



---

# Complete Software Guide for Junos<sup>®</sup> OS for the QFX Series, Release 14.1X53-D10

Release

14.1X53-D10



---

Published: 2014-10-16

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

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

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

*Complete Software Guide for Junos® OS for the QFX Series, Release 14.1X53-D10*  
Release 14.1X53-D10  
Copyright © 2014, Juniper Networks, Inc.  
All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

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



# Table of Contents

	About the Documentation . . . . .	Cxxxv
	Documentation and Release Notes . . . . .	Cxxxv
	Supported Platforms . . . . .	Cxxxv
	Using the Examples in This Manual . . . . .	Cxxxv
	Merging a Full Example . . . . .	Cxxxvi
	Merging a Snippet . . . . .	Cxxxvi
	Documentation Conventions . . . . .	Cxxxvii
	Documentation Feedback . . . . .	Cxxxix
	Requesting Technical Support . . . . .	Cxxxix
	Self-Help Online Tools and Resources . . . . .	Cxxxix
	Opening a Case with JTAC . . . . .	Cxl
<b>Part 1</b>	<b>QFX5100 Switch Overview</b>	
<b>Chapter 1</b>	<b>QFX5100 Switch Overview . . . . .</b>	<b>3</b>
	QFX5100 Device Hardware Overview . . . . .	3
	QFX5100 Hardware . . . . .	3
	System Software . . . . .	7
<b>Part 2</b>	<b>Junos OS Basics</b>	
<b>Chapter 2</b>	<b>Overview . . . . .</b>	<b>11</b>
	Software Overview . . . . .	11
	Configuration File Terms . . . . .	11
	Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements . . . . .	12
	In-Service Software Upgrade (ISSU) System Requirements . . . . .	13
	In-Service Software Upgrade (ISSU) Protocol and Process Support . . . . .	13
	Junos OS Commit Model for Router or Switch Configuration . . . . .	14
	Junos OS Package Names . . . . .	15
	Understanding NTP Time Servers . . . . .	16
	Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release) . . . . .	17
	Understanding Autoinstallation of Configuration Files . . . . .	19
	Typical Uses for Autoinstallation . . . . .	19
	Autoinstallation Configuration Files and IP Addresses . . . . .	19
	Typical Autoinstallation Process on a New Switch . . . . .	20
	Understanding DHCP Services for Switches . . . . .	21
	DHCP Client/Server Model . . . . .	21
	Using DHCP . . . . .	22
	DHCP Relay Servers and DHCP Servers . . . . .	22

Legacy DHCP and Extended DHCP for Server Versions . . . . .	22
Configuring DHCP on a Switch . . . . .	23
How DHCP Works . . . . .	24
Understanding In-Service Software Upgrade (ISSU) . . . . .	25
In-Service Software Upgrade Process . . . . .	25
Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric . . . . .	26
Requirements for Performing an NSSU . . . . .	27
How an NSSU Works . . . . .	27
NSSU Limitations . . . . .	28
NSSU and Junos OS Release Support . . . . .	28
Overview of NSSU Configuration and Operation . . . . .	29
Understanding Software Infrastructure and Processes . . . . .	29
Routing Engine and Packet Forwarding Engine . . . . .	30
Junos OS Processes . . . . .	30
Understanding System Snapshot . . . . .	32
Understanding Zero Touch Provisioning . . . . .	32
Understanding Zero Touch Provisioning . . . . .	33
Zero Touch Provisioning Process . . . . .	34
Zero Touch Provisioning Restart Process Triggers . . . . .	37
User Interfaces . . . . .	39
CLI User Interface Overview . . . . .	39
CLI Overview . . . . .	39
CLI Key Features . . . . .	39
CLI Command Modes . . . . .	40
Configuring Login Tips . . . . .	41
Format for Specifying Filenames and URLs in Junos OS CLI Commands . . . . .	42
Getting Started with Enhanced Layer 2 Software . . . . .	43
Understanding Enhanced Layer 2 Software Support . . . . .	43
Using the ELS Translator Tool . . . . .	44
Configuring a VLAN . . . . .	45
Configuring the Native VLAN Identifier . . . . .	46
Configuring Layer 2 Interfaces . . . . .	46
Configuring Layer 3 Interfaces . . . . .	46
Configuring an IRB Interface . . . . .	47
Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface . . . . .	47
Enhanced Layer 2 CLI Configuration Statement and Command Changes . . . . .	48
Junos OS Operational Mode Commands That Combine Other Commands . . . . .	57
Overview of Junos OS CLI Operational Mode Commands . . . . .	58
CLI Command Categories . . . . .	58
Commonly Used Operational Mode Commands . . . . .	59
Overview of Navigating the CLI . . . . .	60
CLI Command Hierarchy . . . . .	61
CLI Configuration Statements . . . . .	61
Moving Among Hierarchy Levels . . . . .	61
Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands . . . . .	62

	Understanding Junos OS CLI Configuration Mode . . . . .	63
	Configuration Mode Commands . . . . .	64
	Configuration Statements and Identifiers . . . . .	65
	Configuration Statement Hierarchy . . . . .	67
	Licenses . . . . .	69
	Junos OS Feature Licenses . . . . .	69
	Software Features That Require Licenses on the QFX Series . . . . .	70
	Junos OS Feature License Keys . . . . .	71
	Release-Tied License Keys and Upgrade Licenses on MX Series	
	Routers . . . . .	71
	Licensable Ports on MX5, MX10, and MX40 Routers . . . . .	73
	Port Activation on MX104 Routers . . . . .	74
	Generating License Keys . . . . .	75
	Adding New Licenses (CLI Procedure) . . . . .	76
	Deleting a License (CLI Procedure) . . . . .	77
	Saving License Keys . . . . .	78
	Verifying Junos OS License Installation . . . . .	79
	Displaying Installed Licenses . . . . .	79
	Displaying License Usage . . . . .	80
<b>Chapter 3</b>	<b>Installation . . . . .</b>	<b>83</b>
	Software Installation . . . . .	83
	Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI	
	Procedure) . . . . .	84
	Configuring Zero Touch Provisioning . . . . .	87
	Junos OS Package Names . . . . .	92
	Launching a Guest Virtual Machine (VM) to Run a Third Party Application	
	on Junos OS Release 13.2X51-D15 . . . . .	93
	Understanding Guest VMs . . . . .	93
	Prerequisites for Setting up a Virtual Build Environment in the JunosV	
	App Engine . . . . .	93
	Setting up the Virtual Build Environment for the JunosV App Engine . . .	96
	Downloading and Installing the JunosV App Engine Software . . . . .	96
	Launching the VNC Server . . . . .	97
	Launching the FreeBSD Virtual Build Environment (VBE) Virtual Machine	
	(VM) . . . . .	98
	Installing the Junos SDK Packages on the Virtual Build Environment . . .	98
	Prerequisites for Using the Virtual Build Environment . . . . .	99
	Obtaining Junos SDK Certificate Request File and Certificate Key File	
	for the Virtual Build Environment . . . . .	99
	Processing and Obtaining the Certificate File . . . . .	100
	Prerequisites for Packaging the Guest VM . . . . .	100
	Launching the Guest VM on the CentOS Server . . . . .	100
	Copying Required Application to Package with the Guest VM . . . . .	101
	Editing Packaging Tool Scripts . . . . .	102
	Executing Packaging Scripts . . . . .	104
	Copying the Third Party Application to the Switch . . . . .	104
	Configure the Provider Name, License Type, and Deployment	
	Scope . . . . .	104

Configure the Guest VM Options . . . . .	105
Launching a Guest Virtual Machine (VM) to Run a Third Party Application	
on Junos OS Release 13.2X51-D20 . . . . .	109
Understanding Guest VMs . . . . .	109
Troubleshooting Tips . . . . .	110
Copying the Third Party Application to the Switch . . . . .	110
Install the Third Party Application on the Switch . . . . .	110
Configure the Guest VM Options to Launch the Guest VM on the	
Host . . . . .	111
Performing a Recovery Installation . . . . .	116
Performing a Recovery Installation . . . . .	118
Performing an In-Service Software Upgrade (ISSU) . . . . .	119
Preparing the Switch for Software Installation . . . . .	119
Upgrading the Software Using ISSU . . . . .	120
Recovering from a Failed Software Installation . . . . .	121
Software Installation Overview . . . . .	122
Upgrading Jloader Software on QFX Series Devices . . . . .	123
Jloader Software Version 1.1.4 Guidelines . . . . .	124
Upgrading Jloader Software on a QFX3500 Switch . . . . .	125
Upgrading Jloader Software on a QFabric System . . . . .	128
Upgrading Software . . . . .	134
Downloading Software Files with a Browser . . . . .	134
Accessing Software Downloaded to a Remote Location . . . . .	135
Connecting to the Console Port . . . . .	135
Backing Up the Current Configuration Files . . . . .	135
Installing a Standard Software Package . . . . .	136
Upgrading to an ELS-Based Software Package . . . . .	137
Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software	
Upgrade . . . . .	139
Preparing the Switch for Software Installation . . . . .	139
Upgrading the Software Using NSSU . . . . .	141
Upgrading Software by Using Automatic Software Download . . . . .	148
<b>Chapter 4 Configuration . . . . .</b>	<b>151</b>
Initial Configuration . . . . .	151
Configuring Autoinstallation of Configuration Files (CLI Procedure) . . . . .	152
Configuring a DHCP Client (CLI Procedure) . . . . .	154
Configuring a DHCP Server on Switches (CLI Procedure) . . . . .	155
Configuring an Extended DHCP Server on a Switch . . . . .	156
Configuring a Legacy DHCP Server on a Switch (CLI Procedure) . . . . .	156
Configuring a DNS Name Server for Resolving a Hostname into	
Addresses . . . . .	158
Reaching a Domain Name System Server . . . . .	158
Configuring the Hostname of the Router or Switch . . . . .	160
Configuring the Junos OS to Determine Conditions That Trigger Alarms on	
Different Interface Types . . . . .	161
Configuring Junos OS to Disable Protocol Redirect Messages on the Router	
or Switch . . . . .	162

Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses . . . . .	162
Configuring the Junos OS to Display a System Login Announcement . . . . .	163
Configuring the Junos OS to Display a System Login Message . . . . .	163
Configuring Junos OS to Extend the Default Port Address Range . . . . .	164
Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages . . . . .	165
Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets . . . . .	165
Configuring NTP Authentication Keys . . . . .	165
Configuring the NTP Time Server and Time Services . . . . .	166
Configuring the Router or Switch to Operate in Client Mode . . . . .	167
Configuring the Router or Switch to Operate in Symmetric Active Mode . . . . .	167
Configuring the Router or Switch to Operate in Broadcast Mode . . . . .	168
Configuring the Router or Switch to Operate in Server Mode . . . . .	168
Specifying the Physical Location of the Switch . . . . .	169
Configuring the Root Password . . . . .	170
Configuring the Router or Switch to Listen for Broadcast Messages Using NTP . . . . .	171
Configuring the Router or Switch to Listen for Multicast Messages Using NTP . . . . .	172
Configuring System Alarms to Appear Automatically Upon Login . . . . .	172
Configuring Time-Based User Access . . . . .	173
Configuring the Timeout Value for Idle Login Sessions . . . . .	174
Configuring a QFX3500 Device as a Standalone Switch . . . . .	175
Creating an Emergency Boot Device . . . . .	176
Creating a Snapshot and Using It to Boot a QFX Series Switch . . . . .	178
Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch . . . . .	178
Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch . . . . .	179
Creating a Snapshot on the Alternate Slice of the Boot Media . . . . .	179
Creating a Snapshot and Using It to Boot QFX5100 and EX4600 Devices . . . . .	180
Creating a Snapshot on an External USB Flash Drive and Using It to Boot the Device . . . . .	180
Including the Year or Millisecond in Timestamps . . . . .	181
Mapping the Hostname of the Switch to IP Addresses . . . . .	182
Methods for Configuring Junos OS . . . . .	183
Junos OS Command-Line Interface . . . . .	184
ASCII File . . . . .	184
J-Web Package . . . . .	184
Junos XML Management Protocol Software . . . . .	185
NETCONF XML Management Protocol Software . . . . .	185
Configuration Commit Scripts . . . . .	185
Modifying the Default Time Zone for a Router or Switch Running Junos OS . . . . .	186
Rebooting and Halting a Device . . . . .	186

Reverting to the Default Factory Configuration . . . . .	188
Reverting to the Default Factory Configuration by Using the request system zeroize Command . . . . .	188
Reverting to the Rescue Configuration . . . . .	189
Saving Core Files Generated by Junos OS Processes . . . . .	189
Updating the IANA Time Zone Database on Junos Devices . . . . .	190
Importing and Installing Time Zone Files . . . . .	190
Configuring a Custom Time Zone . . . . .	191
Setting the Date and Time . . . . .	192
Specifying Access Privileges for Junos OS Operational Mode Commands . . . . .	192
Synchronizing and Coordinating Time Distribution Using NTP . . . . .	194
Configuring NTP . . . . .	194
Configuring the NTP Boot Server . . . . .	194
Specifying a Source Address for an NTP Server . . . . .	195
Viewing Core Files from Junos OS Processes . . . . .	196
Configuration Examples . . . . .	196
Example: Changing the Requirements for Junos OS Plain-Text Passwords . . . . .	196
Reaching a Domain Name System Server . . . . .	198
Example: Configuring the Name of the Switch, IP Address, and System ID . . . . .	200
Example: Configuring NTP . . . . .	200
Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization . . . . .	203
Configuration Statements . . . . .	204
QFX Series CLI Hierarchy . . . . .	207
[edit access] Hierarchy . . . . .	207
[edit accounting-options] Hierarchy . . . . .	208
[edit chassis] Hierarchy . . . . .	209
[edit class-of-service] Hierarchy . . . . .	211
[edit ethernet-switching-options] Hierarchy . . . . .	213
[edit fabric] Hierarchy . . . . .	215
[edit fc-fabrics] Hierarchy . . . . .	216
[edit fc-options] Hierarchy . . . . .	217
[edit firewall] Hierarchy . . . . .	217
[edit groups] Hierarchy . . . . .	218
[edit interfaces] Hierarchy . . . . .	218
[edit policy-options] Hierarchy . . . . .	224
[edit protocols] Hierarchy . . . . .	224
[edit security] Hierarchy . . . . .	237
[edit snmp] Hierarchy . . . . .	237
[edit system] Hierarchy . . . . .	241
[edit vlans] Hierarchy . . . . .	246
access-end . . . . .	246
access-start . . . . .	247
accounting . . . . .	248
accounting-port . . . . .	249
allow-commands . . . . .	249

allow-configuration	250
allowed-days	250
allow-transients	251
announcement	251
archival	252
arp (System)	253
authentication (Login)	254
authentication-key	255
authentication-order	256
auxiliary	257
boot-server (NTP)	258
broadcast	259
broadcast-client	260
change-type	260
checksum	261
class (Defining Login Classes)	262
class (Assigning a Class to an Individual User)	263
commit	264
compress-configuration-files (System)	265
console (Physical Port)	266
default-address-selection	267
deny-commands	268
deny-configuration	269
destination (Accounting)	270
destination-override	271
direct-access	271
domain-name	272
domain-search	272
explicit-priority	273
events	274
format	274
host-name	275
icmpv4-rate-limit	275
idle-timeout	276
internet-options	276
l2-learning	277
load-key-file	278
location	279
login	280
login-alarms	281
login-tip	281
max-configurations-on-flash	282
maximum-length	282
message	283
minimum-changes	283
minimum-length	284
minimum-lower-cases	285
minimum-numeric	286
minimum-punctuations	287

minimum-upper-cases	288
multicast-client	288
name-server	289
no-multicast-echo	290
no-ping-record-route	291
no-ping-time-stamp	291
no-redirects (IPv4 Traffic)	292
no-split-detection	293
ntp	294
optional	294
password (Login)	295
peer	296
permissions	297
port (TACACS+ Server)	297
ports	298
radius (System)	299
refresh (Commit Scripts)	300
refresh-from (Commit Scripts)	300
retry	301
retry-options	302
root-authentication	303
saved-core-context	304
saved-core-files	304
secret	305
server (TACACS+ Accounting)	305
server (NTP)	306
server (RADIUS Accounting)	307
single-connection	307
source (Commit Scripts)	308
source-address (NTP, RADIUS, System Logging, or TACACS+)	308
source-port (Port Addresses)	309
ssh-dsa	309
ssh-rsa	310
static-host-mapping	311
structured-data	312
syslog (System)	313
system	315
tacplus	320
tacplus-server	321
timeout	322
time-format	323
time-zone	324
traceoptions (Commit Scripts)	326
traceoptions (Layer 2 Learning)	328
tracing	330
trusted-key	331
uid	331
use-imported-time-zones	332
user (Access)	332



<b>Chapter 5</b>	<b>Administration . . . . .</b>	<b>333</b>
	Routine Monitoring . . . . .	333
	Monitoring System Process Information . . . . .	333
	Monitoring System Properties . . . . .	334
	Monitoring Interface Status and Traffic . . . . .	335
	Monitoring Zero Touch Provisioning . . . . .	336
	Using the Console to Monitor Zero Touch Provisioning . . . . .	336
	Using System Log Alerts to Monitor Zero Touch Provisioning . . . . .	336
	Using Error Messages to Monitor Zero Touch Provisioning . . . . .	337
	Using System Log Files to Monitor Zero Touch Provisioning . . . . .	337
	Using the show dhcp client binding Command . . . . .	338
	Using the show dhcp client statistics Command . . . . .	338
	Other Tools to Configure and Monitor Devices Running Junos OS . . . . .	339
	Verifying a Unified In-Service Software Upgrade . . . . .	339
	Verifying Autoinstallation Status . . . . .	340
	Verifying That Automatic Software Download Is Working Correctly . . . . .	341
	Operational Commands . . . . .	342
	commit . . . . .	345
	clear log . . . . .	350
	clear chassis display message . . . . .	351
	clear system commit . . . . .	354
	clear system reboot . . . . .	355
	file . . . . .	359
	file archive . . . . .	361
	file checksum md5 . . . . .	363
	file checksum sha1 . . . . .	364
	file checksum sha-256 . . . . .	365
	file compare . . . . .	366
	file delete . . . . .	369
	file list . . . . .	370
	file rename . . . . .	372
	file show . . . . .	374
	load . . . . .	376
	ping . . . . .	378
	request chassis beacon . . . . .	382
	request chassis fpc . . . . .	384
	request chassis pic . . . . .	388
	request chassis routing-engine master . . . . .	392
	request message . . . . .	397
	request system configuration rescue delete . . . . .	398
	request system configuration rescue save . . . . .	399
	request system halt . . . . .	400
	request system license add . . . . .	406
	request system license delete . . . . .	407
	request system license save . . . . .	408
	request system logout . . . . .	409
	request system power-off . . . . .	410
	request system reboot . . . . .	415
	request system snapshot . . . . .	419

request system software add . . . . .	421
request system software delete . . . . .	430
request system software download . . . . .	434
request system software in-service-upgrade . . . . .	436
request system software nonstop-upgrade . . . . .	449
request system software rollback . . . . .	459
request system software validate . . . . .	463
request system storage cleanup . . . . .	466
request system zeroize . . . . .	476
restart . . . . .	481
rollback . . . . .	492
save . . . . .	493
show chassis alarms . . . . .	495
show chassis beacon . . . . .	509
show chassis environment . . . . .	511
show chassis environment fpc . . . . .	575
show chassis environment pem . . . . .	601
show chassis environment routing-engine . . . . .	610
show chassis fan . . . . .	615
show chassis firmware . . . . .	628
show chassis fpc . . . . .	639
show chassis hardware . . . . .	676
show chassis in-service-upgrade . . . . .	851
show chassis lcd . . . . .	855
show chassis led . . . . .	868
show chassis location . . . . .	878
show chassis mac-addresses . . . . .	882
show chassis nonstop-upgrade . . . . .	887
show chassis pic . . . . .	889
show chassis routing-engine . . . . .	905
show chassis zones . . . . .	927
show cli . . . . .	933
show cli authorization . . . . .	935
show cli directory . . . . .	939
show cli history . . . . .	940
show host . . . . .	941
show interfaces diagnostics optics . . . . .	942
show log . . . . .	948
show ntp associations . . . . .	951
show ntp status . . . . .	953
show subscribers . . . . .	956
show system alarms . . . . .	974
show system audit . . . . .	977
show system boot-messages . . . . .	985
show system buffers . . . . .	992
show system certificate . . . . .	999
show system commit . . . . .	1001
show system configuration archival . . . . .	1004
show system configuration rescue . . . . .	1005

	show system connections . . . . .	1007
	show system core-dumps . . . . .	1026
	show system directory-usage . . . . .	1040
	show system license . . . . .	1044
	show system processes . . . . .	1051
	show system reboot . . . . .	1078
	show system resource-cleanup processes . . . . .	1081
	show system rollback . . . . .	1083
	show system services service-deployment . . . . .	1085
	show system software . . . . .	1086
	show system statistics . . . . .	1094
	show system storage . . . . .	1129
	show system uptime . . . . .	1137
	show system users . . . . .	1142
	show system virtual-memory . . . . .	1147
	show version . . . . .	1205
	start shell . . . . .	1218
	test configuration . . . . .	1220
	traceroute . . . . .	1221
	traceroute monitor . . . . .	1225
<b>Chapter 6</b>	<b>Troubleshooting . . . . .</b>	<b>1227</b>
	Troubleshooting Procedures . . . . .	1227
	Creating an Emergency Boot Device . . . . .	1227
	Performing a Recovery Installation . . . . .	1229
	Rebooting and Halting a Device . . . . .	1230
	Recovering from a Failed Software Installation . . . . .	1232
	Recovering the Root Password . . . . .	1233
	Troubleshooting Network Interfaces . . . . .	1234
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down . . . . .	1234
	Troubleshooting an Aggregated Ethernet Interface . . . . .	1234
<b>Part 3</b>	<b>Configuration and File Management</b>	
<b>Chapter 7</b>	<b>Overview . . . . .</b>	<b>1239</b>
	Configuration Files Overview . . . . .	1239
	Configuration File Terms . . . . .	1239
	Software Overview . . . . .	1240
	Forms of the configure Command . . . . .	1240
	Junos OS Commit Model for Router or Switch Configuration . . . . .	1241
	Understanding Configuration Files . . . . .	1242
	Understanding How the Junos OS Configuration Is Stored . . . . .	1243
<b>Chapter 8</b>	<b>Configuration . . . . .</b>	<b>1245</b>
	Configuration Tasks . . . . .	1245
	Comparing Configuration Changes with a Prior Version . . . . .	1245
	Compressing the Current Configuration File . . . . .	1247
	Creating and Returning to a Rescue Configuration . . . . .	1248
	Loading a Configuration from a File . . . . .	1249

Loading a Previous Configuration File . . . . .	1252
Returning to the Most Recently Committed Junos OS Configuration . . . . .	1252
Returning to a Previously Committed Junos OS Configuration . . . . .	1253
Returning to a Configuration Prior to the One Most Recently	
Committed . . . . .	1253
Displaying Previous Configurations . . . . .	1253
Comparing Configuration Changes with a Prior Version . . . . .	1254
Creating and Returning to a Rescue Configuration . . . . .	1256
Saving a Configuration to a File . . . . .	1257
Reverting to the Default Factory Configuration . . . . .	1258
Reverting to the Rescue Configuration . . . . .	1258
Rolling Back Junos OS Configuration Changes . . . . .	1259
Saving a Configuration to a File . . . . .	1260
Setting or Deleting the Rescue Configuration . . . . .	1261
Uploading a Configuration File . . . . .	1261
Using Junos OS to Configure a Router or Switch to Transfer Its Configuration	
to an Archive Site . . . . .	1263
Configuring the Router or Switch to Transfer Its Currently Active	
Configuration to an Archive . . . . .	1263
Configuring the Transfer Interval for Periodic Transfer of the Active	
Configuration to an Archive Site . . . . .	1263
Configuring Transfer of the Current Active Configuration When a	
Configuration Is Committed . . . . .	1264
Configuring Archive Sites for Transfer of Active Configuration Files . .	1264
Configuration Statements . . . . .	1265
archival . . . . .	1266
archive-sites (Configuration File) . . . . .	1267
configuration . . . . .	1269
transfer-interval (Configuration) . . . . .	1270
transfer-on-commit . . . . .	1271
Default Configurations . . . . .	1271
QFX3500 Switch Default Configuration . . . . .	1271
Configuration Examples . . . . .	1277
Examples: Loading a Configuration from a File . . . . .	1277
<b>Chapter 9 Administration . . . . .</b>	<b>1281</b>
Operational Commands . . . . .	1281
clear log . . . . .	1282
clear system commit . . . . .	1283
file archive . . . . .	1284
file checksum md5 . . . . .	1286
file checksum sha1 . . . . .	1287
file checksum sha-256 . . . . .	1288
file compare . . . . .	1289
file delete . . . . .	1292
file list . . . . .	1293
file rename . . . . .	1295
file show . . . . .	1297
request system configuration rescue delete . . . . .	1299

	request system configuration rescue save . . . . .	1300
	show system commit . . . . .	1301
	show system configuration archival . . . . .	1304
	show system configuration rescue . . . . .	1305
	show system rollback . . . . .	1307
	test configuration . . . . .	1309
<b>Chapter 10</b>	<b>Troubleshooting . . . . .</b>	<b>1311</b>
	Troubleshooting Procedures . . . . .	1311
	Loading a Previous Configuration File . . . . .	1311
	Reverting to the Default Factory Configuration . . . . .	1312
	Reverting to the Rescue Configuration . . . . .	1312
<b>Part 4</b>	<b>User and Access Management</b>	
<b>Chapter 11</b>	<b>Overview . . . . .</b>	<b>1315</b>
	Software Overview . . . . .	1315
	Understanding Software Infrastructure and Processes . . . . .	1315
	Routing Engine and Packet Forwarding Engine . . . . .	1315
	Junos OS Processes . . . . .	1316
	Access Control Overview . . . . .	1317
	Overview of Template Accounts for RADIUS and TACACS+ . . . . .	
	Authentication . . . . .	1318
	Understanding Login Authentication . . . . .	1318
	MAC RADIUS Authentication . . . . .	1319
	Understanding LLDP . . . . .	1319
	Understanding RADIUS Accounting . . . . .	1320
	Understanding VSAs . . . . .	1321
	Juniper Networks Vendor-Specific RADIUS Attributes . . . . .	1321
	Juniper Networks Vendor-Specific TACACS+ Attributes . . . . .	1324
	Understanding Junos OS Access Privilege Levels . . . . .	1325
	Junos OS Login Class Permission Flags . . . . .	1326
	Allowing or Denying Individual Commands for Junos OS Login . . . . .	
	Classes . . . . .	1329
	Junos OS Authentication Order for RADIUS, TACACS+, and Password . . . . .	
	Authentication . . . . .	1330
	Using RADIUS or TACACS+ Authentication . . . . .	1330
	Using Local Password Authentication . . . . .	1331
	Order of Authentication Attempts . . . . .	1331
	Junos OS User Authentication Methods . . . . .	1334
	Junos OS User Accounts Overview . . . . .	1335
	Junos OS Login Classes Overview . . . . .	1337
	Regular Expressions for Allowing and Denying Junos OS Configuration Mode . . . . .	
	Hierarchies . . . . .	1338
	Regular Expressions for Allowing and Denying Junos OS Operational Mode . . . . .	
	Commands . . . . .	1339
	Special Requirements for Junos OS Plain-Text Passwords . . . . .	1339

<b>Chapter 12</b>	<b>Configuration</b>	<b>1343</b>
	Configuration Tasks	1343
	Configuring Access Privilege Levels	1344
	Configuring Login Tips	1344
	Configuring Junos OS User Accounts	1344
	Configuring LLDP	1345
	Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication	1346
	Configuring Local User Template Accounts for User Authentication	1347
	Configuring Management Access	1349
	Configuring RADIUS System Accounting	1349
	Configuring Auditing of User Events on a RADIUS Server	1349
	Specifying RADIUS Server Accounting and Auditing Events	1350
	Configuring RADIUS Server Accounting	1350
	Configuring RADIUS Authentication (QFX Series)	1351
	Configuring RADIUS Server Details	1351
	Configuring MS-CHAPv2 for Password-Change Support	1352
	Specifying a Source Address for the Junos OS to Access External RADIUS Servers	1353
	Configuring Remote Template Accounts for User Authentication	1354
	Configuring the Root Password	1354
	Configuring SNMP	1356
	Configuring SSH Host Keys for Secure Copying of Data	1359
	Configuring SSH Known Hosts	1360
	Configuring Support for SCP File Transfer	1360
	Updating SSH Host Key Information	1361
	Configuring SSH Service for Remote Access to the Router or Switch	1361
	Configuring the Root Login Through SSH	1363
	Configuring the SSH Protocol Version	1363
	Configuring the Client Alive Mechanism	1363
	Configuring TACACS+ Authentication (QFX Series)	1364
	Configuring TACACS+ Server Details	1364
	Specifying a Source Address for the Junos OS to Access External TACACS+ Servers	1365
	Configuring the Same Authentication Service for Multiple TACACS+ Servers	1365
	Configuring Juniper Networks Vendor-Specific TACACS+ Attributes	1366
	Configuring TACACS+ System Accounting	1366
	Specifying TACACS+ Auditing and Accounting Events	1367
	Configuring TACACS+ Server Accounting	1367
	Defining Junos OS Login Classes	1368
	Limiting the Number of User Login Attempts for SSH and Telnet Sessions	1369
	Recovering the Root Password	1370
	Specifying Access Privileges for Junos OS Configuration Mode Hierarchies	1371
	Specifying Access Privileges for Junos OS Operational Mode Commands	1372

Using Junos OS to Configure Logical System Administrators . . . . .	1374
Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands . . . . .	1375
VSA Match Conditions and Actions . . . . .	1376
Configuration Examples . . . . .	1378
Example: Changing the Requirements for Junos OS Plain-Text Passwords . . . . .	1379
Example: Configuring Access Privilege Levels . . . . .	1381
Example: Configuring Access Privileges for Operational Mode Commands . . . . .	1381
Example: Configuring a Plain-Text Password for Root Logins . . . . .	1382
Example: Configuring RADIUS Authentication . . . . .	1384
Example: Configuring RADIUS System Accounting . . . . .	1385
Example: Configuring the Root Password . . . . .	1385
Example: Configuring SSH Authentication for Root Logins . . . . .	1386
Example: Configuring User Accounts . . . . .	1386
Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication . . . . .	1387
Example: Creating Login Classes with Specific Privileges . . . . .	1389
Example: Configuring User Login Accounts . . . . .	1389
Example: Configuring RADIUS Template Accounts . . . . .	1390
Defining Access Privileges Using allow/deny-configuration Statements . .	1390
Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions . . . . .	1391
Configuration Statements . . . . .	1392
access . . . . .	1394
accounting (Access Profile) . . . . .	1395
accounting-options . . . . .	1396
accounting-server . . . . .	1398
accounting-stop-on-access-deny . . . . .	1399
accounting-stop-on-failure . . . . .	1400
advertisement-interval . . . . .	1401
agent-address . . . . .	1402
archival . . . . .	1403
archive-sites (Configuration File) . . . . .	1404
authentication-order . . . . .	1405
authentication-server . . . . .	1406
authorization . . . . .	1407
categories . . . . .	1408
client-list . . . . .	1408
client-list-name . . . . .	1409
clients . . . . .	1409
commit-delay . . . . .	1410
community (SNMP) . . . . .	1411
configuration . . . . .	1412
connection-limit . . . . .	1413
contact . . . . .	1414
disable (LLDP) . . . . .	1414
falling-threshold (Health Monitor) . . . . .	1415

filter-duplicates	1415
full-name	1416
health-monitor	1416
hold-multiplier	1417
idle-timeout (Access)	1418
interface (LLDP)	1419
interval (Health Monitor)	1420
lldp	1421
lldp-configuration-notification-interval	1422
location	1423
management-address	1424
name	1425
nas-ip-address	1425
nonvolatile	1426
oid	1426
order	1427
port (RADIUS Server)	1428
profile	1429
protocols	1430
protocol-version	1443
ptopo-configuration-maximum-hold-time	1443
ptopo-configuration-trap-interval	1444
radius	1445
radius-options (edit system)	1446
radius-server	1447
rate-limit	1448
remote-debug-permission	1449
retry	1450
rising-threshold (Health Monitor)	1451
root-login	1452
services (Switches)	1453
snmp	1454
ssh	1458
system	1459
tacplus-options	1465
targets	1466
traceoptions (LLDP)	1467
transfer-interval (Configuration)	1469
transfer-on-commit	1470
trap-group	1471
trap-options	1472
user (Access)	1473
version	1474



<b>Chapter 13</b>	<b>Administration</b>	<b>1475</b>
	Routine Monitoring	1475
	Monitoring SNMP	1475
	Monitoring Commands	1476
	clear lldp neighbors	1478
	clear lldp statistics	1479
	request component login	1480
	show ethernet-switching interfaces	1482
	show lldp	1486
	show lldp local-information	1491
	show lldp neighbors	1493
	show lldp statistics	1497
	show route instance	1499
	show snmp statistics	1503
	ssh	1507
<b>Part 5</b>	<b>Ethernet Features</b>	
<b>Chapter 14</b>	<b>Overview</b>	<b>1511</b>
	Enhanced Layer 2 Software (ELS) CLI	1511
	Getting Started with Enhanced Layer 2 Software	1511
	Understanding Enhanced Layer 2 Software Support	1511
	Using the ELS Translator Tool	1512
	Configuring a VLAN	1513
	Configuring the Native VLAN Identifier	1514
	Configuring Layer 2 Interfaces	1514
	Configuring Layer 3 Interfaces	1514
	Configuring an IRB Interface	1515
	Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface	1515
	Enhanced Layer 2 CLI Configuration Statement and Command Changes	1516
	Bridging and VLANs	1525
	Ethernet Ring Protection Switching Overview	1525
	Layer 2 Learning and Forwarding for VLANs Overview	1526
	Understanding Bridging and VLANs	1527
	History of VLANs	1527
	How Bridging of VLAN Traffic Works	1527
	Packets Are Either Tagged or Untagged	1529
	Switch Interface Modes—Access, Trunk, or Tagged Access	1529
	Additional Advantages of Using VLANs	1531
	Maximum VLANs and VLAN Members Per Switch	1532
	A Default VLAN Is Configured on Most Switches	1532
	Assigning Traffic to VLANs	1533
	Forwarding VLAN Traffic	1533
	VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces	1534

Understanding Ethernet Ring Protection Switching Functionality . . . . .	1534
Acronyms . . . . .	1535
Ring Nodes . . . . .	1535
Ring Node States . . . . .	1535
Failure Detection . . . . .	1535
Logical Ring . . . . .	1536
FDB Flush . . . . .	1536
Traffic Blocking and Forwarding . . . . .	1536
RAPS Message Blocking and Forwarding . . . . .	1536
Dedicated Signaling Control Channel . . . . .	1537
RAPS Message Termination . . . . .	1538
Multiple Rings . . . . .	1538
Node ID . . . . .	1538
Bridge Domains with the Ring Port (MX Series Routers Only) . . . . .	1538
Understanding Integrated Routing and Bridging . . . . .	1539
Understanding MAC Learning . . . . .	1540
Layer 2 Networking . . . . .	1540
Introduction to the Media Access Control (MAC) Layer 2 Sublayer . . . . .	1541
Overview of Layer 2 Networking . . . . .	1542
Understanding Layer 2 Broadcasting . . . . .	1544
Understanding Unicast . . . . .	1545
Understanding the Unified Forwarding Table . . . . .	1545
Using the Unified Forwarding Table to Optimize Address Storage . . . . .	1545
MAC Address and Host Address Memory Allocation . . . . .	1545
LPM Table Memory Allocation . . . . .	1546
Understanding Q-in-Q Tunneling . . . . .	1547
Understanding Q-in-Q Tunneling . . . . .	1547
How Q-in-Q Tunneling Works . . . . .	1547
How VLAN Translation Works . . . . .	1548
Sending and Receiving Untagged Packets . . . . .	1548
Disabling MAC Address Learning . . . . .	1549
Mapping C-VLANs to S-VLANs . . . . .	1549
Constraints for Q-in-Q Tunneling and VLAN Translation . . . . .	1550
Proxy ARP . . . . .	1551
Understanding Proxy ARP . . . . .	1551
What Is ARP? . . . . .	1552
Proxy ARP Overview . . . . .	1552
Best Practices for Proxy ARP . . . . .	1552
Reflective Relay . . . . .	1553
Understanding Reflective Relay for Use with VEPA Technology . . . . .	1553
VEPA . . . . .	1553
Reflective Relay . . . . .	1553
Spanning Trees . . . . .	1554
Overview of Spanning-Tree Protocols . . . . .	1554
Understanding Spanning Tree Protocols on a QFabric System . . . . .	1555
Understanding MSTP . . . . .	1555
Understanding RSTP . . . . .	1556
Understanding VSTP . . . . .	1557
Understanding BPDU Protection for STP, RSTP, and MSTP . . . . .	1558

	Understanding Loop Protection for STP, RSTP, VSTP, and MSTP . . . . .	1559
	Understanding Root Protection for STP, RSTP, VSTP, and MSTP . . . . .	1560
<b>Chapter 15</b>	<b>Configuration . . . . .</b>	<b>1563</b>
	Bridging and VLAN Configuration Examples . . . . .	1563
	Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches . . . . .	1563
	Example: Configuring Routing Between VLANs on One Switch . . . . .	1576
	Example: Disabling MAC Learning . . . . .	1582
	Example: Setting Up Bridging with Multiple VLANs . . . . .	1583
	Example: Setting Up Basic Bridging and a VLAN on the QFX Series . . . . .	1588
	Reflective Relay Configuration Example . . . . .	1605
	Example: Configuring Reflective Relay for Use with VEPA Technology . . . . .	1605
	STP Configuration Examples . . . . .	1609
	Example: Configuring Faster Convergence and Improving Network Stability with RSTP . . . . .	1610
	Example: Configuring Network Regions for VLANs with MSTP . . . . .	1624
	Example: Connecting an Access Switch to a Distribution Switch . . . . .	1647
	Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations . . . . .	1656
	Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree . . . . .	1660
	Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees . . . . .	1664
	Bridging and VLAN Configuration Tasks . . . . .	1669
	Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure) . . . . .	1669
	Configuring Ethernet Ring Protection Switching (CLI Procedure) . . . . .	1670
	Configuring MAC Limiting (CLI Procedure) . . . . .	1672
	Limiting the Number of MAC Addresses Learned by an Interface . . . . .	1673
	Limiting the Number of MAC Addresses Learned by a VLAN . . . . .	1673
	Configuring MAC Table Aging . . . . .	1674
	Configuring IRB Interfaces . . . . .	1675
	Configuring Static ARP Entries . . . . .	1676
	Configuring the Native VLAN Identifier (CLI Procedure) . . . . .	1677
	Configuring VLANs . . . . .	1678
	Creating a Series of Tagged VLANs . . . . .	1680
	Disabling MAC Learning . . . . .	1681
	Configuring MAC Notification (CLI Procedure) . . . . .	1682
	Enabling MAC Notification . . . . .	1682
	Disabling MAC Notification . . . . .	1683
	Setting the MAC Notification Interval . . . . .	1683
	Q-in-Q Tunneling Configuration Tasks . . . . .	1683
	Configuring Q-in-Q Tunneling . . . . .	1683
	Using the Different Mapping Methods . . . . .	1684
	Configuring All-in-One Bundling . . . . .	1684
	Configuring Many-to-Many Bundling . . . . .	1686

Configuring a Specific Interface Mapping with VLAN ID Translation	
Option . . . . .	1689
Configuring All-in-One Bundling . . . . .	1691
Configuring Many-to-Many Bundling . . . . .	1692
Configuring a Specific Interface Mapping with VLAN ID Translation	
Option . . . . .	1695
Unified Forwarding Table Configuration Task . . . . .	1697
Configuring the Unified Forwarding Table . . . . .	1697
Configuring an Address-Storage Profile . . . . .	1698
Configuring the LPM Allocation . . . . .	1699
Forwarding Mode Configuration Task . . . . .	1702
Configuring the Forwarding Mode . . . . .	1702
Proxy ARP Configuration Task . . . . .	1702
Configuring Proxy ARP (CLI Procedure) . . . . .	1702
Reflective Relay Configuration Tasks . . . . .	1703
Configuring Reflective Relay . . . . .	1703
STP Configuration Tasks . . . . .	1704
Configuring STP . . . . .	1704
Configuring VLAN Spanning-Tree Protocol . . . . .	1705
Unblocking an Interface That Receives BPDUs in Error . . . . .	1708
Protocols Configuration Statement . . . . .	1709
protocols . . . . .	1710
Unified Forwarding Table Configuration Statements . . . . .	1723
forwarding-options (chassis) . . . . .	1724
num-65-127-prefix . . . . .	1725
prefix-65-127-disable . . . . .	1725
Reflective Relay Configuration Statements . . . . .	1726
reflective-relay . . . . .	1726
STP Configuration Statements . . . . .	1726
alarm (STP) . . . . .	1727
block . . . . .	1728
bpdu-block . . . . .	1729
bpdu-block-on-edge . . . . .	1730
bpdu-timeout-action . . . . .	1731
bridge-priority . . . . .	1732
configuration-name (MSTP) . . . . .	1733
cost (STP) . . . . .	1734
disable (STP) . . . . .	1735
disable-timeout (BPDU) . . . . .	1736
edge (STP) . . . . .	1737
forward-delay . . . . .	1738
force-version . . . . .	1739
hello-time . . . . .	1740
interface (Spanning Trees) . . . . .	1741
interface (BPDU) . . . . .	1742
interface (STP) . . . . .	1743
max-age . . . . .	1744
max-hops . . . . .	1745
mode (STP) . . . . .	1746

msti	1747
mstp	1748
no-root-port	1749
priority (STP)	1750
revision-level	1751
rstp	1752
stp	1753
traceoptions (STP)	1754
vlan (STP)	1758
vstp	1759
VLAN Configuration Statements	1760
[edit vlans] Configuration Statement Hierarchy on the QFX Series	1761
Supported Statements in the [edit vlans] Hierarchy Level	1761
Unsupported Statements in the [edit vlans] Hierarchy Level	1763
control-channel	1764
control-vlan	1765
data-channel	1766
description (VLAN)	1767
dhcp-relay	1768
east-interface	1773
ethernet-ring	1774
filter (VLANs)	1775
forwarding-options	1776
guard-interval	1781
hold-interval (Protection Group)	1782
interface (VLANs)	1782
interface-mac-limit	1783
interface-mode	1785
irb (Interfaces)	1787
l3-interface (VLAN)	1790
mac (Static MAC-Based VLANs)	1791
mac-limit	1791
mac-notification	1792
mac-statistics	1793
mac-table-aging-time	1794
mac-table-size	1795
members	1797
native-vlan-id	1798
notification-interval	1799
packet-action	1800
port-mode	1803
protection-group	1804
restore-interval	1805
ring-protection-link-end	1806
ring-protection-link-owner	1806
service-id	1807
switch-options	1808
static (Static MAC-Based VLANs)	1809
static-mac	1810

	vlan-id (VLANs) .....	1811
	vlan-id-list .....	1812
	vlan-rewrite .....	1813
	vlan-tagging .....	1813
	vlangs .....	1814
	west-interface .....	1817
	Q-in-Q Configuration Statements .....	1817
	flexible-vlan-tagging .....	1818
	input-vlan-map .....	1819
	native-vlan-id .....	1820
	output-vlan-map .....	1821
	pop .....	1822
	push .....	1823
	swap .....	1824
	vlan-id-list .....	1825
<b>Chapter 16</b>	<b>Administration .....</b>	<b>1827</b>
	Routine Monitoring .....	1827
	Verifying That MAC Notification Is Working Properly .....	1827
	Verifying That a Series of Tagged VLANs Has Been Created .....	1827
	Verifying That a Private VLAN Is Working .....	1829
	Verifying That Proxy ARP Is Working Correctly .....	1834
	Monitoring Commands .....	1835
	clear ethernet-switching bpdu-error .....	1837
	clear ethernet-switching layer2-protocol-tunneling error .....	1838
	clear ethernet-switching layer2-protocol-tunneling statistics .....	1839
	clear ethernet-switching table .....	1840
	clear spanning-tree statistics .....	1842
	show ethernet-switching interfaces .....	1843
	show ethernet-switching layer2-protocol-tunneling interface .....	1847
	show ethernet-switching layer2-protocol-tunneling statistics .....	1849
	show ethernet-switching layer2-protocol-tunneling vlan .....	1852
	show ethernet-switching mac-learning-log .....	1854
	show ethernet-switching mac-notification .....	1856
	show ethernet-switching statistics aging .....	1858
	show ethernet-switching statistics mac-learning .....	1860
	show ethernet-switching table .....	1864
	show spanning-tree bridge .....	1870
	show spanning-tree interface .....	1875
	show spanning-tree mstp configuration .....	1881
	show spanning-tree statistics .....	1883
	show system statistics arp .....	1885
	show vlans .....	1886
<b>Chapter 17</b>	<b>Troubleshooting .....</b>	<b>1895</b>
	Troubleshooting Procedures .....	1895
	Troubleshooting Ethernet Switching .....	1895

<b>Part 6</b>	<b>OVSDB and VXLAN</b>	
<b>Chapter 18</b>	<b>Overview</b>	<b>1899</b>
	OVSDB Overview	1899
	Open vSwitch Database Support on Juniper Networks Devices	1899
	Understanding the Junos OS Implementation of VXLAN and OVSDB in a VMware NSX for Multi-Hypervisor Environment for the Data Center	1900
	Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices	1902
	Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers	1903
	Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDB	1904
	Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment	1905
	Understanding How to Set Up OVSDB-Managed VXLANs On All Juniper Networks Devices Except QFX5100 Switches	1906
	Understanding How to Set Up OVSDB-Managed VXLANs On QFX5100 Switches	1907
	Understanding Automatically Created OVSDB-Managed VXLANs on a QFX5100 Switch	1908
	Understanding How to Determine the State of an OVSDB-Managed VXLAN	1909
	Open vSwitch Database Schema For Physical Devices	1910
	VXLAN Overview	1912
	Understanding VXLANs	1912
	VXLAN Benefits	1912
	What is a VXLAN?	1913
	Using a QFX5100 Switch with VXLANs	1913
	Using an MX Series Routers as a VTEP	1914
	Manual VXLANs Require PIM	1914
	Load Balancing VXLAN Traffic	1915
<b>Chapter 19</b>	<b>Configuration</b>	<b>1917</b>
	Configuration Examples	1917
	Example: Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections Between Virtual and Physical Entities in a Data Center	1917
	Example: Setting Up Inter-VXLAN Routing and OVSDB Connections in a Data Center	1925
	Example: Configuring VXLAN on MX Series Routers	1934
	Examples: Configuring VXLANs on QFX Series Switches	1944
	Example: Configuring a VXLAN Transit Switch	1944
	Example: Configuring a VXLAN Layer 2 Gateway	1945

Example: Configuring VXLAN to VPLS Stitching with OVSDB . . . . .	1952
Configuration Tasks . . . . .	1983
Installing Open vSwitch Database Components on Juniper Networks	
Devices . . . . .	1983
Creating and Installing an SSL Key and Certificate on a Juniper Networks	
Device for Connection with VMware NSX Controllers . . . . .	1984
Setting Up the Open vSwitch Database Management Protocol on Juniper	
Networks Devices . . . . .	1985
VMware NSX Configuration for Juniper Networks Devices That Function as	
Virtual Tunnel Endpoints . . . . .	1986
Creating a Gateway . . . . .	1987
Creating a Gateway Service . . . . .	1987
Creating a Logical Switch Port . . . . .	1988
Configuring OVSDB-Managed VXLANs . . . . .	1989
Configuring VXLANs on a QFX5100 Switch . . . . .	1991
Configuring a Source IP Address . . . . .	1991
Configuring PIM for VXLANs . . . . .	1991
Configuring VXLANs . . . . .	1991
OVSDB Configuration Statements . . . . .	1992
controller (OVSDB) . . . . .	1993
inactivity-probe-duration . . . . .	1994
ingress-node-replication . . . . .	1995
interfaces (OVSDB) . . . . .	1996
maximum-backoff-duration . . . . .	1996
ovsdb . . . . .	1997
ovsdb-managed . . . . .	1998
port (OVSDB) . . . . .	1999
protocol (OVSDB) . . . . .	2000
traceoptions (OVSDB) . . . . .	2001
VXLAN Configuration Statements . . . . .	2002
decapsulate-accept-inner-vlan . . . . .	2003
encapsulate-inner-vlan . . . . .	2003
multicast-group . . . . .	2004
ovsdb-managed . . . . .	2005
unreachable-vtep-aging-timer . . . . .	2006
vni . . . . .	2006
vtep-source-interface . . . . .	2007
vxlan . . . . .	2007
<b>Chapter 20 Administration . . . . .</b>	<b>2009</b>
OVSDB Monitoring Commands . . . . .	2009
show bridge mac-table . . . . .	2010
show ovsdb controller . . . . .	2014
show ovsdb interface . . . . .	2017
show ovsdb logical-switch . . . . .	2019
show ovsdb mac . . . . .	2021
show ovsdb statistics interface . . . . .	2025
show ovsdb virtual-tunnel-end-point . . . . .	2027



	show vpls mac-table . . . . .	2029
	VXLAN Monitoring Commands . . . . .	2033
	show bridge mac-table . . . . .	2034
	show vpls mac-table . . . . .	2038
	Verifying VXLAN Reachability . . . . .	2042
	Verifying That a Local VXLAN VTEP is Configured Correctly . . . . .	2042
	Verifying MAC Learning from a Remote VTEP . . . . .	2042
	Monitor a Remote VTEP Interface . . . . .	2043
<b>Chapter 21</b>	<b>Troubleshooting . . . . .</b>	<b>2045</b>
	Troubleshooting Procedures . . . . .	2045
	Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSDB-Managed VXLAN . . . . .	2045
<b>Part 7</b>	<b>OpenFlow</b>	
<b>Chapter 22</b>	<b>Overview . . . . .</b>	<b>2049</b>
	Understanding Support for OpenFlow on Devices Running Junos OS . . . . .	2049
	OpenFlow Overview . . . . .	2049
	OpenFlow Virtual Switches . . . . .	2050
	OpenFlow Interfaces . . . . .	2050
	Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS . . . . .	2051
	OpenFlow Operation and Support . . . . .	2051
	OpenFlow Forwarding Actions . . . . .	2054
	Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS . . . . .	2056
	Understanding the OpenFlow Version Negotiation Between the Controller and Junos OS Devices . . . . .	2057
	Understanding OpenFlow Flows and Filters on Devices Running Junos OS . . . . .	2058
	Understanding OpenFlow Flow Instructions on Devices Running Junos OS . . . . .	2060
	Understanding How the OpenFlow Group Action Works . . . . .	2061
	Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS . . . . .	2062
	OpenFlow Flow Entry Timer Overview . . . . .	2062
	Idle Timeout and Hard Timeout . . . . .	2062
	Purge Flow Timer . . . . .	2063
	Understanding OpenFlow Barrier Messages on Devices Running Junos OS . . . . .	2064
	Understanding OpenFlow Multipart Messages on Devices Running Junos OS . . . . .	2064
	OpenFlow v1.0 Compliance Matrix for Devices Running Junos OS . . . . .	2065
	OpenFlow v1.0 Compliance Matrix for QFX5100 Switches . . . . .	2071
	OpenFlow v1.0 Compliance Matrix for EX4550 Switches . . . . .	2078
	OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS . . . . .	2087
<b>Chapter 23</b>	<b>Installing Support for OpenFlow . . . . .</b>	<b>2097</b>
	Installing Support for OpenFlow on Devices Running Junos OS . . . . .	2097
<b>Chapter 24</b>	<b>OpenFlow Basic Configuration . . . . .</b>	<b>2099</b>
	Configuring Support for OpenFlow on MX Series Routers . . . . .	2099
	Configuring the OpenFlow Interfaces . . . . .	2100
	Configuring the OpenFlow Protocol . . . . .	2100

	Configuring the OpenFlow Routing Instance . . . . .	2101
	Example: Enabling OpenFlow on MX Series Routers . . . . .	2102
	Configuring Support for OpenFlow on EX9200 Switches . . . . .	2106
	Configuring the OpenFlow Interfaces . . . . .	2107
	Configuring the OpenFlow Protocol . . . . .	2107
	Configuring the OpenFlow Routing Instance . . . . .	2108
	Example: Enabling OpenFlow on EX9200 Switches . . . . .	2109
	Configuring Support for OpenFlow on QFX5100 Switches . . . . .	2113
	Configuring the OpenFlow Interfaces . . . . .	2114
	Configuring the OpenFlow Protocol . . . . .	2114
	Example: Enabling OpenFlow on QFX5100 Switches . . . . .	2115
	Configuring Support for OpenFlow on EX4550 Switches . . . . .	2120
	Configuring the OpenFlow Interfaces . . . . .	2120
	Configuring the OpenFlow Protocol . . . . .	2120
	Example: Enabling OpenFlow on EX4550 Switches . . . . .	2121
<b>Chapter 25</b>	<b>Configuring OpenFlow Hybrid Interfaces . . . . .</b>	<b>2127</b>
	Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS . . . .	2127
	Configuring OpenFlow Hybrid Interfaces on MX Series Routers . . . . .	2128
	Configuring the Hybrid Physical Interface . . . . .	2129
	Configuring the Hybrid Interface Logical Units . . . . .	2129
	Configuring the Non-Hybrid Interfaces . . . . .	2129
	Configuring OpenFlow . . . . .	2130
	Configuring the Virtual Switch Routing Instances . . . . .	2130
	Example: Configuring OpenFlow Hybrid Interfaces on MX Series Routers . . . .	2131
	Configuring OpenFlow Hybrid Interfaces on EX9200 Switches . . . . .	2138
	Configuring the Hybrid Physical Interface . . . . .	2138
	Configuring the Hybrid Interface Logical Units . . . . .	2139
	Configuring the Non-Hybrid Interfaces . . . . .	2139
	Configuring OpenFlow . . . . .	2140
	Configuring the Virtual Switch Routing Instances . . . . .	2140
	Example: Configuring OpenFlow Hybrid Interfaces on EX9200 Switches . . . .	2140
<b>Chapter 26</b>	<b>Configuring OpenFlow Traffic Steering Across MPLS Networks . . . . .</b>	<b>2149</b>
	Understanding OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects . . . . .	2149
	Example: OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects . . . . .	2150
<b>Chapter 27</b>	<b>Configuration Statements . . . . .</b>	<b>2169</b>
	[edit protocols openflow] Hierarchy Level . . . . .	2169
	address (Protocols OpenFlow) . . . . .	2170
	controller (Protocols OpenFlow) . . . . .	2171
	default-action (Protocols OpenFlow) . . . . .	2172
	id (Protocols OpenFlow) . . . . .	2173
	interfaces (Protocols OpenFlow) . . . . .	2174
	openflow (Protocols OpenFlow) . . . . .	2175
	port (Protocols OpenFlow) . . . . .	2176
	protocol (Protocols OpenFlow) . . . . .	2177
	purge-flow-timer (Protocols OpenFlow) . . . . .	2178

	role (Protocols OpenFlow) . . . . .	2179
	switch (Protocols OpenFlow) . . . . .	2180
	traceoptions (Protocols OpenFlow) . . . . .	2181
<b>Chapter 28</b>	<b>Operational Commands . . . . .</b>	<b>2183</b>
	OpenFlow Operational Mode Commands . . . . .	2183
	show openflow capability . . . . .	2185
	show openflow controller . . . . .	2191
	show openflow filters . . . . .	2194
	show openflow flows . . . . .	2197
	show openflow groups . . . . .	2201
	show openflow interfaces . . . . .	2204
	show openflow statistics flows . . . . .	2208
	show openflow statistics groups . . . . .	2211
	show openflow statistics interfaces . . . . .	2213
	show openflow statistics packet . . . . .	2216
	show openflow statistics queue . . . . .	2218
	show openflow statistics summary . . . . .	2221
	show openflow statistics tables . . . . .	2223
	show openflow summary . . . . .	2225
	show openflow switch . . . . .	2226
<b>Part 8</b>	<b>High Availability</b>	
<b>Chapter 29</b>	<b>Overview . . . . .</b>	<b>2231</b>
	Software Feature Overview . . . . .	2231
	Understanding Graceful Routing Engine Switchover . . . . .	2231
	Graceful Routing Engine Switchover Concepts . . . . .	2231
	Effects of a Routing Engine Switchover . . . . .	2235
	Graceful Routing Engine Switchover System Requirements . . . . .	2237
	Graceful Routing Engine Switchover Platform Support . . . . .	2237
	Graceful Routing Engine Switchover Feature Support . . . . .	2238
	Graceful Routing Engine Switchover DPC Support . . . . .	2239
	Graceful Routing Engine Switchover and Subscriber Access . . . . .	2239
	Graceful Routing Engine Switchover PIC Support . . . . .	2239
	Nonstop Active Routing Concepts . . . . .	2240
	Nonstop Active Routing System Requirements . . . . .	2243
	Nonstop Active Routing Platform and Switching Platform Support . . . . .	2243
	Nonstop Active Routing Protocol and Feature Support . . . . .	2244
	Nonstop Active Routing BFD Support . . . . .	2247
	Nonstop Active Routing BGP Support . . . . .	2248
	Nonstop Active Routing Layer 2 Circuit and VPLS Support . . . . .	2249
	Nonstop Active Routing PIM Support . . . . .	2249
	Nonstop Active Routing MSDP Support . . . . .	2252
	Nonstop Active Routing Support for RSVP-TE LSPs . . . . .	2252
	Nonstop Bridging Concepts . . . . .	2254
	Nonstop Bridging System Requirements . . . . .	2256
	Platform Support . . . . .	2256
	Protocol Support . . . . .	2257
	Graceful Restart Concepts . . . . .	2257

	Understanding VRRP .....	2258
	Overview of VRRP .....	2259
	Sample VRRP Topology .....	2259
<b>Chapter 30</b>	<b>Configuration .....</b>	<b>2261</b>
	Configuration Tasks for Graceful Restart .....	2261
	Configuring Routing Protocols Graceful Restart .....	2261
	Enabling Graceful Restart .....	2262
	Configuring Graceful Restart Options for BGP .....	2262
	Configuring Graceful Restart Options for ES-IS .....	2263
	Configuring Graceful Restart Options for IS-IS .....	2263
	Configuring Graceful Restart Options for OSPF and OSPFv3 .....	2264
	Configuring Graceful Restart Options for RIP and RIPng .....	2266
	Configuring Graceful Restart Options for PIM Sparse Mode .....	2266
	Tracking Graceful Restart Events .....	2267
	Configuration Tasks for Graceful Switchover .....	2268
	Configuring Graceful Routing Engine Switchover .....	2268
	Enabling Graceful Routing Engine Switchover .....	2268
	Configuring Graceful Routing Engine Switchover with Graceful Restart .....	2268
	Synchronizing the Routing Engine Configuration .....	2269
	Verifying Graceful Routing Engine Switchover Operation .....	2269
	Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) .....	2270
	Resetting Local Statistics .....	2270
	Configuration Tasks for Nonstop Bridging .....	2271
	Configuring Nonstop Bridging on Switches (CLI Procedure) .....	2272
	Resetting Local Statistics .....	2273
	Configuration Example for Nonstop Active Routing .....	2274
	Example: Configuring Nonstop Active Routing on Switches .....	2274
	Configuration Tasks for Nonstop Active Routing .....	2277
	Configuring Nonstop Active Routing on Switches .....	2277
	Tracing Nonstop Active Routing Synchronization Events .....	2278
	Configuration Example for VRRP .....	2279
	Example: Configuring VRRP for Load Sharing .....	2279
	Configuration Tasks for VRRP .....	2284
	Configuring Basic VRRP Support .....	2285
	Configuring VRRP Authentication (IPv4 Only) .....	2286
	Configuring the Startup Period for VRRP Operations .....	2287
	Configuring the Advertisement Interval for the VRRP Master .....	2287
	Modifying the Advertisement Interval in Seconds .....	2287
	Modifying the Advertisement Interval in Milliseconds .....	2288
	Configuring VRRP Preemption and Hold Time .....	2288
	Configuring VRRP Preemption .....	2288
	Configuring the Preemption Hold Time .....	2289
	Overriding the Hold Time .....	2289
	Configuring a Route to Be Tracked .....	2289
	Configuring a Logical Interface to Be Tracked .....	2290

Configuring a Backup to Accept Packets Destined for the Virtual IP Address	2292
Configuring Passive ARP Learning for VRRP Backups	2292
Configuring the Silent Period	2293
Configuring Inheritance for a VRRP Group	2293
Configuration Statements for Graceful Restart	2294
disable	2295
disable (BGP Graceful Restart)	2296
graceful-restart (Enabling Globally)	2297
graceful-restart (Protocols BGP)	2299
graceful-restart (Protocols OSPF)	2300
helper-disable (OSPF)	2302
no-strict-lsa-checking	2303
notify-duration	2304
redundancy (Graceful Switchover)	2305
restart-duration	2306
restart-time (BGP Graceful Restart)	2307
stale-routes-time	2308
Configuration Statement for Graceful Switchover	2308
graceful-switchover	2309
redundancy (Graceful Switchover)	2310
Configuration Statement for Nonstop Bridging	2310
nonstop-bridging	2311
Configuration Statements for Nonstop Routing	2311
nonstop-routing	2312
synchronize	2313
traceoptions (Routing Options)	2315
Configuration Statements for VRRP	2317
accept-data	2318
advertise-interval	2319
asymmetric-hold-time	2320
authentication-key	2321
authentication-type	2322
bandwidth-threshold	2323
failover-delay	2324
fast-interval	2325
hold-time (VRRP)	2326
interface (VRRP Group)	2327
preempt (VRRP)	2328
priority (Protocols VRRP)	2329
priority-cost (VRRP)	2330
priority-hold-time	2331
route (Interfaces)	2332
startup-silent-period	2333
traceoptions	2334
track (VRRP)	2336
virtual-address	2337
vrrp-group	2338

<b>Chapter 31</b>	<b>Administration</b>	<b>2341</b>
	Operational Mode Commands for Graceful Restart	2341
	Verifying Graceful Restart Operation	2341
	Graceful Restart Operational Mode Commands	2341
	Verifying BGP Graceful Restart	2342
	Verifying IS-IS and OSPF Graceful Restart	2342
	Verifying CCC and TCC Graceful Restart	2343
	show bgp neighbor	2344
	show log	2358
	show (ospf   ospf3) overview	2361
	Operational Mode Command for Graceful Switchover	2365
	show system switchover	2366
	show task replication	2370
	Operational Mode Command for Nonstop Routing	2371
	show task replication	2372
	Operational Mode Commands for VRRP	2373
	show vrrp	2374
<b>Chapter 32</b>	<b>Troubleshooting</b>	<b>2385</b>
	Troubleshooting Procedures	2385
	Troubleshooting VRRP	2385
<b>Part 9</b>	<b>Interfaces</b>	
<b>Chapter 33</b>	<b>Overview</b>	<b>2389</b>
	Interfaces Overview	2389
	Interfaces Overview	2389
	Network Interfaces	2390
	Special Interfaces	2391
	Overview of Uplink Failure Detection	2392
	Uplink Failure Detection Configuration	2392
	Failure Detection Pair	2393
	Understanding Aggregated Ethernet Interfaces and LACP	2393
	Link Aggregation Group	2394
	Link Aggregation Control Protocol (LACP)	2395
	Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop	
	ECMP Traffic	2396
	Understanding the Hashing Algorithm	2397
	IP (IPv4 and IPv6)	2398
	MPLS	2399
	MAC-in-MAC Packet Hashing	2400
	Layer 2 Header Hashing	2400
	Understanding Interface Naming Conventions	2401
	Physical Part of an Interface Name	2401
	Logical Part of an Interface Name on a Switch Running QFabric	
	Software Package	2405
	Logical Part of a Channelized Interface Name on a Switch Running	
	Enhanced Layer 2 Software	2406
	Wildcard Characters in Interface Names	2406

Understanding Interface Ranges . . . . .	2406
Understanding Layer 3 Logical Interfaces . . . . .	2408
Understanding Local Link Bias . . . . .	2408
Understanding Management Interfaces . . . . .	2410
Understanding Multichassis Link Aggregation . . . . .	2411
Active-Active Mode . . . . .	2412
ICCP and ICL-PL . . . . .	2413
Failure Handling . . . . .	2413
Multichassis Link Protection . . . . .	2414
MC-LAG Packet Forwarding . . . . .	2414
Layer 3 Routing . . . . .	2414
Spanning Tree Protocol (STP) Guidelines . . . . .	2414
MC-LAG Upgrade Guidelines . . . . .	2415
Layer 2 Unicast Features Supported . . . . .	2415
Layer 2 Multicast Features Supported . . . . .	2416
IGMP Snooping on an Active-Active MC-LAG . . . . .	2416
Layer 3 Unicast Features Supported . . . . .	2417
VRRP Active-Standby Support . . . . .	2417
Routed VLAN Interface (RVI) MAC Address Synchronization . . . . .	2417
Address Resolution Protocol (ARP) . . . . .	2418
DHCP Relay with Option 82 . . . . .	2418
Private VLAN (PVLAN) . . . . .	2419
Layer 3 Multicast . . . . .	2419
Understanding Port Ranges and System Modes . . . . .	2421
Port Ranges for Different Media Types . . . . .	2421
Supported System Modes . . . . .	2444
Understanding Redundant Trunk Links . . . . .	2447
Understanding Generic Routing Encapsulation . . . . .	2449
Overview of GRE . . . . .	2449
GRE Tunneling . . . . .	2449
Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch . . . . .	2451
Configuration Limitations . . . . .	2452
Understanding Ethernet OAM Link Fault Management . . . . .	2452
Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups . . . . .	2454
Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing . . . . .	2454
Limitations and Caveats for Resilient Hashing . . . . .	2455
Resilient Hashing on LAGs . . . . .	2455
Resilient Hashing on ECMP . . . . .	2456
<b>Chapter 34 Configuration . . . . .</b>	<b>2457</b>
Configuration Examples . . . . .	2457
Example: Configuring Interfaces for Uplink Failure Detection . . . . .	2457
Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch . . . . .	2462
Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch . . . . .	2466

Example: Configuring Multichassis Link Aggregation . . . . .	2471
Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP . . . . .	2493
Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization . . . . .	2530
Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) . . . . .	2551
Example: Configuring Redundant Trunk Links for Faster Recovery . . . . .	2578
Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches . . . . .	2583
Configuration Tasks . . . . .	2585
Configuring Gigabit and 10-Gigabit Ethernet Interfaces . . . . .	2586
Configuring Port Mode . . . . .	2586
Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces . . . . .	2587
Configuring the Speed of Gigabit Ethernet Copper SFP Interfaces . . . . .	2588
Configuring the IP Options . . . . .	2588
Configuring Aggregated Ethernet LACP . . . . .	2589
Configuring Ethernet Loopback Capability . . . . .	2589
Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) . . . . .	2590
Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing . . . . .	2590
Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing . . . . .	2591
Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing . . . . .	2591
Configuring Interfaces for Uplink Failure Detection . . . . .	2592
Configuring a Layer 3 Logical Interface . . . . .	2593
Configuring Link Aggregation . . . . .	2593
Creating an Aggregated Ethernet Interface . . . . .	2594
Configuring the VLAN Name and VLAN ID Number . . . . .	2594
Configuring Aggregated Ethernet LACP . . . . .	2594
Configuring Local Link Bias (CLI Procedure) . . . . .	2596
Configuring Multichassis Link Aggregation . . . . .	2597
Configuring Generic Routing Encapsulation Tunneling . . . . .	2600
Configuring a GRE Tunnel . . . . .	2601
Configuring the LPM Table With Junos OS 13.2x51-D10 . . . . .	2602
Configuring Ethernet OAM Link Fault Management (CLI Procedure) . . . . .	2604
Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up . . . . .	2606
Configuring Resilient Hashing for Trunk/ECMP Groups . . . . .	2607
Configuring Resilient Hashing on LAGs . . . . .	2607
Configuring Resilient Hashing on ECMP Groups . . . . .	2607
Configuration Tasks . . . . .	2607
Channelizing Interfaces . . . . .	2608
Configuring the System Mode . . . . .	2610



Configuration Statements . . . . .	2612
[edit interfaces et] Configuration Statement Hierarchy on the QFX	
Series . . . . .	2616
Supported Statements in the [edit interfaces et] Hierarchy Level . . .	2616
Unsupported Statements in the [edit interfaces et] Hierarchy Level . .	2620
802.3ad . . . . .	2622
action (OAM LFM) . . . . .	2623
action-profile . . . . .	2624
address . . . . .	2625
aggregated-devices . . . . .	2627
aggregated-ether-options . . . . .	2628
alarm (chassis) . . . . .	2630
allow-remote-loopback . . . . .	2631
authentication-key (ICCP) . . . . .	2631
auto-negotiation . . . . .	2632
backup-liveness-detection . . . . .	2633
backup-peer-ip . . . . .	2633
channel-speed . . . . .	2634
chassis . . . . .	2635
chassis-id . . . . .	2636
configured-flow-control . . . . .	2637
container-devices . . . . .	2638
craft-lockout . . . . .	2639
description (Interfaces) . . . . .	2640
destination (Tunnels) . . . . .	2641
detection-time (Liveness Detection) . . . . .	2641
device-count . . . . .	2642
disk-failure-action . . . . .	2642
ecmp-resilient-hash . . . . .	2643
enhanced-hash-key . . . . .	2644
ethernet . . . . .	2645
ethernet (OAM LFM) . . . . .	2646
ethernet (Alarm) . . . . .	2648
ethernet-switching . . . . .	2649
ether-options . . . . .	2650
eui-64 . . . . .	2651
event (OAM LFM) . . . . .	2651
event-thresholds . . . . .	2652
family . . . . .	2653
fibre-channel (Alarm) . . . . .	2656
filter . . . . .	2657
flow-control . . . . .	2659
force-up . . . . .	2660
fpc . . . . .	2661
frame-error . . . . .	2662
frame-period . . . . .	2662
frame-period-summary . . . . .	2663
gratuitous-arp-reply . . . . .	2663
group . . . . .	2664

group (Redundant Trunk Groups) .....	2665
hash-mode .....	2666
hold-time (Physical Interface) .....	2668
iccp .....	2670
irb (Interfaces) .....	2672
inet (interfaces) .....	2675
inet (enhanced-hash-key) .....	2676
inet6 (interfaces) .....	2677
inet6 (enhanced-hash-key) .....	2678
interface (Multichassis Protection) .....	2679
interface (OAM LFM) .....	2680
interface (Redundant Trunk Groups) .....	2681
interface-mode .....	2682
interface-range .....	2684
interfaces .....	2686
lACP (802.3ad) .....	2693
lACP (Aggregated Ethernet) .....	2694
layer2 (enhanced-hash-key) .....	2695
link-adjacency-loss .....	2696
link-discovery .....	2697
link-down .....	2697
link-event-rate .....	2698
link-fault-management .....	2699
link-to-disable .....	2700
link-to-monitor .....	2700
link-down .....	2701
link-mode .....	2702
link-speed .....	2703
liveness-detection .....	2704
local-bias .....	2705
local-ip-addr (ICCP) .....	2705
loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet) .....	2706
management-ethernet (Alarm) .....	2706
member .....	2707
member-range .....	2707
mc-ae .....	2708
mc-ae-id .....	2709
minimum-interval (Liveness Detection) .....	2709
minimum-links .....	2710
minimum-receive-interval (Liveness Detection) .....	2710
mode (QFX Series) .....	2711
multi-chassis .....	2711
multi-chassis-protection .....	2712
multiplier (Liveness Detection) .....	2712
mtu .....	2713
negotiation-options .....	2714
no-adaptation (Liveness Detection) .....	2714
no-allow-link-events .....	2715

no-gratuitous-arp-request	2715
oam	2716
on-disk-failure	2718
on-loss-of-keepalives	2719
pdu-interval	2720
pdu-threshold	2720
peer (ICCP)	2721
peer (Multichassis)	2722
periodic	2722
pic	2723
preempt-cutover-timer	2724
redundancy (Graceful Switchover)	2725
redundant-trunk-group	2726
remote-loopback	2727
resilient-hash	2727
rx-buffers	2728
routing-engine	2729
service-id	2730
session-establishment-hold-time	2730
source	2731
speed	2732
status-control	2732
symbol-period	2733
syslog (OAM LFM)	2733
targeted-broadcast	2734
threshold (Detection Time)	2734
traceoptions (ICCP)	2735
transmit-interval (Liveness Detection)	2737
traceoptions (Individual Interfaces)	2738
traceoptions (OAM LFM)	2739
traps	2740
tunnel	2741
tunnel-port	2741
tx-buffers	2742
unit	2744
uplink-failure-detection	2745
version (Liveness Detection)	2745
vlan-id	2746
vlan-tagging	2746
<b>Chapter 35 Administration</b>	<b>2747</b>
Routine Monitoring	2747
Monitoring System Process Information	2747
Monitoring System Properties	2748
Monitoring Interface Status and Traffic	2749
Verifying That Layer 3 Logical Interfaces Are Working	2750

Verifying the Status of a LAG Interface .....	2750
Verifying That LACP Is Configured Correctly and Bundle Members Are	
Exchanging LACP Protocol Packets .....	2751
Verifying the LACP Setup .....	2751
Verifying That LACP Packets Are Being Exchanged .....	2751
Verifying That Generic Routing Encapsulation Tunneling Is Working	
Correctly .....	2752
Monitoring Commands .....	2753
monitor interface .....	2754
show forwarding-options enhanced-hash-key .....	2763
show iccp .....	2766
show interfaces diagnostics optics .....	2768
show interfaces ge .....	2782
show interfaces (GRE) .....	2797
show interfaces irb .....	2804
show interfaces mc-ae .....	2810
show interfaces queue .....	2812
show interfaces xe .....	2852
show lacp interfaces .....	2870
show lacp statistics interfaces (View) .....	2875
show oam ethernet link-fault-management .....	2877
show redundant-trunk-group .....	2882
show uplink-failure-detection .....	2884
<b>Chapter 36 Troubleshooting .....</b>	<b>2887</b>
Troubleshooting Procedures .....	2887
Troubleshooting an Aggregated Ethernet Interface .....	2887
Troubleshooting Multichassis Link Aggregation .....	2887
MAC Addresses Learned on MC-AE Interfaces Are Not Removed from	
the MAC Address Table .....	2888
MC-LAG Peer Does Not Go into Standby Mode .....	2889
Secondary MC-LAG Peer with Status Control Set to Standby Becomes	
Inactive .....	2889
Redirect Filters Take Priority over User-Defined Filters .....	2889
Operational Command Output Is Wrong .....	2889
ICCP Connection Might Take Up to 60 Seconds to Become Active ..	2890
MAC Address Age Learned on an MC-AE Interface Is Reset to Zero ..	2890
MAC Address Is Not Learned Remotely in a Default VLAN .....	2890
Snooping Entries Learned on MC-AE Interfaces Are Not Removed ..	2890
ICCP Does Not Come Up After You Add or Delete an Authentication	
Key .....	2891
Local Status Is Standby When It Should Be Active .....	2891
Packets Loop on the Server When ICCP Fails .....	2891
Both MC-LAG Peers Use the Default System ID After a Reboot or an	
ICCP Configuration Change .....	2891
No Commit Checks Are Done for ICL-PL Interfaces .....	2892
Double Failover Scenario .....	2892
Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes	
Down and Up .....	2892

	Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer . . . . .	2892
	AE Interfaces Go Down . . . . .	2892
	Flooding of Upstream Traffic . . . . .	2893
	Troubleshooting Network Interfaces . . . . .	2893
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down . . . . .	2893
<b>Part 10</b>	<b>Routing Options</b>	
<b>Chapter 37</b>	<b>Overview . . . . .</b>	<b>2897</b>
	Routing Options Overview . . . . .	2897
	Overview of Routing Options . . . . .	2897
	Understanding Virtual Router Routing Instances . . . . .	2898
	Understanding Distributed Periodic Packet Management . . . . .	2898
	Understanding Bidirectional Forwarding Detection (BFD) . . . . .	2899
	Understanding the Unified Forwarding Table . . . . .	2899
	Using the Unified Forwarding Table to Optimize Address Storage . . . . .	2899
	MAC Address and Host Address Memory Allocation . . . . .	2900
	LPM Table Memory Allocation . . . . .	2901
<b>Chapter 38</b>	<b>Configuration . . . . .</b>	<b>2903</b>
	Configuration Tasks . . . . .	2903
	Configuring Static Routing . . . . .	2904
	Configuring Per-Packet Load Balancing . . . . .	2904
	Configuring Distributed Periodic Packet Management . . . . .	2906
	Disabling or Enabling Distributed Periodic Packet Management Globally . . . . .	2907
	Disabling or Enabling Distributed Periodic Packet Management for LACP Packets . . . . .	2907
	Configuring Virtual Router Routing Instances . . . . .	2908
	Configuring the Unified Forwarding Table . . . . .	2909
	Configuring an Address-Storage Profile . . . . .	2909
	Configuring the LPM Allocation . . . . .	2910
	Configuration Examples . . . . .	2913
	Examples: Configuring Per-Packet Load Balancing . . . . .	2913
	Examples: Configuring BFD for Static Routes . . . . .	2914
	Understanding BFD for Static Routes . . . . .	2914
	Example: Configuring BFD for Static Routes . . . . .	2918
	Example: Enabling BFD on Qualified Next Hops in Static Routes . . . . .	2923
	Example: Configuring BFD Authentication for Static Routes . . . . .	2929
	Understanding BFD Authentication for Static Routes . . . . .	2929
	Example: Configuring BFD Authentication for Static Routes . . . . .	2931
	Configuration Statements . . . . .	2937
	active . . . . .	2940
	aggregate (Routing) . . . . .	2941
	as-path (Routing Options) . . . . .	2943
	autonomous-system . . . . .	2945
	backup-pe-group . . . . .	2947
	backups . . . . .	2948

bandwidth (Multicast Flow Map) .....	2949
bfd-liveness-detection (Routing Options Static Route) .....	2950
bgp-orf-cisco-mode .....	2954
bmp .....	2956
brief .....	2958
centralized .....	2959
community (Routing Options) .....	2960
confederation .....	2962
disable (Routing Options) .....	2963
description (Routing Instances) .....	2963
discard .....	2964
export (Routing Options) .....	2965
export-rib .....	2966
fate-sharing .....	2968
flow .....	2969
flow-map .....	2970
forwarding-cache (Flow Maps) .....	2971
forwarding-cache (Multicast) .....	2972
forwarding-options (chassis) .....	2974
forwarding-table .....	2975
generate .....	2976
graceful-restart (Enabling Globally) .....	2978
import (Routing Options) .....	2979
import-policy .....	2980
import-rib .....	2981
indirect-next-hop .....	2982
install (Routing Options) .....	2983
instance-export .....	2984
instance-import .....	2984
instance-type .....	2985
interface (Multicast Static Routes) .....	2986
interface (Routing Instances) .....	2987
interface (Routing Options) .....	2988
interface-routes .....	2989
local-address (Routing Options) .....	2990
martians .....	2991
maximum-bandwidth (Routing Options) .....	2992
maximum-paths .....	2993
maximum-prefixes .....	2995
med-igp-update-interval .....	2996
metric (Aggregate, Generated, or Static Route) .....	2997
multicast (Routing Options) .....	2998
no-qos-adjust .....	2999
num-65-127-prefix .....	3000
options (Routing Options) .....	3001
pim-to-igmp-proxy .....	3002
pim-to-mld-proxy .....	3003
policy (Aggregate and Generated Routes) .....	3004
policy (Flow Maps) .....	3005

policy-options	3006
policy-statement	3007
ppm	3011
ppm (Ethernet Switching)	3012
preference (Routing Options)	3013
prefix	3014
prefix-65-127-disable	3014
protocols	3015
qualified-next-hop (Static Routes)	3017
readvertise	3019
redundant-sources	3020
resolution	3021
resolution-ribs	3022
resolve	3023
restart-duration (Routing Options)	3024
retain	3025
reverse-oif-mapping	3026
rpf-check-policy (Routing Options RPF)	3027
rib (General)	3028
rib (Route Resolution)	3030
rib-group (Routing Options)	3031
rib-groups	3032
route-distinguisher-id	3034
route-record	3035
router-id	3036
routing-instances	3037
routing-options	3038
scope	3038
scope-policy	3039
source (Source-Specific Multicast)	3040
source-routing	3041
ssm-groups	3042
ssm-map (Routing Options Multicast)	3043
static (Routes)	3044
subscriber-leave-timer	3046
tag (Routing Options)	3047
threshold (Multicast Forwarding Cache)	3048
timeout (Flow Maps)	3049
timeout (Multicast)	3050
traceoptions (Routing Options)	3051
upstream-interface	3054
<b>Chapter 39 Administration</b>	<b>3055</b>
Routine Monitoring	3055
Monitoring Routing Information	3055
Verifying That Virtual Router Routing Instances Are Working	3056
Operational Commands	3057
clear ipv6 neighbors	3059
show as-path	3060

	show as-path domain . . . . .	3064
	show as-path summary . . . . .	3067
	show ipv6 neighbors . . . . .	3069
	show ipv6 router-advertisement . . . . .	3071
	show route . . . . .	3074
	show route active-path . . . . .	3080
	show route all . . . . .	3085
	show route aspath-regex . . . . .	3087
	show route best . . . . .	3089
	show route brief . . . . .	3092
	show route community . . . . .	3094
	show route community-name . . . . .	3096
	show route damping . . . . .	3098
	show route detail . . . . .	3103
	show route exact . . . . .	3120
	show route export . . . . .	3122
	show route extensive . . . . .	3125
	show route flow validation . . . . .	3142
	show route forwarding-table . . . . .	3144
	show route inactive-path . . . . .	3158
	show route inactive-prefix . . . . .	3161
	show route instance . . . . .	3163
	show route label . . . . .	3171
	show route label-switched-path . . . . .	3174
	show route martians . . . . .	3176
	show route next-hop . . . . .	3178
	show route no-community . . . . .	3184
	show route protocol . . . . .	3187
	show route range . . . . .	3199
	show route receive-protocol . . . . .	3203
	show route resolution . . . . .	3211
	show route snooping . . . . .	3214
	show route source-gateway . . . . .	3222
	show route summary . . . . .	3228
	show route table . . . . .	3232
	show route terse . . . . .	3246
<b>Chapter 40</b>	<b>Troubleshooting . . . . .</b>	<b>3249</b>
	Troubleshooting Procedures . . . . .	3249
	Troubleshooting Virtual Routing Instances . . . . .	3249
	Direct Routes Not Leaked Between Routing Instances . . . . .	3249
<b>Part 11</b>	<b>Border Gateway Protocol</b>	
<b>Chapter 41</b>	<b>Overview . . . . .</b>	<b>3253</b>
	BGP Overview . . . . .	3253
	Understanding BGP . . . . .	3254
	Autonomous Systems . . . . .	3254
	AS Paths and Attributes . . . . .	3254
	External and Internal BGP . . . . .	3255



	Multiple Instances of BGP .....	3255
	BGP Routes Overview .....	3256
	BGP Messages Overview .....	3257
	Open Messages .....	3257
	Update Messages .....	3258
	Keepalive Messages .....	3258
	Notification Messages .....	3258
	Understanding the Advertisement of Multiple Paths to a Single Destination in BGP .....	3258
<b>Chapter 42</b>	<b>Configuration .....</b>	<b>3261</b>
	Basic BGP Configuration .....	3261
	Examples: Configuring External BGP Peering .....	3261
	Understanding External BGP Peering Sessions .....	3261
	Example: Configuring External BGP Point-to-Point Peer Sessions . . .	3262
	Example: Configuring External BGP on Logical Systems with IPv6 Interfaces .....	3269
	Examples: Configuring Internal BGP Peering .....	3284
	Understanding Internal BGP Peering Sessions .....	3284
	Example: Configuring Internal BGP Peer Sessions .....	3285
	Example: Configuring Internal BGP Peering Sessions on Logical Systems .....	3296
	Configuring BGP Monitoring Protocol Version 3 .....	3307
	BGP Path Attribute Configuration .....	3310
	Example: Configuring BGP Local Preference .....	3310
	Understanding the BGP Local Preference .....	3310
	Example: Configuring the Local Preference Value for BGP Routes . . .	3310
	Examples: Configuring BGP MED .....	3323
	Understanding the MED Attribute .....	3323
	Example: Configuring the MED Attribute Directly .....	3325
	Example: Configuring the MED Using Route Filters .....	3338
	Example: Configuring the MED Using Communities .....	3351
	Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates .....	3352
	Examples: Configuring BGP Local AS .....	3362
	Understanding the BGP Local AS Attribute .....	3362
	Example: Configuring a Local AS for EBGp Sessions .....	3367
	Example: Configuring a Private Local AS for EBGp Sessions .....	3377
	Example: Configuring the Accumulated IGP Attribute for BGP .....	3382
	Understanding the Accumulated IGP Attribute for BGP .....	3383
	Example: Configuring the Accumulated IGP Attribute for BGP .....	3383
	BGP Policy Configuration .....	3421
	Example: Configuring BGP Interactions with IGP .....	3421
	Understanding Routing Policies .....	3421
	Example: Injecting OSPF Routes into the BGP Routing Table .....	3422
	Example: Configuring BGP Route Advertisement .....	3425
	Understanding Route Advertisement .....	3425
	Example: Configuring BGP Prefix-Based Outbound Route Filtering . .	3429

Example: Configuring EBGP Multihop . . . . .	3433
Understanding BGP Multihop . . . . .	3433
Example: Configuring EBGP Multihop Sessions . . . . .	3433
Example: Configuring BGP Route Preference (Administrative Distance) . . . . .	3442
Understanding Route Preference Values . . . . .	3442
Example: Configuring the Preference Value for BGP Routes . . . . .	3443
Example: Configuring BGP Path Selection . . . . .	3449
Example: Ignoring the AS Path Attribute When Selecting the Best Path . . . . .	3449
Example: Removing Private AS Numbers . . . . .	3456
Understanding Private AS Number Removal from AS Paths . . . . .	3456
Example: Removing Private AS Numbers from AS Paths . . . . .	3457
BGP BFD Configuration . . . . .	3462
Example: Configuring BFD for BGP . . . . .	3462
Understanding BFD for BGP . . . . .	3462
Example: Configuring BFD on Internal BGP Peer Sessions . . . . .	3463
Example: Configuring BFD Authentication for BGP . . . . .	3471
Understanding BFD Authentication for BGP . . . . .	3471
Example: Configuring BFD Authentication for BGP . . . . .	3473
BGP Load Balancing Configuration . . . . .	3476
Examples: Configuring BGP Multipath . . . . .	3476
Understanding BGP Multipath . . . . .	3476
Example: Load Balancing BGP Traffic . . . . .	3477
Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops . . . . .	3482
Example: Advertising Multiple BGP Paths to a Destination . . . . .	3494
Understanding the Advertisement of Multiple Paths to a Single Destination in BGP . . . . .	3494
Example: Advertising Multiple Paths in BGP . . . . .	3495
Example: Advertising Multiple Paths in BGP . . . . .	3520
Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing . . . . .	3545
IBGP Scaling Configuration . . . . .	3547
Example: Configuring BGP Route Reflectors . . . . .	3547
Understanding BGP Route Reflectors . . . . .	3547
Example: Configuring a Route Reflector . . . . .	3549
Example: Configuring BGP Confederations . . . . .	3564
Understanding BGP Confederations . . . . .	3564
Example: Configuring BGP Confederations . . . . .	3565
BGP Security Configuration . . . . .	3570
Example: Configuring BGP Route Authentication . . . . .	3571
Understanding Route Authentication . . . . .	3571
Example: Configuring Route Authentication for BGP . . . . .	3572
Examples: Configuring TCP and BGP Security . . . . .	3577
Understanding Security Options for BGP with TCP . . . . .	3577
Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers . . . . .	3578
Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List . . . . .	3583

Example: Limiting TCP Segment Size for BGP .....	3586
BGP Flap Configuration .....	3591
Example: Preventing BGP Session Resets .....	3591
Understanding BGP Session Resets .....	3591
Example: Preventing BGP Session Flaps When VPN Families Are Configured .....	3591
Examples: Configuring BGP Flap Damping .....	3599
Understanding BGP Route Flap Damping Parameters .....	3599
Example: Configuring BGP Route Flap Damping Parameters .....	3600
Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family .....	3609
BGP Monitoring Configuration .....	3619
Example: Configuring BGP Trace Operations .....	3619
Understanding Trace Operations for BGP Protocol Traffic .....	3619
Example: Viewing BGP Trace Files on Logical Systems .....	3621
Tracing BMP Operations .....	3625
Configuration Statements .....	3627
accept-remote-nexthop .....	3630
advertise-external .....	3631
advertise-inactive .....	3633
advertise-peer-as .....	3634
algorithm (BGP BFD Authentication) .....	3635
apply-groups .....	3637
apply-groups-except .....	3637
authentication (BGP BFD Liveness Detection) .....	3638
authentication-algorithm .....	3640
authentication-key (Protocols BGP and BMP) .....	3641
authentication-key-chain (Protocols BGP and BMP) .....	3642
bfd-liveness-detection (Protocols BGP) .....	3643
bgp .....	3647
bgp-orf-cisco-mode .....	3648
cluster .....	3650
connection-mode .....	3651
damping (Protocols BGP) .....	3652
description (Protocols BGP) .....	3654
detection-time (BFD Liveness Detection) .....	3655
disable (Protocols BGP) .....	3656
disable (BGP Graceful Restart) .....	3657
export (Protocols BGP) .....	3658
family (Protocols BGP) .....	3659
graceful-restart (Protocols BGP) .....	3663
group (Protocols BGP) .....	3664
hold-down .....	3667
hold-down-interval (BGP BFD Liveness Detection) .....	3669
hold-time (Protocols BGP) .....	3671
import (Protocols BGP) .....	3673
include-mp-next-hop .....	3675
initiation-message .....	3676
keep .....	3677

key-chain (BGP BFD Authentication) .....	3679
local-address (Protocols BGP) .....	3681
local-address (Protocols BMP) .....	3683
local-as .....	3684
local-port .....	3686
local-preference .....	3687
log-updown (Protocols BGP) .....	3688
loops .....	3689
loose-check (BGP BFD Authentication) .....	3691
maximum-ecmp .....	3692
metric-out (Protocols BGP) .....	3693
minimum-interval (BFD Liveness Detection) .....	3695
minimum-interval (transmit-interval) .....	3697
minimum-receive-interval (BFD Liveness Detection) .....	3699
monitor (Protocols BMP) .....	3700
mtu-discovery .....	3701
multihop .....	3703
multiplier (BFD Liveness Detection) .....	3705
neighbor (Protocols BGP) .....	3707
no-adaptation (BFD Liveness Detection) .....	3710
no-advertise-peer-as .....	3711
no-aggregator-id .....	3712
no-client-reflect .....	3713
out-delay .....	3714
outbound-route-filter .....	3716
passive (Protocols BGP) .....	3717
path-selection .....	3718
peer-as (Protocols BGP) .....	3720
post-policy .....	3721
pre-policy .....	3722
preference (Protocols BGP) .....	3723
priority (Protocols BMP) .....	3724
remove-private .....	3725
restart-time (BGP Graceful Restart) .....	3727
route-monitoring .....	3728
session-mode .....	3729
stale-routes-time .....	3730
station .....	3731
station-address .....	3732
station-port .....	3733
statistics-timeout .....	3734
tcp-mss (Protocols BGP) .....	3735
threshold (detection-time) .....	3736
threshold (transmit-interval) .....	3738
traceoptions (Protocols BGP) .....	3740
traceoptions (Protocols BMP) .....	3743
transmit-interval (BFD Liveness Detection) .....	3745
version (BFD Liveness Detection) .....	3747

<b>Chapter 43</b>	<b>Administration</b>	<b>3749</b>
	Routine Monitoring	3749
	Monitoring BGP Routing Information	3749
	Operational Commands	3749
	clear bgp damping	3750
	clear bgp neighbor	3751
	clear bgp table	3753
	show bgp bmp	3755
	show bgp group	3757
	show bgp neighbor	3764
	show bgp summary	3778
	show policy damping	3784
	show route damping	3786
	show route detail	3791
<b>Part 12</b>	<b>Intermediate System to Intermediate System</b>	
<b>Chapter 44</b>	<b>Overview</b>	<b>3811</b>
	IS-IS Overview	3811
	IS-IS Overview	3812
	IS-IS Terminology	3812
	ISO Network Addresses	3813
	IS-IS Packets	3814
	Persistent Route Reachability	3815
	IS-IS Support for Multipoint Network Clouds	3815
	Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels	3816
	Understanding BFD Authentication for IS-IS	3816
	BFD Authentication Algorithms	3817
	Security Authentication Keychains	3817
	Strict Versus Loose Authentication	3818
	Understanding Hitless Authentication Key Rollover for IS-IS	3818
	Understanding Loop-Free Alternate Routes for IS-IS	3819
	Configuring Link Protection for IS-IS	3820
	Configuring Node-Link Protection for IS-IS	3821
	Excluding an IS-IS Interface as a Backup for Protected Interfaces	3821
	Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS	3822
	Using Operational Mode Commands to Monitor Protected IS-IS Routes	3822
<b>Chapter 45</b>	<b>Configuration</b>	<b>3823</b>
	Configuration Guidelines	3823
	Example: Configuring IS-IS	3823
	Configuration Examples	3828
	Example: Configuring Multi-Level IS-IS	3829
	Example: Configuring Hitless Authentication Key Rollover for IS-IS	3837
	Example: Redistributing OSPF Routes into IS-IS	3842
	Example: Configuring BFD for IS-IS	3850

Example: Configuring BFD Authentication for IS-IS . . . . .	3855
Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies . . . . .	3859
Understanding IS-IS IPv4 and IPv6 Unicast Topologies . . . . .	3859
Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies . . . . .	3859
Example: Configuring IS-IS Multicast Topology . . . . .	3867
IS-IS Multicast Topologies Overview . . . . .	3868
Example: Configuring IS-IS Multicast Topology . . . . .	3869
Example: Configuring Link and Node Protection for IS-IS Routes . . . . .	3883
Understanding Loop-Free Alternate Routes for IS-IS . . . . .	3883
Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN . . . . .	3887
Example: Configuring an IS-IS Default Route Policy on Logical Systems . .	3897
Example: Configuring IS-IS for CLNS . . . . .	3903
Understanding IS-IS for CLNS . . . . .	3903
Example: Configuring IS-IS for CLNS . . . . .	3903
Example: Configuring IS-IS Designated Routers . . . . .	3905
Understanding IS-IS Designated Routers . . . . .	3905
Example: Configuring Designated Router Election Priority for IS-IS . .	3906
Example: Enabling Packet Checksums on IS-IS Interfaces . . . . .	3906
Configuration Tasks . . . . .	3908
Configuring IS-IS Authentication . . . . .	3909
Configuring Authentication Without Network-Wide Deployment . . . . .	3910
Configuration Statements . . . . .	3910
authentication-key (Protocols IS-IS) . . . . .	3913
authentication-key-chain (Protocols IS-IS) . . . . .	3914
authentication-type (Protocols IS-IS) . . . . .	3915
bfd-liveness-detection (Protocols IS-IS) . . . . .	3916
checksum (Protocols IS-IS) . . . . .	3918
csnp-interval . . . . .	3919
disable (Protocols IS-IS) . . . . .	3920
export (Protocols IS-IS) . . . . .	3921
external-preference (Protocols IS-IS) . . . . .	3922
family (Protocols IS-IS) . . . . .	3923
graceful-restart (Protocols IS-IS) . . . . .	3924
hello-authentication-key . . . . .	3925
hello-authentication-key-chain . . . . .	3926
hello-authentication-type . . . . .	3927
hello-interval (Protocols IS-IS) . . . . .	3928
hello-padding . . . . .	3929
hold-time (Protocols IS-IS) . . . . .	3931
ignore-attached-bit . . . . .	3932
interface (Protocols IS-IS) . . . . .	3933
ipv4-multicast . . . . .	3935
ipv4-multicast-metric . . . . .	3936
ipv6-multicast . . . . .	3936
ipv6-multicast-metric . . . . .	3937
ipv6-unicast . . . . .	3938
ipv6-unicast-metric . . . . .	3939
isis . . . . .	3940

level (Global IS-IS) . . . . .	3941
link-protection (Protocols IS-IS) . . . . .	3942
loose-authentication-check . . . . .	3942
lsp-interval . . . . .	3943
lsp-lifetime . . . . .	3944
max-areas . . . . .	3945
mesh-group (Protocols IS-IS) . . . . .	3946
metric (Protocols IS-IS) . . . . .	3947
no-adjacency-holddown . . . . .	3948
no-authentication-check . . . . .	3949
no-csnp-authentication . . . . .	3949
no-eligible-backup (Protocols IS-IS) . . . . .	3950
no-hello-authentication . . . . .	3950
no-ipv4-multicast . . . . .	3951
no-ipv4-routing . . . . .	3952
no-ipv6-multicast . . . . .	3953
no-ipv6-routing . . . . .	3954
no-ipv6-unicast . . . . .	3955
no-psnp-authentication . . . . .	3955
no-unicast-topology . . . . .	3956
node-link-protection (Protocols IS-IS) . . . . .	3956
overload (Protocols IS-IS) . . . . .	3957
passive (Protocols IS-IS) . . . . .	3960
point-to-point . . . . .	3961
preference (Protocols IS-IS) . . . . .	3962
prefix-export-limit (Protocols IS-IS) . . . . .	3963
priority (Protocols IS-IS) . . . . .	3964
reference-bandwidth (Protocols IS-IS) . . . . .	3965
rib-group (Protocols IS-IS) . . . . .	3966
spf-options (Protocols IS-IS) . . . . .	3967
topologies (Protocols IS-IS) . . . . .	3968
traceoptions (Protocols IS-IS) . . . . .	3969
traffic-engineering (Protocols IS-IS) . . . . .	3972
wide-metrics-only . . . . .	3975
<b>Chapter 46 Administration . . . . .</b>	<b>3977</b>
Operational Commands . . . . .	3977
clear isis adjacency . . . . .	3978
clear isis database . . . . .	3980
clear isis overload . . . . .	3982
clear isis statistics . . . . .	3984
show isis adjacency . . . . .	3986
show isis authentication . . . . .	3990
show isis backup coverage . . . . .	3992
show isis backup label-switched-path . . . . .	3994
show isis backup spf results . . . . .	3996
show isis database . . . . .	4000
show isis hostname . . . . .	4012
show isis interface . . . . .	4014

show isis overview . . . . .	4018
show isis route . . . . .	4021
show isis spf . . . . .	4025
show isis statistics . . . . .	4030

## Part 13

### Chapter 47

## Open Shortest Path First

<b>Overview . . . . .</b>	<b>4035</b>
OSPF Overview . . . . .	4035
OSPF Overview . . . . .	4036
OSPF Default Route Preference Values . . . . .	4038
OSPF Routing Algorithm . . . . .	4038
OSPF Three-Way Handshake . . . . .	4039
OSPF Version 3 . . . . .	4040
OSPF Areas and Router Functionality Overview . . . . .	4041
Areas . . . . .	4041
Area Border Routers . . . . .	4041
Backbone Areas . . . . .	4041
AS Boundary Routers . . . . .	4042
Backbone Router . . . . .	4042
Internal Router . . . . .	4042
Stub Areas . . . . .	4042
Not-So-Stubby Areas . . . . .	4042
Transit Areas . . . . .	4043
Packets Overview . . . . .	4043
OSPF Packet Header . . . . .	4043
Hello Packets . . . . .	4044
Database Description Packets . . . . .	4044
Link-State Request Packets . . . . .	4044
Link-State Update Packets . . . . .	4044
Link-State Acknowledgment Packets . . . . .	4045
Link-State Advertisement Packet Types . . . . .	4045
OSPF External Metrics Overview . . . . .	4046



<b>Chapter 48</b>	<b>Configuration</b>	<b>4047</b>
	Basic OSPF Area Configuration	4047
	Examples: Configuring OSPF Designated Routers	4047
	OSPF Designated Router Overview	4047
	Example: Configuring an OSPF Router Identifier	4048
	Example: Controlling OSPF Designated Router Election	4050
	Examples: Configuring OSPF Areas	4052
	Understanding OSPF Areas and Backbone Areas	4052
	Example: Configuring a Single-Area OSPF Network	4053
	Example: Configuring a Multiarea OSPF Network	4055
	Advanced OSPF Area Configuration	4058
	Examples: Configuring OSPF Stub and Not-So-Stubby Areas	4059
	Understanding OSPF Stub Areas, Totally Stubby Areas, and	
	Not-So-Stubby Areas	4059
	Example: Configuring OSPF Stub and Totally Stubby Areas	4060
	Example: Configuring OSPF Not-So-Stubby Areas	4064
	Example: Configuring OSPF Multiarea Adjacency	4069
	Multiarea Adjacency for OSPF	4069
	Example: Configuring Multiarea Adjacency for OSPF	4070
	Example: Disabling OSPFv2 Compatibility with RFC 1583	4073
	OSPFv2 Compatibility with RFC 1583 Overview	4074
	Example: Disabling OSPFv2 Compatibility with RFC 1583	4074
	OSPF Interface Configuration	4075
	Examples: Configuring OSPF Interfaces	4075
	About OSPF Interfaces	4076
	Example: Configuring an Interface on a Broadcast or Point-to-Point	
	Network	4077
	Example: Configuring an OSPFv2 Interface on a Nonbroadcast	
	Multiaccess Network	4079
	Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint	
	Network	4082
	Example: Configuring OSPF Demand Circuits	4084
	Example: Configuring a Passive OSPF Interface	4086
	Example: Configuring OSPFv2 Peer interfaces	4088
	Example: Configuring Multiple Address Families for OSPFv3	4090
	Understanding Multiple Address Families for OSPFv3	4090
	Example: Configuring Multiple Address Families for OSPFv3	4091

OSPF Route Control Configuration . . . . .	4094
Examples: Configuring OSPF Route Summarization . . . . .	4094
Understanding OSPF Route Summarization . . . . .	4094
Example: Summarizing Ranges of Routes in OSPF Link-State	
Advertisements . . . . .	4094
Example: Limiting the Number of Prefixes Exported to OSPF . . . . .	4100
Configuring OSPF Refresh and Flooding Reduction in Stable	
Topologies . . . . .	4102
Examples: Configuring OSPF Traffic Control . . . . .	4103
Understanding OSPF Traffic Control . . . . .	4103
Example: Controlling the Cost of Individual OSPF Network	
Segments . . . . .	4105
Example: Dynamically Adjusting OSPF Interface Metrics Based on	
Bandwidth . . . . .	4109
Example: Controlling OSPF Route Preferences . . . . .	4111
Example: Configuring OSPF Overload Mode . . . . .	4113
OSPF Overload Function Overview . . . . .	4113
Example: Configuring OSPF to Make Routing Devices Appear	
Overloaded . . . . .	4114
OSPF Fault Detection Configuration . . . . .	4117
Example: Configuring OSPF Timers . . . . .	4117
OSPF Timers Overview . . . . .	4117
Example: Configuring OSPF Timers . . . . .	4118
Example: Configuring BFD for OSPF . . . . .	4123
Understanding BFD for OSPF . . . . .	4123
Example: Configuring BFD for OSPF . . . . .	4126
Example: Configuring BFD Authentication for OSPF . . . . .	4129
BFD Authentication for OSPF Overview . . . . .	4130
Configuring BFD Authentication for OSPF . . . . .	4131
OSPF Redundancy Features Configuration . . . . .	4134
Examples: Configuring Graceful Restart for OSPF . . . . .	4134
Graceful Restart for OSPF Overview . . . . .	4135
Example: Configuring Graceful Restart for OSPF . . . . .	4136
Example: Configuring the Helper Capability Mode for OSPFv2 Graceful	
Restart . . . . .	4140
Example: Configuring the Helper Capability Mode for OSPFv3 Graceful	
Restart . . . . .	4144
Example: Disabling Strict LSA Checking for OSPF Graceful Restart . .	4147
OSPF Traffic Engineering Configuration . . . . .	4150
Examples: Configuring OSPF Traffic Engineering . . . . .	4150
OSPF Support for Traffic Engineering . . . . .	4150
Example: Enabling OSPF Traffic Engineering Support . . . . .	4153
Example: Configuring the Traffic Engineering Metric for a Specific OSPF	
Interface . . . . .	4157
Example: Configuring OSPF Passive Traffic Engineering Mode . . . . .	4159
OSPF Passive Traffic Engineering Mode . . . . .	4159
Example: Configuring OSPF Passive Traffic Engineering Mode . . . . .	4159

OSPF Database Protection Configuration . . . . .	4162
Example: Configuring OSPF Database Protection . . . . .	4162
OSPF Database Protection Overview . . . . .	4162
Configuring OSPF Database Protection . . . . .	4163
OSPF Policy Configuration . . . . .	4164
Examples: Configuring OSPF Routing Policy . . . . .	4164
Understanding OSPF Routing Policy . . . . .	4164
Example: Injecting OSPF Routes into the BGP Routing Table . . . . .	4166
Example: Redistributing Static Routes into OSPF . . . . .	4169
Example: Configuring an OSPF Import Policy . . . . .	4172
Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF . . . . .	4176
Examples: Configuring Routing Policy for Network Summaries . . . . .	4180
Import and Export Policies for Network Summaries Overview . . . . .	4180
Example: Configuring an OSPF Export Policy for Network Summaries . . . . .	4181
Example: Configuring an OSPF Import Policy for Network Summaries . . . . .	4190
OSPF Monitoring Configuration . . . . .	4198
Example: Configuring OSPF Trace Options . . . . .	4198
Tracing OSPF Protocol Traffic . . . . .	4198
Example: Tracing OSPF Protocol Traffic . . . . .	4200
Configuration Statements . . . . .	4205
area . . . . .	4207
area-range . . . . .	4209
authentication (Protocols OSPF) . . . . .	4211
context-identifier (Protocols OSPF) . . . . .	4212
bfd-liveness-detection (Protocols OSPF) . . . . .	4213
database-protection . . . . .	4217
disable (OSPF) . . . . .	4219
export (Protocols OSPF) . . . . .	4221
external-preference (Protocols OSPF) . . . . .	4222
graceful-restart (Protocols OSPF) . . . . .	4223
import (Protocols OSPF) . . . . .	4225
interface (Protocols OSPF) . . . . .	4226
no-nssa-abr . . . . .	4228
no-rfc-1583 . . . . .	4229
ospf . . . . .	4230
ospf3 . . . . .	4230
overload (Protocols OSPF) . . . . .	4231
preference (Protocols OSPF) . . . . .	4232
prefix-export-limit (Protocols OSPF) . . . . .	4233
reference-bandwidth (Protocols OSPF) . . . . .	4234
rib-group (Protocols OSPF) . . . . .	4235
topology (OSPF) . . . . .	4236
traceoptions (Protocols OSPF) . . . . .	4237
traffic-engineering (OSPF) . . . . .	4240

<b>Chapter 49</b>	<b>Administration</b>	<b>4243</b>
	Routine Monitoring	4243
	Monitoring OSPF Routing Information	4243
	Operational Commands	4243
	clear (ospf   ospf3) database	4245
	clear (ospf   ospf3) database-protection	4248
	clear (ospf   ospf3) io-statistics	4249
	clear (ospf   ospf3) neighbor	4250
	clear (ospf   ospf3) statistics	4252
	clear (ospf   ospf3) overload	4254
	show (ospf   ospf3) backup coverage	4255
	show (ospf   ospf3) backup neighbor	4258
	show ospf context-identifier	4260
	show ospf database	4262
	show (ospf   ospf3) interface	4270
	show (ospf   ospf3) io-statistics	4276
	show (ospf   ospf3) log	4278
	show (ospf   ospf3) neighbor	4281
	show (ospf   ospf3) overview	4287
	show (ospf   ospf3) route	4292
	show (ospf   ospf3) statistics	4298
<b>Part 14</b>	<b>Routing Information Protocol</b>	
<b>Chapter 50</b>	<b>Overview</b>	<b>4305</b>
	RIP Overview	4305
	RIP Overview	4305
	Distance-Vector Routing Protocols	4305
	RIP Protocol Overview	4306
	RIP Packets	4307
	Maximizing Hop Count	4308
	Split Horizon and Poison Reverse Efficiency Techniques	4308
	Limitations of Unidirectional Connectivity	4309
<b>Chapter 51</b>	<b>Configuration</b>	<b>4311</b>
	RIP Configuration Tasks	4311
	Example: Configuring RIP	4311
	Understanding Basic RIP Routing	4311
	Example: Configuring a Basic RIP Network	4312
	Example: Configuring Authentication for RIP Routes	4318
	Understanding RIP Authentication	4318
	Example: Configuring Route Authentication for RIP	4318
	Enabling Authentication with Plain-Text Passwords (CLI Procedure)	4323
	Enabling Authentication with MD5 Authentication (CLI Procedure)	4323
	Example: Configuring BFD for RIP	4324
	Understanding BFD for RIP	4324
	Example: Configuring BFD for RIP	4325

Example: Configuring BFD Authentication for RIP . . . . .	4330
Understanding BFD Authentication for RIP . . . . .	4330
Example: Configuring BFD Authentication for RIP . . . . .	4332
Example: Applying Policies to RIP Routes Imported from Neighbors . . . . .	4338
Understanding RIP Import Policy . . . . .	4338
Example: Applying Policies to RIP Routes Imported from Neighbors . . . . .	4338
Examples: Controlling Traffic with Metrics in a RIP Network . . . . .	4344
Understanding Traffic Control with Metrics in a RIP Network . . . . .	4344
Example: Controlling Traffic in a RIP Network with an Incoming Metric . . . . .	4345
Example: Controlling Traffic in a RIP Network with an Outgoing Metric . . . . .	4346
Example: Configuring the Metric Value Added to Imported RIP Routes . . . . .	4348
Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets . . . . .	4352
Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets . . . . .	4352
Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets . . . . .	4352
Example: Redistributing Routes Among RIP Instances . . . . .	4356
Understanding Route Redistribution Among RIP instances . . . . .	4356
Example: Redistributing Routes Between Two RIP Instances . . . . .	4357
Example: Configuring RIP Timers . . . . .	4361
Understanding RIP Timers . . . . .	4362
Example: Configuring RIP Timers . . . . .	4362
Example: Tracing RIP Protocol Traffic . . . . .	4368
Understanding RIP Trace Operations . . . . .	4368
Example: Tracing RIP Protocol Traffic . . . . .	4369
RIP Configuration Statements . . . . .	4373
any-sender . . . . .	4374
authentication-key (Protocols RIP) . . . . .	4375
authentication-type (Protocols RIP) . . . . .	4376
bfd-liveness-detection (Protocols RIP) . . . . .	4377
check-zero . . . . .	4380
export (Protocols RIP) . . . . .	4381
group (Protocols RIP) . . . . .	4382
holddown (Protocols RIP) . . . . .	4384
import (Protocols RIP) . . . . .	4385
message-size . . . . .	4386
metric-in (Protocols RIP) . . . . .	4387
metric-out (Protocols RIP) . . . . .	4388
neighbor (Protocols RIP) . . . . .	4389
preference (Protocols RIP) . . . . .	4390
receive (Protocols RIP) . . . . .	4391
rib-group (Protocols RIP) . . . . .	4392
rip . . . . .	4392
route-timeout (Protocols RIP) . . . . .	4393
send (Protocols RIP) . . . . .	4394

	traceoptions (Protocols RIP) . . . . .	4395
	update-interval (Protocols RIP) . . . . .	4398
<b>Chapter 52</b>	<b>Administration . . . . .</b>	<b>4399</b>
	Routine Monitoring . . . . .	4399
	Monitoring RIP Routing Information . . . . .	4399
	RIP Operational Commands . . . . .	4399
	clear rip general-statistics . . . . .	4400
	clear rip statistics . . . . .	4401
	show rip general-statistics . . . . .	4402
	show rip neighbor . . . . .	4404
	show rip statistics . . . . .	4406
<b>Part 15</b>	<b>MPLS Applications</b>	
<b>Chapter 53</b>	<b>Overview . . . . .</b>	<b>4411</b>
	MPLS Overview . . . . .	4411
	MPLS Overview . . . . .	4411
	Understanding MPLS Components . . . . .	4412
	Provider Edge Switches . . . . .	4412
	Provider Switch . . . . .	4413
	Components Required for All Switches in the MPLS Network . . . . .	4413
	Understanding MPLS Label Operations . . . . .	4415
	MPLS Label-Switched Paths and MPLS Labels . . . . .	4415
	Reserved Labels . . . . .	4416
	MPLS Label Operations . . . . .	4416
	Penultimate-Hop Popping and Ultimate-Hop Popping . . . . .	4418
	Understanding CoS MPLS EXP Classifiers and Rewrite Rules . . . . .	4419
	EXP Classifiers . . . . .	4419
	EXP Rewrite Rules . . . . .	4420
	Schedulers . . . . .	4421
	Understanding Using MPLS-Based Layer 3 VPNs . . . . .	4422
	MPLS-Based Layer 3 VPNs . . . . .	4422
	MPLS Features . . . . .	4423
	MPLS Feature Support on the QFX Series and EX4600 Switch Overview . . . . .	4423
	Supported MPLS Features . . . . .	4423
	Unsupported MPLS Features . . . . .	4425
	Supported MPLS Scaling Values . . . . .	4426
	Introduction to LDP for QFX5100 . . . . .	4426
	LDP Introduction . . . . .	4427
	Junos OS LDP Protocol Implementation . . . . .	4427
	LDP Operation . . . . .	4427
	Tunneling LDP LSPs in RSVP LSPs . . . . .	4428
	Tunneling LDP LSPs in RSVP LSPs Overview . . . . .	4428
	Label Operations . . . . .	4428
	LDP Message Types . . . . .	4430
	Discovery Messages . . . . .	4430
	Session Messages . . . . .	4430
	Advertisement Messages . . . . .	4430
	Notification Messages . . . . .	4431

	LDP Session Protection . . . . .	4431
	LDP Graceful Restart . . . . .	4431
<b>Chapter 54</b>	<b>Configuration . . . . .</b>	<b>4433</b>
	Configuration Guidelines . . . . .	4433
	MPLS Configuration Guidelines . . . . .	4433
	LDP Configuration Guidelines for QFX5100 . . . . .	4434
	Minimum LDP Configuration . . . . .	4434
	Enabling and Disabling LDP . . . . .	4435
	Enabling Strict Targeted Hello Messages for LDP . . . . .	4435
	Filtering Inbound LDP Label Bindings . . . . .	4435
	Examples: Filtering Inbound LDP Label Bindings . . . . .	4437
	Filtering Outbound LDP Label Bindings . . . . .	4437
	Examples: Filtering Outbound LDP Label Bindings . . . . .	4438
	Specifying the Transport Address Used by LDP . . . . .	4439
	Collecting LDP Statistics . . . . .	4440
	LDP Statistics Output . . . . .	4440
	Disabling LDP Statistics on the Penultimate-Hop Router . . . . .	4441
	LDP Statistics Limitations . . . . .	4442
	Tracing LDP Protocol Traffic . . . . .	4442
	Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels . . . . .	4442
	Tracing LDP Protocol Traffic Within FECs . . . . .	4443
	Examples: Tracing LDP Protocol Traffic . . . . .	4444
	Configuration Examples . . . . .	4445
	Example: Configuring MPLS-Based Layer 3 VPNs . . . . .	4445
	Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks . . . . .	4454
	Example: Configuring LDP Downstream on Demand . . . . .	4462
	Configuration Tasks . . . . .	4467
	Configuring MPLS on Provider Edge Switches . . . . .	4468
	Configuring the Ingress PE Switch . . . . .	4468
	Configuring the Egress PE Switch . . . . .	4470
	Configuring MPLS on Provider Switches . . . . .	4471
	Configuring Static Label Switched Paths for MPLS . . . . .	4472
	Configuring the Ingress PE Switch . . . . .	4473
	Configuring the Provider and the Egress PE Switch . . . . .	4474
	Configuring MPLS Firewall Filters and Policers . . . . .	4474
	Configuring an MPLS Firewall Filter . . . . .	4476
	Applying an MPLS Firewall Filter to an MPLS Interface . . . . .	4476
	Configuring Policers for LSPs . . . . .	4477
	Configuring CoS Bits for an MPLS Network . . . . .	4478
	Configuring a Global MPLS EXP Classifier . . . . .	4479
	Configuring Rewrite Rules for MPLS EXP Classifiers . . . . .	4480
	Configuring MPLS to Gather Statistics . . . . .	4481
	Configuring Automatic Bandwidth Allocation for LSPs . . . . .	4482
	Configuring Automatic Bandwidth Allocation on LSPs . . . . .	4483
	Requesting Automatic Bandwidth Allocation Adjustment . . . . .	4488
	Configuring Reporting of Automatic Bandwidth Allocation Statistics . . . . .	4489

Configuring the LDP Timer for Hello Messages . . . . .	4492
Configuring the LDP Timer for Link Hello Messages . . . . .	4493
Configuring the LDP Timer for Targeted Hello Messages . . . . .	4493
Configuring the Delay Before LDP Neighbors Are Considered Down . . . . .	4493
Configuring the LDP Hold Time for Link Hello Messages . . . . .	4494
Configuring the LDP Hold Time for Targeted Hello Messages . . . . .	4494
Configuring the Interval for LDP Keepalive Messages . . . . .	4494
Configuring the LDP Keepalive Timeout . . . . .	4495
Configuring LDP Route Preferences . . . . .	4495
Configuring LDP Graceful Restart . . . . .	4495
Enabling Graceful Restart . . . . .	4496
Disabling LDP Graceful Restart or Helper Mode . . . . .	4496
Configuring Reconnect Time . . . . .	4497
Configuring Recovery Time and Maximum Recovery Time . . . . .	4497
Configuring the Prefixes Advertised into LDP from the Routing Table . . . . .	4498
Example: Configuring the Prefixes Advertised into LDP . . . . .	4498
Configuring LDP LSP Traceroute . . . . .	4498
Configuring Miscellaneous LDP Properties . . . . .	4500
Configuring LDP to Use the IGP Route Metric . . . . .	4500
Preventing Addition of Ingress Routes to the inet.0 Routing Table . . . . .	4500
Multiple-Instance LDP and Carrier-of-Carriers VPNs . . . . .	4501
Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router . . . . .	4501
Enabling LDP over RSVP-Established LSPs . . . . .	4501
Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks . . . . .	4502
Configuring the TCP MD5 Signature for LDP Sessions . . . . .	4502
Configuring LDP Session Protection . . . . .	4503
Disabling SNMP Traps for LDP . . . . .	4504
Configuring LDP Synchronization with the IGP on LDP Links . . . . .	4504
Configuring LDP Synchronization with the IGP on the Router . . . . .	4505
Configuring the Label Withdrawal Timer . . . . .	4505
Ignoring the LDP Subnet Check . . . . .	4505
Configuring Ethernet over MPLS (L2 Circuit) . . . . .	4506
Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire) . . . . .	4506
Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire) . . . . .	4507
Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit . . . . .	4507
Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit . . . . .	4508
Configuration Statements . . . . .	4508
[edit protocols mpls] Hierarchy Level . . . . .	4509
[edit protocols rsdp] Hierarchy Level . . . . .	4513
auto-bandwidth (MPLS Statistics) . . . . .	4515
auto-bandwidth . . . . .	4516
adjust-interval . . . . .	4517
adjust-threshold . . . . .	4517
adjust-threshold-overflow-limit . . . . .	4518
adjust-threshold-underflow-limit . . . . .	4518



exp	4519
maximum-bandwidth (Protocols MPLS)	4520
minimum-bandwidth	4520
minimum-bandwidth-adjust-interval	4521
minimum-bandwidth-adjust-threshold-change	4521
minimum-bandwidth-adjust-threshold-value	4522
monitor-bandwidth	4522
system-defaults	4523
LDP Configuration Statements for QFX5100	4523
allow-subnet-mismatch	4525
authentication-algorithm	4526
authentication-key (Protocols LDP)	4527
authentication-key-chain (Protocols LDP)	4528
deaggregate	4529
disable (Protocols LDP)	4530
dod-request-policy	4531
downstream-on-demand	4531
egress-policy	4532
explicit-null (Protocols LDP)	4532
export (Protocols LDP)	4533
fec	4534
graceful-restart (Protocols LDP)	4535
hello-interval (Protocols LDP)	4536
helper-disable (LDP)	4537
hold-time (Protocols LDP)	4538
ignore-lsp-metrics	4539
igp-synchronization	4539
import (Protocols LDP)	4540
interface (Protocols LDP)	4541
keepalive-interval	4542
keepalive-timeout	4543
l2-smart-policy	4543
label-withdrawal-delay	4544
ldp	4545
ldp-synchronization	4548
log-updown (Protocols LDP)	4549
maximum-neighbor-recovery-time	4550
no-forwarding	4551
policing (Protocols LDP)	4552
preference (Protocols LDP)	4553
reconnect-time	4554
recovery-time	4554
session (ldp)	4555
session-protection	4556
strict-targeted-hellos	4556
targeted-hello	4557
traceoptions (Protocols LDP)	4558
track-igp-metric	4560
traffic-statistics (Protocols LDP)	4561

	transport-address .....	4563
<b>Chapter 55</b>	<b>Administration .....</b>	<b>4565</b>
	Routine Monitoring .....	4565
	Verifying That MPLS Is Working Correctly .....	4565
	Verifying the Physical Layer on the Switches .....	4565
	Verifying the Routing Protocol .....	4566
	Verifying the Core Interfaces Being Used for the MPLS Traffic .....	4566
	Verifying RSVP .....	4566
	Operational Mode Commands .....	4567
	clear ldp neighbor .....	4569
	clear ldp session .....	4570
	clear ldp statistics .....	4571
	clear mpls lsp .....	4572
	clear rsvp session .....	4574
	clear rsvp statistics .....	4576
	monitor label-switched-path .....	4577
	ping mpls bgp .....	4580
	ping mpls l2circuit .....	4582
	ping mpls l3vpn .....	4585
	ping mpls ldp .....	4588
	ping mpls lsp-end-point .....	4591
	ping mpls rsvp .....	4593
	request mpls lsp adjust-autobandwidth .....	4598
	show ldp database .....	4600
	show ldp fec-filters .....	4608
	show ldp interface .....	4609
	show ldp neighbor .....	4611
	show ldp path .....	4613
	show ldp route .....	4615
	show ldp session .....	4619
	show ldp statistics .....	4625
	show ldp traffic-statistics .....	4629
	show security keychain .....	4633
	show link-management .....	4636
	show link-management peer .....	4640
	show link-management routing .....	4642
	show link-management statistics .....	4645
	show link-management te-link .....	4647
	show mpls call-admission-control .....	4649
	show mpls cspf .....	4651
	show mpls diffserv-te .....	4653
	show route forwarding-table .....	4655
	show mpls interface .....	4663
	show mpls lsp .....	4665
	show mpls lsp autobandwidth .....	4681
	show mpls path .....	4684
	show mpls static-lsp .....	4685
	show rsvp interface .....	4688

	show rsvp neighbor . . . . .	4693
	show rsvp session . . . . .	4698
	show rsvp statistics . . . . .	4707
	show rsvp version . . . . .	4711
	show ted database . . . . .	4714
	show ted link . . . . .	4721
	show ted protocol . . . . .	4724
	traceroute mpls ldp . . . . .	4726
	traceroute mpls rsvp . . . . .	4729
<b>Chapter 56</b>	<b>Troubleshooting . . . . .</b>	<b>4733</b>
	Troubleshooting Procedures . . . . .	4733
	Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch . . . . .	4733
<b>Part 16</b>	<b>Multicast</b>	
<b>Chapter 57</b>	<b>Overview . . . . .</b>	<b>4737</b>
	Introduction to PIM Basics . . . . .	4737
	PIM Overview . . . . .	4737
	Basic PIM Network Components . . . . .	4740
	PIM on Aggregated Interfaces . . . . .	4740
	Introduction to PIM Sparse Mode . . . . .	4741
	Understanding PIM Sparse Mode . . . . .	4741
	Rendezvous Point . . . . .	4743
	RP Mapping Options . . . . .	4743
	Designated Router . . . . .	4744
	Introduction to Static RP . . . . .	4744
	Understanding Static RP . . . . .	4744
	Introduction to Anycast RP . . . . .	4745
	Understanding RP Mapping with Anycast RP . . . . .	4745
	Introduction to PIM Bootstrap Router . . . . .	4745
	Understanding the PIM Bootstrap Router . . . . .	4745
	Introduction to PIM Filtering . . . . .	4746
	Understanding Multicast Message Filters . . . . .	4746
	Filtering MAC Addresses . . . . .	4747
	Filtering RP and DR Register Messages . . . . .	4747
	Introduction to PIM RPT and SPT Cutover . . . . .	4748
	Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees . . . . .	4748
	Building an RPT Between the RP and Receivers . . . . .	4749
	PIM Sparse Mode Source Registration . . . . .	4750
	Multicast Shortest-Path Tree . . . . .	4753
	SPT Cutover . . . . .	4754
	SPT Cutover Control . . . . .	4757
	Introduction to IGMP . . . . .	4757
	Understanding Group Membership Protocols . . . . .	4757
	Understanding IGMP . . . . .	4759

	Introduction to IGMP Snooping . . . . .	4761
	IGMP Snooping Overview . . . . .	4761
	How IGMP Snooping Works . . . . .	4761
	How IGMP Snooping Works with Routed VLAN Interfaces . . . . .	4762
	How Hosts Join and Leave Multicast Groups . . . . .	4762
	IGMP Snooping and Forwarding Interfaces . . . . .	4762
	General Forwarding Rules . . . . .	4763
	Using a Switch as an IGMP Querier . . . . .	4763
	Introduction to MLD . . . . .	4764
	Understanding MLD . . . . .	4764
	Introduction to MSDP . . . . .	4767
	Understanding MSDP . . . . .	4767
	Filtering MSDP SA Messages . . . . .	4769
	Introduction to Source-Specific Multicast . . . . .	4769
	Source-Specific Multicast Groups Overview . . . . .	4769
	Understanding PIM Source-Specific Mode . . . . .	4770
	PIM SSM . . . . .	4771
	Introduction to Multicast VLAN Registration . . . . .	4773
	Understanding Multicast VLAN Registration . . . . .	4773
	How MVR Works . . . . .	4773
<b>Chapter 58</b>	<b>Configuration . . . . .</b>	<b>4777</b>
	PIM Basics . . . . .	4777
	Changing the PIM Version . . . . .	4778
	Modifying the PIM Hello Interval . . . . .	4778
	Preserving Multicast Performance by Disabling Response to the ping Utility . . . . .	4779
	Configuring PIM Trace Options . . . . .	4780
	Disabling PIM . . . . .	4782
	Disabling the PIM Protocol . . . . .	4782
	Disabling PIM On an Interface . . . . .	4783
	Disabling PIM for a Family . . . . .	4783
	Disabling PIM for a Rendezvous Point . . . . .	4784
	PIM Designated Router . . . . .	4784
	Configuring Interface Priority for PIM Designated Router Selection . . . . .	4784
	Configuring PIM Designated Router Election on Point-to-Point Links . . . . .	4785
	PIM Sparse Mode . . . . .	4786
	Enabling PIM Sparse Mode . . . . .	4786
	Configuring PIM Join Load Balancing . . . . .	4787
	Modifying the Join State Timeout . . . . .	4790
	Example: Enabling Join Suppression . . . . .	4791
	Static RP . . . . .	4795
	Configuring Local PIM RPs . . . . .	4795
	Configuring the Static PIM RP Address on the Non-RP Routing Device . . . . .	4797
	Anycast RP . . . . .	4798
	Example: Configuring PIM Anycast With or Without MSDP . . . . .	4799
	Configuring a PIM Anycast RP Router with MSDP . . . . .	4802
	Configuring a PIM Anycast RP Router Using Only PIM . . . . .	4803

Configuring All PIM Anycast Non-RP Routers . . . . .	4804
Example: Configuring Multiple RPs in a Domain with Anycast RP . . . . .	4805
PIM Bootstrap Router . . . . .	4807
Configuring PIM Bootstrap Properties for IPv4 or IPv6 . . . . .	4807
Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain . . . . .	4809
Example: Configuring PIM BSR Filters . . . . .	4809
PIM Filtering . . . . .	4810
Configuring Interface-Level PIM Neighbor Policies . . . . .	4810
Filtering Outgoing PIM Join Messages . . . . .	4811
Filtering Incoming PIM Join Messages . . . . .	4812
Configuring Register Message Filters on a PIM RP and DR . . . . .	4813
PIM RPT and SPT Cutover . . . . .	4815
Example: Configuring the PIM Assert Timeout . . . . .	4815
Example: Configuring the PIM SPT Threshold Policy . . . . .	4818
PIM and the BFD Protocol . . . . .	4821
Configuring BFD for PIM . . . . .	4821
Configuring BFD Authentication for PIM . . . . .	4823
Configuring BFD Authentication Parameters . . . . .	4823
Viewing Authentication Information for BFD Sessions . . . . .	4824
IGMP . . . . .	4826
Configuring IGMP . . . . .	4826
Enabling IGMP . . . . .	4828
Changing the IGMP Version . . . . .	4829
Modifying the IGMP Host-Query Message Interval . . . . .	4830
Modifying the IGMP Last-Member Query Interval . . . . .	4831
Specifying Immediate-Leave Host Removal for IGMP . . . . .	4831
Filtering Unwanted IGMP Reports at the IGMP Interface Level . . . . .	4832
Accepting IGMP Messages from Remote Subnetworks . . . . .	4833
Modifying the IGMP Query Response Interval . . . . .	4834
Modifying the IGMP Robustness Variable . . . . .	4835
Limiting the Maximum IGMP Message Rate . . . . .	4836
Enabling IGMP Static Group Membership . . . . .	4836
Recording IGMP Join and Leave Events . . . . .	4843
Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces . . . . .	4844
Tracing IGMP Protocol Traffic . . . . .	4845
Disabling IGMP . . . . .	4847
IGMP Snooping . . . . .	4847
Configuring IGMP Snooping . . . . .	4848
Configuring VLAN-Specific IGMP Snooping Parameters . . . . .	4849
Example: Configuring IGMP Snooping . . . . .	4849
Using a Switch as an IGMP Querier . . . . .	4851
MLD . . . . .	4852
Examples: Configuring MLD . . . . .	4852
Understanding MLD . . . . .	4853
Configuring MLD . . . . .	4855
Enabling MLD . . . . .	4856
Modifying the MLD Version . . . . .	4857

Modifying the MLD Host-Query Message Interval . . . . .	4858
Modifying the MLD Query Response Interval . . . . .	4858
Modifying the MLD Last-Member Query Interval . . . . .	4859
Specifying Immediate-Leave Host Removal for MLD . . . . .	4860
Filtering Unwanted MLD Reports at the MLD Interface Level . . . . .	4861
Example: Modifying the MLD Robustness Variable . . . . .	4861
Limiting the Maximum MLD Message Rate . . . . .	4863
Enabling MLD Static Group Membership . . . . .	4863
Example: Recording MLD Join and Leave Events . . . . .	4870
Configuring the Number of MLD Multicast Group Joins on Logical Interfaces . . . . .	4872
Tracing MLD Protocol Traffic . . . . .	4873
Disabling MLD . . . . .	4875
MSDP . . . . .	4875
Configuring MSDP . . . . .	4876
Tracing MSDP Protocol Traffic . . . . .	4877
Configuring the Interface to Accept Traffic from a Remote Source . . . . .	4879
Example: Configuring MSDP . . . . .	4879
Example: Configuring MSDP with Active Source Limits and Mesh Groups . . . . .	4880
Example: Configuring PIM Anycast With or Without MSDP . . . . .	4886
Configuring a PIM Anycast RP Router with MSDP . . . . .	4890
Source-Specific Multicast . . . . .	4891
Example: Configuring PIM SSM on a Network . . . . .	4891
Example: Configuring an SSM-Only Domain . . . . .	4892
Example: Configuring SSM Mapping . . . . .	4893
Example: Configuring Source-Specific Multicast Groups with Any-Source Override . . . . .	4895
Example: Configuring SSM Maps for Different Groups to Different Sources . . . . .	4899
Multiple SSM Maps and Groups for Interfaces . . . . .	4899
Example: Configuring Multiple SSM Maps Per Interface . . . . .	4899
PIM Configuration Statements . . . . .	4902
address (Anycast RPs) . . . . .	4905
address (Local RPs) . . . . .	4905
address (Static RPs) . . . . .	4906
algorithm . . . . .	4907
anycast-pim . . . . .	4908
assert-timeout . . . . .	4909
authentication (Protocols PIM) . . . . .	4910
bfd-liveness-detection (Protocols PIM) . . . . .	4911
bootstrap . . . . .	4912
bootstrap-export . . . . .	4913
bootstrap-import . . . . .	4914
bootstrap-priority . . . . .	4915
detection-time (BFD for PIM) . . . . .	4916
disable (PIM) . . . . .	4917
dr-election-on-p2p . . . . .	4918
dr-register-policy . . . . .	4918
embedded-rp . . . . .	4919

export (Protocols PIM Bootstrap) .....	4920
export (Protocols PIM) .....	4920
family (Bootstrap) .....	4921
family (Protocols PIM) .....	4922
family (Local RP) .....	4923
group (RPF Selection) .....	4924
group-ranges .....	4925
hello-interval (Protocols PIM) .....	4926
hold-time (Protocols PIM) .....	4927
import (Protocols PIM Bootstrap) .....	4928
import (Protocols PIM) .....	4929
infinity .....	4930
interface .....	4931
join-load-balance .....	4932
join-prune-timeout .....	4933
key-chain (Protocols PIM) .....	4934
local .....	4935
local-address (Protocols PIM) .....	4936
loose-check .....	4937
maximum-rps .....	4938
minimum-interval (PIM BFD Liveness Detection) .....	4939
minimum-interval (PIM BFD Transmit Interval) .....	4940
minimum-receive-interval .....	4941
mode (Protocols PIM) .....	4942
multiplier .....	4942
neighbor-policy .....	4943
next-hop (PIM RPF Selection) .....	4943
no-adaptation (PIM BFD Liveness Detection) .....	4944
override-interval .....	4945
pim .....	4946
prefix-list (PIM RPF Selection) .....	4949
priority (Bootstrap) .....	4950
priority (PIM Interfaces) .....	4951
priority (PIM RPs) .....	4952
propagation-delay .....	4953
reset-tracking-bit .....	4954
rib-group (Protocols PIM) .....	4955
rp .....	4956
rp-register-policy .....	4958
rp-set .....	4959
rpf-selection .....	4960
source (PIM RPF Selection) .....	4961
spt-threshold .....	4962
static (Protocols PIM) .....	4963
threshold (PIM BFD Detection Time) .....	4964
threshold (PIM BFD Transmit Interval) .....	4965
transmit-interval (PIM BFD Liveness Detection) .....	4966
traceoptions (Protocols PIM) .....	4967
version (BFD) .....	4970

version (PIM) .....	4971
wildcard-source (PIM RPF Selection) .....	4972
IGMP Configuration Statements .....	4972
accounting (Protocols IGMP) .....	4973
accounting (Protocols IGMP Interface) .....	4973
asm-override-ssm .....	4974
disable (Protocols IGMP) .....	4974
exclude (Protocols IGMP) .....	4975
group (Protocols IGMP) .....	4976
group-count (Protocols IGMP) .....	4977
group-increment (Protocols IGMP) .....	4977
group-limit (IGMP) .....	4978
group-policy (Protocols IGMP) .....	4979
igmp .....	4980
immediate-leave (Protocols IGMP) .....	4982
interface (Protocols IGMP) .....	4983
maximum-transmit-rate (Protocols IGMP) .....	4984
oif-map (IGMP Interface) .....	4984
passive (IGMP) .....	4985
promiscuous-mode (Protocols IGMP) .....	4986
query-interval (Protocols IGMP) .....	4986
query-last-member-interval (Protocols IGMP) .....	4987
query-response-interval (Protocols IGMP) .....	4988
robust-count (Protocols IGMP) .....	4989
source (Protocols IGMP) .....	4990
source-count (Protocols IGMP) .....	4991
source-increment (Protocols IGMP) .....	4991
static (Protocols IGMP) .....	4992
traceoptions (Protocols IGMP) .....	4993
version (Protocols IGMP) .....	4995
IGMP Snooping Configuration Statements .....	4995
data-forwarding .....	4997
disable (IGMP Snooping) .....	4997
group (IGMP Snooping) .....	4998
group-limit (IGMP and MLD Snooping) .....	4999
groups (Multicast VLAN Registration) .....	5000
host-only-interface .....	5001
igmp-querier .....	5002
igmp-snooping .....	5003
immediate-leave (Bridge Domains) .....	5004
install (Multicast VLAN Registration) .....	5005
interface (Bridge Domains) .....	5006
interface (IGMP Snooping) .....	5007
l2-querier .....	5007
multicast-router-interface (IGMP Snooping) .....	5008
proxy (Multicast VLAN Registration) .....	5008
query-interval (Bridge Domains) .....	5009
query-last-member-interval (Bridge Domains) .....	5010
query-response-interval (Bridge Domains) .....	5011



receiver	5012
robust-count (IGMP Snooping)	5012
source (Multicast VLAN Registration)	5013
source-address	5013
source-address (IGMP Querier)	5014
source-vlans	5014
static (IGMP Snooping)	5015
traceoptions (IGMP Snooping)	5016
version (IGMP Snooping)	5018
vlan (IGMP Snooping)	5019
MSDP Configuration Statements	5019
active-source-limit	5021
authentication-key	5022
data-encapsulation	5023
default-peer	5024
disable (Protocols MSDP)	5025
export (Protocols MSDP)	5026
group (Protocols MSDP)	5027
import (Protocols MSDP)	5028
local-address (Protocols MSDP)	5029
maximum (MSDP Active Source Messages)	5030
mode (Protocols MSDP)	5031
msdp	5032
peer (Protocols MSDP)	5034
rib-group (Protocols MSDP)	5035
source (Protocols MSDP)	5036
threshold (MSDP Active Source Messages)	5037
traceoptions (Protocols MSDP)	5038
Source-Specific Multicast Configuration Statements	5040
asm-override-ssm	5041
policy (SSM Maps)	5042
ssm-groups	5043
ssm-map (Protocols IGMP)	5044
ssm-map (Routing Options Multicast)	5044
ssm-map-policy (IGMP)	5045
<b>Chapter 59 Administration</b>	<b>5047</b>
Routine Monitoring	5047
Monitoring IGMP Snooping	5047
Verifying the IGMP Snooping Group Timeout Value	5048
Monitoring Commands for Multicast Protocols	5048
clear igmp membership	5051
clear igmp-snooping membership	5054
clear igmp statistics	5055
clear igmp-snooping statistics	5057
clear msdp cache	5058
clear msdp statistics	5059
clear multicast bandwidth-admission	5060
clear multicast scope	5062

clear multicast sessions . . . . .	5063
clear multicast statistics . . . . .	5064
clear pim join . . . . .	5065
clear pim register . . . . .	5067
clear pim statistics . . . . .	5069
mtrace . . . . .	5072
mtrace from-source . . . . .	5075
mtrace monitor . . . . .	5078
mtrace to-gateway . . . . .	5080
show configuration protocols igmp . . . . .	5083
show igmp group . . . . .	5085
show igmp interface . . . . .	5089
show igmp statistics . . . . .	5093
show igmp-snooping membership . . . . .	5096
show igmp-snooping route . . . . .	5099
show igmp-snooping statistics . . . . .	5101
show igmp-snooping vlans . . . . .	5103
show msdp . . . . .	5105
show msdp source . . . . .	5107
show msdp source-active . . . . .	5109
show msdp statistics . . . . .	5112
show multicast flow-map . . . . .	5116
show multicast interface . . . . .	5118
show multicast minfo . . . . .	5120
show multicast next-hops . . . . .	5122
show multicast pim-to-igmp-proxy . . . . .	5125
show multicast pim-to-mld-proxy . . . . .	5127
show multicast route . . . . .	5129
show multicast rpf . . . . .	5136
show multicast scope . . . . .	5140
show multicast sessions . . . . .	5142
show multicast usage . . . . .	5145
show pim bootstrap . . . . .	5148
show pim interfaces . . . . .	5150
show pim join . . . . .	5153
show pim neighbors . . . . .	5175
show pim rps . . . . .	5179
show pim source . . . . .	5186
show pim statistics . . . . .	5189
show system statistics igmp . . . . .	5202
test msdp . . . . .	5206

## Part 17

### Chapter 60

## Security

Overview . . . . .	5209
Firewall Filters . . . . .	5209
Overview of Firewall Filters . . . . .	5209
Firewall Filter Types . . . . .	5210
Firewall Filter Components . . . . .	5211

Firewall Filter Processing . . . . .	5211
Understanding Filter-Based Forwarding . . . . .	5212
Understanding How Firewall Filters Are Evaluated . . . . .	5212
Understanding How Firewall Filters Control Packet Flows . . . . .	5214
Understanding Firewall Filter Match Conditions . . . . .	5215
Filter Match Conditions . . . . .	5215
Numeric Filter Match Conditions . . . . .	5216
Interface Filter Match Conditions . . . . .	5216
IP Address Filter Match Conditions . . . . .	5217
MAC Address Filter Match Conditions . . . . .	5217
Bit-Field Filter Match Conditions . . . . .	5218
Firewall Filter Match Conditions and Actions . . . . .	5219
Understanding How a Firewall Filter Tests a Protocol . . . . .	5233
Understanding Firewall Filter Planning . . . . .	5234
Planning the Number of Firewall Filters to Create . . . . .	5236
Understanding How Many Firewall Filters Are Supported . . . . .	5236
Egress Filters . . . . .	5237
Avoid Configuring too Many Filters . . . . .	5237
Policers can Limit Egress Filters . . . . .	5238
Planning for Filter-Specific Policers . . . . .	5239
Planning for Filter-Based Forwarding . . . . .	5239
Understanding Firewall Filter Processing Points for Bridged and Routed Packets . . . . .	5239
Applying Firewall Filters to Interfaces . . . . .	5240
Policers . . . . .	5241
Overview of Policers . . . . .	5241
Policer Overview . . . . .	5241
Policer Types . . . . .	5242
Policer Actions . . . . .	5243
Policer Colors . . . . .	5244
Filter-Specific Policers . . . . .	5244
Suggested Naming Convention for Policers . . . . .	5244
Policer Counters . . . . .	5245
Policer Algorithms . . . . .	5245
How Many Policers are Supported? . . . . .	5245
Policers can Limit Egress Firewall Filters . . . . .	5245
Understanding Policers with Link Aggregation Groups . . . . .	5246
Understanding Color-Blind Mode for Single-Rate Tricolor Marking . . . . .	5247
Understanding Color-Aware Mode for Single-Rate Tricolor Marking . . . . .	5247
Summary of PLP Changes . . . . .	5247
Understanding Color-Blind Mode for Two-Rate Tricolor Marking . . . . .	5249
Understanding Color-Aware Mode for Two-Rate Tricolor Marking . . . . .	5249
Summary of PLP Changes . . . . .	5249
Effect on Green Packets (Low PLP) . . . . .	5250
Effect on Yellow Packets (Medium PLP) . . . . .	5250
Effect on Red Packets (High PLP) . . . . .	5251

Port Security . . . . .	5251
Overview of Access Port Protection . . . . .	5251
Mitigation of Ethernet Switching Table Overflow Attacks . . . . .	5252
Mitigation of Rogue DHCP Server Attacks . . . . .	5252
Protection Against ARP Spoofing Attacks . . . . .	5253
Protection Against DHCP Snooping Database Alteration Attacks . . . . .	5253
Protection Against DHCP Starvation Attacks . . . . .	5253
Understanding Port Security . . . . .	5254
Understanding DHCP Snooping for Port Security . . . . .	5256
DHCP Snooping Basics . . . . .	5256
DHCP Snooping Process . . . . .	5257
DHCPv6 Snooping . . . . .	5258
Rapid Commit for DHCPv6 . . . . .	5259
DHCP Server Access . . . . .	5259
Static IP Address Additions to the DHCP Snooping Database . . . . .	5262
Snooping DHCP Packets That Have Invalid IP Addresses . . . . .	5262
Prioritizing Snooped Packets . . . . .	5263
Understanding DAI for Port Security . . . . .	5263
Address Resolution Protocol . . . . .	5263
ARP Spoofing . . . . .	5264
Dynamic ARP Inspection . . . . .	5264
Prioritizing Inspected Packets . . . . .	5265
Understanding MAC Limiting and MAC Move Limiting for Port Security . . . . .	5266
MAC Limiting . . . . .	5266
MAC Move Limiting . . . . .	5266
Actions for MAC Limiting . . . . .	5267
MAC Addresses That Exceed the MAC Limit or MAC Move Limit . . . . .	5267
Understanding Trusted and Untrusted Ports . . . . .	5268
Understanding Trusted DHCP Servers for Port Security . . . . .	5268
Understanding DHCP Option 82 for Port Security . . . . .	5269
DHCP Option 82 Processing . . . . .	5269
Suboption Components of Option 82 . . . . .	5270
Configurations That Support Option 82 . . . . .	5270
Understanding Static ARP Entries . . . . .	5271
Device Security . . . . .	5272
Understanding Storm Control . . . . .	5272
Understanding Unicast RPF . . . . .	5273
Unicast RPF for Switches Overview . . . . .	5274
Unicast RPF Implementation . . . . .	5274
When to Enable Unicast RPF . . . . .	5275
When Not to Enable Unicast RPF . . . . .	5276
Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches . . . . .	5276
Understanding Unknown Unicast Forwarding . . . . .	5277

<b>Chapter 61</b>	<b>Configuration</b>	<b>5279</b>
	Firewall and Policer Configuration Examples	5279
	Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device	5279
	Example: Using Two-Color Policers and Prefix Lists	5283
	Example: Using Policers to Manage Oversubscription	5286
	Device Security Configuration Example	5288
	Example: Configuring Storm Control to Prevent Network Outages	5288
	Firewall and Policer Configuration Tasks	5290
	Configuring Firewall Filters	5290
	Configuring a Firewall Filter	5290
	Applying a Firewall Filter to a Port	5292
	Applying a Firewall Filter to a VLAN	5292
	Applying a Firewall Filter to a Layer 3 (Routed) Interface	5293
	Applying Firewall Filters to Interfaces	5293
	Assigning Forwarding Classes and Loss Priority	5294
	Configuring Color-Blind Egress Policers for Medium-Low PLP	5296
	Configuring Two-Color and Three-Color Policers to Control Traffic Rates	5296
	Configuring Two-Color Policers	5297
	Configuring Three-Color Policers	5297
	Specifying Policers in a Firewall Filter Configuration	5298
	Applying a Firewall Filter That Includes a Policer	5298
	Configuring MPLS Firewall Filters and Policers	5299
	Configuring MPLS Firewall Filters	5299
	Examples: Configuring MPLS Firewall Filters	5300
	Configuring Policers for LSPs	5300
	Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch	5301
	Configuring a Filter to De-encapsulate GRE Traffic	5301
	Applying the Filter to an Interface	5302
	Device Security Configuration Tasks	5303
	Configuring Unicast RPF (CLI Procedure)	5303
	Disabling Unicast RPF (CLI Procedure)	5304
	Configuring Unknown Unicast Forwarding (CLI Procedure)	5305
	Configuration Statements for Firewall Filters	5306
	family	5307
	filter	5308
	filter (Layer 2 and Layer 3 Interfaces)	5309
	filter (VLANs)	5310
	firewall	5311
	from	5312
	interface-specific	5313
	term	5313
	then (Filters)	5314
	Configuration Statements for Policers	5314
	action	5315
	bandwidth-limit	5315
	burst-size-limit	5316

color-aware .....	5317
color-blind .....	5318
committed-burst-size .....	5319
committed-information-rate .....	5320
excess-burst-size .....	5321
filter-specific .....	5322
firewall .....	5323
if-exceeding .....	5324
loss-priority high then discard (Three-Color Policer) .....	5325
peak-burst-size .....	5326
peak-information-rate .....	5327
policer .....	5328
single-rate .....	5329
then (Policers) .....	5330
three-color-policer .....	5331
two-rate .....	5332
Configuration Statements for Port Security .....	5332
circuit-id .....	5333
dhcp-snooping-file .....	5334
fc-map .....	5335
fcoe-trusted .....	5337
mac-move-limit .....	5338
no-allowed-mac-log .....	5339
no-gratuitous-arp-request .....	5340
persistent-learning .....	5340
port-error-disable .....	5341
vendor-id .....	5343
write-interval .....	5344
Configuration Statements for Port Security .....	5344
accept-source-mac .....	5346
arp-inspection .....	5348
dhcp-security .....	5350
dhcp-service .....	5352
group (DHCP Security) .....	5353
interface (DHCP Security) .....	5354
interface-mac-limit .....	5355
no-dhcp-snooping .....	5357
no-option-82 .....	5358
option-82 .....	5359
overrides (DHCP Security) .....	5360
recovery-timeout .....	5361
static-ip .....	5362
switch-options .....	5363
trusted .....	5364
untrusted .....	5364
Configuration Statements for Device Security .....	5364
action-shutdown .....	5365
bandwidth-level .....	5366
bandwidth-percentage .....	5367

	interface (Unknown Unicast Forwarding) . . . . .	5368
	no-broadcast . . . . .	5369
	no-multicast . . . . .	5370
	no-registered-multicast . . . . .	5371
	no-unknown-unicast . . . . .	5372
	no-unregistered-multicast . . . . .	5373
	rpf-check . . . . .	5374
	storm-control . . . . .	5375
	storm-control-profiles . . . . .	5376
	unknown-unicast-forwarding . . . . .	5377
<b>Chapter 62</b>	<b>Administration . . . . .</b>	<b>5379</b>
	Routine Monitoring . . . . .	5379
	Monitoring Firewall Filter Traffic . . . . .	5379
	Monitoring Traffic for All Firewall Filters and Policers That Are Configured . . . . .	5379
	Monitoring Traffic for a Specific Firewall Filter . . . . .	5380
	Monitoring Traffic for a Specific Policer . . . . .	5380
	Monitoring Port Security . . . . .	5381
	Verifying That Firewall Filters Are Operational . . . . .	5382
	Verifying That DAI Is Working Correctly . . . . .	5383
	Verifying That DHCP Snooping Is Working Correctly . . . . .	5384
	Verifying That MAC Limiting Is Working Correctly . . . . .	5385
	Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly . . . . .	5385
	Verifying That Allowed MAC Addresses Are Working Correctly . . . . .	5386
	Verifying That Interfaces Are Shut Down . . . . .	5386
	Customizing the Ethernet Switching Table Display to View Information for a Specific Interface . . . . .	5387
	Verifying That MAC Move Limiting Is Working Correctly . . . . .	5388
	Verifying That the Port Error Disable Setting Is Working Correctly . . . . .	5388
	Verifying Unicast RPF Status . . . . .	5389
	Verifying That a Trusted DHCP Server Is Working Correctly . . . . .	5392
	Verifying That Three-Color Policers Are Operational . . . . .	5392
	Verifying That Two-Color Policers Are Operational . . . . .	5393
	Monitoring Commands . . . . .	5393
	clear arp inspection statistics . . . . .	5395
	clear dhcp snooping binding . . . . .	5396
	clear ethernet-switching port-error . . . . .	5397
	clear firewall . . . . .	5398
	show arp inspection statistics . . . . .	5399
	show dhcp snooping binding . . . . .	5400
	show firewall . . . . .	5402
	show firewall policer . . . . .	5406
	show interfaces filters . . . . .	5408

<b>Chapter 63</b>	<b>Troubleshooting</b> . . . . .	<b>5411</b>
	Troubleshooting Procedures . . . . .	5411
	Troubleshooting Firewall Filter Configuration . . . . .	5411
	Firewall Filter Configuration Returns a No Space Available in TCAM Message . . . . .	5411
	Filter Counts Previously Dropped Packet . . . . .	5413
	Matching Packets Not Counted . . . . .	5414
	Counter Reset When Editing Filter . . . . .	5414
	Cannot Include loss-priority and policer Actions in Same Term . . . . .	5414
	Cannot Egress Filter Certain Traffic Originating on QFX Switch . . . . .	5415
	Firewall Filter Match Condition Not Working with Q-in-Q Tunneling . . . . .	5415
	Egress Firewall Filters with Private VLANs . . . . .	5415
	Egress Filtering of L2PT Traffic Not Supported . . . . .	5416
	Cannot Drop BGP Packets in Certain Circumstances . . . . .	5416
	Invalid Statistics for Policer . . . . .	5416
	Policers can Limit Egress Filters . . . . .	5416
	Troubleshooting Policer Configuration . . . . .	5418
	Incomplete Count of Packet Drops . . . . .	5418
	Counter Reset When Editing Filter . . . . .	5418
	Invalid Statistics for Policer . . . . .	5418
	Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured . . . . .	5419
	Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured . . . . .	5420
	Policers Can Limit Egress Filters . . . . .	5420
<b>Part 18</b>	<b>Services</b>	
<b>Chapter 64</b>	<b>Overview</b> . . . . .	<b>5425</b>
	Port Mirroring . . . . .	5425
	Understanding Port Mirroring . . . . .	5425
	Port Mirroring Overview . . . . .	5425
	Port Mirroring Instance Types . . . . .	5426
	Port-Mirroring Terminology . . . . .	5426
	Port Mirroring and STP . . . . .	5428
	Port Mirroring Constraints and Limitations . . . . .	5428
	Understanding Layer 3 Logical Interfaces . . . . .	5430
	DHCP Relay . . . . .	5431
	DHCP and BOOTP Relay Overview . . . . .	5431
<b>Chapter 65</b>	<b>Configuration</b> . . . . .	<b>5433</b>
	Configuration Examples . . . . .	5433
	Examples: Configuring Port Mirroring for Local Analysis . . . . .	5433
	Example: Mirroring Employee Web Traffic with a Firewall Filter . . . . .	5435
	Example: Configuring Port Mirroring for Remote Analysis . . . . .	5438



Example: Mirroring Employee Web Traffic with a Firewall Filter . . . . .	5443
Configuration Tasks . . . . .	5446
Configuring Port Mirroring . . . . .	5446
Configuring Port Mirroring for Local Analysis . . . . .	5447
Configuring Port Mirroring for Remote Analysis . . . . .	5448
Filtering the Traffic Entering an Analyzer . . . . .	5448
Configuring DHCP and BOOTP . . . . .	5449
Configuration Statements for Port Mirroring . . . . .	5450
analyzer . . . . .	5451
egress . . . . .	5452
ethernet-switching (Port Mirroring) . . . . .	5453
family (Port Mirroring) . . . . .	5454
inet (Port Mirroring) . . . . .	5455
ingress (Port Mirroring) . . . . .	5456
input . . . . .	5457
instance (Port Mirroring) . . . . .	5458
interface (Port Mirroring) . . . . .	5459
ip-address (Port Mirroring) . . . . .	5460
output . . . . .	5461
port-mirroring . . . . .	5462
routing-instance (Port Mirroring) . . . . .	5463
vlan (Port Mirroring) . . . . .	5464
Configuration Statements for Encryption . . . . .	5464
authentication-key-chains . . . . .	5466
cache-size . . . . .	5467
cache-timeout-negative . . . . .	5468
ca-name . . . . .	5468
certificates . . . . .	5469
certification-authority . . . . .	5470
crl (Encryption Interface) . . . . .	5470
encoding . . . . .	5471
enrollment-retry . . . . .	5471
enrollment-url . . . . .	5472
file . . . . .	5472
key (Authentication Keychain) . . . . .	5473
key-chain (Security) . . . . .	5474
ldap-url . . . . .	5475
local . . . . .	5476
maximum-certificates . . . . .	5477
path-length . . . . .	5477
secret . . . . .	5478
security . . . . .	5479
ssh-known-hosts . . . . .	5480
start-time (Authentication Key Transmission) . . . . .	5481
traceoptions . . . . .	5483
Configuration Statements for DHCP . . . . .	5484
dhcp-local-server . . . . .	5485
dhcp-relay . . . . .	5490

<b>Chapter 66</b>	<b>Administration</b> .....	<b>5497</b>
	Monitoring Commands for Port Mirroring .....	5497
	show analyzer .....	5498
<b>Chapter 67</b>	<b>Troubleshooting</b> .....	<b>5501</b>
	Troubleshooting Procedures .....	5501
	Troubleshooting Port Mirroring .....	5501
	Port Mirroring Constraints and Limitations .....	5501
	Egress Port Mirroring with VLAN Translation .....	5503
	Egress Port Mirroring with Private VLANs .....	5503
<b>Part 19</b>	<b>Storage</b>	
<b>Chapter 68</b>	<b>Overview</b> .....	<b>5507</b>
	Software Features Overview .....	5507
	Overview of Fibre Channel .....	5508
	Fibre Channel Transport Protocol .....	5509
	How FC Works on the Switch .....	5509
	Supported FC Features and Functions .....	5512
	Lossless Transport Support .....	5512
	Overview of FIP .....	5513
	FCoE and FIP Snooping .....	5513
	Understanding DCB Features and Requirements .....	5515
	Lossless Transport .....	5515
	ETS .....	5516
	DCBX .....	5517
	Understanding FCoE .....	5518
	FCoE Devices .....	5519
	FCoE Frames .....	5520
	Virtual Links .....	5521
	FCoE VLANs .....	5521
	Understanding FCoE Transit Switch Functionality .....	5524
	Understanding FCoE and FIP Session High Availability .....	5528
	High Availability for Fibre Channel Process Termination (FCoE-FC Gateway Mode) .....	5528
	High Availability for FIP Snooping .....	5528
	Nonstop Software Upgrade (QFabric Systems) .....	5529
	Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches .....	5530
	Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch .....	5531
	FC Network Security .....	5532
	VN2VF_Port FIP Snooping Functions .....	5533
	FIP Snooping Firewall Filters .....	5534
	FIP Snooping Session Scalability .....	5534
	VN2VF_Port FIP Snooping Implementation .....	5534
	T11 VN2VF_Port FIP Snooping Specification .....	5538

Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch	5539
VN2VN_Port FIP Snooping and FIP Snooping Virtual Links	5539
VN2VN_Port Communication Modes	5540
Network Security	5541
VN2VN_Port FIP Snooping Functions	5541
Scalability	5541
VN2VN_Port FIP Snooping Implementation	5541
ENode-Facing Interfaces	5542
Network-Facing Interfaces (Connecting to Another Transit Switch)	5543
Beacon Period (VN2VN_Port FIP Snooping Link Maintenance)	5544
QFabric System Differences in VN2VN_Port FIP Snooping Traffic Handling	5544
Understanding FIP Snooping, FBF, and MVR Filter Scalability	5546
VFP TCAM Architecture and Allocation	5546
VFP TCAM Entry Consumption	5547
Rejected Filter Configurations (No Available VFP TCAM Space)	5550
VFP TCAM Allocation and Consumption (Scaling) Examples	5551
Filter Configuration Recommendations	5553
Understanding MC-LAGs on an FCoE Transit Switch	5555
Supported Topology	5555
FIP Snooping and FCoE Trusted Ports	5557
CoS and Data Center Bridging (DCB)	5558
Understanding CoS Flow Control (Ethernet PAUSE and PFC)	5559
Ethernet PAUSE	5560
PFC	5564
Lossless Transport Support Summary	5567
Understanding Fibre Channel Terminology	5569
DCBX	5580
Understanding DCBX	5580
DCBX Basics	5580
DCBX Modes and Support	5581
DCBX Attribute Types	5584
DCBX Application Protocol TLV Exchange	5585
DCBX and PFC	5586
DCBX and ETS	5587
Understanding DCBX Application Protocol TLV Exchange	5589
Applications	5590
Application Maps	5590
Classifying and Prioritizing Application Traffic	5591
Enabling Interfaces to Exchange Application Protocol Information	5592
Disabling DCBX Application Protocol Exchange	5592
Learn About Technology	5593
Data Center Technology Overview Videos	5593
Learn About Video: Why Do We Need an IP Fabric?	5593
Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?	5593
Learn About Video: Why Use an Overlay Network in a Data Center?	5593
Conceptual Documents That Contain Technology Overview Videos	5594

<b>Chapter 69</b>	<b>Configuration</b>	<b>5595</b>
	Configuration Examples	5595
	Example: Configuring DCBX Application Protocol TLV Exchange	5595
	Example: Configuring CoS PFC for FCoE Traffic	5606
	Configuration Examples	5614
	Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG	5614
	Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)	5636
	Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)	5641
	Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)	5647
	FCoE and FIP Snooping Configuration Tasks	5655
	Enabling and Disabling CoS OxID Hash Control on Standalone Switches	5656
	Configuring VLANs for FCoE Traffic on an FCoE Transit Switch	5657
	Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch	5662
	Enabling VN2VN_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch	5665
	DCBX Configuration Tasks	5667
	Configuring the DCBX Mode	5668
	Configuring DCBX Autonegotiation	5669
	Disabling the ETS Recommendation TLV	5672
	Defining an Application for DCBX Application Protocol TLV Exchange	5672
	Configuring an Application Map for DCBX Application Protocol TLV Exchange	5674
	Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange	5675
	Configuration Statements	5675
	application (Application Maps)	5677
	application (Applications)	5678
	application-map	5679
	application-maps	5680
	applications (Applications)	5681
	applications (DCBX)	5682
	beacon-period	5683
	code-points (Application Maps)	5684
	dcbx	5685
	dcbx-version	5686
	destination-port (Applications)	5687
	disable (DCBX)	5688
	enhanced-transmission-selection	5689
	ether-type	5690
	examine-vn2vf	5691
	examine-vn2vn	5692
	family fcoe	5693
	fc-map	5694
	fip-security	5696

	fcoe-trusted .....	5697
	interface (DCBX) .....	5698
	interface (FIP Snooping) .....	5699
	no-recommendation-tlv .....	5700
	oxid .....	5701
	policy-options .....	5702
	priority-flow-control .....	5703
	protocol (Applications) .....	5704
	recommendation-tlv .....	5705
<b>Chapter 70</b>	<b>Administration .....</b>	<b>5707</b>
	Operational Commands .....	5707
	clear fip snooping enode .....	5708
	clear fip snooping statistics .....	5709
	clear fip snooping vlan .....	5710
	clear fip vlan-discovery statistics .....	5711
	restart .....	5712
	show dcbx .....	5723
	show dcbx neighbors .....	5724
	show fip snooping .....	5746
	show fip snooping enode .....	5751
	show fip snooping fcf .....	5755
	show fip snooping interface .....	5758
	show fip snooping statistics .....	5761
	show fip snooping vlan .....	5764
	show fip vlan-discovery .....	5768
<b>Chapter 71</b>	<b>Troubleshooting .....</b>	<b>5771</b>
	Troubleshooting Procedures .....	5771
	Troubleshooting Dropped FCoE Traffic .....	5771
	Troubleshooting Dropped FIP Traffic .....	5774
<b>Part 20</b>	<b>Traffic Management</b>	
<b>Chapter 72</b>	<b>Overview .....</b>	<b>5779</b>
	CoS Overview .....	5779
	Overview of Junos OS CoS for the QFX Series and EX4600 Switch .....	5781
	CoS Standards .....	5781
	How Junos CoS Works .....	5782
	Default CoS Behavior .....	5783
	Overview of Policers .....	5783
	Policer Overview .....	5784
	Policer Types .....	5784
	Policer Actions .....	5785
	Policer Colors .....	5786
	Filter-Specific Policers .....	5786
	Suggested Naming Convention for Policers .....	5787
	Policer Counters .....	5787
	Policer Algorithms .....	5787
	How Many Policers are Supported? .....	5787

Policers can Limit Egress Firewall Filters . . . . .	5788
Understanding Junos CoS Components . . . . .	5789
Code-Point Aliases . . . . .	5789
Policers . . . . .	5789
Classifiers . . . . .	5789
Forwarding Classes . . . . .	5790
Forwarding Class Sets . . . . .	5790
Flow Control (Ethernet PAUSE, PFC, and ECN) . . . . .	5790
WRED Profiles . . . . .	5791
Schedulers . . . . .	5791
Rewrite Rules . . . . .	5792
Understanding CoS Packet Flow . . . . .	5793
CoS Inputs and Outputs Overview . . . . .	5795
Understanding Default CoS Settings . . . . .	5796
Default Forwarding Classes and Queue Mapping . . . . .	5796
Default Forwarding Class Sets (Priority Groups) . . . . .	5797
Default Code-Point Aliases . . . . .	5797
Default Classifiers . . . . .	5799
Default Rewrite Rules . . . . .	5801
Default Drop Profile . . . . .	5801
Default Schedulers . . . . .	5802
Default Scheduler Maps . . . . .	5804
Default Shared Buffer Configuration . . . . .	5804
Understanding Host Inbound Traffic Classification . . . . .	5805
Understanding Host Routing Engine Outbound Traffic Queues and Defaults . . . . .	5806
Understanding CoS Code-Point Aliases . . . . .	5808
Default Code-Point Aliases . . . . .	5808
Understanding CoS Classifiers . . . . .	5810
Interfaces and Output Queues . . . . .	5810
Behavior Aggregate Classifiers . . . . .	5811
Fixed Classifiers on Ethernet Interfaces . . . . .	5814
Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports) . . . . .	5815
Multifield Classifiers . . . . .	5815
Packet Classification for Routed VLAN Interfaces (RVIs) . . . . .	5816
Understanding CoS MPLS EXP Classifiers and Rewrite Rules . . . . .	5817
EXP Classifiers . . . . .	5817
EXP Rewrite Rules . . . . .	5818
Schedulers . . . . .	5819
Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces . . . . .	5820
Supported Classifier and Rewrite Rule Types . . . . .	5820
Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration . . . . .	5821
Default Classifiers . . . . .	5823
Default Rewrite Rules . . . . .	5824
Classifier Precedence . . . . .	5824
Classifier Behavior and Limitations . . . . .	5825
Rewrite Rule Precedence and Behavior . . . . .	5826

Classifier and Rewrite Rule Configuration Interaction with Ethernet	
Interface Configuration . . . . .	5827
Understanding CoS Forwarding Classes . . . . .	5830
Default Forwarding Classes . . . . .	5831
Forwarding Class Configuration Rules . . . . .	5832
Lossless Transport Support . . . . .	5833
Understanding CoS Forwarding Class Sets (Priority Groups) . . . . .	5835
Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows . . . . .	5837
Lossless Transport Features Introduced in Junos OS Release 12.3 . . . . .	5838
Default Lossless Priority Configuration . . . . .	5838
Configuring Lossless Priorities . . . . .	5841
Backward Compatibility with Junos OS Releases Earlier Than Release	
12.3 . . . . .	5854
Configuration Rules and Recommendations . . . . .	5855
Understanding Default CoS Scheduling and Classification . . . . .	5856
Default Classification . . . . .	5856
Default Scheduling . . . . .	5859
Default DCBX Advertisement . . . . .	5861
Default Scheduling and Classification Summary . . . . .	5862
Understanding CoS Hierarchical Port Scheduling (ETS) . . . . .	5862
Hierarchical Scheduling Tiers . . . . .	5863
Hierarchical Scheduling and ETS . . . . .	5863
ETS Advertisement in DCBX . . . . .	5864
Hierarchical Scheduling Process . . . . .	5865
Strict-High Priority Queues and Hierarchical Scheduling . . . . .	5866
Default Hierarchical Scheduling . . . . .	5866
Understanding CoS Output Queue Schedulers . . . . .	5868
Output Queue Scheduling Components . . . . .	5868
Default Schedulers . . . . .	5869
Transmit Rate (Minimum Guaranteed Bandwidth) . . . . .	5872
Sharing Extra Bandwidth . . . . .	5872
Shaping Rate (Maximum Bandwidth) . . . . .	5873
Scheduling Priority . . . . .	5873
Scheduler Drop-Profile Maps . . . . .	5874
Buffer Size . . . . .	5874
Explicit Congestion Notification . . . . .	5875
Scheduler Maps . . . . .	5876
Understanding CoS Priority Group Scheduling . . . . .	5877
Priority Group Scheduling Components . . . . .	5877
Default Traffic Control Profile . . . . .	5878
Guaranteed Rate (Minimum Guaranteed Bandwidth) . . . . .	5878
Sharing Extra Bandwidth . . . . .	5878
Shaping Rate (Maximum Bandwidth) . . . . .	5879
Scheduler Maps . . . . .	5879
Understanding CoS Traffic Control Profiles . . . . .	5880
Understanding CoS Priority Group and Queue Guaranteed Rates (Minimum	
Bandwidth) . . . . .	5881
Guaranteeing Bandwidth Using Hierarchical Scheduling . . . . .	5881
Priority Group Guaranteed Rate (Minimum Bandwidth) . . . . .	5883

Queue Transmit Rate (Minimum Bandwidth) . . . . .	5883
Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth) . . . . .	5884
Priority Group Shaping . . . . .	5884
Queue Shaping . . . . .	5884
Shaping Maximum Bandwidth Using Hierarchical Scheduling . . . . .	5885
Understanding CoS Scheduling Behavior and Configuration	
Considerations . . . . .	5886
Understanding CoS Buffer Configuration . . . . .	5891
Buffer Pools . . . . .	5892
Default Buffer Pool Values . . . . .	5900
Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios . . . . .	5903
Optimizing Buffer Configuration . . . . .	5906
General Buffer Configuration Rules and Considerations . . . . .	5908
Understanding CoS WRED Drop Profiles . . . . .	5909
Drop Profile Parameters . . . . .	5910
Default Drop Profile . . . . .	5911
Packet Drop Method . . . . .	5911
Drop Profile Maps . . . . .	5912
Congestion Prevention . . . . .	5912
Configuring a WRED Drop Profile and Applying it to an Output Queue . . . . .	5913
Understanding CoS Rewrite Rules . . . . .	5914
Understanding CoS Flow Control (Ethernet PAUSE and PFC) . . . . .	5916
Ethernet PAUSE . . . . .	5917
PFC . . . . .	5921
Lossless Transport Support Summary . . . . .	5924
Understanding CoS Explicit Congestion Notification . . . . .	5926
How ECN Works . . . . .	5926
WRED Drop Profile Control of ECN Thresholds . . . . .	5931
Support, Limitations, and Notes . . . . .	5932
Understanding DCBX Features and Requirements . . . . .	5934
Lossless Transport . . . . .	5934
ETS . . . . .	5935
DCBX . . . . .	5936
Understanding DCBX . . . . .	5937
DCBX Basics . . . . .	5937
DCBX Modes and Support . . . . .	5938
DCBX Attribute Types . . . . .	5941
DCBX Application Protocol TLV Exchange . . . . .	5942
DCBX and PFC . . . . .	5943
DCBX and ETS . . . . .	5943
Understanding DCBX Application Protocol TLV Exchange . . . . .	5946
Applications . . . . .	5946
Application Maps . . . . .	5947
Classifying and Prioritizing Application Traffic . . . . .	5948
Enabling Interfaces to Exchange Application Protocol Information . . . . .	5949



Disabling DCBX Application Protocol Exchange . . . . .	5949
QFX5100 Switches Only . . . . .	5950
Understanding PFC Functionality Across Layer 3 Interfaces . . . . .	5950
QFX3500 and QFX3600 Virtual Chassis Only . . . . .	5952
CoS on Virtual Chassis Switch Ports . . . . .	5953
Access Interface CoS Support . . . . .	5953
VCP Interface CoS Support . . . . .	5955
CPU-Generated Host Outbound Traffic . . . . .	5956
Virtual Chassis Fabric Only . . . . .	5957
CoS on Virtual Chassis Fabric (VCF) EX4300 Leaf Devices (Mixed Mode) . . . . .	5958
VCF CoS in Mixed Mode with an EX4300 Leaf Device . . . . .	5958
Scheduling on an EX4300 VCF Leaf Device . . . . .	5960
Learn About Technology . . . . .	5963
Data Center Technology Overview Videos . . . . .	5963
Learn About Video: Why Do We Need an IP Fabric? . . . . .	5963
Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric? . . . . .	5964
Learn About Video: Why Use an Overlay Network in a Data Center? . . . . .	5964
Conceptual Documents That Contain Technology Overview Videos . . . . .	5964
<b>Chapter 73 Configuration . . . . .</b>	<b>5965</b>
Configuration Examples . . . . .	5965
Example: Configuring CoS Hierarchical Port Scheduling (ETS) . . . . .	5966
Example: Configuring CoS PFC for FCoE Traffic . . . . .	5987
Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG . . . . .	5995
Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) . . . . .	6019
Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface . . . . .	6028
Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces . . . . .	6036
Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) . . . . .	6050
Example: Configuring Unicast Classifiers . . . . .	6066
Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers . . . . .	6069
Example: Configuring WRED Drop Profiles . . . . .	6071
Example: Configuring Drop Profile Maps . . . . .	6073
Example: Configuring Forwarding Classes . . . . .	6075
Example: Configuring Forwarding Class Sets . . . . .	6078
Example: Configuring Queue Schedulers . . . . .	6081
Example: Configuring Queue Scheduling Priority . . . . .	6087
Example: Configuring ECN . . . . .	6090
Example: Configuring Traffic Control Profiles (Priority Group Scheduling) . . . . .	6094
Example: Configuring Minimum Guaranteed Output Bandwidth . . . . .	6096

Example: Configuring Maximum Output Bandwidth . . . . .	6101
Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic . . . . .	6104
Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled . . . . .	6110
Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic . . . . .	6116
Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic . . . . .	6122
Example: Configuring DCBX Application Protocol TLV Exchange . . . . .	6128
Configuration Examples (QFX5100 Switches Only) . . . . .	6138
Example: Configuring PFC Across Layer 3 Interfaces . . . . .	6138
Configuration Tasks . . . . .	6156
Configuring CoS . . . . .	6157
Defining CoS Code-Point Aliases . . . . .	6159
Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p) . . . .	6160
Configuring a Global MPLS EXP Classifier . . . . .	6161
Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers . . . . .	6162
Configuring CoS WRED Drop Profiles . . . . .	6163
Configuring CoS Drop Profile Maps . . . . .	6164
Defining CoS Forwarding Classes . . . . .	6164
Defining CoS Forwarding Class Sets . . . . .	6166
Disabling the ETS Recommendation TLV . . . . .	6167
Defining CoS Queue Schedulers . . . . .	6167
Defining CoS Queue Scheduling Priority . . . . .	6171
Changing the Host Outbound Traffic Default Queue Mapping . . . . .	6172
Defining CoS Traffic Control Profiles (Priority Group Scheduling) . . . . .	6172
Configuring CoS PFC (Congestion Notification Profiles) . . . . .	6174
Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control . . . .	6177
Configuring CoS Asymmetric Ethernet PAUSE Flow Control . . . . .	6178
Configuring Global Ingress and Egress Shared Buffers . . . . .	6179
Defining CoS Rewrite Rules . . . . .	6182
Configuring Rewrite Rules for MPLS EXP Classifiers . . . . .	6184
Assigning CoS Components to Interfaces . . . . .	6185
Configuring the DCBX Mode . . . . .	6186
Configuring DCBX Autonegotiation . . . . .	6187
Defining an Application for DCBX Application Protocol TLV Exchange . . . .	6190
Configuring an Application Map for DCBX Application Protocol TLV Exchange . . . . .	6191
Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange . . . . .	6192
Configuration Statements . . . . .	6192
application (Application Maps) . . . . .	6195
application (Applications) . . . . .	6196
application-map . . . . .	6197
application-maps . . . . .	6198
applications (Applications) . . . . .	6199

applications (DCBX) .....	6200
buffer-partition (Egress) .....	6201
buffer-partition (Ingress) .....	6203
buffer-size .....	6205
cable-length (Congestion Notification) .....	6207
class-of-service .....	6208
class (Forwarding Classes) .....	6212
class (Forwarding Class Sets) .....	6213
classifiers .....	6214
code-point (Input Congestion Notification) .....	6215
code-point (Output Congestion Notification) .....	6216
code-point (Rewrite Rules) .....	6217
code-point-aliases .....	6217
code-points (Application Maps) .....	6218
code-points (CoS) .....	6218
configured-flow-control .....	6219
congestion-notification-profile .....	6220
dcbx .....	6222
dcbx-version .....	6223
destination-port (Applications) .....	6224
disable (DCBX) .....	6225
drop-probability .....	6226
drop-profile .....	6227
drop-profile-map .....	6227
drop-profiles .....	6228
dscp .....	6229
dscp-ipv6 .....	6231
dscp-code-point .....	6232
egress (Buffer Configuration) .....	6233
enhanced-transmission-selection .....	6234
ether-type .....	6235
exp .....	6236
explicit-congestion-notification .....	6237
fill-level .....	6238
flow-control .....	6239
flow-control-queue (Output Congestion Notification) .....	6240
forwarding-class .....	6242
forwarding-class (Host Outbound Traffic) .....	6243
forwarding-class-set .....	6243
forwarding-class-sets .....	6244
forwarding-classes .....	6245
guaranteed-rate .....	6247
host-outbound-traffic .....	6248
ieee-802.1 .....	6249
ieee-802.1 (Input Congestion Notification) .....	6250
ieee-802.1 (Output Congestion Notification) .....	6251
import .....	6252
ingress (Buffer Configuration) .....	6253
input (Congestion Notification) .....	6254

interface (DCBX) .....	6255
interfaces (Class of Service) .....	6256
interpolate .....	6257
loss-priority (Classifiers) .....	6258
loss-priority (Drop Profiles) .....	6259
loss-priority (Rewrite Rules) .....	6260
multi-destination .....	6261
mru .....	6262
output (Congestion Notification) .....	6263
output-traffic-control-profile .....	6264
pfc (Input Congestion Notification) .....	6265
policy-options .....	6266
priority (Schedulers) .....	6267
priority-flow-control .....	6268
protocol (Applications) .....	6269
protocol (Drop Profile Map) .....	6270
queue-num .....	6271
recommendation-tlv .....	6272
rewrite-rules .....	6273
rx-buffers .....	6274
scheduler .....	6275
scheduler-map .....	6275
scheduler-maps .....	6276
schedulers .....	6277
shaping-rate .....	6278
shared-buffer .....	6280
system-defaults .....	6281
traceoptions (Class of Service) .....	6282
traffic-control-profiles .....	6284
transmit-rate .....	6285
tx-buffers .....	6287
unit .....	6288
<b>Chapter 74 Administration .....</b>	<b>6289</b>
Routine Monitoring .....	6289
Monitoring CoS Classifiers .....	6289
Monitoring CoS Forwarding Classes .....	6290
Monitoring Interfaces That Have CoS Components .....	6291
Monitoring CoS Rewrite Rules .....	6292
Monitoring CoS Scheduler Maps .....	6293
Monitoring CoS Value Aliases .....	6294
Operational Commands .....	6295
show class-of-service .....	6297
show class-of-service classifier .....	6301
show class-of-service code-point-aliases .....	6303
show class-of-service congestion-notification .....	6305
show class-of-service drop-profile .....	6308
show class-of-service forwarding-class .....	6311
show class-of-service forwarding-class-set .....	6313

	show class-of-service forwarding-table . . . . .	6315
	show class-of-service forwarding-table classifier . . . . .	6319
	show class-of-service forwarding-table classifier mapping . . . . .	6321
	show class-of-service forwarding-table drop-profile . . . . .	6323
	show class-of-service forwarding-table rewrite-rule . . . . .	6325
	show class-of-service forwarding-table rewrite-rule mapping . . . . .	6327
	show class-of-service forwarding-table scheduler-map . . . . .	6328
	show class-of-service interface . . . . .	6330
	show class-of-service multi-destination . . . . .	6358
	show class-of-service rewrite-rule . . . . .	6359
	show class-of-service scheduler-map . . . . .	6361
	show class-of-service shared-buffer . . . . .	6363
	show class-of-service traffic-control-profile . . . . .	6365
	show dcbx . . . . .	6369
	show dcbx neighbors . . . . .	6370
	show interfaces queue . . . . .	6392
	show pfe filter hw summary . . . . .	6432
	show pfe next-hop . . . . .	6434
	show pfe route . . . . .	6439
	show pfe terse . . . . .	6448
	show pfe version . . . . .	6450
<b>Chapter 75</b>	<b>Troubleshooting . . . . .</b>	<b>6451</b>
	Troubleshooting Procedures . . . . .	6451
	Troubleshooting Dropped FCoE Traffic . . . . .	6451
	Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth . . . . .	6454
	Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth . . . . .	6455
	Troubleshooting Egress Queue Bandwidth Impacted by Congestion . . . .	6456
	Troubleshooting an Unexpected Rewrite Value . . . . .	6457
	Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic . . . . .	6458
<b>Part 21</b>	<b>Network Management and Monitoring</b>	
<b>Chapter 76</b>	<b>Overview . . . . .</b>	<b>6463</b>
	Network Management . . . . .	6463
	Understanding Device and Network Management Features . . . . .	6463
	Understanding Network Management Implementation on the QFabric System . . . . .	6466
	Understanding Telnet on the QFabric System . . . . .	6467
	Understanding Tracing and Logging Operations . . . . .	6468
	Automation . . . . .	6469
	Overview of QFX5100 Switch Automation Enhancements . . . . .	6470
	Features of the QFX5100 Switch Automation Enhancements . . . . .	6470
	Overview of Python with QFX5100 Switch Automation Enhancements . .	6471
	Understanding Automation Scripts Support . . . . .	6473

How Commit Scripts Work . . . . .	6474
Commit Script Input . . . . .	6475
Commit Script Output . . . . .	6476
Commit Scripts and the Junos OS Commit Model . . . . .	6477
Avoiding Potential Conflicts When Using Multiple Commit Scripts . . . . .	6480
Overview of Generating Persistent or Transient Configuration Changes . . . . .	6481
Differences Between Persistent and Transient Changes . . . . .	6481
Interaction of Configuration Changes and Configuration Groups . . . . .	6484
Tag Elements and Templates for Generating Changes . . . . .	6484
Required Boilerplate for Commit Scripts . . . . .	6485
How Op Scripts Work . . . . .	6486
Required Boilerplate for Op Scripts . . . . .	6487
Junos Space . . . . .	6489
Understanding Junos Space Support . . . . .	6489
Network Analytics . . . . .	6490
Network Analytics Overview . . . . .	6490
Analytics Feature Overview . . . . .	6491
Network Analytics Enhancements Overview . . . . .	6492
Summary of CLI Changes . . . . .	6493
Understanding Network Analytics Configuration and Status . . . . .	6497
Understanding Network Analytics Streaming Data . . . . .	6499
Understanding Enhanced Network Analytics Streaming Data . . . . .	6501
Google Protocol Buffer (GPB) . . . . .	6501
JavaScript Object Notation (JSON) . . . . .	6504
Comma-separated Values (CSV) . . . . .	6504
Tab-separated Values (TSV) . . . . .	6504
Queue Statistics Output for JSON, CSV, and TSV . . . . .	6505
Traffic Statistics Output for JSON, CSV, and TSV . . . . .	6505
Understanding Enhanced Analytics Local File Output . . . . .	6506
Prototype File for the Google Protocol Buffer Stream Format . . . . .	6508
sFlow Technology . . . . .	6509
Understanding How to Use sFlow Technology for Network Monitoring on a Switch . . . . .	6509
Sampling Mechanism and Architecture of sFlow Technology on Switches . . . . .	6509
Adaptive Sampling . . . . .	6511
sFlow Agent Address Assignment . . . . .	6511
sFlow Limitations on Switches . . . . .	6512
SNMP . . . . .	6513
Understanding the Implementation of SNMP . . . . .	6513
Understanding the Implementation of SNMP on the QFabric System . . . . .	6516
Fabric Chassis MIB . . . . .	6518
Utility MIB . . . . .	6522
SNMPv3 Overview . . . . .	6523
Minimum SNMPv3 Configuration on a Device Running Junos OS . . . . .	6524
Understanding RMON . . . . .	6525
RMON Overview . . . . .	6525
Alarm Thresholds and Events . . . . .	6526
RMON MIB Event, Alarm, Log, and History Control Tables . . . . .	6527

Understanding Health Monitoring . . . . .	6529
SNMP MIBs Support . . . . .	6530
MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6530
MIBs Supported on QFabric Systems . . . . .	6539
SNMP Traps Support . . . . .	6546
SNMP Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6546
SNMP Traps Supported on QFabric Systems . . . . .	6554
MIB Objects for the QFX Series . . . . .	6558
QFX Series Standalone Switches . . . . .	6558
QFabric Systems . . . . .	6558
QFabric System QFX3100 Director Device . . . . .	6559
QFabric System QFX3008-I Interconnect Device . . . . .	6559
QFabric System QFX3600-I Interconnect Device . . . . .	6559
QFabric System Node Devices . . . . .	6560
System Logging . . . . .	6560
Overview of Junos OS System Log Messages . . . . .	6560
Overview of Single-Chassis System Logging Configuration . . . . .	6561
Understanding the Implementation of System Log Messages on the QFabric System . . . . .	6562
<b>Chapter 77 Configuration . . . . .</b>	<b>6565</b>
Configuration Examples . . . . .	6565
Examples: Configuring System Logging . . . . .	6565
Examples: Assigning an Alternative Facility . . . . .	6567
Example: Configuring System Log Messages . . . . .	6568
Example: Monitoring Network Traffic Using sFlow Technology . . . . .	6571
Example: Configuring SNMP . . . . .	6575
Example: Configuring Network Analytics . . . . .	6577
Configuration Tasks for Network Management . . . . .	6583
Configuring Console and Auxiliary Port Properties . . . . .	6583
Configuring SSH Service for Remote Access to the Router or Switch . . . . .	6584
Configuring the Root Login Through SSH . . . . .	6585
Configuring the SSH Protocol Version . . . . .	6586
Configuring the Client Alive Mechanism . . . . .	6586
Configuring Telnet Service for Remote Access to a Switch . . . . .	6586
Configuration Tasks for Automation . . . . .	6587
Invoking the Python Interpreter . . . . .	6587
Controlling the Execution of Commit Scripts . . . . .	6588
Enabling Commit Scripts to Execute . . . . .	6588
Removing Commit Scripts from the Configuration . . . . .	6589
Deactivating Commit Scripts . . . . .	6590
Activating Inactive Commit Scripts . . . . .	6590
Configuration Tasks for Network Analytics . . . . .	6590
Configuring Queue Monitoring . . . . .	6591
Configuring Traffic Monitoring . . . . .	6593
Configuring a Local File for Network Analytics Data . . . . .	6594
Configuring a Remote Collector for Streaming Analytics Data . . . . .	6595

Configuration Tasks for sFlow Technology .....	6596
Configuring sFlow Technology .....	6596
Configuration Tasks for SNMP .....	6597
Configuring SNMP .....	6598
Configuring the SNMP Community String .....	6601
Configuring SNMP Trap Groups .....	6602
Adding a Group of Clients to an SNMP Community .....	6603
Configuring the Interfaces on Which SNMP Requests Can Be Accepted ..	6604
Configuring MIB Views .....	6605
Configuring RMON Alarms and Events .....	6606
Configuring SNMP .....	6607
Configuring an Event .....	6607
Configuring an Alarm .....	6608
Configuring Health Monitoring .....	6609
Creating SNMPv3 Users .....	6609
Configuring Access Privileges for a Group .....	6611
Assigning a Security Name to a Group .....	6612
Configuring SNMPv3 Traps on a Device Running Junos OS .....	6613
Configuring SNMP Informs .....	6614
Configuration Tasks for System Log Messages .....	6615
Junos OS Minimum System Logging Configuration .....	6616
Junos OS System Log Configuration Statements .....	6616
Adding a Text String to System Log Messages Directed to a Remote	
Destination .....	6617
Directing System Log Messages to a Log File .....	6618
Directing System Log Messages to a Remote Machine .....	6619
Directing System Log Messages to a User Terminal .....	6620
Directing System Log Messages to the Console .....	6620
Disabling the System Logging of a Facility .....	6620
Displaying a Log File from a Single-Chassis System .....	6621
Including Priority Information in System Log Messages .....	6622
Including the Year or Millisecond in Timestamps .....	6623
Logging Messages in Structured-Data Format .....	6624
Interpreting Messages Generated in Structured-Data Format .....	6625
Interpreting Messages Generated in Standard Format .....	6628
Specifying Log File Size, Number, and Archiving Properties .....	6629
Specifying the Facility and Severity of Messages to Include in the Log ..	6631
Junos OS System Logging Facilities and Message Severity Levels .....	6633
Default Facilities for System Log Messages Directed to a Remote	
Destination .....	6634
Alternate Facilities for System Log Messages Directed to a Remote	
Destination .....	6635
Changing the Alternative Facility Name for System Log Messages Directed	
to a Remote Destination .....	6636
Using Regular Expressions to Refine the Set of Logged Messages .....	6637
Configuration Statements for Network Management .....	6639
connection-limit .....	6641
destination-override .....	6642
no-remote-trace .....	6642



protocol-version . . . . .	6643
rate-limit . . . . .	6644
ssh . . . . .	6645
telnet . . . . .	6646
tracing . . . . .	6647
Configuration Statements for Automation . . . . .	6647
allow-transients . . . . .	6648
apply-macro . . . . .	6649
checksum . . . . .	6650
command . . . . .	6651
commit . . . . .	6652
description . . . . .	6653
direct-access . . . . .	6653
file (Commit Scripts) . . . . .	6654
file (Op Scripts) . . . . .	6655
no-allow-url . . . . .	6656
op . . . . .	6657
optional . . . . .	6658
refresh (Commit Scripts) . . . . .	6659
refresh (Op Scripts) . . . . .	6660
refresh-from (Commit Scripts) . . . . .	6661
refresh-from (Op Scripts) . . . . .	6662
scripts . . . . .	6663
source (Commit Scripts) . . . . .	6665
source (Op Scripts) . . . . .	6666
Configuration Statements for Network Analytics . . . . .	6666
analytics . . . . .	6667
depth-threshold . . . . .	6671
interfaces (Analytics) . . . . .	6672
latency-threshold . . . . .	6674
queue-statistics . . . . .	6676
streaming-servers . . . . .	6678
traceoptions (Analytics) . . . . .	6679
traffic-statistics . . . . .	6680
Configuration Statements for sFlow Technology . . . . .	6681
agent-id . . . . .	6682
collector (sFlow Technology) . . . . .	6682
interfaces (sFlow) . . . . .	6683
polling-interval . . . . .	6684
sample-rate . . . . .	6685
sfow . . . . .	6686
source-ip . . . . .	6687
traceoptions (sFlow Technology) . . . . .	6688
udp-port . . . . .	6689
Configuration Statements for SNMP . . . . .	6689
access (SNMP) . . . . .	6693
address (SNMP) . . . . .	6693
address-mask . . . . .	6694
agent-address . . . . .	6694

alarm (SNMP RMON) .....	6695
authentication-md5 .....	6696
authentication-none .....	6697
authentication-password .....	6698
authentication-sha .....	6699
authorization .....	6700
bucket-size .....	6701
categories .....	6701
client-list .....	6702
client-list-name .....	6702
clients .....	6703
commit-delay .....	6703
community (SNMP) .....	6704
community (RMON) .....	6705
community-name (SNMP) .....	6706
contact .....	6707
description (SNMP) .....	6707
description (RMON) .....	6708
destination-port (SNMP) .....	6708
engine-id .....	6709
event .....	6710
falling-event-index (RMON) .....	6711
falling-threshold (Health Monitor) .....	6712
falling-threshold (RMON) .....	6713
falling-threshold-interval .....	6714
filter-duplicates .....	6714
filter-interfaces .....	6715
group (Associating a Security Name) .....	6715
group (Configuring Access Privileges) .....	6716
health-monitor .....	6717
history .....	6718
interface (SNMP) .....	6719
interface (RMON) .....	6720
interval (Health Monitor) .....	6720
interval (RMON) .....	6721
local-engine .....	6722
location .....	6723
message-processing-model .....	6723
name .....	6724
nonvolatile .....	6724
notify .....	6725
notify-filter (Applying to the Management Target) .....	6726
notify-filter (Configuring the Profile Name) .....	6726
notify-view .....	6727
oid .....	6727
oid (SNMPv3) .....	6728
owner .....	6729
parameters .....	6729
port (SNMP) .....	6730

privacy-3des	6731
privacy-aes128	6732
privacy-des	6733
privacy-none	6733
privacy-password	6734
read-view	6735
remote-engine	6736
request-type	6737
retry-count (SNMPv3)	6738
rising-event-index	6739
rising-threshold (Health Monitor)	6740
rising-threshold (RMON)	6741
rmon	6742
sample-type	6743
security-level (Defining Access Privileges)	6744
security-level (Generating SNMP Notifications)	6745
security-model (Access Privileges)	6746
security-model (Group)	6747
security-model (SNMP Notifications)	6748
security-name (Community String)	6749
security-name (Security Group)	6750
security-name (SNMP Notifications)	6751
security-to-group	6752
snmp	6753
snmp-community	6757
source-address (SNMP)	6757
startup-alarm	6758
syslog-subtag	6759
tag (Configuring Notification Targets)	6759
tag (Configuring the SNMP Community)	6760
tag-list	6760
target-address	6761
target-parameters	6762
targets	6763
timeout	6763
traceoptions (SNMP)	6764
trap-group	6766
trap-options	6767
type (RMON Notification)	6768
type (SNMPv3)	6769
user	6769
usm	6770
v3	6772
vacm	6774
variable	6775
version	6776
view (Configuring a MIB View)	6777
view (Associating MIB View with a Community)	6778
write-view	6778

	Configuration Statements for System Log Messages . . . . .	6778
	archive (All System Log Files) . . . . .	6780
	archive (Individual System Log File) . . . . .	6782
	archive (QFabric System) . . . . .	6783
	console (System Logging) . . . . .	6784
	explicit-priority . . . . .	6785
	facility-override . . . . .	6785
	file (QFabric System) . . . . .	6786
	file (System Logging) . . . . .	6787
	files . . . . .	6788
	host (System) . . . . .	6789
	log-prefix (System) . . . . .	6791
	match . . . . .	6791
	size (System) . . . . .	6792
	structured-data . . . . .	6793
	syslog (System) . . . . .	6794
	syslog (QFabric System) . . . . .	6796
	time-format . . . . .	6797
	user (System Logging) . . . . .	6798
<b>Chapter 78</b>	<b>Administration . . . . .</b>	<b>6799</b>
	Monitoring Tasks . . . . .	6799
	Displaying a Log File from a Single-Chassis System . . . . .	6799
	Monitoring Traffic Through the Router or Switch . . . . .	6800
	Displaying Real-Time Statistics About All Interfaces on the Router or Switch . . . . .	6800
	Displaying Real-Time Statistics About an Interface on the Router or Switch . . . . .	6801
	Monitoring RMON MIB Tables . . . . .	6803
	Monitoring SNMP . . . . .	6804
	Monitoring System Log Messages . . . . .	6805
	Pinging Hosts . . . . .	6806
	Tracing SNMP Activity on a Device Running Junos OS . . . . .	6807
	Configuring the Number and Size of SNMP Log Files . . . . .	6808
	Configuring Access to the Log File . . . . .	6808
	Configuring a Regular Expression for Lines to Be Logged . . . . .	6809
	Configuring the Trace Operations . . . . .	6809
	Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage . . . . .	6810
	Displaying Commit Script Output . . . . .	6812
	Commands for General Monitoring . . . . .	6814
	monitor traffic . . . . .	6815
	ping . . . . .	6825
	Commands for Network Analytics . . . . .	6828
	monitor start (Analytics) . . . . .	6830
	show analytics collector . . . . .	6833
	show analytics configuration . . . . .	6835
	show analytics queue-statistics . . . . .	6839
	show analytics status . . . . .	6841
	show analytics streaming-servers . . . . .	6845

show analytics traffic-statistics .....	6847
Commands for sFlow Technology .....	6849
clear sflow collector statistics .....	6850
show sflow .....	6851
show sflow collector .....	6853
show sflow interface .....	6854
Commands for SNMP .....	6855
clear snmp history .....	6856
clear snmp statistics .....	6857
request snmp spoof-trap .....	6859
request snmp utility-mib clear instance .....	6865
request snmp utility-mib set instance .....	6866
show snmp health-monitor .....	6867
show snmp inform-statistics .....	6872
show snmp mib .....	6874
show snmp rmon .....	6877
show snmp rmon history .....	6881
show snmp statistics .....	6882
show snmp v3 .....	6886
Commands for Syslog .....	6888
show log .....	6889
<b>Chapter 79</b>	
<b>Troubleshooting .....</b>	<b>6893</b>
Troubleshooting Overview .....	6893
Understanding Troubleshooting Resources .....	6893
Troubleshooting Overview .....	6895
QFX5100 Switch with Automation Enhancements Frequently Asked	
Questions .....	6897
Who Should You Contact If You Have Problems with Loading, Installing	
or Updating Libraries? .....	6898
Who Should You Contact If You Have Problems with Puppet for Junos	
OS? .....	6898
Who Should You Contact If You Have Problems with Chef for Junos	
OS? .....	6898
What Happens to the User Partition If You Downgrade a QFX5100	
Switch That Is Running the jinstall-qfx-5-flex-x.tgz Software Bundle	
to a QFX Switch That Is Running a Different QFX5100 Software	
Bundle? .....	6898
How Do You Recover Junos OS Binaries That You Have Deleted? ...	6898
How Do You Recover from a System Crash? .....	6898
How Can You Verify That a QFX5100 Switch Is Running a	
jinstall-qfx-5-flex-x.tgz Software Bundle? .....	6898
Troubleshooting Procedures .....	6899
Recovering from a Failed Software Installation .....	6899
Loading a Previous Configuration File .....	6900
Reverting to the Default Factory Configuration .....	6900
Reverting to the Rescue Configuration .....	6901
Recovering the Root Password .....	6901
Troubleshooting a Deprecated Network Analytics Configuration .....	6903

**Part 22****Chapter 80****Virtual Chassis**

<b>Overview</b>	<b>6907</b>
Virtual Chassis Overview	6907
Understanding QFX Series Virtual Chassis	6907
QFX Virtual Chassis Overview	6908
QFX5100 Switches in a Virtual Chassis	6908
QFX3500 and QFX3600 Switches in a Virtual Chassis	6909
EX4300 Switches in a QFX Series Virtual Chassis	6909
Understanding QFX Series Virtual Chassis Components	6909
Virtual Chassis Ports (VCPs)	6909
Maximum Switch Support	6910
Master Role	6910
Backup Role	6911
Linecard Role	6911
Member Switch and Member ID	6912
Mastership Priority	6913
Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric	6913
Virtual Chassis Fabric Summary	6914
Understanding Mixed Virtual Chassis Fabric	6914
Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches	6915
Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches	6915
Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis	6916
Understanding Mixed EX4300 and EX4600 Virtual Chassis	6916
Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis	6916
Understanding How the Master in a Virtual Chassis Is Elected	6917
Understanding Software Upgrades in a QFX Series Virtual Chassis	6918
Understanding Global Management of a Virtual Chassis	6919
Understanding Nonvolatile Storage in a Virtual Chassis	6921
Nonvolatile Memory Features	6921
Understanding QFX Series Virtual Chassis Port Link Aggregation	6921
Understanding Split and Merge in a Virtual Chassis	6922
What Happens When a Virtual Chassis Configuration Splits	6923
Merging Virtual Chassis Configurations	6924
Understanding Automatic Software Update on Virtual Chassis Member Switches	6925
Automatic Software Update Basics	6926
Automatic Software Update Restrictions	6926
Understanding MAC Address Assignment on a Virtual Chassis	6928

<b>Chapter 81</b>	<b>Configuration</b>	<b>6931</b>
	Configuration Tasks	6931
	Configuring a QFX Series Virtual Chassis (CLI Procedure)	6931
	Configuring a QFX Series Virtual Chassis with a Preprovisioned Configuration File	6932
	Configuring a QFX Series Virtual Chassis with a Nonprovisioned Configuration File	6934
	Adding a New Switch to an Existing QFX Series Virtual Chassis (CLI Procedure)	6936
	Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure)	6938
	Remove, Repair, and Reinstall the Same Switch	6939
	Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration	6939
	Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch	6940
	Configuring Mastership of a Virtual Chassis (CLI Procedure)	6941
	Configuring Mastership Using a Preprovisioned Configuration File	6941
	Configuring Mastership Using a Configuration File That Is Not Preprovisioned	6942
	Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure)	6943
	Disabling Split and Merge in a Virtual Chassis (CLI Procedure)	6944
	Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)	6944
	Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure)	6946
	Configuration Statements	6946
	[edit virtual-chassis] Configuration Statement Hierarchy	6947
	Supported Statements in the [edit virtual-chassis] Hierarchy Level	6947
	Unsupported Statements in the [edit virtual-chassis] Hierarchy Level	6948
	aliases (Virtual Chassis)	6949
	alias-name (Virtual Chassis aliases)	6950
	auto-sw-update	6951
	id	6953
	location (Virtual Chassis)	6954
	mac-persistence-timer	6955
	mastership-priority	6956
	member	6958
	no-management-vlan	6959
	no-split-detection	6960
	package-name	6961
	preprovisioned	6962
	role	6963
	serial-number	6966
	serial-number (Virtual Chassis aliases)	6967
	traceoptions (Virtual Chassis)	6968
	vcp-no-hold-time	6971

	virtual-chassis . . . . .	6973
<b>Chapter 82</b>	<b>Administration . . . . .</b>	<b>6975</b>
	Routine Monitoring . . . . .	6975
	Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member . . . . .	6975
	Operational Commands . . . . .	6976
	clear virtual-chassis vc-port statistics . . . . .	6978
	request session member . . . . .	6980
	request virtual-chassis recycle . . . . .	6981
	request virtual-chassis renumber . . . . .	6982
	request virtual-chassis vc-port . . . . .	6983
	show virtual-chassis active-topology . . . . .	6985
	show virtual-chassis device-topology . . . . .	6990
	show virtual-chassis protocol adjacency . . . . .	6996
	show virtual-chassis protocol database . . . . .	7000
	show virtual-chassis protocol interface . . . . .	7004
	show virtual-chassis protocol route . . . . .	7007
	show virtual-chassis protocol statistics . . . . .	7010
	show virtual-chassis login . . . . .	7013
	show virtual-chassis . . . . .	7014
	show virtual-chassis vc-path . . . . .	7018
	show virtual-chassis vc-port . . . . .	7020
	show virtual-chassis vc-port statistics . . . . .	7024
<b>Part 23</b>	<b>Virtual Chassis Fabric</b>	
<b>Chapter 83</b>	<b>Overview . . . . .</b>	<b>7033</b>
	Virtual Chassis Fabric Overview . . . . .	7033
	Virtual Chassis Fabric Overview . . . . .	7033
	Understanding Virtual Chassis Fabric Components . . . . .	7035
	Spine-and-Leaf Topology . . . . .	7035
	Spine Devices . . . . .	7036
	Leaf Devices . . . . .	7036
	Routing Engine Role . . . . .	7037
	Linecard Role . . . . .	7038
	Master Routing Engine Election Process . . . . .	7038
	Virtual Chassis Ports (VCPs) . . . . .	7039
	Automatic Virtual Chassis Port (VCP) Conversion . . . . .	7039
	VCF Configuration Options . . . . .	7040
	Fabric Mode . . . . .	7040
	Mixed Mode . . . . .	7041
	Virtual Management Ethernet Interface . . . . .	7041
	Virtual Chassis Fabric Port Link Aggregation Group Bundles . . . . .	7041
	Virtual Chassis Fabric License Requirements . . . . .	7042
	Hardware Requirements for a Virtual Chassis Fabric . . . . .	7042
	Software Requirements in a Virtual Chassis Fabric . . . . .	7042
	Understanding Virtual Chassis Fabric Configuration . . . . .	7043
	Virtual Chassis Fabric Setup . . . . .	7043
	Configuration File Management in a VCF . . . . .	7045



Logging into a Virtual Chassis Fabric . . . . .	7045
Understanding Interface Numbering . . . . .	7045
Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric . . . . .	7046
Virtual Chassis Fabric Summary . . . . .	7046
Understanding Mixed Virtual Chassis Fabric . . . . .	7047
Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches . . . . .	7047
Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches . . . . .	7048
Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis . . . . .	7049
Understanding Mixed EX4300 and EX4600 Virtual Chassis . . . . .	7049
Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis . . . . .	7049
Understanding Traffic Flow Through a Virtual Chassis Fabric . . . . .	7050
Understanding Software Upgrades in a Virtual Chassis Fabric . . . . .	7050
Virtual Chassis Fabric Software Basics . . . . .	7051
Nonstop Software Upgrade (NSSU) . . . . .	7051
Automatic Software Update . . . . .	7051
Traditional Software Upgrade . . . . .	7051
<b>Chapter 84 Configuration . . . . .</b>	<b>7053</b>
Configuration Tasks . . . . .	7053
Autoprovisioning a Virtual Chassis Fabric . . . . .	7053
Preprovisioning a Virtual Chassis Fabric . . . . .	7056
Configuring a Nonprovisioned Virtual Chassis Fabric . . . . .	7059
Adding a Device to a Virtual Chassis Fabric . . . . .	7062
Adding a Leaf Device to an Autoprovisioned Virtual Chassis Fabric . . . . .	7063
Adding a Spine Device to an Autoprovisioned Virtual Chassis Fabric . . . . .	7064
Adding a Spine or Leaf Device to a Preprovisioned Virtual Chassis Fabric . . . . .	7065
Adding a Spine or Leaf Device to a Nonprovisioned Virtual Chassis Fabric . . . . .	7067
Removing a Device From a Virtual Chassis Fabric . . . . .	7069
Upgrading Software for a Virtual Chassis Fabric . . . . .	7070
NSSU . . . . .	7071
Automatic Software Update . . . . .	7071
Standard Upgrade . . . . .	7071
Configuration Statements . . . . .	7072
[edit virtual-chassis] Configuration Statement Hierarchy . . . . .	7072
Supported Statements in the [edit virtual-chassis] Hierarchy Level . . . . .	7073
Unsupported Statements in the [edit virtual-chassis] Hierarchy Level . . . . .	7074
aliases (Virtual Chassis) . . . . .	7075
alias-name (Virtual Chassis aliases) . . . . .	7076
auto-provisioned . . . . .	7077
auto-sw-update . . . . .	7078

	enhanced-hash-key . . . . .	7080
	fabric-load-balance . . . . .	7082
	id . . . . .	7083
	inactivity-interval (Fabric Load Balance) . . . . .	7084
	location (Virtual Chassis) . . . . .	7085
	mac-persistence-timer . . . . .	7086
	mastership-priority . . . . .	7087
	member . . . . .	7089
	no-management-vlan . . . . .	7090
	no-split-detection . . . . .	7091
	package-name . . . . .	7092
	preprovisioned . . . . .	7093
	role . . . . .	7094
	serial-number . . . . .	7097
	serial-number (Virtual Chassis aliases) . . . . .	7098
	traceoptions (Virtual Chassis) . . . . .	7099
	virtual-chassis . . . . .	7102
<b>Chapter 85</b>	<b>Administration . . . . .</b>	<b>7105</b>
	Routine Monitoring . . . . .	7105
	Verifying the Member ID, Role, Status, and Neighbor Member Connections of a Virtual Chassis Fabric Member Device . . . . .	7105
	Verifying Virtual Chassis Port Connections in a Virtual Chassis Fabric . . . .	7106
	Verifying the Virtual Chassis Fabric Mode Settings . . . . .	7107
	Operational Commands . . . . .	7107
	clear virtual-chassis vc-port statistics . . . . .	7109
	request session member . . . . .	7111
	request virtual-chassis mode . . . . .	7112
	request virtual-chassis reactivate . . . . .	7115
	request virtual-chassis vc-port . . . . .	7116
	request virtual-chassis vc-port diagnostics optics . . . . .	7118
	show forwarding-options enhanced-hash-key . . . . .	7119
	show virtual-chassis active-topology . . . . .	7122
	show virtual-chassis device-topology . . . . .	7127
	show virtual-chassis login . . . . .	7133
	show virtual-chassis mode . . . . .	7134
	show virtual-chassis protocol adjacency . . . . .	7137
	show virtual-chassis protocol database . . . . .	7141
	show virtual-chassis protocol interface . . . . .	7145
	show virtual-chassis protocol route . . . . .	7148
	show virtual-chassis protocol statistics . . . . .	7151
	show virtual-chassis . . . . .	7154
	show virtual-chassis vc-port . . . . .	7158
	show virtual-chassis vc-port diagnostics optics . . . . .	7162
	show virtual-chassis vc-port statistics . . . . .	7176
<b>Chapter 86</b>	<b>Troubleshooting Procedures . . . . .</b>	<b>7183</b>
	Troubleshooting Virtual Chassis Fabric . . . . .	7183
	Virtual Chassis Port Link Does Not Form . . . . .	7183
	QFX5100 Leaf Device Assumes Routing Engine Role . . . . .	7184

<b>Part 24</b>	<b>Troubleshooting</b>	
<b>Chapter 87</b>	<b>Overview</b>	<b>7187</b>
	General Troubleshooting	7187
	Understanding Troubleshooting Resources	7187
	Troubleshooting Overview	7189
	Alarms	7191
	Understanding Alarms	7191
	Chassis Alarm Messages on a QFX3500 Device	7192
	Interface Alarm Messages	7195
	System Utilization Alarms	7195
<b>Chapter 88</b>	<b>Administration</b>	<b>7197</b>
	Routine Monitoring Using the CLI	7197
	Monitoring SNMP	7197
	Tracing SNMP Activity on a Device Running Junos OS	7199
	Configuring the Number and Size of SNMP Log Files	7200
	Configuring Access to the Log File	7200
	Configuring a Regular Expression for Lines to Be Logged	7200
	Configuring the Trace Operations	7200
	Monitoring RMON MIB Tables	7202
	Displaying a Log File from a Single-Chassis System	7203
	Monitoring System Log Messages	7204
	Monitoring Traffic Through the Router or Switch	7205
	Displaying Real-Time Statistics About All Interfaces on the Router or Switch	7205
	Displaying Real-Time Statistics About an Interface on the Router or Switch	7206
	Pinging Hosts	7207
<b>Chapter 89</b>	<b>Troubleshooting</b>	<b>7209</b>
	Configuration and File Management	7209
	Loading a Previous Configuration File	7209
	Reverting to the Default Factory Configuration	7210
	Reverting to the Rescue Configuration	7211
	Cleaning Up the System File Storage Space	7211
	Ethernet Switching	7212
	Troubleshooting Ethernet Switching	7212
	Troubleshooting Layer 2 Protocol Tunneling	7213
	Drop Threshold Statistics Might Be Incorrect	7213
	Egress Filtering of L2PT Traffic Not Supported	7213
	Troubleshooting Private VLANs	7214
	Limitations of Private VLANs	7214
	Forwarding with Private VLANs	7214
	Egress Firewall Filters with Private VLANs	7215
	Egress Port Mirroring with Private VLANs	7216
	Troubleshooting Q-in-Q and VLAN Translation Configuration	7217
	Firewall Filter Match Condition Not Working with Q-in-Q Tunneling	7217
	Egress Port Mirroring with VLAN Translation	7217

Hardware .....	7217
Troubleshooting QFX3100 Director Device Isolation .....	7218
High Availability .....	7219
Troubleshooting VRRP .....	7219
Interfaces .....	7220
Troubleshooting an Aggregated Ethernet Interface .....	7220
Troubleshooting Network Interfaces .....	7220
The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down .....	7220
Troubleshooting Multichassis Link Aggregation .....	7221
MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table .....	7221
MC-LAG Peer Does Not Go into Standby Mode .....	7222
Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive .....	7222
Redirect Filters Take Priority over User-Defined Filters .....	7222
Operational Command Output Is Wrong .....	7223
ICCP Connection Might Take Up to 60 Seconds to Become Active ..	7223
MAC Address Age Learned on an MC-AE Interface Is Reset to Zero ..	7223
MAC Address Is Not Learned Remotely in a Default VLAN .....	7224
Snooping Entries Learned on MC-AE Interfaces Are Not Removed ..	7224
ICCP Does Not Come Up After You Add or Delete an Authentication Key .....	7224
Local Status Is Standby When It Should Be Active .....	7224
Packets Loop on the Server When ICCP Fails .....	7224
Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change .....	7225
No Commit Checks Are Done for ICL-PL Interfaces .....	7225
Double Failover Scenario .....	7225
Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up .....	7225
Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer .....	7226
AE Interfaces Go Down .....	7226
Flooding of Upstream Traffic .....	7226

Junos OS Basics .....	7226
Rebooting and Halting a Device .....	7227
Recovering from a Failed Software Installation .....	7228
Recovering the Root Password .....	7229
Creating an Emergency Boot Device .....	7230
Performing a Recovery Installation .....	7232
Performing a QFabric System Recovery Installation on the Director Group .....	7233
(Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive .....	7234
Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software .....	7236
Troubleshooting Network Interfaces .....	7240
The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down .....	7240
Troubleshooting an Aggregated Ethernet Interface .....	7240
Layer 3 Protocols .....	7241
Troubleshooting Virtual Routing Instances .....	7241
Direct Routes Not Leaked Between Routing Instances .....	7241
MPLS .....	7242
Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch .....	7242
Network Management .....	7242
Understanding Troubleshooting Resources .....	7242
Troubleshooting Overview .....	7244
QFX5100 Switch with Automation Enhancements Frequently Asked Questions .....	7247
Who Should You Contact If You Have Problems with Loading, Installing or Updating Libraries? .....	7247
Who Should You Contact If You Have Problems with Puppet for Junos OS? .....	7247
Who Should You Contact If You Have Problems with Chef for Junos OS? .....	7247
What Happens to the User Partition If You Downgrade a QFX5100 Switch That Is Running the jinstall-qfx-5-flex-x.tgz Software Bundle to a QFX Switch That Is Running a Different QFX5100 Software Bundle? .....	7247
How Do You Recover Junos OS Binaries That You Have Deleted? . . .	7247
How Do You Recover from a System Crash? .....	7247
How Can You Verify That a QFX5100 Switch Is Running a jinstall-qfx-5-flex-x.tgz Software Bundle? .....	7248
Recovering from a Failed Software Installation .....	7248
Loading a Previous Configuration File .....	7249
Reverting to the Default Factory Configuration .....	7250
Reverting to the Rescue Configuration .....	7250
Recovering the Root Password .....	7251
Troubleshooting a Deprecated Network Analytics Configuration .....	7252

Security .....	7253
Troubleshooting Firewall Filter Configuration .....	7253
Firewall Filter Configuration Returns a No Space Available in TCAM Message .....	7253
Filter Counts Previously Dropped Packet .....	7255
Matching Packets Not Counted .....	7255
Counter Reset When Editing Filter .....	7256
Cannot Include loss-priority and policer Actions in Same Term .....	7256
Cannot Egress Filter Certain Traffic Originating on QFX Switch .....	7256
Firewall Filter Match Condition Not Working with Q-in-Q Tunneling ..	7257
Egress Firewall Filters with Private VLANs .....	7257
Egress Filtering of L2PT Traffic Not Supported .....	7258
Cannot Drop BGP Packets in Certain Circumstances .....	7258
Invalid Statistics for Policer .....	7258
Policers can Limit Egress Filters .....	7258
Troubleshooting Policer Configuration .....	7259
Incomplete Count of Packet Drops .....	7260
Counter Reset When Editing Filter .....	7260
Invalid Statistics for Policer .....	7260
Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured .....	7260
Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured .....	7261
Policers Can Limit Egress Filters .....	7262
Services .....	7263
Troubleshooting Port Mirroring .....	7263
Port Mirroring Constraints and Limitations .....	7263
Egress Port Mirroring with VLAN Translation .....	7265
Egress Port Mirroring with Private VLANs .....	7265
Traffic Management .....	7266
Troubleshooting Dropped FCoE Traffic .....	7266
Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth .....	7269
Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth .....	7270
Troubleshooting Egress Queue Bandwidth Impacted by Congestion .....	7271
Troubleshooting an Unexpected Rewrite Value .....	7272
Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic .....	7273
Virtual Chassis Fabric .....	7275
Troubleshooting Virtual Chassis Fabric .....	7275
Virtual Chassis Port Link Does Not Form .....	7276
QFX5100 Leaf Device Assumes Routing Engine Role .....	7276

# List of Figures

<b>Part 1</b>	<b>QFX5100 Switch Overview</b>	
<b>Chapter 1</b>	<b>QFX5100 Switch Overview</b>	<b>3</b>
	Figure 1: QFX5100-48S Switch	4
	Figure 2: QFX5100-48T Switch	4
	Figure 3: QFX5100-24Q Switch	5
	Figure 4: QFX-EM-4Q Expansion Module	6
	Figure 5: EX4600-EM-8F Expansion Module	6
	Figure 6: QFX5100-96S Switch	7
<b>Part 2</b>	<b>Junos OS Basics</b>	
<b>Chapter 2</b>	<b>Overview</b>	<b>11</b>
	Figure 7: DHCP Client/Server Model	21
	Figure 8: DHCP Four-Step Transfer	24
	Figure 9: Commands That Combine Other Commands	57
	Figure 10: CLI Command Hierarchy	61
	Figure 11: Command Output Options	62
	Figure 12: Configuration Mode Hierarchy of Statements	67
<b>Part 3</b>	<b>Configuration and File Management</b>	
<b>Chapter 8</b>	<b>Configuration</b>	<b>1245</b>
	Figure 13: Overriding the Current Configuration	1277
	Figure 14: Using the replace Option	1278
	Figure 15: Using the merge Option	1278
	Figure 16: Using a Patch File	1279
	Figure 17: Using the set Option	1279
<b>Part 5</b>	<b>Ethernet Features</b>	
<b>Chapter 14</b>	<b>Overview</b>	<b>1511</b>
	Figure 18: Protocol Packets from the Network to the Router	1536
	Figure 19: Protocol Packets from the Router or Switch to the Network	1536
<b>Chapter 15</b>	<b>Configuration</b>	<b>1563</b>
	Figure 20: Ethernet Ring Protection Switching Example	1565
	Figure 21: IRB with One Switch	1577
	Figure 22: Reflective Relay Topology	1607
	Figure 23: Network Topology for RSTP	1611
	Figure 24: Network Topology for MSTP	1625
	Figure 25: BPDU Protection Topology	1657

	Figure 26: Network Topology for Loop Protection . . . . .	1662
	Figure 27: Network Topology for Root Protection . . . . .	1666
<b>Part 6</b>	<b>OVSDB and VXLAN</b>	
<b>Chapter 18</b>	<b>Overview . . . . .</b>	<b>1899</b>
	Figure 28: High-Level NSX for Multi-Hypervisor Architecture . . . . .	1901
	Figure 29: Integration of Juniper Networks Device That Implements VXLAN and OVSDB into NSX for Multi-Hypervisor Environment . . . . .	1901
	Figure 30: VXLAN Packet Format . . . . .	1913
<b>Chapter 19</b>	<b>Configuration . . . . .</b>	<b>1917</b>
	Figure 31: VXLAN/OVSDB Layer 2 Gateway Topology . . . . .	1919
	Figure 32: Inter-VXLAN Routing and OVSDB Topology . . . . .	1927
	Figure 33: QFX5100 Acting as a VXLAN Transit Switch . . . . .	1945
	Figure 34: QFX5100 Acting as a VTEP . . . . .	1947
<b>Part 7</b>	<b>OpenFlow</b>	
<b>Chapter 26</b>	<b>Configuring OpenFlow Traffic Steering Across MPLS Networks . . . . .</b>	<b>2149</b>
	Figure 35: Connecting OpenFlow Networks Using MPLS LSP Tunnel Cross-Connects . . . . .	2149
	Figure 36: Connecting OpenFlow Networks Using MPLS Tunnel Cross-Connects . . . . .	2152
<b>Part 8</b>	<b>High Availability</b>	
<b>Chapter 29</b>	<b>Overview . . . . .</b>	<b>2231</b>
	Figure 37: Preparing for a Graceful Routing Engine Switchover . . . . .	2233
	Figure 38: Graceful Routing Engine Switchover Process . . . . .	2234
	Figure 39: Nonstop Active Routing Switchover Preparation Process . . . . .	2241
	Figure 40: Nonstop Active Routing During a Switchover . . . . .	2242
	Figure 41: Nonstop Bridging Switchover Preparation Process . . . . .	2255
	Figure 42: Nonstop Bridging During a Switchover . . . . .	2256
	Figure 43: Basic VRRP Topology . . . . .	2260
<b>Chapter 30</b>	<b>Configuration . . . . .</b>	<b>2261</b>
	Figure 44: VRRP Load-Sharing Configuration . . . . .	2280
<b>Part 9</b>	<b>Interfaces</b>	
<b>Chapter 33</b>	<b>Overview . . . . .</b>	<b>2389</b>
	Figure 45: Uplink Failure Detection Configuration on Switches . . . . .	2392
	Figure 46: Egress Traffic Flow with Local Link Bias . . . . .	2409
	Figure 47: Egress Traffic Flow without Local Link Bias . . . . .	2409
	Figure 48: Redundant Trunk Group, Link 1 Active . . . . .	2448
	Figure 49: Redundant Trunk Group, Link 2 Active . . . . .	2448
<b>Chapter 34</b>	<b>Configuration . . . . .</b>	<b>2457</b>
	Figure 50: Uplink Failure Detection Configuration on Switches . . . . .	2459
	Figure 51: Configuring a Multichassis LAG Between Switch A and Switch B . . . . .	2472



	Figure 52: Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP . . .	2495
	Figure 53: Configuring a Multichassis LAG Between Switch A and Switch B . . .	2531
	Figure 54: Configuring a Multichassis LAG Between Switch A and Switch B . . .	2553
	Figure 55: Topology for Configuring the Redundant Trunk Links . . . . .	2580
<b>Part 10</b>	<b>Routing Options</b>	
<b>Chapter 38</b>	<b>Configuration . . . . .</b>	<b>2903</b>
	Figure 56: Customer Routes Connected to a Service Provider . . . . .	2919
	Figure 57: BFD Enabled on Qualified Next Hops . . . . .	2924
	Figure 58: Customer Routes Connected to a Service Provider . . . . .	2932
<b>Part 11</b>	<b>Border Gateway Protocol</b>	
<b>Chapter 41</b>	<b>Overview . . . . .</b>	<b>3253</b>
	Figure 59: ASs, EBGp, and IBGP . . . . .	3255
<b>Chapter 42</b>	<b>Configuration . . . . .</b>	<b>3261</b>
	Figure 60: BGP Peering Session . . . . .	3262
	Figure 61: Typical Network with BGP Peer Sessions . . . . .	3263
	Figure 62: Typical Network with BGP Peer Sessions . . . . .	3270
	Figure 63: Internal and External BGP . . . . .	3284
	Figure 64: Typical Network with IBGP Sessions . . . . .	3287
	Figure 65: Typical Network with IBGP Sessions . . . . .	3297
	Figure 66: Typical Network with IBGP Sessions and Multiple Exit Points . . . . .	3311
	Figure 67: Default MED Example . . . . .	3324
	Figure 68: Typical Network with IBGP Sessions and Multiple Exit Points . . . . .	3327
	Figure 69: Typical Network with IBGP Sessions and Multiple Exit Points . . . . .	3339
	Figure 70: Topology for Delaying the MED Update . . . . .	3354
	Figure 71: Local AS Configuration . . . . .	3365
	Figure 72: Topology for Configuring the Local AS . . . . .	3368
	Figure 73: Topology for Configuring a Private Local AS . . . . .	3378
	Figure 74: Advertisement of Multiple Paths in BGP . . . . .	3385
	Figure 75: BGP Prefix-Based Outbound Route Filtering . . . . .	3430
	Figure 76: Typical Network with EBGp Multihop Sessions . . . . .	3434
	Figure 77: BGP Preference Value Topology . . . . .	3445
	Figure 78: Topology for Ignoring the AS-Path Length . . . . .	3450
	Figure 79: Topology for Removing a Private AS from the Advertised AS Path . . . . .	3457
	Figure 80: Typical Network with IBGP Sessions . . . . .	3464
	Figure 81: BGP Load Balancing . . . . .	3479
	Figure 82: Topology for Accepting a Remote Next Hop . . . . .	3484
	Figure 83: Advertisement of Multiple Paths in BGP . . . . .	3496
	Figure 84: Advertisement of Multiple Paths in BGP . . . . .	3521
	Figure 85: Simple Route Reflector Topology (One Cluster) . . . . .	3548
	Figure 86: Basic Route Reflection (Multiple Clusters) . . . . .	3548
	Figure 87: Hierarchical Route Reflection (Clusters of Clusters) . . . . .	3549
	Figure 88: IBGP Network Using a Route Reflector . . . . .	3551
	Figure 89: BGP Confederations . . . . .	3565
	Figure 90: Typical Network Using BGP Confederations . . . . .	3566
	Figure 91: Authentication for BGP . . . . .	3573

	Figure 92: Typical Network with BGP Peer Sessions . . . . .	3578
	Figure 93: TCP Maximum Segment Size for BGP . . . . .	3587
	Figure 94: Topology for the EBGP Case . . . . .	3595
	Figure 95: Topology for the RR Case . . . . .	3595
	Figure 96: BGP Flap Damping Topology . . . . .	3601
	Figure 97: MBGP MVPN with BGP Route Flap Damping . . . . .	3609
<b>Part 12</b>	<b>Intermediate System to Intermediate System</b>	
<b>Chapter 44</b>	<b>Overview . . . . .</b>	<b>3811</b>
	Figure 98: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2 . . . . .	3816
	Figure 99: Link Protection and Node-Link Protection Comparison for IS-IS Routes . . . . .	3820
<b>Chapter 45</b>	<b>Configuration . . . . .</b>	<b>3823</b>
	Figure 100: Simple IS-IS Topology . . . . .	3824
	Figure 101: IS-IS Multi-Level Topology . . . . .	3830
	Figure 102: Hitless Authentication Key Rollover for IS-IS . . . . .	3839
	Figure 103: IS-IS Route Redistribution Topology . . . . .	3843
	Figure 104: Configuring BFD for IS-IS . . . . .	3850
	Figure 105: IS-IS BFD Authentication Topology . . . . .	3856
	Figure 106: IS-IS IPv4 and IPv6 Unicast Topologies . . . . .	3861
	Figure 107: Configuring IS-IS Multicast Topology . . . . .	3870
	Figure 108: Link Protection and Node-Link Protection Comparison for IS-IS Routes . . . . .	3885
	Figure 109: IS-IS Node-Link Protection Topology . . . . .	3888
	Figure 110: IS-IS Logical Systems with a Default Route to an ISP . . . . .	3898
	Figure 111: IS-IS Checksum Topology . . . . .	3907
<b>Part 13</b>	<b>Open Shortest Path First</b>	
<b>Chapter 47</b>	<b>Overview . . . . .</b>	<b>4035</b>
	Figure 112: OSPF Three-Way Handshake . . . . .	4039
<b>Chapter 48</b>	<b>Configuration . . . . .</b>	<b>4047</b>
	Figure 113: Multiarea OSPF Topology . . . . .	4052
	Figure 114: Typical Single-Area OSPF Network Topology . . . . .	4054
	Figure 115: Typical Multiarea OSPF Network Topology . . . . .	4056
	Figure 116: OSPF AS Network with Stub Areas and NSSAs . . . . .	4059
	Figure 117: OSPF Network Topology with Stub Areas and NSSAs . . . . .	4062
	Figure 118: OSPF Network Topology with Stub Areas and NSSAs . . . . .	4066
	Figure 119: IPv4 Unicast Realm . . . . .	4092
	Figure 120: Summarizing Ranges of Routes in OSPF . . . . .	4095
	Figure 121: OSPF Metric Configuration . . . . .	4106
	Figure 122: Sample Topology Used for an OSPF Export Network Summary Policy . . . . .	4182
	Figure 123: Sample Topology Used for an OSPF Import Network Summary Policy . . . . .	4191

<b>Part 14</b>	<b>Routing Information Protocol</b>	
<b>Chapter 50</b>	<b>Overview</b>	<b>4305</b>
	Figure 124: Distance-Vector Protocol	4306
	Figure 125: Split Horizon Example	4308
	Figure 126: Poison Reverse Example	4309
	Figure 127: Limitations of Unidirectional Connectivity	4310
<b>Chapter 51</b>	<b>Configuration</b>	<b>4311</b>
	Figure 128: Sample RIP Network Topology	4312
	Figure 129: RIP Authentication Network Topology	4319
	Figure 130: RIP BFD Network Topology	4327
	Figure 131: RIP BFD Authentication Network Topology	4333
	Figure 132: RIP Import Policy Network Topology	4339
	Figure 133: Controlling Traffic in a RIP Network with the Incoming Metric	4345
	Figure 134: Controlling Traffic in a RIP Network with the Outgoing Metric	4347
	Figure 135: RIP Incoming Metrics Network Topology	4348
	Figure 136: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology	4353
	Figure 137: Redistributing Routes Between RIP Instances Network Topology	4357
	Figure 138: RIP Timers Network Topology	4363
	Figure 139: RIP Trace Operations Network Topology	4370
<b>Part 15</b>	<b>MPLS Applications</b>	
<b>Chapter 53</b>	<b>Overview</b>	<b>4411</b>
	Figure 140: Label Encoding	4416
	Figure 141: MPLS Label Swapping	4417
	Figure 142: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs	4429
	Figure 143: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs	4429
<b>Chapter 54</b>	<b>Configuration</b>	<b>4433</b>
	Figure 144: MPLS-Based Layer 3 VPN	4446
	Figure 145: IPv6 Networks Linked by MPLS IPv4 Tunnels	4455
<b>Part 16</b>	<b>Multicast</b>	
<b>Chapter 57</b>	<b>Overview</b>	<b>4737</b>
	Figure 146: Rendezvous Point as Part of the RPT and SPT	4743
	Figure 147: Building an RPT Between the RP and the Receiver	4750
	Figure 148: PIM Register Message and PIM Join Message Exchanged	4751
	Figure 149: Traffic Sent from the Source to the RP Router	4752
	Figure 150: Traffic Sent from the RP Router Toward the Receiver	4752
	Figure 151: Receiver DR Sends a PIM Join Message to the Source	4754
	Figure 152: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router	4755
	Figure 153: RP Router Receives PIM Prune Message	4755
	Figure 154: RP Router Sends a PIM Prune Message to the Source DR	4756

	Figure 155: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router . . . . .	4756
	Figure 156: Routing Devices Start Up on a Subnet . . . . .	4765
	Figure 157: Querier Routing Device Is Determined . . . . .	4766
	Figure 158: General Query Message Is Issued . . . . .	4766
	Figure 159: Reports Are Received by the Querier Routing Device . . . . .	4766
	Figure 160: Host Has No Interested Receivers and Sends a Done Message to Routing Device . . . . .	4767
	Figure 161: Host Address Timer Expires and Address Is Removed from Multicast Address List . . . . .	4767
	Figure 162: Receiver Announces Desire to Join Group G and Source S . . . . .	4772
	Figure 163: Router 3 (Last-Hop Router) Joins the Source Tree . . . . .	4772
	Figure 164: (S,G) State Is Built Between the Source and the Receiver . . . . .	4773
<b>Chapter 58</b>	<b>Configuration . . . . .</b>	<b>4777</b>
	Figure 165: Join Suppression . . . . .	4792
	Figure 166: PIM Assert Topology . . . . .	4817
	Figure 167: Routing Devices Start Up on a Subnet . . . . .	4853
	Figure 168: Querier Routing Device Is Determined . . . . .	4854
	Figure 169: General Query Message Is Issued . . . . .	4854
	Figure 170: Reports Are Received by the Querier Routing Device . . . . .	4855
	Figure 171: Host Has No Interested Receivers and Sends a Done Message to Routing Device . . . . .	4855
	Figure 172: Host Address Timer Expires and Address Is Removed from Multicast Address List . . . . .	4855
	Figure 173: Source-Active Message Flooding . . . . .	4883
	Figure 174: Network on Which to Configure PIM SSM . . . . .	4891
	Figure 175: Receiver Sends Messages to Join Group G and Source S . . . . .	4896
	Figure 176: Router 3 (Last-Hop Router) Joins the Source Tree . . . . .	4896
	Figure 177: (S,G) State Is Built Between the Source and the Receiver . . . . .	4897
	Figure 178: Simple RPF Topology . . . . .	4897
<b>Part 17</b>	<b>Security</b>	
<b>Chapter 60</b>	<b>Overview . . . . .</b>	<b>5209</b>
	Figure 179: Evaluation of Terms Within a Firewall Filter . . . . .	5213
	Figure 180: Application of Firewall Filters to Control Packet Flow . . . . .	5215
	Figure 181: Flow of Tricolor Marking Policer Operation . . . . .	5242
	Figure 182: DHCP Server Connected Directly to a Switching Device . . . . .	5260
	Figure 183: DHCP Server Connected Directly to Switching Device 2, with Switching Device 2 Connected to Switching Device 1 Through a Trusted Trunk Port . . . . .	5260
	Figure 184: Switching Device Is the DHCP Server . . . . .	5261
	Figure 185: Switching Device Acting as Relay Agent Through Router to DHCP Server . . . . .	5262
	Figure 186: Switch Relays DHCP Requests to Server . . . . .	5271
	Figure 187: Symmetrically Routed Interfaces . . . . .	5275
	Figure 188: Asymmetrically Routed Interfaces . . . . .	5276

<b>Part 18</b>	<b>Services</b>	
<b>Chapter 65</b>	<b>Configuration</b>	<b>5433</b>
	Figure 189: Network Topology for Local Port Mirroring Example	5434
<b>Part 19</b>	<b>Storage</b>	
<b>Chapter 68</b>	<b>Overview</b>	<b>5507</b>
	Figure 190: ENode Components	5520
	Figure 191: FCoE Transit Switch Connecting FCoE Devices to an FC Switch	5526
	Figure 192: FCoE Transit Switch Performs VN2VF_Port FIP Snooping	5533
	Figure 193: VN2VN_Port Traffic Across a QFabric Interconnect Device	5545
	Figure 194: Supported Topology for an MC-LAG on an FCoE Transit Switch	5556
<b>Chapter 69</b>	<b>Configuration</b>	<b>5595</b>
	Figure 195: PFC for FCoE Traffic Configuration Components Block Diagram	5608
	Figure 196: Supported Topology for an MC-LAG on an FCoE Transit Switch	5616
	Figure 197: VN2VN_Port FIP Snooping (FCoE Hosts Connected to Same Transit Switch) Topology	5638
	Figure 198: VN2VN_Port FIP Snooping (FCoE Hosts Connected to Different Transit Switches) Topology	5643
	Figure 199: VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected) Topology	5650
<b>Part 20</b>	<b>Traffic Management</b>	
<b>Chapter 72</b>	<b>Overview</b>	<b>5779</b>
	Figure 200: Packet Flow Across the Network	5783
	Figure 201: Flow of Tricolor Marking Policer Operation	5784
	Figure 202: CoS Classifier, Queues, and Scheduler	5794
	Figure 203: Packet Flow Through Configurable CoS Components	5794
	Figure 204: Hierarchical Scheduling Tiers	5864
	Figure 205: Hierarchical Scheduling Packet Flow	5866
	Figure 206: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling	5882
	Figure 207: Setting Maximum Bandwidth Using Hierarchical Scheduling	5886
	Figure 208: WRED-Drop Profile Packet Drop	5910
	Figure 209: Explicit Congestion Notification	5928
	Figure 210: Enabling PFC Across Layer 3 Interface Hops	5951
<b>Chapter 73</b>	<b>Configuration</b>	<b>5965</b>
	Figure 211: Hierarchical Port Scheduling Components Block Diagram	5970
	Figure 212: Hierarchical Port Scheduling Packet Flow Block Diagram	5970
	Figure 213: PFC for FCoE Traffic Configuration Components Block Diagram	5990
	Figure 214: Supported Topology for an MC-LAG on an FCoE Transit Switch	5998
	Figure 215: Topology of the Two Lossless FCoE Priorities Example	6038
	Figure 216: Topology of the Lossless FCoE and iSCSI Priorities Example	6052
	Figure 217: WRED Drop Profile Packet Drop Example	6072
	Figure 218: Enabling PFC Across Layer 3 Interface Hops	6140

<b>Part 21</b>	<b>Network Management and Monitoring</b>	
<b>Chapter 76</b>	<b>Overview</b>	<b>6463</b>
	Figure 219: Commit Script Input and Output	6475
	Figure 220: Standard Commit Model	6478
	Figure 221: Commit Model with Commit Scripts Added	6478
	Figure 222: Configuration Evaluation by Multiple Commit Scripts	6480
	Figure 223: Op Script Input and Output	6487
	Figure 224: SNMP Communication Flow	6515
	Figure 225: Setting Thresholds	6526
<b>Chapter 77</b>	<b>Configuration</b>	<b>6565</b>
	Figure 226: sFlow Technology Monitoring System	6572
	Figure 227: Inform Request and Response	6615
<b>Part 22</b>	<b>Virtual Chassis</b>	
<b>Chapter 80</b>	<b>Overview</b>	<b>6907</b>
	Figure 228: Console Session Redirection (EX4200 Virtual Chassis Pictured)	6919
	Figure 229: Management Ethernet Port Redirection to the VME Interface	6920
<b>Part 23</b>	<b>Virtual Chassis Fabric</b>	
<b>Chapter 83</b>	<b>Overview</b>	<b>7033</b>
	Figure 230: VCF Spine-and-Leaf Architecture	7034
	Figure 231: VCF Spine-and-Leaf Architecture	7036

# List of Tables

	<b>About the Documentation</b> . . . . .	<b>CXXXV</b>
	Table 1: Notice Icons . . . . .	CXXXvii
	Table 2: Text and Syntax Conventions . . . . .	CXXXvii
<b>Part 2</b>	<b>Junos OS Basics</b>	
<b>Chapter 2</b>	<b>Overview</b> . . . . .	<b>11</b>
	Table 3: Configuration File Terms . . . . .	11
	Table 4: ISSU Protocol Support . . . . .	13
	Table 5: Legacy DHCP and Extended DHCP Server Hierarchy Levels . . . . .	23
	Table 6: Platform and Release Support for NSSU on a Virtual Chassis Fabric . . . . .	29
	Table 7: Junos OS Processes . . . . .	30
	Table 8: ELS Support . . . . .	44
	Table 9: Enhanced Layer 2 CLI Changes . . . . .	49
	Table 10: Commonly Used Operational Mode Commands . . . . .	59
	Table 11: Summary of Configuration Mode Commands . . . . .	64
	Table 12: Configuration Mode Top-Level Statements . . . . .	66
	Table 13: Junos OS Feature Licenses and Model Numbers for QFX Series Devices . . . . .	70
	Table 14: Upgrade Licenses for Enhancing Port Capacity . . . . .	73
	Table 15: Port Activation License Model for MX104 Routers . . . . .	74
<b>Chapter 3</b>	<b>Installation</b> . . . . .	<b>83</b>
	Table 16: Junos OS and Jloader Software Compatibility Matrix for the QFX3500 Switch and QFX3500 Node Device . . . . .	123
	Table 17: Junos OS and Jloader Software Compatibility Matrix for the QFX3008-I Interconnect Device . . . . .	123
	Table 18: Junos OS and Jloader Software Compatibility Matrix for the QFX3600-I Interconnect Device and QFX3600 Node Device . . . . .	124
	Table 19: Uboot Software Release and Jloader Software Compatibility Matrix . . . . .	124
<b>Chapter 4</b>	<b>Configuration</b> . . . . .	<b>151</b>
	Table 20: DHCP Client Settings . . . . .	154
	Table 21: Methods for Configuring Junos OS . . . . .	183
<b>Chapter 5</b>	<b>Administration</b> . . . . .	<b>333</b>
	Table 22: Summary of System Process Information Output Fields . . . . .	333
	Table 23: Summary of Key System Properties Output Fields . . . . .	334
	Table 24: request system storage cleanup Output Fields . . . . .	468
	Table 25: show chassis alarms Output Fields . . . . .	501
	Table 26: show chassis led Output Fields . . . . .	509

Table 27: show chassis environment Output Fields . . . . .	518
Table 28: show chassis environment fpc Output Fields . . . . .	578
Table 29: show chassis environment pem Output Fields . . . . .	603
Table 30: show chassis environment routing-engine Output Fields . . . . .	612
Table 31: show chassis fan Output Fields . . . . .	617
Table 32: show chassis firmware Output Fields . . . . .	631
Table 33: show chassis fpc Output Fields . . . . .	646
Table 34: Routing Engines Displaying DIMM Information . . . . .	679
Table 35: show chassis hardware Output Fields . . . . .	683
Table 36: show chassis in-service-upgrade Output Fields . . . . .	851
Table 37: show chassis lcd Output Fields . . . . .	857
Table 38: show chassis led Output Fields . . . . .	869
Table 39: show chassis location Output Fields . . . . .	880
Table 40: show chassis mac-addresses Output Fields . . . . .	884
Table 41: show chassis nonstop-upgrade Output Fields . . . . .	887
Table 42: show chassis pic Output Fields . . . . .	893
Table 43: show chassis routing-engine Output Fields . . . . .	908
Table 44: show chassis zones Output Fields . . . . .	928
Table 45: show cli Output Fields . . . . .	933
Table 46: show cli authorization Output Fields . . . . .	935
Table 47: show cli directory Output Fields . . . . .	939
Table 48: show cli history Output Fields . . . . .	940
Table 49: show interfaces diagnostics optics Output Fields . . . . .	942
Table 50: show ntp associations Output Fields . . . . .	951
Table 51: show ntp status Output Fields . . . . .	953
Table 52: show subscribers Output Fields . . . . .	959
Table 53: show system alarms Output Fields . . . . .	974
Table 54: show system buffers Output Fields . . . . .	995
Table 55: show system certificate Output Fields . . . . .	999
Table 56: show system commit Output Fields . . . . .	1001
Table 57: show system connections Output Fields . . . . .	1009
Table 58: show system core-dumps Output Fields . . . . .	1029
Table 59: show system directory-usage Output Fields . . . . .	1042
Table 60: show system license Output Fields . . . . .	1044
Table 61: show system processes Output Fields . . . . .	1058
Table 62: show system resource-cleanup processes Output Fields . . . . .	1081
Table 63: show system services service-deployment Output Fields . . . . .	1085
Table 64: show system storage Output Fields . . . . .	1132
Table 65: show system uptime Output Fields . . . . .	1139
Table 66: show system users Output Fields . . . . .	1144
Table 67: show system virtual-memory Output Fields . . . . .	1150
Table 68: traceroute Output Fields . . . . .	1223
Table 69: traceroute monitor Output Fields . . . . .	1226

## Part 3

### Chapter 7

## Configuration and File Management

Overview . . . . .	1239
--------------------	------

Table 70: Configuration File Terms . . . . .	1239
Table 71: Forms of the configure Command . . . . .	1240



<b>Chapter 8</b>	<b>Configuration</b> . . . . .	<b>1245</b>
	Table 72: Options for the load Command . . . . .	1262
<b>Chapter 9</b>	<b>Administration</b> . . . . .	<b>1281</b>
	Table 73: show system commit Output Fields . . . . .	1301
<b>Part 4</b>	<b>User and Access Management</b>	
<b>Chapter 11</b>	<b>Overview</b> . . . . .	<b>1315</b>
	Table 74: Junos OS Processes . . . . .	1316
	Table 75: Juniper Networks Vendor-Specific RADIUS Attributes . . . . .	1322
	Table 76: Juniper Networks Vendor-Specific TACACS+ Attributes . . . . .	1324
	Table 77: Login Class Permission Flags . . . . .	1326
	Table 78: Order of Authentication Attempts . . . . .	1331
	Table 79: Predefined System Login Classes . . . . .	1337
	Table 80: Configuration Mode Hierarchies—Common Regular Expression Operators . . . . .	1338
	Table 81: Common Regular Expression Operators to Allow or Deny Operational Mode Commands . . . . .	1339
	Table 82: Special Requirements for Plain-Text Passwords . . . . .	1340
<b>Chapter 12</b>	<b>Configuration</b> . . . . .	<b>1343</b>
	Table 83: Match Conditions . . . . .	1377
	Table 84: Actions for VSAs . . . . .	1378
<b>Chapter 13</b>	<b>Administration</b> . . . . .	<b>1475</b>
	Table 85: show ethernet-switching interfaces Output Fields . . . . .	1482
	Table 86: show lldp Output Fields . . . . .	1486
	Table 87: show lldp local-information Output Fields . . . . .	1491
	Table 88: show lldp neighbors Output Fields . . . . .	1493
	Table 89: show lldp statistics Output Fields . . . . .	1497
	Table 90: show route instance Output Fields . . . . .	1499
	Table 91: show snmp statistics Output Fields . . . . .	1503
<b>Part 5</b>	<b>Ethernet Features</b>	
<b>Chapter 14</b>	<b>Overview</b> . . . . .	<b>1511</b>
	Table 92: ELS Support . . . . .	1512
	Table 93: Enhanced Layer 2 CLI Changes . . . . .	1517
	Table 94: Sample IRB Values . . . . .	1539
	Table 95: Number of Supported IRBs/RVIs by Platform . . . . .	1540
	Table 96: Unified Forwarding Table Profiles . . . . .	1546
	Table 97: Example Host Table Combinations Using l2-profile-one . . . . .	1546
<b>Chapter 15</b>	<b>Configuration</b> . . . . .	<b>1563</b>
	Table 98: Components to Configure for This Example . . . . .	1565
	Table 99: Components of the Multiple VLAN Topology . . . . .	1577
	Table 100: Components of the Multiple VLAN Topology . . . . .	1584
	Table 101: Components of the Basic Bridging Configuration Topology . . . . .	1589
	Table 102: Components of the Topology for Configuring Reflective Relay . . . . .	1607
	Table 103: Topology for Configuring RSTP on the QFX Series . . . . .	1611

	Table 104: Topology for Configuring MSTP on the QFX Series . . . . .	1626
	Table 105: Components of the Topology for Connecting an Access Switch to a Distribution Switch . . . . .	1648
	Table 106: Components of the Topology for Configuring BPDU Protection on the QFX Series . . . . .	1658
	Table 107: Topology for Configuring Loop Protection on the QFX Series . . . . .	1662
	Table 108: Topology for Configuring Root Protection on the QFX Series . . . . .	1666
	Table 109: Unified Forwarding Table Profiles . . . . .	1698
	Table 110: Example LPM Table Combinations Using I2-and I3 Profiles With Junos OS 13.2X51-D10 and 13.2X52-D10 . . . . .	1700
	Table 111: LPM Table Combinations for I2 and I3 profiles With Junos OS 13.2X51-D15 . . . . .	1701
	Table 112: Unified Forwarding Table Profiles . . . . .	1724
	Table 113: Unsupported [edit vlans] Configuration Statements on EX Series Switches . . . . .	1763
<b>Chapter 16</b>	<b>Administration . . . . .</b>	<b>1827</b>
	Table 114: show ethernet-switching interfaces Output Fields . . . . .	1843
	Table 115: show ethernet-switching layer2-protocol-tunneling interface Output Fields . . . . .	1847
	Table 116: show ethernet-switching layer2-protocol-tunneling statistics Output Fields . . . . .	1850
	Table 117: show ethernet-switching layer2-protocol-tunneling vlan Output Fields . . . . .	1852
	Table 118: show ethernet-switching mac-learning-log Output Fields . . . . .	1854
	Table 119: show ethernet-switching mac-notification Output Fields . . . . .	1856
	Table 120: show ethernet-switching statistics aging Output Fields . . . . .	1858
	Table 121: show ethernet-switching statistics mac-learning Output Fields . . . .	1861
	Table 122: show ethernet-switching table Output Fields . . . . .	1864
	Table 123: show spanning-tree bridge Output Fields . . . . .	1870
	Table 124: show spanning-tree Interface Output Fields . . . . .	1875
	Table 125: show spanning-tree mstp configuration Output Fields . . . . .	1881
	Table 126: show spanning-tree statistics Output Fields . . . . .	1883
	Table 127: show vlans Output Fields . . . . .	1887
<b>Part 6</b>	<b>OVSDB and VXLAN</b>	
<b>Chapter 18</b>	<b>Overview . . . . .</b>	<b>1899</b>
	Table 128: OVSDB Support on Junos OS Devices . . . . .	1899
	Table 129: NSX Multi-Hypervisor Components and Products That Can Be Implemented . . . . .	1900
	Table 130: Summary of Configuration Tasks for Setting Up An OVSDB-Managed VXLAN on All Juniper Networks Devices Except QFX5100 Switches . . . . .	1907
	Table 131: Summary of Configuration Tasks for Setting Up An OVSDB-Managed VXLAN on QFX5100 Switches . . . . .	1908
	Table 132: OVSDB Schema Tables . . . . .	1910
<b>Chapter 19</b>	<b>Configuration . . . . .</b>	<b>1917</b>
	Table 133: Components of the Topology for Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections . . . . .	1921

	Table 134: Components of the Topology for Setting Up Inter-VXLAN Routing and OVSDb Connections in a Data Center . . . . .	1928
	Table 135: Create a Gateway in NSX Manager: Key Configurations . . . . .	1987
	Table 136: Create a Gateway Service in NSX Manager: Key Configurations . . . . .	1988
	Table 137: Create a Logical Switch Port in NSX Manager: Key Configurations . . . . .	1988
<b>Chapter 20</b>	<b>Administration . . . . .</b>	<b>2009</b>
	Table 138: show bridge mac-table Output fields . . . . .	2011
	Table 139: show ovssdb controller Output Fields . . . . .	2014
	Table 140: show ovssdb interface Output Fields . . . . .	2017
	Table 141: show ovssdb logical-switch Output Fields . . . . .	2020
	Table 142: show ovssdb mac Output Fields . . . . .	2022
	Table 143: show ovssdb statistics interface Output Fields . . . . .	2025
	Table 144: show ovssdb virtual-tunnel-end-point Output Fields . . . . .	2027
	Table 145: show vpls mac-table Output fields . . . . .	2029
	Table 146: show bridge mac-table Output fields . . . . .	2035
	Table 147: show vpls mac-table Output fields . . . . .	2038
<b>Part 7</b>	<b>OpenFlow</b>	
<b>Chapter 22</b>	<b>Overview . . . . .</b>	<b>2049</b>
	Table 148: OpenFlow v1.0 Support on Devices Running Junos OS . . . . .	2052
	Table 149: OpenFlow v1.3.1 Support on Devices Running Junos OS . . . . .	2053
	Table 150: OpenFlow Versions Negotiated Between the Controller and a Junos OS Device and the Numerical Value Associated with Each Version . . . . .	2057
	Table 151: OpenFlow Flow Elements . . . . .	2058
	Table 152: OpenFlow Flow Entry Timers . . . . .	2062
	Table 153: Junos OS Support for OpenFlow v1.0 Message Types . . . . .	2065
	Table 154: Junos OS Support for OpenFlow v1.0 Port Structure Flags . . . . .	2067
	Table 155: Junos OS Support for OpenFlow v1.0 Match Conditions . . . . .	2068
	Table 156: Junos OS Support for OpenFlow v1.0 Wildcards . . . . .	2069
	Table 157: Junos OS Support for OpenFlow v1.0 Flow Actions . . . . .	2069
	Table 158: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT_PACKET_OUT) . . . . .	2070
	Table 159: Junos OS Support for OpenFlow v1.0 Statistics . . . . .	2071
	Table 160: Junos OS Support for OpenFlow v1.0 Features . . . . .	2071
	Table 161: Junos OS Support for OpenFlow v1.0 Message Types . . . . .	2072
	Table 162: Junos OS Support for OpenFlow v1.0 Port Structure Flags . . . . .	2073
	Table 163: Junos OS Support for OpenFlow v1.0 Match Conditions . . . . .	2074
	Table 164: Junos OS Support for OpenFlow v1.0 Wildcards . . . . .	2075
	Table 165: Junos OS Support for OpenFlow v1.0 Flow Actions . . . . .	2076
	Table 166: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT_PACKET_OUT) . . . . .	2077
	Table 167: Junos OS Support for OpenFlow v1.0 Statistics . . . . .	2077
	Table 168: Junos OS Support for OpenFlow v1.0 Features . . . . .	2078
	Table 169: Junos OS Support for OpenFlow v1.0 Message Types . . . . .	2079
	Table 170: Junos OS Support for OpenFlow v1.0 Port Structure Flags . . . . .	2080
	Table 171: Junos OS Support for OpenFlow v1.0 Match Conditions . . . . .	2081
	Table 172: Junos OS Support for OpenFlow v1.0 Wildcards . . . . .	2084
	Table 173: Junos OS Support for OpenFlow v1.0 Flow Actions . . . . .	2085

	Table 174: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT_PACKET_OUT) . . . . .	2085
	Table 175: Junos OS Support for OpenFlow v1.0 Statistics . . . . .	2086
	Table 176: Junos OS Support for OpenFlow v1.0 Features . . . . .	2087
	Table 177: Junos OS Support for OpenFlow v1.3.1 Message Types . . . . .	2087
	Table 178: Junos OS Support for OpenFlow v1.3.1 Features Reply Messages . . . . .	2090
	Table 179: Junos OS Support for OpenFlow v1.3.1 Port Structure Flags . . . . .	2090
	Table 180: Junos OS Support for OpenFlow v1.3.1 Port Numbering . . . . .	2091
	Table 181: Junos OS Support for OpenFlow v1.3.1 Match Conditions . . . . .	2092
	Table 182: Junos OS Support for OpenFlow v1.3.1 Flow Actions . . . . .	2094
	Table 183: Junos OS Support for OpenFlow v1.3.1 Multipart Messages . . . . .	2094
	Table 184: Junos OS Support for OpenFlow v1.3.1 Flow Instructions . . . . .	2095
	Table 185: Junos OS Support for OpenFlow v1.3.1 Group Types . . . . .	2096
<b>Chapter 25</b>	<b>Configuring OpenFlow Hybrid Interfaces . . . . .</b>	<b>2127</b>
	Table 186: Summary of Logical Interfaces . . . . .	2132
	Table 187: Summary of Logical Interfaces . . . . .	2142
<b>Chapter 28</b>	<b>Operational Commands . . . . .</b>	<b>2183</b>
	Table 188: OpenFlow Operational Mode Commands . . . . .	2183
	Table 189: show openflow capability Output Fields . . . . .	2185
	Table 190: show openflow controller Output Fields . . . . .	2191
	Table 191: show openflow filters Output Fields . . . . .	2194
	Table 192: show openflow flows Output Fields . . . . .	2197
	Table 193: show openflow groups Output Fields . . . . .	2201
	Table 194: show openflow interfaces Output Fields . . . . .	2204
	Table 195: show openflow statistics flows Output Fields . . . . .	2208
	Table 196: show openflow statistics groups Output Fields . . . . .	2211
	Table 197: show openflow statistics interfaces Output Fields . . . . .	2213
	Table 198: show openflow statistics packet Output Fields . . . . .	2216
	Table 199: show openflow statistics queue Output Fields . . . . .	2218
	Table 200: show openflow statistics summary Output Fields . . . . .	2221
	Table 201: show openflow statistics tables Output Fields . . . . .	2223
	Table 202: show openflow summary Output Fields . . . . .	2225
	Table 203: show openflow switch Output Fields . . . . .	2226
<b>Part 8</b>	<b>High Availability</b>	
<b>Chapter 29</b>	<b>Overview . . . . .</b>	<b>2231</b>
	Table 204: Effects of a Routing Engine Switchover . . . . .	2236
	Table 205: Graceful Routing Engine Switchover Feature Support . . . . .	2238
	Table 206: Nonstop Active Routing Platform Support . . . . .	2243
	Table 207: Nonstop Active Routing Protocol and Feature Support . . . . .	2245
<b>Chapter 30</b>	<b>Configuration . . . . .</b>	<b>2261</b>
	Table 208: Settings for VRRP Load-Sharing Example . . . . .	2281
	Table 209: Interface State and Priority Cost Usage . . . . .	2291
<b>Chapter 31</b>	<b>Administration . . . . .</b>	<b>2341</b>
	Table 210: show bgp neighbor Output Fields . . . . .	2345
	Table 211: show ospf overview Output Fields . . . . .	2362

	Table 212: show system switchover Output Fields . . . . .	2368
	Table 213: show task replication Output Fields . . . . .	2370
	Table 214: show task replication Output Fields . . . . .	2372
	Table 215: show vrrp Output Fields . . . . .	2374
<b>Part 9</b>	<b>Interfaces</b>	
<b>Chapter 33</b>	<b>Overview . . . . .</b>	<b>2389</b>
	Table 216: Network Interface Types and Purposes . . . . .	2390
	Table 217: Special Interface Types and Purposes . . . . .	2391
	Table 218: IPv4 and IPv6 Hashing Fields . . . . .	2398
	Table 219: MPLS Hashing Fields . . . . .	2399
	Table 220: MAC-in-MAC Hashing Fields . . . . .	2400
	Table 221: Layer 2 Header Hashing Fields . . . . .	2401
	Table 222: ICCP Failure Scenarios . . . . .	2413
	Table 223: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package . . . . .	2423
	Table 224: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software . . . . .	2427
	Table 225: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package . . . . .	2430
	Table 226: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software . . . . .	2433
	Table 227: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package . . . . .	2436
	Table 228: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software . . . . .	2438
	Table 229: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package . . . . .	2441
	Table 230: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software . . . . .	2445
	Table 231: Firewall Filter Application Points for Tunneled Packets . . . . .	2451
	Table 232: Features Not Supported with GRE . . . . .	2452
	Table 233: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from Trunk Groups . . . . .	2455
<b>Chapter 34</b>	<b>Configuration . . . . .</b>	<b>2457</b>
	Table 234: Settings for Uplink Failure Protection Example . . . . .	2459
	Table 235: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch . . . . .	2463
	Table 236: Components of the Topology for Configuring a Multichassis LAG Between Two Switches . . . . .	2472
	Table 237: Components of the Topology for Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP . . . . .	2495
	Table 238: Components of the Topology for Configuring a Multichassis LAG Between Two Switches . . . . .	2532
	Table 239: Components of the Topology for Configuring a Multichassis LAG Between Two Switches . . . . .	2553
	Table 240: Components of the Redundant Trunk Link Topology . . . . .	2580

	Table 241: Example LPM Table Combinations Using l2-profile-one With Junos OS 13.2X51-D10 . . . . .	2603
	Table 242: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10 . . . . .	2604
	Table 243: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software . . . . .	2611
	Table 244: Unsupported [edit interfaces et] Configuration Statements for the QFX Series . . . . .	2620
<b>Chapter 35</b>	<b>Administration . . . . .</b>	<b>2747</b>
	Table 245: Summary of System Process Information Output Fields . . . . .	2747
	Table 246: Summary of Key System Properties Output Fields . . . . .	2748
	Table 247: Output Control Keys for the monitor interface Command . . . . .	2754
	Table 248: Output Control Keys for the monitor interface traffic Command . . . . .	2755
	Table 249: monitor interface Output Fields . . . . .	2756
	Table 250: show forwarding-options enhanced-hash-key Output Fields . . . . .	2763
	Table 251: show iccp . . . . .	2766
	Table 252: show interfaces diagnostics optics Output Fields . . . . .	2768
	Table 253: show interfaces ge Output Fields . . . . .	2783
	Table 254: GRE show interfaces Output Fields . . . . .	2798
	Table 255: show interfaces irb Output Fields . . . . .	2804
	Table 256: show interfaces mc-ae Output Fields . . . . .	2810
	Table 257: Layer 2 Overhead, Transmitted Packets/Bytes . . . . .	2813
	Table 258: show interfaces queue Output Fields . . . . .	2816
	Table 259: Byte Count by Interface Hardware . . . . .	2819
	Table 260: show interfaces xe Output Fields . . . . .	2853
	Table 261: show lacp interfaces Output Fields . . . . .	2871
	Table 262: show lacp statistics interfaces Output Fields . . . . .	2875
	Table 263: show oam ethernet link-fault-management Output Fields . . . . .	2877
	Table 264: show redundant-trunk-group Output Fields . . . . .	2882
	Table 265: show uplink-failure-detection Output Fields . . . . .	2884
<b>Part 10</b>	<b>Routing Options</b>	
<b>Chapter 37</b>	<b>Overview . . . . .</b>	<b>2897</b>
	Table 266: Unified Forwarding Table Profiles . . . . .	2900
	Table 267: Example Host Table Combinations Using l2-profile-one . . . . .	2900
<b>Chapter 38</b>	<b>Configuration . . . . .</b>	<b>2903</b>
	Table 268: Unified Forwarding Table Profiles . . . . .	2909
	Table 269: Example LPM Table Combinations Using l2-and l3 Profiles With Junos OS 13.2X51-D10 and 13.2X52-D10 . . . . .	2911
	Table 270: LPM Table Combinations for l2 and l3 profiles With Junos OS 13.2X51-D15 . . . . .	2912
	Table 271: Unified Forwarding Table Profiles . . . . .	2974
<b>Chapter 39</b>	<b>Administration . . . . .</b>	<b>3055</b>
	Table 272: Filtering Route Messages . . . . .	3055
	Table 273: Summary of Key Routing Information Output Fields . . . . .	3056
	Table 274: show as-path Output Fields . . . . .	3061

	Table 275: show as-path domain Output Fields . . . . .	3064
	Table 276: show as-path summary Output Fields . . . . .	3067
	Table 277: show ipv6 neighbors Output Fields . . . . .	3069
	Table 278: show ipv6 router-advertisement Output Fields . . . . .	3071
	Table 279: show route Output Fields . . . . .	3075
	Table 280: show route damping Output Fields . . . . .	3099
	Table 281: show route detail Output Fields . . . . .	3103
	Table 282: Next-hop Types Output Field Values . . . . .	3108
	Table 283: State Output Field Values . . . . .	3109
	Table 284: Communities Output Field Values . . . . .	3111
	Table 285: show route export Output Fields . . . . .	3122
	Table 286: show route extensive Output Fields . . . . .	3125
	Table 287: show route flow validation Output Fields . . . . .	3142
	Table 288: show route forwarding-table Output Fields . . . . .	3147
	Table 289: show route instance Output Fields . . . . .	3164
	Table 290: show route martians Output Fields . . . . .	3176
	Table 291: show route receive-protocol Output Fields . . . . .	3203
	Table 292: show route resolution Output Fields . . . . .	3212
	Table 293: show route summary Output Fields . . . . .	3228
	Table 294: show route terse Output Fields . . . . .	3246
<b>Part 11</b>	<b>Border Gateway Protocol</b>	
<b>Chapter 42</b>	<b>Configuration . . . . .</b>	<b>3261</b>
	Table 295: MED Options for Routing Table Path Selection . . . . .	3325
	Table 296: Default Route Preference Values . . . . .	3442
	Table 297: Damping Parameters . . . . .	3600
<b>Chapter 43</b>	<b>Administration . . . . .</b>	<b>3749</b>
	Table 298: show bgp bmp Output Fields . . . . .	3755
	Table 299: show bgp group Output Fields . . . . .	3758
	Table 300: show bgp neighbor Output Fields . . . . .	3765
	Table 301: show bgp summary Output Fields . . . . .	3779
	Table 302: show policy damping Output Fields . . . . .	3785
	Table 303: show route damping Output Fields . . . . .	3787
	Table 304: show route detail Output Fields . . . . .	3791
	Table 305: Next-hop Types Output Field Values . . . . .	3796
	Table 306: State Output Field Values . . . . .	3797
	Table 307: Communities Output Field Values . . . . .	3799
<b>Part 12</b>	<b>Intermediate System to Intermediate System</b>	
<b>Chapter 45</b>	<b>Configuration . . . . .</b>	<b>3823</b>
	Table 308: IPv4 Statements . . . . .	3868
	Table 309: IPv6 Statements . . . . .	3868
	Table 310: Default Metric Values for Routes Exported into IS-IS . . . . .	3947
<b>Chapter 46</b>	<b>Administration . . . . .</b>	<b>3977</b>
	Table 311: show isis adjacency Output Fields . . . . .	3987
	Table 312: show isis authentication Output Fields . . . . .	3990

	Table 313: show isis backup coverage Output Fields . . . . .	3992
	Table 314: show isis backup label-switched-path Output Fields . . . . .	3994
	Table 315: show isis backup spf results Output Fields . . . . .	3997
	Table 316: show isis database Output Fields . . . . .	4001
	Table 317: show isis hostname Output Fields . . . . .	4012
	Table 318: show isis interface Output Fields . . . . .	4015
	Table 319: show isis overview Output Fields . . . . .	4018
	Table 320: show isis route Output Fields . . . . .	4022
	Table 321: show isis spf Output Fields . . . . .	4025
	Table 322: show isis statistics Output Fields . . . . .	4031
<b>Part 13</b>	<b>Open Shortest Path First</b>	
<b>Chapter 47</b>	<b>Overview . . . . .</b>	<b>4035</b>
	Table 323: Default Route Preference Values for OSPF . . . . .	4038
<b>Chapter 49</b>	<b>Administration . . . . .</b>	<b>4243</b>
	Table 324: show (ospf   ospf3) backup coverage Output Fields . . . . .	4255
	Table 325: show (ospf   ospf3) backup neighbor Output Fields . . . . .	4258
	Table 326: show ospf context-identifier Output Fields . . . . .	4261
	Table 327: show ospf database Output Fields . . . . .	4263
	Table 328: show (ospf   ospf3) interface Output Fields . . . . .	4271
	Table 329: show (ospf   ospf3) io-statistics Output Fields . . . . .	4276
	Table 330: show (ospf   ospf3) log Output Fields . . . . .	4278
	Table 331: show (ospf   ospf3) neighbor Output Fields . . . . .	4282
	Table 332: show ospf overview Output Fields . . . . .	4288
	Table 333: show (ospf   ospf3) route Output Fields . . . . .	4293
	Table 334: show (ospf   ospf3) statistics Output Fields . . . . .	4298
<b>Part 14</b>	<b>Routing Information Protocol</b>	
<b>Chapter 51</b>	<b>Configuration . . . . .</b>	<b>4311</b>
	Table 335: Configuring Simple RIP Authentication . . . . .	4323
	Table 336: Configuring MD5 RIP Authentication . . . . .	4324
<b>Chapter 52</b>	<b>Administration . . . . .</b>	<b>4399</b>
	Table 337: show rip general-statistics Output Fields . . . . .	4402
	Table 338: show rip neighbor Output Fields . . . . .	4405
	Table 339: show rip statistics Output Fields . . . . .	4407
<b>Part 15</b>	<b>MPLS Applications</b>	
<b>Chapter 53</b>	<b>Overview . . . . .</b>	<b>4411</b>
	Table 340: MPLS Features on the QFX Series and on the EX4600 Switch . . . . .	4423
	Table 341: MPLS Scaling Values . . . . .	4426
<b>Chapter 54</b>	<b>Configuration . . . . .</b>	<b>4433</b>
	Table 342: from Operators That Apply to LDP Received-Label Filtering . . . . .	4436
	Table 343: to Operators for LDP Outbound-Label Filtering . . . . .	4438
	Table 344: Local CE Switch in the MPLS-Based Layer 3 VPN Topology . . . . .	4447
	Table 345: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology . . . . .	4447



	Table 346: Layer 3 VPN Components of the Local PE Switch . . . . .	4447
	Table 347: Layer 3 VPN Components of the Remote PE Switch . . . . .	4448
	Table 348: Supported Match Conditions for MPLS Firewall Filters . . . . .	4475
	Table 349: Supported Actions for MPLS Firewall Filters . . . . .	4475
<b>Chapter 55</b>	<b>Administration . . . . .</b>	<b>4565</b>
	Table 350: Output Control Keys for the monitor label-switched-path Command . . . . .	4577
	Table 351: monitor label-switched-path Output Fields . . . . .	4578
	Table 352: show ldp database Output Fields . . . . .	4601
	Table 353: show ldp fec-filters Output Fields . . . . .	4608
	Table 354: show ldp interface Output Fields . . . . .	4609
	Table 355: show ldp neighbor Output Fields . . . . .	4611
	Table 356: show ldp path Output Fields . . . . .	4613
	Table 357: show ldp route Output Fields . . . . .	4615
	Table 358: show ldp session Output Fields . . . . .	4619
	Table 359: show ldp statistics Output Fields . . . . .	4625
	Table 360: show ldp traffic-statistics Output Fields . . . . .	4630
	Table 361: show security keychain Output Fields . . . . .	4633
	Table 362: show link-management Output Fields . . . . .	4636
	Table 363: show link-management peer Output Fields . . . . .	4640
	Table 364: show link-management routing Output Fields . . . . .	4642
	Table 365: show link-management statistics Output Fields . . . . .	4645
	Table 366: show link-management te-link Output Fields . . . . .	4647
	Table 367: show mpls call-admission-control Output Fields . . . . .	4649
	Table 368: show mpls cspf Output Fields . . . . .	4651
	Table 369: show mpls diffserv-te Output Fields . . . . .	4653
	Table 370: show route forwarding-table Output Fields . . . . .	4656
	Table 371: show mpls interface Output Fields . . . . .	4663
	Table 372: show mpls lsp Output Fields . . . . .	4667
	Table 373: show mpls lsp autobandwidth Output Fields . . . . .	4681
	Table 374: show mpls path Output Fields . . . . .	4684
	Table 375: show mpls static-lsp Output Fields . . . . .	4686
	Table 376: show rsvp interface Output Fields . . . . .	4688
	Table 377: show rsvp neighbor Output Fields . . . . .	4693
	Table 378: show rsvp session Output Fields . . . . .	4700
	Table 379: show rsvp statistics Output Fields . . . . .	4707
	Table 380: show rsvp version Output Fields . . . . .	4711
	Table 381: show ted database Output Fields . . . . .	4714
	Table 382: show ted link Output Fields . . . . .	4721
	Table 383: show ted protocol Output Fields . . . . .	4724
	Table 384: traceroute mpls ldp Output Fields . . . . .	4727
	Table 385: traceroute mpls rsvp Output Fields . . . . .	4730
<b>Part 16</b>	<b>Multicast</b>	
<b>Chapter 57</b>	<b>Overview . . . . .</b>	<b>4737</b>
	Table 386: ASM and SSM Terminology . . . . .	4771
<b>Chapter 58</b>	<b>Configuration . . . . .</b>	<b>4777</b>

	Table 387: PIM Join Filter Match Conditions . . . . .	4812
	Table 388: IGMP Event Messages . . . . .	4843
	Table 389: Components of the IGMP Snooping Topology . . . . .	4850
	Table 390: MLD Event Messages . . . . .	4870
	Table 391: Source-Active Message Flooding Explanation . . . . .	4883
<b>Chapter 59</b>	<b>Administration . . . . .</b>	<b>5047</b>
	Table 392: Summary of IGMP Snooping Output Fields . . . . .	5047
	Table 393: mtrace Output Fields . . . . .	5072
	Table 394: mtrace from-source Output Fields . . . . .	5076
	Table 395: mtrace monitor Output Fields . . . . .	5078
	Table 396: mtrace to-gateway Output Fields . . . . .	5081
	Table 397: show igmp group Output Fields . . . . .	5083
	Table 398: show igmp group Output Fields . . . . .	5085
	Table 399: show igmp interface Output Fields . . . . .	5089
	Table 400: show igmp statistics Output Fields . . . . .	5093
	Table 401: show igmp-snooping membership Output Fields . . . . .	5096
	Table 402: show igmp-snooping route Output Fields . . . . .	5099
	Table 403: show igmp-snooping statistics Output Fields . . . . .	5101
	Table 404: show igmp-snooping vlans Output Fields . . . . .	5103
	Table 405: show msdp Output Fields . . . . .	5105
	Table 406: show msdp source Output Fields . . . . .	5108
	Table 407: show msdp source-active Output Fields . . . . .	5110
	Table 408: show msdp statistics Output Fields . . . . .	5112
	Table 409: show multicast flow-map Output Fields . . . . .	5116
	Table 410: show multicast interface Output Fields . . . . .	5118
	Table 411: show multicast minfo Output Fields . . . . .	5120
	Table 412: show multicast next-hops Output Fields . . . . .	5123
	Table 413: show multicast pim-to-igmp-proxy Output Fields . . . . .	5125
	Table 414: show multicast pim-to-mln-proxy Output Fields . . . . .	5127
	Table 415: show multicast route Output Fields . . . . .	5130
	Table 416: show multicast rpf Output Fields . . . . .	5137
	Table 417: show multicast scope Output Fields . . . . .	5140
	Table 418: show multicast sessions Output Fields . . . . .	5142
	Table 419: show multicast usage Output Fields . . . . .	5146
	Table 420: show pim bootstrap Output Fields . . . . .	5148
	Table 421: show pim interfaces Output Fields . . . . .	5150
	Table 422: show pim join Output Fields . . . . .	5155
	Table 423: show pim neighbors Output Fields . . . . .	5176
	Table 424: show pim rps Output Fields . . . . .	5180
	Table 425: show pim source Output Fields . . . . .	5187
	Table 426: show pim statistics Output Fields . . . . .	5190
<b>Part 17</b>	<b>Security</b>	
<b>Chapter 60</b>	<b>Overview . . . . .</b>	<b>5209</b>
	Table 427: Actions for Firewall Filters . . . . .	5218
	Table 428: Supported Match Conditions for Firewall Filters . . . . .	5219
	Table 429: Actions for Firewall Filters . . . . .	5230
	Table 430: Action Modifiers for Firewall Filters . . . . .	5231

	Table 431: Supported Firewall Filter Numbers . . . . .	5236
	Table 432: Policer Actions . . . . .	5243
	Table 433: Color-Blind Mode TCM Color-to-PLP Mapping . . . . .	5247
	Table 434: Color-Aware Mode Single-Rate PLP Mapping . . . . .	5247
	Table 435: Color-Blind Mode TCM Color-to-PLP Mapping . . . . .	5249
	Table 436: Color-Aware Mode Two-Rate PLP Mapping . . . . .	5250
	Table 437: DHCPv6 Messages and Equivalent DHCPv4 Messages . . . . .	5258
<b>Chapter 61</b>	<b>Configuration . . . . .</b>	<b>5279</b>
	Table 438: Servers Connected to Switch . . . . .	5286
	Table 439: Unicast Forwarding Classes . . . . .	5294
<b>Chapter 62</b>	<b>Administration . . . . .</b>	<b>5379</b>
	Table 440: show arp inspection statistics Output Fields . . . . .	5399
	Table 441: show dhcp snooping binding Output Fields . . . . .	5400
	Table 442: show firewall Output Fields . . . . .	5402
	Table 443: show firewall policer Output Fields . . . . .	5406
	Table 444: show interfaces filters Output Fields . . . . .	5408
<b>Part 18</b>	<b>Services</b>	
<b>Chapter 64</b>	<b>Overview . . . . .</b>	<b>5425</b>
	Table 445: Port Mirroring Terms and Definitions . . . . .	5426
<b>Chapter 66</b>	<b>Administration . . . . .</b>	<b>5497</b>
	Table 446: show analyzer Output Fields . . . . .	5498
<b>Part 19</b>	<b>Storage</b>	
<b>Chapter 68</b>	<b>Overview . . . . .</b>	<b>5507</b>
	Table 447: Fibre Channel Protocol Layers . . . . .	5509
	Table 448: VFP TCAM Entry Consumption Summary . . . . .	5550
	Table 449: Asymmetric Ethernet PAUSE Flow Control Configuration . . . . .	5561
	Table 450: Flow Control State Advertised to the Connected Peer (Autonegotiation) . . . . .	5562
	Table 451: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces . . . . .	5563
	Table 452: Default PFC Priority to Queue and Forwarding Class Mapping . . . . .	5565
	Table 453: Fibre Channel Terms . . . . .	5569
	Table 454: Summary of Differences Between IEEE DCBX and DCBX Version 1.01 . . . . .	5582
<b>Chapter 69</b>	<b>Configuration . . . . .</b>	<b>5595</b>
	Table 455: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier) . . . . .	5597
	Table 456: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier) . . . . .	5597
	Table 457: Components of DCBX Application Protocol Exchange Configuration Topology . . . . .	5598
	Table 458: Components of the PFC for FCoE Traffic Configuration Topology . . . . .	5607

	Table 459: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology . . . . .	5616
	Table 460: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to the Same FCoE Transit Switch) . . . . .	5638
	Table 461: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to Different FCoE Transit Switches) . . . .	5643
	Table 462: Components of the VN2VN_Port FIP Snooping Configuration Topology (FCoE Hosts Indirectly Connected Across an Aggregation Layer FCoE Transit Switch) . . . . .	5649
<b>Chapter 70</b>	<b>Administration . . . . .</b>	<b>5707</b>
	Table 463: show dcbx output fields . . . . .	5723
	Table 464: show dcbx neighbors Output Fields . . . . .	5724
	Table 465: show fip snooping Output Fields . . . . .	5746
	Table 466: show fip snooping enode Output Fields . . . . .	5751
	Table 467: show fip snooping fcf Output Fields . . . . .	5755
	Table 468: show fip snooping interface Output Fields . . . . .	5758
	Table 469: show fip snooping statistics Output Fields . . . . .	5761
	Table 470: show fip snooping vlan Output Fields . . . . .	5764
	Table 471: show fip vlan-discovery Output Fields . . . . .	5768
<b>Part 20</b>	<b>Traffic Management</b>	
<b>Chapter 72</b>	<b>Overview . . . . .</b>	<b>5779</b>
	Table 472: Policer Actions . . . . .	5785
	Table 473: CoS Mappings—Inputs and Outputs . . . . .	5795
	Table 474: Default Forwarding Classes and Queue Mapping . . . . .	5796
	Table 475: Default IEEE 802.1 Code-Point Aliases . . . . .	5797
	Table 476: Default DSCP and DCSP IPv6 Code-Point Aliases . . . . .	5798
	Table 477: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier) . . . . .	5799
	Table 478: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier) . . . . .	5799
	Table 479: Default IEEE 802.1 Multidestination Classifiers . . . . .	5800
	Table 480: Default DSCP IP and IPv6 Unicast Classifiers . . . . .	5800
	Table 481: Default Drop Profile . . . . .	5802
	Table 482: Default Schedulers . . . . .	5802
	Table 483: Default Scheduler Maps . . . . .	5804
	Table 484: Default Ingress Shared Buffer Configuration . . . . .	5804
	Table 485: Default Egress Shared Buffer Configuration . . . . .	5804
	Table 486: Routing Engine Protocol Default Queue Mapping . . . . .	5806
	Table 487: Default IEEE 802.1 Code-Point Aliases . . . . .	5808
	Table 488: Default DSCP and DSCP IPv6 Code-Point Aliases . . . . .	5809
	Table 489: Default BA Classification . . . . .	5812
	Table 490: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping . . . . .	5813
	Table 491: Supported Classifiers and Rewrite Rules . . . . .	5820
	Table 492: Ethernet Interface Support for Classifier and Rewrite Rule Configuration . . . . .	5823

Table 493: Default Forwarding Classes for Unicast Packets . . . . .	5831
Table 494: Default Forwarding Classes for Multicast Packets . . . . .	5832
Table 495: Mapping of Default Unicast Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute . . . . .	5840
Table 496: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3 . . . . .	5842
Table 497: Default Output Flow Control Profile . . . . .	5847
Table 498: User-Configured Output Flow Control Profile . . . . .	5848
Table 499: Results of Lossless Priority Configuration . . . . .	5852
Table 500: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier) . . . . .	5857
Table 501: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier) . . . . .	5857
Table 502: Default IEEE 802.1 Multidestination Classifiers . . . . .	5858
Table 503: Default DSCP IP and IPv6 Unicast Classifiers . . . . .	5858
Table 504: Default Scheduler Configuration . . . . .	5859
Table 505: Hierarchical Scheduling Tiers . . . . .	5863
Table 506: Output Queue Scheduler Components . . . . .	5869
Table 507: Other Scheduling Components . . . . .	5869
Table 508: Default Schedulers . . . . .	5870
Table 509: Priority Group Scheduler Components . . . . .	5877
Table 510: Other Scheduling Components . . . . .	5878
Table 511: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler) . . . . .	5896
Table 512: Egress Queue Dedicated Buffer Allocation (Example 1) . . . . .	5898
Table 513: Egress Queue Dedicated Buffer Allocation with Another Remainder Queue (Example 2) . . . . .	5898
Table 514: QFX5100 and EX4600 Switch Default Shared Ingress Buffer Values (KB) . . . . .	5901
Table 515: QFX3500 and QFX3600 Switch Default Shared Ingress Buffer Values (KB) . . . . .	5901
Table 516: Default Shared Ingress Buffer Values (Percentage) . . . . .	5901
Table 517: QFX5100 and EX4600 Switch Default Shared Egress Buffer Values (KB) . . . . .	5902
Table 518: QFX3500 and QFX3600 Switch Default Shared Egress Buffer Values (KB) . . . . .	5902
Table 519: Default Shared Egress Buffer Values (Percentage) . . . . .	5902
Table 520: Default Ingress Shared Buffer Configuration . . . . .	5904
Table 521: Default Egress Shared Buffer Configuration . . . . .	5904
Table 522: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic . . . . .	5904
Table 523: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic . . . . .	5904
Table 524: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled . . . . .	5905
Table 525: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled . . . . .	5905
Table 526: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic . . . . .	5905

	Table 527: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic . . . . .	5906
	Table 528: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Lossless Traffic . . . . .	5906
	Table 529: Recommended Egress Shared Buffer Configuration for Networks with Mostly Lossless Traffic . . . . .	5906
	Table 530: Asymmetric Ethernet PAUSE Flow Control Configuration . . . . .	5918
	Table 531: Flow Control State Advertised to the Connected Peer (Autonegotiation) . . . . .	5919
	Table 532: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces . . . . .	5920
	Table 533: Default PFC Priority to Queue and Forwarding Class Mapping . . . . .	5922
	Table 534: ECN Bit Codes . . . . .	5927
	Table 535: Traffic Behavior on ECN-Enabled Queues . . . . .	5929
	Table 536: Summary of Differences Between IEEE DCBX and DCBX Version 1.01 . . . . .	5939
	Table 537: Default Forwarding Class Configuration . . . . .	5954
	Table 538: Support of QFX CoS Features on a VCF in Mixed Mode with an EX4300 Leaf Device . . . . .	5958
<b>Chapter 73</b>	<b>Configuration . . . . .</b>	<b>5965</b>
	Table 539: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology . . . . .	5968
	Table 540: Components of the PFC for FCoE Traffic Configuration Topology . . . . .	5989
	Table 541: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology . . . . .	5998
	Table 542: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3 . . . . .	6021
	Table 543: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology . . . . .	6029
	Table 544: Components of the Two Lossless FCoE Priorities Configuration Topology . . . . .	6038
	Table 545: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology . . . . .	6052
	Table 546: ba-ucast-classifier Loss Priority Assignments . . . . .	6067
	Table 547: BA-mcast-classifier Loss Priority Assignments . . . . .	6070
	Table 548: Forwarding-Class-to-Queue Example Configuration . . . . .	6077
	Table 549: Components of the Forwarding Class Sets Configuration Example . . . . .	6079
	Table 550: Components of the Queue Scheduler Configuration Example . . . . .	6084
	Table 551: Components of the Queue Scheduler Priority Configuration Example . . . . .	6088
	Table 552: Components of the ECN Configuration Example . . . . .	6091
	Table 553: Components of the Minimum Guaranteed Output Bandwidth Configuration Example . . . . .	6098
	Table 554: Components of the Maximum Output Bandwidth Configuration Example . . . . .	6102
	Table 555: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies . . . . .	6106

Table 556: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE . . .	6112
Table 557: Components of the Recommended Shared Buffer Configuration for Multicast Network Topologies . . . . .	6118
Table 558: Components of the Recommended Shared Buffer Configuration for Lossless Network Topologies . . . . .	6124
Table 559: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier) . . . . .	6129
Table 560: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier) . . . . .	6130
Table 561: Components of DCBX Application Protocol Exchange Configuration Topology . . . . .	6131
Table 562: Components of the PFC Across Layer 3 Interfaces Topology . . . . .	6140
Table 563: Default Egress Shared Buffer Partitioning . . . . .	6201
Table 564: Default Ingress Shared Buffer Partitioning . . . . .	6203
Table 565: Default Output Queue Buffer Sizes . . . . .	6206
<b>Chapter 74 Administration . . . . .</b>	<b>6289</b>
Table 566: Summary of Key CoS Classifier Output Fields . . . . .	6289
Table 567: Summary of Key CoS Forwarding Class Output Fields . . . . .	6291
Table 568: Summary of Key CoS Interfaces Output Fields . . . . .	6292
Table 569: Summary of Key CoS Rewrite Rule Output Fields . . . . .	6293
Table 570: Summary of Key CoS Scheduler Maps Output Fields . . . . .	6293
Table 571: Summary of Key CoS Value Alias Output Fields . . . . .	6295
Table 572: show class-of-service Output Fields . . . . .	6297
Table 573: show class-of-service classifier Output Fields . . . . .	6301
Table 574: show class-of-service code-point-aliases Output Fields . . . . .	6303
Table 575: show class-of-service congestion-notification Output Fields . . . . .	6305
Table 576: show class-of-service drop-profile Output Fields . . . . .	6308
Table 577: show class-of-service forwarding-class Output Fields . . . . .	6311
Table 578: show class-of-service forwarding-class-set Output Fields . . . . .	6313
Table 579: show class-of-service forwarding-table classifier Output Fields . . .	6319
Table 580: show class-of-service forwarding-table classifier mapping Output Fields . . . . .	6321
Table 581: show class-of-service forwarding-table drop-profile Output Fields . . . . .	6323
Table 582: show class-of-service forwarding-table rewrite-rule Output Fields . . . . .	6325
Table 583: show class-of-service forwarding-table rewrite-rule mapping Output Fields . . . . .	6327
Table 584: show class-of-service forwarding-table scheduler-map Output Fields . . . . .	6328
Table 585: show class-of-service interface Output Fields . . . . .	6331
Table 586: show class-of-service multi-destination Output Fields . . . . .	6358
Table 587: show class-of-service rewrite-rule Output Fields . . . . .	6359
Table 588: show class-of-service scheduler-map Output Fields . . . . .	6361
Table 589: show class-of-service shared-buffer Output Fields . . . . .	6363
Table 590: show class-of-service traffic-control-profile Output Fields . . . . .	6365
Table 591: show dcbx output fields . . . . .	6369

	Table 592: show dcbx neighbors Output Fields . . . . .	6370
	Table 593: Layer 2 Overhead, Transmitted Packets/Bytes . . . . .	6393
	Table 594: show interfaces queue Output Fields . . . . .	6396
	Table 595: Byte Count by Interface Hardware . . . . .	6399
	Table 596: show pfe filter Output Fields . . . . .	6432
	Table 597: show pfe next-hop Output Fields . . . . .	6436
	Table 598: show pfe route Output Fields . . . . .	6442
	Table 599: QFX Series show pfe route Hardware Table Output Fields . . . . .	6442
<b>Chapter 75</b>	<b>Troubleshooting . . . . .</b>	<b>6451</b>
	Table 600: Components of the Rate Shaping Troubleshooting Example . . . .	6459
<b>Part 21</b>	<b>Network Management and Monitoring</b>	
<b>Chapter 76</b>	<b>Overview . . . . .</b>	<b>6463</b>
	Table 601: Device and Network Management Features on the QFX Series and EX4600 . . . . .	6463
	Table 602: Differences Between Persistent and Transient Changes . . . . .	6482
	Table 603: Network Analytics CLI Changes . . . . .	6494
	Table 604: Configuration and Status Output in Junos OS Release 13.2X51-D10 and 13.2X50-D15 . . . . .	6498
	Table 605: Streamed Queue Statistics Data Output Fields . . . . .	6499
	Table 606: Streamed Traffic Statistics Data Output Fields . . . . .	6500
	Table 607: GPB Stream Format Message Header Information . . . . .	6502
	Table 608: Streamed Queue Statistics Data Output Fields . . . . .	6505
	Table 609: Streamed Traffic Statistics Data Output Fields . . . . .	6505
	Table 610: Output Fields for Queue Statistics in Local Analytics File . . . . .	6507
	Table 611: Output Fields for Traffic Statistics in Local Analytics File . . . . .	6507
	Table 612: Fabric Chassis MIB Tables and Objects . . . . .	6519
	Table 613: Fabric Chassis MIB SNMPv2 Traps . . . . .	6521
	Table 614: RMON Event Table . . . . .	6527
	Table 615: RMON Alarm Table . . . . .	6527
	Table 616: jnxRmon Alarm Table . . . . .	6528
	Table 617: RMON History Control Table . . . . .	6528
	Table 618: Monitored Object Instances . . . . .	6530
	Table 619: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6531
	Table 620: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6536
	Table 621: Standard MIBs Supported on QFabric Systems . . . . .	6540
	Table 622: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems . . . . .	6543
	Table 623: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6546
	Table 624: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6549
	Table 625: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6551
	Table 626: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis . . . . .	6553



	Table 627: Standard SNMPv2 Traps Supported on QFabric Systems . . . . .	6555
	Table 628: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems . . . . .	6556
<b>Chapter 77</b>	<b>Configuration . . . . .</b>	<b>6565</b>
	Table 629: Minimum Configuration Statements for System Logging . . . . .	6616
	Table 630: Fields in Structured-Data Messages . . . . .	6625
	Table 631: Facility and Severity Codes in the priority-code Field . . . . .	6627
	Table 632: Fields in Standard-Format Messages . . . . .	6628
	Table 633: Junos OS System Logging Facilities . . . . .	6631
	Table 634: System Log Message Severity Levels . . . . .	6632
	Table 635: Junos OS System Logging Facilities . . . . .	6633
	Table 636: System Log Message Severity Levels . . . . .	6633
	Table 637: Default Facilities for Messages Directed to a Remote Destination . .	6634
	Table 638: Facilities for the facility-override Statement . . . . .	6635
	Table 639: Regular Expression Operators for the match Statement . . . . .	6638
<b>Chapter 78</b>	<b>Administration . . . . .</b>	<b>6799</b>
	Table 640: Output Control Keys for the monitor interface Command . . . . .	6802
	Table 641: SNMP Tracing Flags . . . . .	6809
	Table 642: Commit Script Configuration and Operational Mode Commands . .	6812
	Table 643: Match Conditions for the monitor traffic Command . . . . .	6817
	Table 644: Logical Operators for the monitor traffic Command . . . . .	6818
	Table 645: Arithmetic and Relational Operators for the monitor traffic Command . . . . .	6820
	Table 646: monitor start Command Output Fields . . . . .	6830
	Table 647: show analytics collector Command Output Fields . . . . .	6833
	Table 648: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D15 and Later) . . . . .	6835
	Table 649: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier) . . . . .	6836
	Table 650: show analytics queue-statistics Command Output Fields . . . . .	6839
	Table 651: show analytics status Command Output Fields . . . . .	6841
	Table 652: show analytics streaming-servers Command Output Fields . . . .	6845
	Table 653: show analytics traffic-statistics Command Output Fields . . . . .	6847
	Table 654: show sflow Output Fields . . . . .	6851
	Table 655: show sflow collector Output Fields . . . . .	6853
	Table 656: show sflow interface Output Fields . . . . .	6854
	Table 657: show snmp health-monitor Output Fields . . . . .	6867
	Table 658: show snmp inform-statistics Output Fields . . . . .	6872
	Table 659: show snmp mib Output Fields . . . . .	6875
	Table 660: show snmp rmon Output Fields . . . . .	6877
	Table 661: show snmp statistics Output Fields . . . . .	6882
	Table 662: show snmp v3 Output Fields . . . . .	6887
<b>Chapter 79</b>	<b>Troubleshooting . . . . .</b>	<b>6893</b>
	Table 663: Troubleshooting Resources on the QFX Series . . . . .	6893
	Table 664: Troubleshooting on the QFX Series . . . . .	6895

<b>Part 22</b>	<b>Virtual Chassis</b>	
<b>Chapter 80</b>	<b>Overview</b>	<b>6907</b>
	Table 665: Virtual Chassis Fabric Summary	6914
	Table 666: Virtual Chassis Summary	6915
	Table 667: Automatic Software Update Support	6927
<b>Chapter 82</b>	<b>Administration</b>	<b>6975</b>
	Table 668: show virtual-chassis active-topology Output Fields	6985
	Table 669: show virtual-chassis device-topology Output Fields	6990
	Table 670: show virtual-chassis protocol adjacency Output Fields	6997
	Table 671: show virtual-chassis protocol database Output Fields	7000
	Table 672: show virtual-chassis protocol interface Output Fields	7005
	Table 673: show virtual-chassis protocol route Output Fields	7007
	Table 674: show virtual-chassis protocol statistics Output Fields	7010
	Table 675: show virtual-chassis Output Fields	7014
	Table 676: show virtual-chassis vc-path Output Fields	7018
	Table 677: show virtual-chassis vc-port Output Fields	7020
	Table 678: show virtual-chassis vc-port statistics Output Fields	7025
<b>Part 23</b>	<b>Virtual Chassis Fabric</b>	
<b>Chapter 83</b>	<b>Overview</b>	<b>7033</b>
	Table 679: Virtual Chassis Fabric Summary	7047
	Table 680: Virtual Chassis Summary	7048
<b>Chapter 85</b>	<b>Administration</b>	<b>7105</b>
	Table 681: show forwarding-options enhanced-hash-key Output Fields	7119
	Table 682: show virtual-chassis active-topology Output Fields	7122
	Table 683: show virtual-chassis device-topology Output Fields	7127
	Table 684: show virtual-chassis mode Output Fields	7134
	Table 685: show virtual-chassis protocol adjacency Output Fields	7138
	Table 686: show virtual-chassis protocol database Output Fields	7141
	Table 687: show virtual-chassis protocol interface Output Fields	7146
	Table 688: show virtual-chassis protocol route Output Fields	7148
	Table 689: show virtual-chassis protocol statistics Output Fields	7151
	Table 690: show virtual-chassis Output Fields	7154
	Table 691: show virtual-chassis vc-port Output Fields	7158
	Table 692: show virtual-chassis vc-port diagnostics optics Output Fields	7163
	Table 693: show virtual-chassis vc-port statistics Output Fields	7177
<b>Part 24</b>	<b>Troubleshooting</b>	
<b>Chapter 87</b>	<b>Overview</b>	<b>7187</b>
	Table 694: Troubleshooting Resources on the QFX Series	7187
	Table 695: Troubleshooting on the QFX Series	7189
	Table 696: Alarm Terms and Definitions	7192
	Table 697: QFX3500 Chassis Alarm Messages	7193
<b>Chapter 88</b>	<b>Administration</b>	<b>7197</b>
	Table 698: SNMP Tracing Flags	7201

<b>Chapter 89</b>	Table 699: Output Control Keys for the monitor interface Command . . . . .	7207
	<b>Troubleshooting . . . . .</b>	<b>7209</b>
	Table 700: Troubleshooting Resources on the QFX Series . . . . .	7243
	Table 701: Troubleshooting on the QFX Series . . . . .	7245
	Table 702: Components of the Rate Shaping Troubleshooting Example . . . .	7274



# About the Documentation

- [Documentation and Release Notes on page cxxxv](#)
- [Supported Platforms on page cxxxv](#)
- [Using the Examples in This Manual on page cxxxv](#)
- [Documentation Conventions on page cxxxvii](#)
- [Documentation Feedback on page cxxxix](#)
- [Requesting Technical Support on page cxxxix](#)

## Documentation and Release Notes

---

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

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

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

## Supported Platforms

---

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

- [QFX Series standalone switches](#)

## Using the Examples in This Manual

---

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

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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

## Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

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



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

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

## Documentation Feedback

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

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

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

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

## Opening a Case with JTAC

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

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

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

## PART 1

# QFX5100 Switch Overview

- [QFX5100 Switch Overview on page 3](#)



## CHAPTER 1

# QFX5100 Switch Overview

- [QFX5100 Device Hardware Overview on page 3](#)

## QFX5100 Device Hardware Overview

---

The QFX5100 line of switches is Juniper Network's second generation of top-of-rack switch solutions for data centers and campus distribution or aggregation environments. The QFX5100 portfolio consists of high-performance fixed-configuration switches that add higher port densities, additional scalability, and improved latency to the QFX Series.

This topic covers:

- [QFX5100 Hardware on page 3](#)
- [System Software on page 7](#)

## QFX5100 Hardware


QFX5100 line of switches offer two compact 1 U models and a 2 U model that provide wire-speed packet performance, very low latency, and rich set of Layer 2 and Layer 3 features. In addition to a high-throughput Packet Forwarding Engine, the performance of the control plane running on all the QFX5100 switches is enhanced by the 1.5 Ghz dual core Intel CPU with 8 GB of memory and 32 GB of solid-state drive (SSD) storage.

The QFX5100 line of switches include both 10GE and 40GE fixed-configurations:

- QFX5100-48S

As shown in [Figure 1 on page 4](#), the QFX5100-48S is a 10-Gigabit Ethernet Enhanced Small Form-Factor Pluggable (SFP+) top-of-rack switch with 48 SFP+ ports and 6 Quad SFP+ (QSFP+) ports. Each SFP+ port can operate as a native 10 Gigabit port or a 1 Gigabit port when 1\_Gigabit optics are inserted. Each QSFP+ port can operate at native 40-Gigabit speed or as 4 independent 10-Gigabit port speeds. The 6 QSFP+ ports can be used as either access ports or as uplinks. The QFX5100-48S provides full duplex throughput of 1.44 Tbps. The QFX5100-48S has a 1 U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

Figure 1: QFX5100-48S Switch



The QFX5100-48S can be used as:

- A standalone switch.
- A Node device in a QFabric system.

The QFX5100-48S is supported on both the QFX3000-G and QFX3000-M QFabric systems.

- A master, backup, or linecard in a QFX Virtual Chassis.

A QFX Series Virtual Chassis allows you to interconnect up to ten QFX3500, QFX3600, or QFX5100 switches into one logical device and manage the device as a single chassis using a ring topology.


- A spine or leaf device in a Virtual Chassis Fabric (VCF).

VCF uses Virtual Chassis technology to interconnect multiple devices into a single logical device and manage that device as a single logical device inside of a fabric architecture. VCF architecture supports up to 32 total devices in a spine and leaf topology. Using all QFX5100 devices in a VCF, you can configure a maximum of 8 spine devices and 24 leaf devices. When mixed with other supported devices in the VCF, you can configure a maximum of 4 spine devices and 28 leaf devices.

- QFX5100-48T

As shown in [Figure 2 on page 4](#), the QFX5100-48T is a tri-speed 100/1000/10GBASE-T top-of-rack switch with 48 10GBASE-T access ports and 6 QSFP+ ports. Each 40-Gigabit QSFP+ port can operate either as a native 40-Gigabit port or be channelized into 4 independent 10 Gigabit ports. The 6 QSFP+ ports can be used as either access ports or as uplinks. The QFX5100-48T provides full duplex throughput of 720 Gbps. The QFX5100-48T has a 1U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

Figure 2: QFX5100-48T Switch



The QFX5100-48T can be used as:

- A standalone switch.

- A master, backup, or linecard in a QFX Virtual Chassis.

A QFX Series Virtual Chassis allows you to interconnect up to ten QFX5100, QFX3500, or QFX3600, switches into one logical device and manage the device as a single chassis using a ring topology.

- A spine device or a leaf device in a Virtual Chassis Fabric (VCF).


VCF uses Virtual Chassis technology to interconnect multiple devices into a single logical device and manage that device as a single logical device inside of a fabric architecture. VCF architecture supports up to 20 total devices in a spine and leaf topology.

- QFX5100-24Q

As shown in [Figure 3 on page 5](#), the QFX5100-24Q is a 40-Gigabit Ethernet QSFP+ switch with 24 high-density QSFP+ ports. Each QSFP+ port can operate as a native 40 Gbps port or as 4 independent 10 Gbps ports. The QFX5100-24Q switch has a 1 U form factor and comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

The QFX5100-24Q switch has two module bays for the optional expansion module, QFX-EM-4Q, which can add a total of 8 additional QSFP+ ports to the chassis. When fully populated with QFX-EM-4Q Expansion Modules, the QFX5100-24Q switch is equivalent to 128 10 Gbps interfaces (96 + 16 + 16). Of these total ports, 104 logical ports are available for 10G port channelization. For full details on the different port channelization modes, see *Port Panel of a QFX5100-24Q Device*. All ports on the QFX5100-24Q and QFX-EM-4Q can be configured as either access ports or as uplinks. The QFX5100-24Q switch provides full duplex throughput of 2.56 Tbps.

**Figure 3: QFX5100-24Q Switch**



The QFX5100-24Q can be used as:

- A standalone switch.
- A master, backup, or linecard in a QFX Virtual Chassis.

A QFX Series Virtual Chassis allows you to interconnect up to ten QFX3500, QFX3600, or QFX5100 switches into one logical device and manage the device as a single chassis in a ring topology.

- A spine or leaf device in a Virtual Chassis Fabric (VCF).

VCF uses Virtual Chassis technology to interconnect multiple devices into a single logical device and manage that device as a single logical device inside of a fabric architecture. VCF architecture supports up to 32 total devices in a spine and leaf topology. Using all QFX5100 devices in a VCF, you can configure a maximum of 32

total devices, of which 8 devices can be configured as spine devices. When mixed with other supported devices in the VCF, you can configure a maximum of 4 spine devices of the total 32 devices.




**BEST PRACTICE:** Use QFX5100-24Q switches as master and backup in a QFX Virtual Chassis; use as spine devices in a VCF.

The QFX5100-24Q switch has two bays on the port panel for optional expansion modules. The QFX5100-24Q supports two expansion modules to increase port density:


- QFX-EM-4Q, which provides 4 additional 40-Gigabit Quad SFP+ (QSFP+) ports. See [Figure 4 on page 6](#).

**Figure 4: QFX-EM-4Q Expansion Module**



- EX4600-EM-8F, which provides 8 additional 10-Gigabit Ethernet Enhanced Small Form-Factor Pluggable (SFP+) ports. See [Figure 5 on page 6](#).

**Figure 5: EX4600-EM-8F Expansion Module**



The QFX5100-24Q is configured for the QFX-EM-4Q by default, but any combination of the two modules is supported. Expansion modules can be hot-inserted or hot-removed. However, when an EX4600-EM-8F is inserted instead of the default QFX-EM-4Q, the new configuration causes the interfaces to temporarily go down. Likewise, when an EX4600-EM-8F is running on the QFX5100-24Q and it is swapped with a QFX-EM-4Q, the interfaces temporarily go down, which can cause a short disruption in traffic.


- QFX5100-96S

As shown in [Figure 6 on page 7](#), the QFX5100-96S switch is a 10-Gigabit Ethernet Enhanced Small Form-Factor Pluggable (SFP+) top-of-rack switch with 96 SFP+ ports and 8 Quad SFP+ (QSFP+) ports. Each SFP+ port can operate as a native 10 Gbps port or as a 1 Gbps port. QSFP+ ports 96 and 100 can operate at native 40 Gbps speed or can be channelized to 4 independent 10 Gbps port speeds. The 8 QSFP+ ports can be used as either access ports or as uplinks. The QFX5100-96S switch has a 2 U form factor and comes standard with redundant fans and redundant power



supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow and with AC or DC power supplies.

**Figure 6: QFX5100-96S Switch**



The QFX5100-96S can be used as:

- A standalone switch.
- A member in a QFX Virtual Chassis.

A QFX Series Virtual Chassis allows you to interconnect up to ten QFX3500, QFX3600, or QFX5100 switches into one logical device and manage the device as a single chassis in a ring topology.

A spine or leaf device in a Virtual Chassis Fabric (VCF).

VCF uses Virtual Chassis technology to interconnect multiple devices into a single logical device and manage that device as a single logical device inside of a fabric architecture. VCF architecture supports up to 32 total devices in a spine and leaf topology. Using all QFX5100 devices in a VCF, you can configure a maximum of 32 total devices, of which 8 devices can be configured as spine devices. When mixed with other supported devices in the VCF, you can configure a maximum of 4 spine devices of the total 32 devices.

## System Software

QFX Series devices use the Junos operating system (OS), which provides Layer 2 and Layer 3 switching, routing, and security services. Junos OS is installed on a QFX5100 switch's 32-gigabyte (GB) internal solid state flash drive. The same Junos OS code base that runs on QFX5100 switches also runs on all Juniper Networks EX Series switches, and J Series, M Series, MX Series, and T Series routers.

For more information about which features are supported on QFX Series devices, see *QFX Series Software Features Overview*.

You manage the switch using the Junos OS command-line interface (CLI), accessible through the console and out-of-band management ports on the device.

### Related Documentation

- *QFX5100 Device Models*
- *QFX Series Software Features Overview*
- *Virtual Chassis Features on the QFX Series*
- *Virtual Chassis Fabric Hardware Documentation*

- [Virtual Chassis Fabric Overview on page 7033](#)
- *QFX Series Software Features Overview*
- *Virtual Chassis Features on the QFX Series*
- *Virtual Chassis Fabric Hardware Documentation*
- [Virtual Chassis Fabric Overview on page 7033](#)

## PART 2

# Junos OS Basics

- [Overview on page 11](#)
- [Installation on page 83](#)
- [Configuration on page 151](#)
- [Administration on page 333](#)
- [Troubleshooting on page 1227](#)



## CHAPTER 2

# Overview

- [Software Overview on page 11](#)
- [User Interfaces on page 39](#)
- [Licenses on page 69](#)

## Software Overview

---

- [Configuration File Terms on page 11](#)
- [Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements on page 12](#)
- [In-Service Software Upgrade \(ISSU\) System Requirements on page 13](#)
- [Junos OS Commit Model for Router or Switch Configuration on page 14](#)
- [Junos OS Package Names on page 15](#)
- [Understanding NTP Time Servers on page 16](#)
- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 17](#)
- [Understanding Autoinstallation of Configuration Files on page 19](#)
- [Understanding DHCP Services for Switches on page 21](#)
- [Understanding In-Service Software Upgrade \(ISSU\) on page 25](#)
- [Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric on page 26](#)
- [Understanding Software Infrastructure and Processes on page 29](#)
- [Understanding System Snapshot on page 32](#)
- [Understanding Zero Touch Provisioning on page 32](#)

## Configuration File Terms

[Table 3 on page 11](#) lists the various configuration file terms and their definitions.

**Table 3: Configuration File Terms**

Term	Definition
active configuration	Current committed configuration of a switch.

Table 3: Configuration File Terms (*continued*)

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

**Related  
Documentation**

- [Loading a Previous Configuration File on page 1252](#)
- [Reverting to the Rescue Configuration on page 189](#)
- [Understanding Configuration Files on page 1242](#)

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

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

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

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

- Related Documentation**
- [Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 42](#)

## In-Service Software Upgrade (ISSU) System Requirements

To perform an in-service software upgrade (ISSU), your device must be running Junos OS Release 13.2X51-D15 or later.



**NOTE:** ISSU does not support extension application packages developed with the Junos SDK.



Video: [How Does ISSU Work on the QFX5100?](#)

- [In-Service Software Upgrade \(ISSU\) Protocol and Process Support on page 13](#)

### In-Service Software Upgrade (ISSU) Protocol and Process Support

Table 4 on page 13 lists the protocols and processes that are supported during an ISSU. Protocols that are not supported might cause packet loss.

**Table 4: ISSU Protocol Support**

Protocol	Junos OS Release
Graceful Routing Engine switchover (GRES)	Junos OS 13.2X51-D15 and later
Internet Group Management Protocol (IGMP)	Junos OS 13.2X51-D15 and later
Layer 2 MAC routes	Junos OS 13.2X51-D15 and later
Layer 3 unicast and multicast routes	Junos OS 13.2X51-D15 and later
Layer 2 multicast routes	Junos OS 13.2X51-D15 and later
Link Aggregation Control Protocol (LACP)	Junos OS 13.2X51-D15 and later
<p><b>NOTE:</b> Configure LACP before you issue an ISSU.</p> <p>The LACP <b>periodic fast</b> mode is not supported. Instead, configure the <b>periodic slow</b> mode. If you configure the <b>periodic fast</b> mode, the configuration can be committed without any commit or system log error messages, but you might experience a larger than expected amount of traffic drops. Traffic drops occur because the LACP links go down during an ISSU.</p> <p>Link changes are processed after an ISSU is complete.</p>	
Multicast Listener Discovery (MLD) snooping	Junos OS 13.2X51-D15 and later
Nonstop bridging	Junos OS 13.2X51-D15 and later

Table 4: ISSU Protocol Support (*continued*)

Protocol	Junos OS Release
Nonstop active routing	Junos OS 13.2X51-D15 and later
Spanning tree protocols:	Junos OS 13.2X51-D15 and later
<ul style="list-style-type: none"> <li>• Multiple Spanning Tree Protocol (MSTP)</li> <li>• Rapid Spanning Tree Protocol (RSTP)</li> <li>• Spanning Tree Protocol (STP)</li> <li>• VLAN Spanning Tree Protocol (VSTP)</li> </ul>	

- Related Documentation**
- [Understanding In-Service Software Upgrade \(ISSU\) on page 25](#)
  - [Performing an In-Service Software Upgrade \(ISSU\) on page 119](#)

## Junos OS Commit Model for Router or Switch Configuration

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

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

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

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



**NOTE:** The QFX3500 switch has only one Routing Engine.

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

- Distribution—Propagates a configuration across the routing plane on a multichassis router or switch. Distribution occurs automatically. There is no user command available to control the distribution process. If a configuration is locked during a distribution of



a configuration, the locked configuration does not receive the distributed configuration file, so the synchronization fails. You need to clear the lock before the configuration and resynchronize the routing planes.



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

- Related Documentation**
- *Configuring Junos OS for the First Time on a Router or Switch with a Single Routing Engine*
  - [commit on page 345](#)

## Junos OS Package Names

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

A software package name is in the following format:



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

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

where:

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

A sample switch software package name is:

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

- Related Documentation**
- [Upgrading Software on page 134](#)
  - *Upgrading Software on a QFabric System*

- [Software Installation Overview on page 122](#)

## Understanding NTP Time Servers

The IETF defined the Network Time Protocol (NTP) to synchronize the clocks of computer systems connected to each other over a network. Most large networks have an NTP server that ensures that time on all devices is synchronized, regardless of the device location. If you use one or more NTP servers on your network, ensure you include the NTP server addresses in your Junos OS configuration.

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

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



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

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



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

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

## Related Documentation

- [Configuring the NTP Time Server and Time Services on page 166](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)

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

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

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



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

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

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

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

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

```
user@switch# deactivate class-of-service
```
2. Follow the upgrade procedure to Junos OS Release 11.3 or later software.
3. Make the following changes to the CoS configuration while the CoS configuration is still deactivated:
  - Remove the **excess-rate** statement from the CoS configuration if you have used it at the **[edit class-of-service schedulers]** or **[edit class-of-service traffic-control-profiles]** hierarchy level.
  - Remove the **strict-high** and **strict** priority queue configurations if you have used them at the **[edit class-of-service schedulers]** hierarchy level.
  - Remove the default multidestination forwarding classes (**mcast-be**, **mcast-af**, **mcast-ef**, and **mcast-nc**) if you have used them at the **[edit class-of-service schedulers]**, **[edit class-of-service rewrite-rules]**, **[edit class-of-service classifiers]**, **[edit class-of-service scheduler-maps]**, or **[edit class-of-service forwarding-class-sets]**

hierarchy level. Alternatively, you can change the mapping of the multdestination traffic to use the new default multdestination forwarding class (**mcast**).

4. If desired, configure **strict-high** priority queues in accordance with the Junos OS Release 11.3 or later configuration rules, and map multdestination traffic to the default multdestination forwarding class (**mcast**).

5. Activate the CoS configuration:

```
user@switch# activate class-of-service
```

6. Commit the CoS configuration:

```
user@switch# commit
```



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

#### Related Documentation

- [Upgrading Software on page 134](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\)](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)

## Understanding Autoinstallation of Configuration Files

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

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

This topic describes:

- [Typical Uses for Autoinstallation on page 19](#)
- [Autoinstallation Configuration Files and IP Addresses on page 19](#)
- [Typical Autoinstallation Process on a New Switch on page 20](#)

---

### Typical Uses for Autoinstallation

Typical uses for autoinstallation of the software include:

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

---

### Autoinstallation Configuration Files and IP Addresses

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

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

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

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

### Typical Autoinstallation Process on a New Switch

---

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

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

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

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

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

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

- Related Documentation**
- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 152](#)
  - [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
  - [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)
  - [Configuration Files Terms](#)

## Understanding DHCP Services for Switches

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


This topic describes:

- [DHCP Client/Server Model on page 21](#)
- [Using DHCP on page 22](#)
- [DHCP Relay Servers and DHCP Servers on page 22](#)
- [Legacy DHCP and Extended DHCP for Server Versions on page 22](#)
- [Configuring DHCP on a Switch on page 23](#)
- [How DHCP Works on page 24](#)

### DHCP Client/Server Model

DHCP IP address allocation works on a client/server model in which the server, in this case a switch, assigns the client reusable IP information from an address pool. A DHCP client might receive offer messages from multiple DHCP servers and can accept any one of the offers; however, the client usually accepts the first offer it receives. See [Figure 7 on page 21](#).

**Figure 7: DHCP Client/Server Model**



g041177

## Using DHCP

---

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

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

Another valuable DHCP feature is automatic software download for installation of software packages on switches. DHCP clients configured for automatic software download receive messages as part of the DHCP message exchange process—when the software package name in the DHCP server message is different from that of the software package that booted the DHCP client switch, the new software is downloaded and installed. See [“Upgrading Software by Using Automatic Software Download” on page 148](#).

## DHCP Relay Servers and DHCP Servers

---

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

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

## Legacy DHCP and Extended DHCP for Server Versions

---

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

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

The extended DHCP server version has the following added features:

- Graceful Routing Engine switchover (GRES), which provides mirroring support for clients. For details, see *High Availability Features for EX Series Switches Overview*.



- Virtual routing and forwarding (VRF), which allows multiple instances of a routing table to simultaneously coexist on the same switch. For details, see *Understanding Virtual Routing Instances on EX Series Switches*.



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

Legacy DHCP and extended DHCP servers can be configured at the hierarchy levels shown in [Table 5 on page 23](#):

**Table 5: Legacy DHCP and Extended DHCP Server Hierarchy Levels**

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

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

### Configuring DHCP on a Switch

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

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

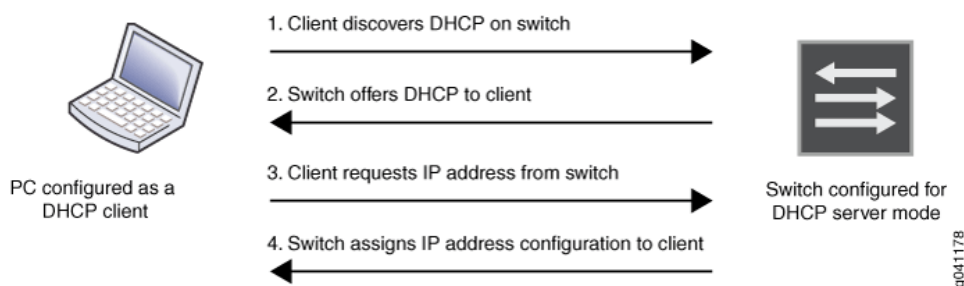
For directions for configuring either a legacy DHCP server or an extended DHCP server, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 155](#).

To configure a DHCP client, set the client’s DHCP interface address in the `[edit interfaces interface-name unit 0 family inet dhcp]` hierarchy. For directions for configuring a DHCP client on a switch, see [“Configuring a DHCP Client \(CLI Procedure\)” on page 154](#).

## How DHCP Works

DHCP consists of a four-step transfer process beginning with a broadcast DHCP discovery message from the client. As the second step, the client receives a DHCP offer message from the server. This message includes the IP address and mask, and some other specific parameters. The client then sends a DHCP request message to accept the IP address and other parameters that it received from the server in the previous step. The DHCP server sends a DHCP response message and removes the now-allocated address from the DHCP address pool. See [Figure 8 on page 24](#).

**Figure 8: DHCP Four-Step Transfer**



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

### Related Documentation

- [Configuring a DHCP Client \(CLI Procedure\) on page 154](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 155](#)
- [Configuring an Extended DHCP Relay Server on EX Series Switches \(CLI Procedure\)](#)
- [Configuring a DHCP SIP Server \(CLI Procedure\)](#)
- [Upgrading Software by Using Automatic Software Download on page 148](#)
- [Monitoring DHCP Services](#)

## Understanding In-Service Software Upgrade (ISSU)

An in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.



Video: [How Does ISSU Work on the QFX5100?](#)



**NOTE:** ISSU is supported in Junos OS Release 13.2X51-D15 and later.

ISSU provides the following benefits:

- Eliminates network downtime during software image upgrades
- Reduces operating costs, while delivering higher service levels
- Allows fast implementation of new features
- [In-Service Software Upgrade Process on page 25](#)

### In-Service Software Upgrade Process

When you request an ISSU on a standalone device:

1. The management process (mgd) verifies that non-stop routing (NSR), graceful Routing Engine switchover (GRES), and non-stop bridging (NSB) are enabled.
2. The switch downloads and validates the software package.
3. The ISSU state machine spawns the backup Routing Engine (RE) with the newer software.
4. The ISSU state machine checks to see if the backup RE has synchronized all of the data with the master RE.
5. The ISSU state machine moves the devices (for example, forwarding ASIC, FPGA, management port and serial console) from the master RE to the backup RE.
6. The mastership is switched between the REs, so the backup RE becomes the master RE.
7. The old master RE is shut down.

#### Related Documentation

- [In-Service Software Upgrade \(ISSU\) System Requirements on page 13](#)
- [Performing an In-Service Software Upgrade \(ISSU\) on page 119](#)

## Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member switches on a Virtual Chassis Fabric with minimal network traffic disruption during the upgrade. A Virtual Chassis Fabric can contain 20 members—up to 2 members can be in the routing engine role, and up to 18 line cards can be configured in the line card role. You can upgrade software for a fixed configuration of switches (QFX3500/QFX3600 and QFX5100 switches, or EX4300 and QFX5100 switches) in a Virtual Chassis Fabric, or for a mixed mode of switches (combination of QFX3500/QFX3600, EX4300, and QFX5100 switches) in a Virtual Chassis Fabric. For information on performing a nonstop software upgrade on a Virtual Chassis, see *Understanding Nonstop Software Upgrade on a Virtual Chassis*.

Performing an NSSU provides these benefits:

- No disruption to the control plane—NSSU uses graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) to ensure no disruption to the control plane. During the upgrade process, interface, kernel, and routing protocol information is preserved.
- Minimal disruption to network traffic—An NSSU minimizes network traffic disruption by:
  - Upgrading line cards one at a time, or in groups, permits traffic to continue to flow through the line cards that are not being upgraded.
  - Upgrading member switches one at a time enables the master and backup to maintain their master and backup roles (although mastership will change) without disruption to traffic.

To achieve minimal disruption to traffic, you must configure link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards or Virtual Chassis Fabric members. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.



**NOTE:** Because NSSU upgrades the software on each line card or on each Virtual Chassis Fabric member one at a time, an upgrade using NSSU can take longer than an upgrade using the `request system software add` command.

You can reduce the amount of time an upgrade takes by configuring line-card upgrade groups. The line cards in an upgrade group are upgraded simultaneously, reducing the amount of time it takes to complete an upgrade. See [“Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\)”](#) on page 84.

---

This topic covers:

- [Requirements for Performing an NSSU on page 27](#)
- [How an NSSU Works on page 27](#)

- [NSSU Limitations on page 28](#)
- [NSSU and Junos OS Release Support on page 28](#)
- [Overview of NSSU Configuration and Operation on page 29](#)

### Requirements for Performing an NSSU

The following requirements apply to Virtual Chassis Fabric:

- Graceful Routing Engine switchover (GRES) must be enabled.
- Nonstop active routing (NSR) and nonstop bridging (NSB) must be enabled.



**NOTE:** Using NSB is recommended for any mode of VCF (pre-provisioned, auto-provisioned, and non-provisioned) to avoid loss of Layer 2 control protocol adjacency during a routing engine switchover..



**NOTE:** Issue the `commit synchronize` command to enable NSB and NSR.

- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis Fabric members.

The following are requirements for Virtual Chassis Fabric members:

- Only two pre-provisioned members in the routing engine role are supported. If more than two routing engines are configured, a warning will be issued, and NSSU will stop.
- The Virtual Chassis Fabric members are connected in a spine and leaf topology. A spine and leaf topology prevents the Virtual Chassis Fabric from splitting during an NSSU. Each leaf device must be connected to both spine devices.
- The Virtual Chassis Fabric must be preprovisioned so that the line card role has been explicitly assigned to member switches acting in a line card role, and that the routing engine role has been explicitly assigned to member switches acting in a routing engine role. During an NSSU, the Virtual Chassis Fabric members must maintain their roles—the master and backup must maintain their master and backup roles (although mastership will change), the member switches must remain their routing engine roles, and the remaining switches must maintain their linecard roles.
- A two-member Virtual Chassis Fabric must have **no-split-detection** configured so that the Virtual Chassis Fabric does not split when an NSSU upgrades a member.

### How an NSSU Works

This section describes what happens when you request an NSSU on a Virtual Chassis Fabric:

- [Virtual Chassis Fabric on page 28](#)

### ***Virtual Chassis Fabric***

When you request an NSSU on a Virtual Chassis Fabric:

1. The Virtual Chassis Fabric master verifies that:
  - The backup is online.
  - Graceful Routing Engine switchover (GRES), nonstop active routing (NSR), and nonstop bridging (NSB) is enabled.
  - The Virtual Chassis Fabric has a preprovisioned configuration.
2. The master installs the new software image on the backup and reboots it.
3. The backup resynchronizes with the master.
4. The master installs the new software image on member switches that are in the linecard role and reboots them, one at a time. The master waits for each member to become online and active before starting the software upgrade on the next member.
5. If you have configured upgrade groups, the line cards in the first upgrade group (or the line card in slot 0, if no upgrade groups are defined) download the new image and then restart. Traffic continues to flow through the line cards in the other upgrade groups during this process.
6. When line cards restarted are online again, the line cards in the next upgrade group download the new image and restart. This process continues until all online line cards have restarted with the new software.
7. The master becomes the backup, and the upgraded backup becomes the master.
8. The software on the original master is upgraded and the original master is automatically rebooted. After the original master has rejoined the Virtual Chassis Fabric, you can optionally return control to it by requesting a graceful Routing Engine switchover.

### **NSSU Limitations**

---

You cannot use an NSSU to downgrade the software—that is, to install an earlier version of the software than is currently running on the switch. To install an earlier software version, use the **request system software add** command.

You cannot roll back to the previous software version after you perform an upgrade using NSSU. If you need to rollback to the previous software version, you can do so by rebooting from the alternate root partition if you have not already copied the new software version into the alternate root partition.

### **NSSU and Junos OS Release Support**

---

A Virtual Chassis Fabric must be running a Junos OS release that supports NSSU before you can perform an NSSU. If a Virtual Chassis Fabric is running a software version that does not support NSSU, use the **request system software add** command.

[Table 6 on page 29](#) lists the QFX Series switches and Virtual Chassis that support NSSU and the Junos OS release at which they began supporting it.

**Table 6: Platform and Release Support for NSSU on a Virtual Chassis Fabric**

Platform	Junos OS Release
Virtual Chassis Fabric	13.2X51-D20 or later

### Overview of NSSU Configuration and Operation

You must ensure that the configuration meets the requirements described in [“Requirements for Performing an NSSU” on page 27](#). NSSU requires no additional configuration.

You perform an NSSU by executing the [request system software nonstop-upgrade](#) command. For detailed instructions on how to perform an NSSU, see the topics in Related Documentation.

#### Related Documentation

- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade on page 139](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2270](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2272](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2274](#)
- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\) on page 84](#)

## Understanding Software Infrastructure and Processes

Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the switch.

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

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

- [Routing Engine and Packet Forwarding Engine on page 30](#)
- [Junos OS Processes on page 30](#)

## Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

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

## Junos OS Processes

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

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

Table 7 on page 30 describes the primary Junos OS processes.

**Table 7: Junos OS Processes**

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



Table 7: Junos OS Processes (*continued*)

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

- Related Documentation**
- [Junos OS Baseline Network Operations Guide](#)
  - [Junos OS Administration Library for Routing Devices](#)

## Understanding System Snapshot



**NOTE:** On QFX3500 and QFX3600 switches running Enhanced Layer 2 Software, all of the directories that reside in the “/” partition are read only.

You can create copies of the software running on a switch using the system snapshot feature. The system snapshot feature takes a “snapshot” of the files currently used to run the switch—the complete contents of the `/config` and `/var` directories, which include the running Junos OS, the active configuration, and the rescue configuration—and copies all of these files into an alternate (internal, meaning internal flash, or an external, meaning USB flash) memory source. You can then use this snapshot to boot the switch at the next boot up or as a backup boot option.

You can only use snapshots to move files to external memory if the switch was booted from internal memory, or to move files to internal memory if the switch was booted from external memory. You cannot create a snapshot in the memory source that booted the switch even if the snapshot is being created on a different partition in the same memory source.

Snapshots are particularly useful for moving files onto USB flash drives. You cannot use the `copy` command or any other file-moving technique to move files from an internal memory source to USB memory on the switch.

System snapshots on the switch have the following limitations:

- You cannot use snapshots to move files to any destination outside of the switch other than an installed external USB flash drive.
- Snapshot commands are always executed on a local switch.

- Related Documentation**
- [Creating a Snapshot and Using It to Boot a QFX Series Switch on page 178](#)

## Understanding Zero Touch Provisioning

- [Understanding Zero Touch Provisioning on page 33](#)
- [Zero Touch Provisioning Process on page 34](#)
- [Zero Touch Provisioning Restart Process Triggers on page 37](#)

## Understanding Zero Touch Provisioning



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

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

The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If the DHCP server does not respond or provide the software image and configuration files, the switch boots with the preinstalled software and default factory configuration. On switches running Enhanced Layer 2 Software, Junos Extended Dynamic Host Configuration Protocol (JDHCP) is used instead of legacy DHCP. JDHCP supports the same functionality as DHCP, and all configuration options remain the same. JDHCP is an enhanced version of legacy DHCP software.



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

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

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



**NOTE:** On QFX3500 and QFX3600 switches running the original CLI, you cannot use ZTP to upgrade from Junos OS Release 12.2 or later to Junos OS Release 13.2X51-D15 or later.

## Zero Touch Provisioning Process

---

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



**NOTE:** If you are performing Zero Touch Provisioning with a Junos OS image that contains enhanced automation for the QFX5100 switch, configure root authentication, and the provider name, license type, and deployment scope for Chef and Puppet at the [edit system] hierarchy in the configuration file that is fetched from the server:

```
{master:0}
root# set root-authentication (encrypted-password password | plain-text-password
password | ssh-dsa public-key | ssh-rsa public-key)
root# set extensions providers juniper license-type customer deployment-scope
commercial
root# set extensions providers chef license-type customer deployment-scope
commercial
```

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



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

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

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

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



**NOTE:** On EX4300 and QFX5100 switches running Enhanced Layer 2 Software, and QFX5100 switches running a Junos OS image that contains enhanced automation, you can specify the name of a script file or a configuration file in suboption 01. ZTP determines if the file is a script file based on the first line that is included in the file. If the first line contains

`#!` characters followed by an interpreter path—for example, `#!/usr/libexec/ui/cscript`—ZTP determines that the file is a script file, and executes the script file with the specified interpreter path. If the script returns an error, ZTP will fetch the script file and execute the script file until the script executes successfully. If the file does not contain special characters or an interpreter path, ZTP determines that the file is a configuration file.

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

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

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



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

.....



**NOTE:** DHCP option 43 suboptions 05 through 255 are reserved.

.....

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



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

.....

7. (Optional) If DHCP option 7 is specified, you can configure one or more system log (syslog) servers.

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

### Zero Touch Provisioning Restart Process Triggers

---

ZTP restarts when any of the following events occur:

- Request for configuration file, script file, or image file fails.
- Configuration file is incorrect, and commit fails.
- No configuration file and no image file is available.
- Image file is corrupted, and installation fails.
- No file server information is available.
- DHCP client does not have valid ZTP parameters configured.
- When none of the DHCP client interfaces goes to a bound state.
- ZTP transaction fails after six attempts to fetch configuration file or image file.

When any of these events occur, ZTP resets the DHCP client state machine on all of the DHCP client-configured interfaces (management and network) and then restarts the state machine. Restarting the state machine enables the DHCP client to get the latest DHCP server-configured parameters.

Before ZTP restarts, approximately 15 to 30 seconds must elapse to allow enough time to build a list of bound and unbound DHCP client interfaces.

The list of bound and unbound DHCP client interfaces can contain:

- No entries.
- Multiple DHCP client interfaces.

Priority is given to the DHCP client interfaces that have received all ZTP parameters (software image file, configuration file, and file server information) from the DHCP server.

After the lists of bound and unbound client interfaces are created, and a DHCP client gets selected for ZTP activity, then any existing default route is deleted, and the DHCP client interface that was selected adds a new default route. In order to add a new default route, only one ZTP instance can be active.

After ZTP restarts, the DHCP client attempts fetching files from the DHCP server for up to six times, with ten to fifteen seconds elapsing between attempts. Every attempt, whether successful or not, is logged and can be seen on the console.

If there is a failure, or the number of attempts exceeds the limit, ZTP stops. ZTP then clears the DHCP client bindings and restarts state machine on the DHCP-configured interfaces.

The ZTP restart process continues until there is either a successful software upgrade, or an operator manually commits a user configuration and deletes the ZTP configuration.

**Related  
Documentation**

- [Configuring Zero Touch Provisioning on page 87](#)
- [Monitoring Zero Touch Provisioning on page 336](#)



---

## User Interfaces

---

- [CLI User Interface Overview on page 39](#)
- [Configuring Login Tips on page 41](#)
- [Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 42](#)
- [Getting Started with Enhanced Layer 2 Software on page 43](#)
- [Junos OS Operational Mode Commands That Combine Other Commands on page 57](#)
- [Overview of Junos OS CLI Operational Mode Commands on page 58](#)
- [Overview of Navigating the CLI on page 60](#)
- [Understanding the Brief, Detail, Extensive, and Terse Options of Junos OS Operational Commands on page 62](#)
- [Understanding Junos OS CLI Configuration Mode on page 63](#)

### CLI User Interface Overview

- [CLI Overview on page 39](#)
- [CLI Key Features on page 39](#)
- [CLI Command Modes on page 40](#)

---

#### CLI Overview

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

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

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

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

---

#### CLI Key Features

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

- Consistent command names—Commands that provide the same type of function have the same name, regardless of the portion of the software on which they are operating.

For example, all **show** commands display software information and statistics, and all **clear** commands erase various types of system information.

- Lists and short descriptions of available commands—Information about available commands is provided at each level of the CLI command hierarchy. If you type a question mark (?) at any level, you see a list of the available commands along with a short description of each command. This means that if you already are familiar with the Junos OS, you can use many of the CLI commands without referring to the documentation.
- Command completion—Command completion for command names (keywords) and for command options is available at each level of the hierarchy. To complete a command or option that you have partially typed, press Tab or the Spacebar. If the partially typed letters begin a string that uniquely identifies a command, the complete command name appears. Otherwise, a beep indicates that you have entered an ambiguous command, and the possible completions are displayed. Completion also applies to other strings, such as filenames, interface names, usernames, and configuration statements.

---

## CLI Command Modes

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

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

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

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

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

When you commit the candidate configuration, you can require an explicit confirmation for the commit to become permanent by using the **commit confirmed** command. This is useful for verifying that a configuration change works correctly and does not prevent management access to the switch. After you issue the **commit confirmed** command, you

must issue another **commit** command within the defined period of time (10 minutes by default), or the system reverts to the previous configuration.

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

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

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

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

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

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

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

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

#### Related Documentation

- [Configuring Login Tips on page 41](#)
- [Overview of Navigating the CLI on page 60](#)
- *CLI User Guide*
- [Other Tools to Configure and Monitor Devices Running Junos OS on page 339](#)

## Configuring Login Tips

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

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

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

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

- Related Documentation**
- [CLI User Interface Overview on page 39](#)
  - [Defining Junos OS Login Classes](#)
  - [login-tip on page 281](#)

## Format for Specifying Filenames and URLs in Junos OS CLI Commands

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

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

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

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



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

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

```
user@host> file delete lcc0-re0:/var/tmp/junk
```
- **a:filename** or **a:path/filename**—File on the local removable media. The default path is **/** (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- **hostname:/path/filename**, **hostname:filename**, **hostname:path/filename**, or **"scp://hostname/path/filename"**—File on an **scp/ssh** client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify **hostname** as **username@hostname**.

- **ftp://hostname/path/filename**—File on an FTP server. You can also specify **hostname** as **username@hostname** or **username:password@hostname**. The default path is the user's home directory. To specify an absolute path, the path must start with **%2F**; for example, **ftp://hostname/%2Fpath/filename**. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required and you do not specify the password or **prompt**, an error message is displayed:

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

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

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

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



**NOTE:** You cannot specify a URL for a file on an HTTP server, because HTTP URLs are not writable.

#### Related Documentation

- [Format for Specifying IP Addresses, Network Masks, and Prefixes in Junos OS Configuration Statements on page 12](#)
- [Default Directories for Junos OS File Storage on the Router or Switch](#)

## Getting Started with Enhanced Layer 2 Software

- [Understanding Enhanced Layer 2 Software Support on page 43](#)
- [Using the ELS Translator Tool on page 44](#)
- [Configuring a VLAN on page 45](#)
- [Configuring the Native VLAN Identifier on page 46](#)
- [Configuring Layer 2 Interfaces on page 46](#)
- [Configuring Layer 3 Interfaces on page 46](#)
- [Configuring an IRB Interface on page 47](#)
- [Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface on page 47](#)
- [Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 48](#)

### Understanding Enhanced Layer 2 Software Support

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

ELS is available on the following EX Series switches and QFX Series devices.

Table 8: ELS Support

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

ELS is supported on the EX4300, EX4600, and EX9200 switches for all Junos OS releases, starting with the initial releases shown in [Table 8 on page 44](#).

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

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



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

### Using the ELS Translator Tool

The ELS Translator is a web-based tool that converts Junos OS Layer 2 configurations to Enhanced Layer 2 Software (ELS) configurations. This conversion tool supports all Juniper Networks EX Series, MX Series, and QFX Series platforms with ELS installed. The ELS Translator is hosted on Juniper Networks Customer Support website for EX Series switches, MX Series Universal Edge routers, and QFX Series switches and is available to registered users, internal users, partners, and premium service contract customers. You need to login using your Juniper Networks user name and password to access the ELS Translator tool.

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

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

1. Log onto your device using the console port.



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

2. Copy your entire existing configuration into another file. Save the file to a remote location. See [“Saving a Configuration to a File” on page 1257](#).
3. Retain the portion of your existing configuration related to management network connectivity (such as `[edit system]`). Delete all other top-level configuration hierarchy levels (such as `[edit interfaces]`, `[edit protocols]`, and `[edit vlans]`). Issue a **commit** operation to remove the deleted configuration hierarchy levels.
4. Perform the software upgrade. Reboot your device to complete the upgrade. See [“Software Installation Overview” on page 122](#)



**NOTE:** Maintain your console port connection during the reboot.

5. [Click](#) to access the ELS translator tool in a web browser. Follow the instructions on the page to update your configuration.
6. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator Tool onto your switch. See [“Uploading a Configuration File” on page 1261](#).
7. Commit the new configuration.



**NOTE:** It is possible a script might not translate correctly, so review translated scripts carefully before loading the converted configuration on your switch or other device.

## Configuring a VLAN

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface. EX Series and QFX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN.

To configure a VLAN:

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

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

2. Assign at least one interface to the VLAN:

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

---

### Configuring the Native VLAN Identifier

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

To configure the native VLAN ID:

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

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

2. Configure the native VLAN ID:

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

3. Assign the interface to the native VLAN ID:

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

---

### Configuring Layer 2 Interfaces

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

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for trunk interface mode:

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

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for access interface mode:

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

---

### Configuring Layer 3 Interfaces

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



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

To specify an IP address for the logical unit using IPv4:

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

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

To specify an IP address for the logical unit using IPv6:

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

### Configuring an IRB Interface

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

To configure an IRB interface:

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

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

2. Create an IRB logical interface:

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

3. Associate the IRB interface with the VLAN:

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

### Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface

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

To configure an aggregated Ethernet interface:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

To configure LACP:

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

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

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

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

---

### Enhanced Layer 2 CLI Configuration Statement and Command Changes

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



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

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

**Table 9: Enhanced Layer 2 CLI Changes**

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   analyzer {     name {       ...     }   } } </pre>	<pre> forwarding-options {   analyzer {     name {       ...     }   } } </pre>	Statements moved to different hierarchy.
<pre> ethernet-switching-options {   authentication-whitelist {     ...   } } </pre>	<pre> switch-options {   ...   authentication-whitelist {     ...   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   bpdu-block {     ...   } } </pre>	<pre> protocols {   layer2-control {     bpdu-block {       ...     }   } } </pre>	Statement moved to different hierarchy.
<pre> ethernet-switching-options {   dot1q-tunneling {     ether-type (0x8100   0x88a8   0x9100);     ...   } } </pre>	<pre> interfaces interface-name {   ether-options {     ethernet-switch-profile {       tag-protocol-id [tpids];     }   } }  interfaces interface-name {   aggregated-ether-options {     ethernet-switch-profile {       tag-protocol-id [tpids];     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.

Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   interfaces <i>interface-name</i> {     no-mac-learning;     ...   } } </pre>	<pre> switch-options {   interfaces <i>interface-name</i> {     no-mac-learning;     ...   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   mac-notification {     notification-interval <i>seconds</i>;     ...   } } </pre>	—	Statements deleted.
<pre> ethernet-switching-options {   mac-table-aging-time <i>seconds</i>;   ... } </pre>	<pre> protocols {   l2-learning {     global-mac-table-aging-time <i>seconds</i>;     ...   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   nonstop-bridging; } </pre>	<pre> protocols {   layer2-control {     nonstop-bridging {     }   } } </pre>	Statement moved to different hierarchy.
<pre> ethernet-switching-options {   port-error-disable {     disable-timeout <i>timeout</i>;     ...   } } </pre>	<pre> interfaces <i>interface-name</i> family   ethernet-switching {     recovery-timeout <i>seconds</i>;   } </pre>	Statement replaced with a new statement.
<pre> ethernet-switching-options {   redundant-trunk-group {     group <i>name</i> {       description;       interface <i>interface-name</i> {         primary;       }       preempt-cutover-timer <i>seconds</i>;       ...     }   } } </pre>	<pre> switch-options {   redundant-trunk-group {     group <i>name</i> {       description;       interface <i>interface-name</i> {         primary;       }       preempt-cutover-timer <i>seconds</i>;       ...     }   } } </pre>	Hierarchy renamed.

Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   secure-access-port {     interface (all   <i>interface-name</i>) {       (dhcp-trusted   no-dhcp-trusted );       static-ip <i>ip-address</i> {         mac <i>mac-address</i>;         vlan <i>vlan-name</i>;       }     }   }   vlan (all   <i>vlan-name</i>) {     (arp-inspection   no-arp-inspection );     dhcp-option82 {       disable;       circuit-id {         prefix <i>hostname</i>;         use-interface-description;         use-vlan-id;       }       remote-id {         prefix (<i>hostname</i>   mac   none);         use-interface-description;         use-string <i>string</i>;       }       vendor-id [<i>string</i>];     }     (examine-dhcp   no-examine-dhcp);   }   (ip-source-guard   no-ip-source-guard); } </pre>	<pre> vlans <i>vlan-name</i> forwarding-options{   dhcp-security {     arp-inspection;     group <i>group-name</i> {       interface <i>interface-name</i> {         static-ip <i>ip-address</i> {           mac <i>mac-address</i>;         }       }     }     overrides {       no-option-82;       trusted;     }   }   ip-source-guard;   no-dhcp-snooping;   option-82 {     circuit-id {       prefix {         host-name;         routing-instance-name;       }       use-interface-description (device           logical);       use-vlan-id;     }     remote-id {       host-name;       use-interface-description (device           logical);       use-string <i>string</i>;     }     vendor-id {       use-string <i>string</i>;     }   } } </pre>	<p>Statements moved to different hierarchy.</p> <p><b>NOTE:</b> The statement <b>examine-dhcp</b> does not exist in the changed hierarchy. Instead, DHCP snooping is enabled automatically when other DHCP security features are enabled on a VLAN. See <i>Configuring Port Security (CLI Procedure)</i> for additional information.</p>
<pre> ethernet-switching-options {   secure-access-port {     dhcp-snooping-file {       location <i>local_pathname</i>   <i>remote_URL</i>;       timeout <i>seconds</i>;       write-interval <i>seconds</i>;     }   } } </pre>	<pre> system [   processes [     dhcp-service     dhcp-snooping-file <i>local_pathname</i>         <i>remote_URL</i>;     write-interval <i>interval</i>;   ] ] </pre>	<p>Statement moved to different hierarchy.</p>
<pre> ethernet-switching-options {   secure-access-port vlan (all   <i>vlan-name</i>{     mac-move-limit   } } </pre>	<pre> vlans <i>vlan-name</i> switch-options {   mac-move-limit } </pre>	<p>Statement moved to different hierarchy.</p>

Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   static {     vlan <i>vlan-id</i> {       mac <i>mac-address</i> next-hop         <i>interface-name</i>;       ...     }   } } </pre>	<pre> vlangs {   <i>vlan-name</i> {     switch-options {       interface <i>interface-name</i> {         static-mac <i>mac-address</i>;         ...       }     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   storm-control {     (...)   } } </pre>	<pre> forwarding-options {   storm-control-profiles <i>profile-name</i> {     (...)   } }  interfaces <i>interface-name</i> unit <i>number</i> family   ethernet-switching {     storm-control <i>storm-control-profile</i>;   } </pre>	Storm control configuration is done in two steps. The first step is to create a storm control profile at the [edit forwarding-options] hierarchy, and the second step is to bind the profile to a logical interface at the [edit interfaces] hierarchy. See <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i> for additional information.
<pre> ethernet-switching-options {   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt;       &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable           no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;;     ...   } } </pre>	—	Statements removed.
<pre> ethernet-switching-options {   unknown-unicast-forwarding {     (...)   } } </pre>	<pre> switch-options {   unknown-unicast-forwarding {     (...)   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	<pre> switch-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	Hierarchy renamed.

Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> interfaces <i>interface-name</i> {   ether-options {     link-mode <i>mode</i>;     speed (auto-negotiation   <i>speed</i>)   } } </pre>	<pre> interfaces <i>interface-name</i> {   link-mode <i>mode</i>;   speed <i>speed</i> } </pre>	Statements moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       native-vlan-id <i>vlan-id</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   native-vlan-id <i>vlan-id</i> } </pre>	Statement moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       port-mode <i>mode</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       interface-mode <i>mode</i>     }   } } </pre>	Statement replaced with a new statement.
interfaces <i>vlan</i>	interfaces <i>irb</i>	Statement replaced with a new statement.

Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> protocols {   igmp-snooping {     traceoptions {       file filename &lt;files number&gt;       &lt;no-stamp&gt; &lt;replace&gt;       &lt;size maximum-file-size&gt;       &lt;world-readable         no-world-readable&gt;;       flag flag &lt;flag-modifier&gt; &lt;disable&gt;;     }     vlan (all   vlan-identifier) {       disable;       data-forwarding {         receiver {           install;           source-vlans vlan-name;         }         source {           groups ip-address;         }       }       immediate-leave;       interface (all   interface-name) {         multicast-router-interface;         static {           group multicast-ip-address;         }         proxy {           source-address ip-address;         }         robust-count number;       }     }   } } </pre>	<pre> protocols {   igmp-snooping {     vlan vlan-name {       immediate-leave;       interface interface-name {         group-limit &lt;1..65535&gt;         host-only-interface         multicast-router-interface;         immediate-leave;         static {           group multicast-ip-address {             source &lt;&gt;           }         }       }     }     l2-querier {       source-address ip-address;     }     proxy {       source-address ip-address;     }     query-interval number;     query-last-member-interval number;     query-response-interval number;     robust-count number;     traceoptions {       file filename &lt;files number&gt;       &lt;no-stamp&gt; &lt;replace&gt;       &lt;size maximum-file-size&gt;       &lt;world-readable         no-world-readable&gt;;       flag flag &lt;flag-modifier&gt;;     }   } } </pre>	IGMP snooping is configured on a VLAN.
<pre> vlans {   vlan-name {     dot1q-tunneling {       customer-vlans (id   native   range);       layer2-protocol-tunneling all         protocol-name {         drop-threshold number;         shutdown-threshold number;         ...       }     }   } } </pre>	<pre> interface interface-name {   encapsulation extended-vlan-bridge;   flexible-vlan-tagging;   native-vlan-id number;   unit logical-unit-number {     input-vlan-map action;     output-vlan-map action;     vlan-id number;     vlan-id-list [vlan-id vlan-id-vlan-id];   } } </pre>	Statements replaced with new statements and moved to different hierarchy



Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> vlsns {   vln-nme {     filter{       input filter-nme       output filter-nme;       ...     }   } } </pre>	<pre> vlsns {   vln-nme {     forwarding-options {       filter{         input filter-nme         output filter-nme;         ...       }     }   } } </pre>	Statements moved to different hierarchy.
<pre> vlsns {   vln-nme {     interface interface-nme {       egress;       ingress;       mapping (native (push   swap)   policy           tag (push   swap));       pvlan-trunk;       ...     }   } } </pre>	—	Statements removed. You can assign interfaces to a VLAN using the [edit interfaces <i>interface-nme</i> unit <i>logical-unit-number</i> family ethernet-switching vln members <i>vln-nme</i> ] hierarchy.
<pre> vlsns {   vln-nme {     isolation-id id-number;     ...   } } </pre>	—	Statement removed.
<pre> vlsns {   vln-nme {     l3-interface vln.logical-interface-number;     ...   } } </pre>	<pre> vlsns {   vln-nme {     l3-interface irb.logical-interface-number;     ...   } } </pre>	Syntax changed.
<pre> vlsns {   vln-nme {     l3-interface-ingress-counting       layer-3-interface-nme;     ...   } } </pre>	—	Statement removed. Ingress traffic is automatically tracked.

Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> vlands {   vlan-name {     mac-limit limit action action;     ...   } } </pre>	<pre> vlands {   vlan-name {     switch-options {       interface-mac-limit limit {         packet-action action;         ...       }     }   } }  vlands {   vlan-name {     switch-options {       interface interface-name {         interface-mac-limit limit {           packet-action action;           ...         }       }     }   } } </pre>	Statements moved to different hierarchies and renamed.
<pre> vlands {   vlan-name {     mac-table-aging-time seconds;     ...   } } </pre>	<pre> protocols {   l2-learning {     global-mac-table-aging-time seconds;     ...   } } </pre>	Statement moved to different hierarchy and renamed.
<pre> vlands {   vlan-name {     no-local-switching;     ...   } } </pre>	—	Statement removed.
<pre> vlands {   vlan-name {     no-mac-learning;     ...   } } </pre>	<pre> vlands {   vlan-name {     switch-options {       no-mac-learning limit       ...     }   } } </pre>	Statement moved to different hierarchy.
<pre> vlands {   vlan-name {     primary-vlan vlan-name;     ...   } } </pre>	—	Statement removed.

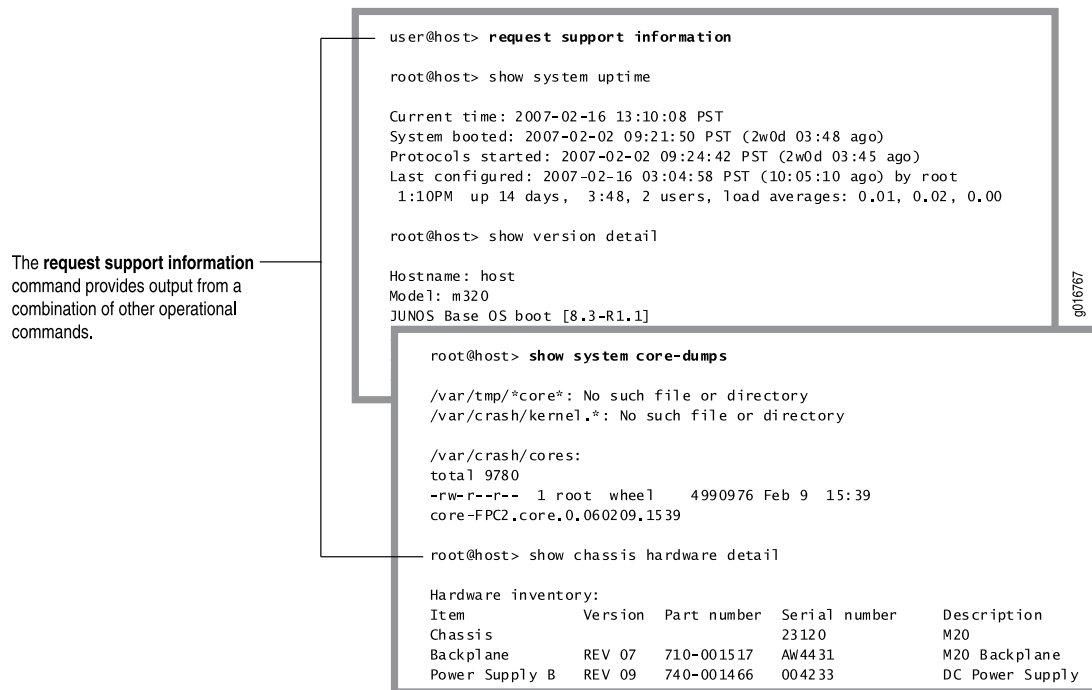
Table 9: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> vlangs {   vlan-name {     vlan-prune;     ...   } } </pre>	—	Statement removed.
<pre> vlangs {   vlan-name {     vlan-range vlan-id-low-vlan-id-high;     ...   } } </pre>	<pre> vlangs {   vlan-name {     vlan-id-list [vlan-id-numbers];     ...   } } </pre>	Statement replaced with new statement.

## Junos OS Operational Mode Commands That Combine Other Commands

In some cases, some Junos OS operational commands are created from a combination of other operational commands. These commands can be useful shortcuts for collecting information about the device, as shown in [Figure 9 on page 57](#).

Figure 9: Commands That Combine Other Commands



### Related Documentation

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

## Overview of Junos OS CLI Operational Mode Commands

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

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

### CLI Command Categories

---

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

- Commands for controlling the CLI environment—Some set commands in the **set** hierarchy configure the CLI display screen. For information about these commands, see *Understanding the Junos OS CLI Modes, Commands, and Statement Hierarchies*.
- Commands for monitoring and troubleshooting—The following commands display information and statistics about the software and test network connectivity. Detailed command descriptions are provided in the *Junos OS Interfaces Command Reference*.
  - **clear**—Clear statistics and protocol database information.
  - **mtrace**—Trace mtrace packets from source to receiver.
  - **monitor**—Perform real-time debugging of various software components, including the routing protocols and interfaces.
  - **ping**—Determine the reachability of a remote network host.
  - **show**—Display the current configuration and information about interfaces, routing protocols, routing tables, routing policy filters, system alarms, and the chassis.
  - **test**—Test the configuration and application of policy filters and autonomous system (AS) path regular expressions.
  - **traceroute**—Trace the route to a remote network host.
- Commands for connecting to other network systems—The **ssh** command opens Secure Shell connections, and the **telnet** command opens telnet sessions to other hosts on the network. For information about these commands, see the [CLI Explorer](#).
- Commands for copying files—The **copy** command copies files from one location on the router or switch to another, from the router or switch to a remote system, or from a remote system to the router or switch. For information about these commands, see the [CLI Explorer](#).
- Commands for restarting software processes—The commands in the **restart** hierarchy restart the various Junos OS processes, including the routing protocol, interface, and SNMP. For information about these commands, see the [CLI Explorer](#).
- A command—**request**—for performing system-level operations, including stopping and rebooting the router or switch and loading Junos OS images. For information about this command, see the [CLI Explorer](#).

- A command—**start**—to exit the CLI and start a UNIX shell. For information about this command, see the [CLI Explorer](#).
- A command—**configure**—for entering configuration mode, which provides a series of commands that configure Junos OS, including the routing protocols, interfaces, network management, and user access. For information about the CLI configuration commands, see “[Understanding Junos OS CLI Configuration Mode](#)” on page 63.
- A command—**quit**—to exit the CLI. For information about this command, see the [CLI Explorer](#).
- For more information about the CLI operational mode commands, see the [CLI Explorer](#).

### Commonly Used Operational Mode Commands

Table 10 on page 59 lists some operational commands you may find useful for monitoring router or switch operation. For a complete description of operational commands, see the Junos OS command references.



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

**Table 10: Commonly Used Operational Mode Commands**

Items to Check	Description	Command
Software version	Versions of software running on the router or switch	<b>show version</b>
Log files	Contents of the log files	<b>monitor</b>
	Log files and their contents and recent user logins	<b>show log</b>
Remote systems	Host reachability and network connectivity	<b>ping</b>
	Route to a network system	<b>traceroute</b>
Configuration	Current system configuration	<b>show configuration</b>
Manipulate files	List of files and directories on the router or switch	<b>file list</b>
	Contents of a file	<b>file show</b>
Interface information	Detailed information about interfaces	<b>show interfaces</b>

Table 10: Commonly Used Operational Mode Commands (*continued*)

Items to Check	Description	Command
Chassis	Chassis alarm status	<b>show chassis alarms</b>
	Information currently on craft display	<b>show chassis craft-interface</b>
	Router or switch environment information	<b>show chassis environment</b>
	Hardware inventory	<b>show chassis hardware</b>
Routing table information	Information about entries in the routing tables	<b>show route</b>
Forwarding table information	Information about data in the kernel's forwarding table	<b>show route forwarding-table</b>
IS-IS	Adjacent routers or switches	<b>show isis adjacency</b>
OSPF	Display standard information about OSPF neighbors	<b>show ospf neighbor</b>
BGP	Display information about BGP neighbors	<b>show bgp neighbor</b>
MPLS	Status of interfaces on which MPLS is running	<b>show mpls interface</b>
	Configured LSPs on the router or switch, as well as all ingress, transit, and egress LSPs	<b>show mpls lsp</b>
	Routes that form a label-switched path	<b>show route label-switched-path</b>
RSVP	Status of interfaces on which RSVP is running	<b>show rsvp interface</b>
	Currently active RSVP sessions	<b>show rsvp session</b>
	RSVP packet and error counters	<b>show rsvp statistics</b>

**Related Documentation**

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

## Overview of Navigating the CLI


This topic describes how to navigate the CLI.

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

## CLI Command Hierarchy

CLI commands are organized in a hierarchy. Commands that perform a similar function are grouped together under the same level of the hierarchy. For example, all commands that display information about the system and the system software are grouped under the **show system** command, and all commands that display information about the routing table are grouped under the **show route** command. [Figure 10 on page 61](#) illustrates a portion of the **show** command hierarchy.

**Figure 10: CLI Command Hierarchy**



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

## CLI Configuration Statements

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

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

## Moving Among Hierarchy Levels

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

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

### Related Documentation

- [CLI User Interface Overview on page 39](#)
- [CLI User Guide](#)

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

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

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

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

2. Choose the option you wish to use with the command. (See [Figure 11 on page 62](#).)

Figure 11: Command Output Options

Command output with the **brief** option.

```
user@host> show interfaces fe-1/1/1 brief
Physical interface: fe-1/1/1, Enabled, Physical link is Down
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback:
  Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
```

Command output with the **terse** option.

```
user@host> show interfaces fe-1/1/1 terse
Interface      Admin Link Proto  Local      Remote
fe-1/1/1       up    down
```

Command output with the **extensive** option.

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

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



## Understanding Junos OS CLI Configuration Mode

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

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

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

This topic covers:

- [Configuration Mode Commands on page 64](#)
- [Configuration Statements and Identifiers on page 65](#)
- [Configuration Statement Hierarchy on page 67](#)

## Configuration Mode Commands

Table 11 on page 64 summarizes each CLI configuration mode command. The commands are organized alphabetically.

**Table 11: Summary of Configuration Mode Commands**

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

Table 11: Summary of Configuration Mode Commands (*continued*)

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

### Configuration Statements and Identifiers

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

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

Table 12 on page 66 describes top-level CLI configuration mode statements.



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

**Table 12: Configuration Mode Top-Level Statements**

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

Table 12: Configuration Mode Top-Level Statements (*continued*)

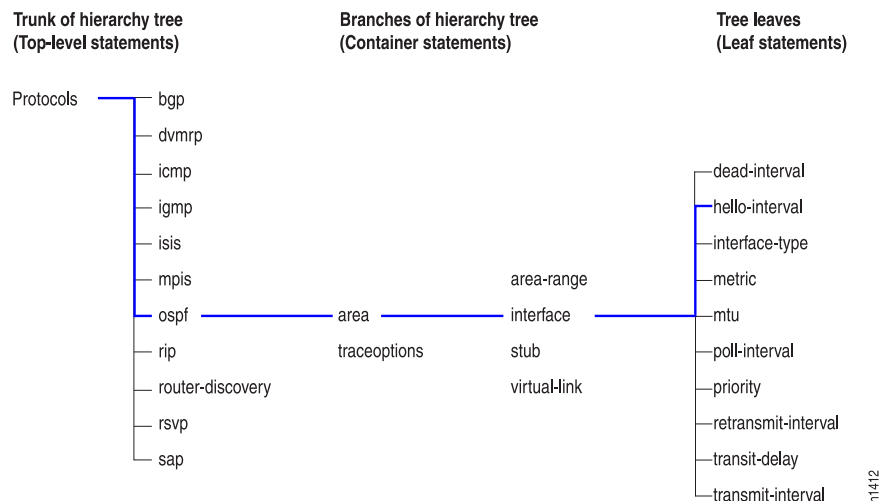
Statement	Description
<b>routing-instances</b>	Configure multiple routing instances. For information about the statements in this hierarchy, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
<b>routing-options</b>	Configure protocol-independent routing options, such as static routes,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

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

### Configuration Statement Hierarchy

The Junos OS configuration consists of a hierarchy of *statements*. There are two types of statements: *container statements*, which are statements that contain other statements, and *leaf statements*, which do not contain other statements (see [Figure 12 on page 67](#)). All of the container and leaf statements together form the *configuration hierarchy*.

Figure 12: Configuration Mode Hierarchy of Statements



Each statement at the top level of the configuration hierarchy resides at the trunk (or root level) of a hierarchy tree. The top-level statements are container statements, containing other statements that form the tree branches. The leaf statements are the leaves of the hierarchy tree. An individual hierarchy of statements, which starts at the trunk of the hierarchy tree, is called a *statement path*. [Figure 12 on page 67](#) illustrates the hierarchy tree, showing a statement path for the portion of the protocol configuration hierarchy that configures the hello interval on an interface in an OSPF area.

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

The CLI represents the statement path shown in [Figure 12 on page 67](#) as **[edit protocols ospf area *area-number* interface *interface-name*]** and displays the configuration as follows:

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

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

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

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

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

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

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

**Related Documentation**

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

## Licenses

- [Junos OS Feature Licenses on page 69](#)
- [Software Features That Require Licenses on the QFX Series on page 70](#)
- [Junos OS Feature License Keys on page 71](#)
- [Generating License Keys on page 75](#)
- [Adding New Licenses \(CLI Procedure\) on page 76](#)
- [Deleting a License \(CLI Procedure\) on page 77](#)
- [Saving License Keys on page 78](#)
- [Verifying Junos OS License Installation on page 79](#)

## Junos OS Feature Licenses

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

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

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

**Related Documentation**

- [License Enforcement](#)
- [Junos OS Feature License Keys on page 71](#)
- [Software Feature Licenses](#)

- [Verifying Junos OS License Installation on page 79](#)

## Software Features That Require Licenses on the QFX Series



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

[Table 13 on page 70](#) lists the licenses you can purchase for each QFX Series software feature.

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

**Table 13: Junos OS Feature Licenses and Model Numbers for QFX Series Devices**

Licensed Software Feature	Supported Devices	Number of Licenses Required	Model Number
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Multi-protocol Label Switching (MPLS)	QFX3500, QFX3600, QFX5100-48S, and QFX5100-48T switches	One per switch	QFX-JSL-EDGE-ADV1
QFX Series advanced feature license for Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Multi-protocol Label Switching (MPLS)	QFX5100-24Q and QFX5100-96S switches	One per switch	QFX5100-HDNSE-LIC
Virtual Extensible Local Area Network (VXLAN) and Open vSwitch Database (OVSDb)	QFX5100-48S and QFX5100-48T switches	One per switch, two per Virtual Chassis and Virtual Chassis Fabric	QFX-JSL-EDGE-ADV1
Virtual Extensible Local Area Network (VXLAN) and Open vSwitch Database (OVSDb)	QFX5100-24Q and QFX5100-96S switches	One per switch, two per Virtual Chassis and Virtual Chassis Fabric	QFX5100-HDNSE-LIC
QFX Series advanced feature license for Border Gateway Protocol (BGP)	QFX3100 Director device	One per Node device in a network Node group	QFX-JSL-DRCTR-ADV1
QFX Series advanced feature license for Fibre Channel	QFX3500 switch	One per switch on which fibre channel ports are configured	QFX-JSL-EDGE-FC



Table 13: Junos OS Feature Licenses and Model Numbers for QFX Series Devices (*continued*)

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

#### Related Documentation

- [Junos OS Feature Licenses on page 69](#)
- [Junos OS Feature License Keys on page 71](#)
- [Generating License Keys on page 75](#)
- [Generating the License Keys for a QFabric System](#)
- [Adding New Licenses \(CLI Procedure\) on page 76](#)
- [Deleting a License \(CLI Procedure\) on page 77](#)
- [Saving License Keys on page 78](#)
- [Verifying Junos OS License Installation on page 79](#)

## Junos OS Feature License Keys

Some Junos OS software features require a license to be activated. To enable each licensed feature, you must purchase, install, manage, and verify a license key that corresponds to the licensed feature.

### Release-Tied License Keys and Upgrade Licenses on MX Series Routers

The Junos OS licensing infrastructure currently associates a license feature with attributes such as date, platform, and validity. In addition to these attributes, for MX Series routers running Junos OS Release 12.2 and later, a licensed feature can be associated with a

release number at the time of generating the license key. This type of release-tied license key is used to validate a particular licensed feature while attempting a software upgrade. The upgrade process aborts if the release number in the license key is earlier than the Junos OS release number to which the system is being upgraded.

Additionally, an upgrade license key can be generated for a release-tied licensed feature. An upgrade license key is used for carrying forward a capacity license to the upgrade release. Although an upgrade license might be an acceptable license on the current release, it does not add to the existing capacity limit. The capacity added in the upgrade license key is valid for the upgrade software release only.

The release number embedded in the license key indicates the maximum release number up to which Junos OS can be upgraded.

As an example, assume that your system is running Junos OS Release 12.2 and is using the **scale-subscriber** licensed feature with a later release-tied upgrade license key installed. If you request a software upgrade to the later release of Junos OS, the software upgrade operation fails and the following error message is displayed:

```
mgd: error: No valid upgrade license found for feature 'scale-subscriber'.  
Aborting Software upgrade.  
Validation failed
```

In this example, to successfully upgrade to the later release of Junos OS, the release number included in the upgrade license key should be greater than or equal to the later release number. Also, you can perform software upgrades up to the previous release without any additional license keys to retain the existing scale limit.

**NOTE:**

When you install a release-tied license, the following apply:

- You can purchase an upgrade capacity license only if a base capacity license for the same scale-tier has already been generated or purchased.
  - You cannot install an upgrade license if the capacity does not match any of the existing base capacity licenses on the system.
  - The license installation fails when you install a lower release number license key on a higher software release number.
  - A release-tied license can be installed on a Junos OS release number that is lower than or equal to the release number included in the license key. For example, a 12.2 license key is valid on Junos OS Release 12.1.
  - An upgrade license is valid only on the target release number specified in the license key, but can be installed on an earlier Junos OS release. For example, a 4 K scale-tier upgrade license for Junos OS Release 12.2 can be installed on an earlier release, and the installed count of licenses remains unaltered.
  - Release-tied licenses of the previous release are not deleted on upgrading Junos OS to a newer release version.
-

### Licensable Ports on MX5, MX10, and MX40 Routers

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

The base capacity of a router is identified by the Ideeprom assembly ID (I2C ID), which defines the board type. However, the Junos OS licensing infrastructure allows the use of restricted ports without a license for a grace period of 30 days. After the grace period expires, the router reverts back to the base capacity if no upgrade license is purchased and installed for the locked ports. The I2C ID along with an upgrade license determine the final capacity of an MX5, MX10, or MX40 router.

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

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

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

[Table 14 on page 73](#) displays the chassis types and their associated port capacity, I2C ID, base capacity, feature ID, feature name, and the final capacity after a license upgrade.

**Table 14: Upgrade Licenses for Enhancing Port Capacity**

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

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

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

### Port Activation on MX104 Routers

Starting with Junos OS Release 13.3, license keys are available to activate the ports on the MX104 router. MX104 routers have four built-in ports. By default, in the absence of valid licenses, all four built-in ports are deactivated. By installing licenses, you can activate any two of the four or all of the four built-in ports. For instance, you can install a license to activate the first two built-in ports (xe-2/0/0 and xe-2/0/1) or you can install a license to activate the next two built-in ports (xe-2/0/2 and xe-2/0/3). You can also install a license to activate all four built-in ports (xe-2/0/0, xe-2/0/1, xe-2/0/2, and xe-2/0/3). If you have already activated two of the built-in ports, you can install an additional license to activate the other two built-in ports on the MX104 router.

A feature ID is assigned to every license for activating the built-in ports on the MX104 router. The port license model with the feature ID is described in [Table 15 on page 74](#).

**Table 15: Port Activation License Model for MX104 Routers**

Feature ID	Feature Name	Functionality
F1	MX104 2X10G Port Activate (0 and 1)	Ability to activate first two built-in ports (xe-2/0/0 and xe-2/0/1)
F2	MX104 2X10G Port Activate (2 and 3)	Ability to activate next two built-in ports (xe-2/0/2 and xe-2/0/3)

Both the features are also provided in a single license key for ease of use. To activate all four ports, you must either install the licenses for both the features listed in [Table 15 on page 74](#) or the single license key for both features. If you install the single license key when feature IDs F1 and F2 are already installed, the license does not get rejected. Also, MX104 routers do not support the graceful license expiry policy. A graceful license expiry policy allows the use of a feature for a certain period of time (usually a grace period of 30 days), and reverts if the license for that feature is not installed after the grace period.

#### Related Documentation

- [Junos OS Feature Licenses on page 69](#)
- *License Enforcement*
- *Software Feature Licenses*

- [Verifying Junos OS License Installation on page 79](#)
- [show system license on page 1044](#)

## Generating License Keys

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

Before generating the license keys for a device:

- Purchase the required licenses for the device. See "[Software Features That Require Licenses on the QFX Series](#)" on page 70.
- Note down the authorization code in the e-mail you received from Juniper Networks when you purchased the license.
- Determine the serial number of the device. For instructions, see *Locating the Serial Number on a QFX3500 Device or Component*.

To generate the license keys for a device:



**NOTE:** This procedure shows you how to generate license keys on a QFX Series device, but you can follow the same procedure for any device.

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

The Manage Product Licenses page appears.



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

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

The Generate Licenses - QFX Series Product page appears.

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

The Generate Licenses - QFX Series Product Devices page appears.

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

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

7. Click **Confirm**.

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

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

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

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

10. Select the delivery method you want to use to obtain your new license keys.

To download the license keys:

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

To e-mail the license keys:

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

#### Related Documentation

- [Software Features That Require Licenses on the QFX Series on page 70](#)
- [Adding New Licenses \(CLI Procedure\) on page 76](#)
- [Locating the Serial Number on a QFX3500 Device or Component](#)

## Adding New Licenses (CLI Procedure)

Before adding new licenses, complete the following tasks:

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



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

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

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

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

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

```
user@host> request system license add terminal
```

2. When prompted, enter the license key, separating multiple license keys with a blank line.

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

3. Go on to [“Verifying Junos OS License Installation” on page 79](#).

On routers that have graceful Routing Engine switchover (GRES) enabled, after successfully adding the new license on the master Routing Engine, the license keys are automatically synchronized on the backup Routing Engine as well. However, in case GRES is not enabled, the new license is added on each Routing Engine separately. This ensures that the license key is enabled on the backup Routing Engine during changeover of mastership between the Routing Engines.

To add a new license key to a router with dual Routing Engines without GRES:

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



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

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

#### Related Documentation

- [Deleting a License \(CLI Procedure\) on page 77](#)
- [Junos OS Feature Licenses on page 69](#)
- [Verifying Junos OS License Installation on page 79](#)
- [request system license add on page 406](#)

## Deleting a License (CLI Procedure)

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

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

1. Display the licenses available to be deleted.

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

```
Possible completions:
```

```
E00468XXX4      License key identifier
JUNOS10XXX1      License key identifier
JUNOS10XXX2      License key identifier
JUNOS10XXX3      License key identifier
JUNOS10XXX4      License key identifier
[               Open a set of values
```

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

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

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

- Delete all license keys from the current device.

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

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

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

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

3. Go on to [“Verifying Junos OS License Installation” on page 79](#).



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

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

#### Related Documentation

- [Adding New Licenses \(CLI Procedure\) on page 76](#)
- [Saving License Keys on page 78](#)
- [Junos OS Feature Licenses on page 69](#)
- [Verifying Junos OS License Installation on page 79](#)
- [request system license delete on page 407](#)

## Saving License Keys

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

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



1. From the CLI operational mode, enter one of the following CLI commands:
  - To save the installed license keys to a file or URL, enter the following command:
 

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

 For example, the following command saves the installed license keys to a file named **license.config**:
  - To save a license key from the terminal, enter the following command:
 

```
user@host> request system license save ftp://user@host/license.config
```
2. Go on to “[Verifying Junos OS License Installation](#)” on page 79.

- Related Documentation**
- [Adding New Licenses \(CLI Procedure\) on page 76](#)
  - [Deleting a License \(CLI Procedure\) on page 77](#)
  - [Junos OS Feature Licenses on page 69](#)
  - [Verifying Junos OS License Installation on page 79](#)

## Verifying Junos OS License Installation

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

- [Displaying Installed Licenses on page 79](#)
- [Displaying License Usage on page 80](#)

### Displaying Installed Licenses

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

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

## Sample Output

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

```

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

```

Licenses installed:
License identifier: E000185416
License version: 2
Features:
  subscriber-acct - Per Subscriber Radius Accounting
                    permanent
  subscriber-auth - Per Subscriber Radius Authentication
                    permanent

```

```

subscriber-addr - Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
permanent
subscriber-ip   - Dynamic and Static IP
permanent

```

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

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



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

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

### Displaying License Usage

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

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

### Sample Output

```

user@host> show system license usage

```

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

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

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

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



## CHAPTER 3

# Installation

- [Software Installation on page 83](#)

### Software Installation

---

- [Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade \(CLI Procedure\) on page 84](#)
- [Configuring Zero Touch Provisioning on page 87](#)
- [Junos OS Package Names on page 92](#)
- [Launching a Guest Virtual Machine \(VM\) to Run a Third Party Application on Junos OS Release 13.2X51-D15 on page 93](#)
- [Launching a Guest Virtual Machine \(VM\) to Run a Third Party Application on Junos OS Release 13.2X51-D20 on page 109](#)
- [Performing a Recovery Installation on page 116](#)
- [Performing a Recovery Installation on page 118](#)
- [Performing an In-Service Software Upgrade \(ISSU\) on page 119](#)
- [Recovering from a Failed Software Installation on page 121](#)
- [Software Installation Overview on page 122](#)
- [Upgrading Jloader Software on QFX Series Devices on page 123](#)
- [Upgrading Software on page 134](#)
- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade on page 139](#)
- [Upgrading Software by Using Automatic Software Download on page 148](#)

## Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)

Nonstop software upgrade (NSSU) enables you to upgrade the software running on an EX Series switch with redundant Routing Engines, on most EX Series Virtual Chassis, QFX3500, QFX3600, and QFX5100 Virtual Chassis, and Virtual Chassis Fabric by using a single command and with minimal disruption to network traffic.

In its default configuration, NSSU upgrades each line card in a switch or Virtual Chassis or Virtual Chassis Fabric one at a time. Traffic continues to flow through the other line cards while a line card is being restarted as part of the upgrade. This behavior allows you to minimize disruption to traffic by configuring link aggregation groups (LAGs) such that the member links of each LAG reside on different line cards. When one member link of a LAG is down, the remaining links are up, and traffic continues to flow through the LAG.

To reduce the time an NSSU takes, you can configure line-card upgrade groups on an EX6200 or EX8200 switch with redundant Routing Engines, EX8200 Virtual Chassis, QFX3500, QFX3600, and QFX5100 Virtual Chassis, and Virtual Chassis Fabric.



**NOTE:** NSSU line-card upgrade groups are not supported for NSSUs on EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4300 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or any mixed Virtual Chassis composed of EX4200, EX4500, or EX4550 switches.

When you define an upgrade group, NSSU upgrades the line cards in the upgrade group at the same time instead of sequentially. To achieve minimal traffic disruption, you must define the line-card upgrade groups such that the member links of the LAGs reside on line cards that are in different upgrade groups. For information on how to configure LAGs, see *Configuring Aggregated Ethernet Links (CLI Procedure)*.

To configure line-card upgrade groups on a standalone EX6200 or EX8200 switch:

- To create an upgrade group and add a line card to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs slot-number
```

For example, to create an upgrade group called **group3** and add the line card in slot 5 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group3 fpcs 5
```

If **group3** already exists, this command adds line card 5 to **group3**.

- To create an upgrade group and add multiple line cards to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name fpcs (NSSU Upgrade Groups)
[list-of-slot-numbers]
```

For example, to create an upgrade group called **primary** and add line cards in slots 1, 4, and 7 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group primary fpcs [1 4 7]
```

If **primary** already exists, this command adds line cards in slots 1, 4, and 7 to **primary**.

To configure line-card upgrade groups on an EX8200 Virtual Chassis:

- To create an upgrade group and add a line card on a Virtual Chassis member to it:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name member (NSSU Upgrade Groups) member-id
fpcs slot-number
```

For example, to create an upgrade group called **primary-ny** and add the line card on member 1 in slot 5 to it:

```
[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 1 fpcs 5
```

If **primary-ny** already exists, this command adds line card 5 on member 1 to **primary-ny**.

- To create an upgrade group that contains multiple line cards on a Virtual Chassis member:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name member member-id fpcs
[list-of-slot-numbers]
```

For example, to create an upgrade group called **primary-ny** that contains the line cards in slots 1 and 2 on member 0 and in slots 3 and 4 on member 1:

```
[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 0 fpcs [1 2]
```

```
[edit chassis]
user@switch# set nssu upgrade-group primary-ny member 1 fpcs [3 4]
```

To configure line-card upgrade groups on a Virtual Chassis Fabric:

- To create an upgrade group and add a line card member on a Virtual Chassis Fabric to the upgrade group:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name primary fpcs value
```

For example, to create an upgrade group called **vcf** and add a line card member:

```
[edit chassis]
user@switch# set nssu upgrade-group vcf primary fpcs 2
```

If **vcf** already exists, this command adds line card 2 to **vcf**.

- To create an upgrade group that contains multiple line cards on a Virtual Chassis Fabric:

```
[edit chassis]
user@switch# set nssu upgrade-group group-name primary fpcs [list-of-slot-numbers]
```

For example, to create an upgrade group called **vcf** that contains line cards 1 and 2:

```
[edit chassis]
user@switch# set nssu upgrade-group vcf primary fpcs [1 2]
```

- Related Documentation**
- *Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches*
  - *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
  - *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
  - [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade on page 139](#)
  - *Understanding Nonstop Software Upgrade on EX Series Switches*



## Configuring Zero Touch Provisioning



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

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

The device uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If the DHCP server does not respond or provide the software image and configuration files, the device continues using the preinstalled Junos OS software and default factory configuration. On switches running Enhanced Layer 2 Software, Junos Extended Dynamic Host Configuration Protocol (JDHCP) is used instead of legacy DHCP. JDHCP supports the same functionality as DHCP, and all configuration options remain the same. JDHCP is an enhanced version of legacy DHCP software. If you are performing Zero Touch Provisioning with a Junos OS image that contains enhanced automation for the QFX5100 switch, you can use DHCP option 43 suboption 01 to run script files, not just load configuration files. Using scripts, you can create device-specific configuration files, and perform HTTP request operations to web servers to download specific configuration files or Junos OS software.



**NOTE:** If the ZTP configuration is enabled, the switch broadcasts DHCP DISCOVER packets on its interfaces. If the DHCP server on the network responds with DHCP vendor options set with the necessary values to initiate ZTP, then ZTP proceeds. To disable broadcasting the DHCP DISCOVER packets without performing the ZTP process, manually delete the **auto-image-upgrade** statement located at the **[edit chassis]** hierarchy. If ZTP completes without errors, the **auto-image-upgrade** statement is automatically deleted.

Before you begin, ensure that the switch has access to the following network resources:

- A DHCP server to lease IP addresses and information on software images and configuration files on the network.

Refer to your DHCP server documentation for configuration instructions.

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



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

- A Domain Name System (DNS) server to perform reverse DNS lookup
- (Optional) An NTP server to perform time synchronization on the network
- (Optional) A system log (syslog) server to manage system log messages and alerts



**CAUTION:** We recommend that you do not commit a user configuration while the device is performing ZTP activity—for example, updating the software image or applying a configuration file.

Perform the following steps to configure ZTP:

1. Boot the device.

The device continues to use the preinstalled Junos OS software and default factory configuration.

2. Issue the **request system zeroize** command on the device.
3. Download the software image file and the configuration file to the FTP, HTTP, TFTP server that the device will download these files from.

You can download either one or both of these files.



**NOTE:** If you are performing Zero Touch Provisioning with a Junos OS image that contains enhanced automation for the QFX5100 device, configure root authentication, and the provider name, license type, and deployment scope for Chef and Puppet at the [edit system] hierarchy in the configuration file that is fetched from the server:

```
{master:0}
root# set root-authentication (encrypted-password password |
plain-text-password password | ssh-dsa public-key | ssh-rsa public-key)
root# set extensions providers juniper license-type customer deployment-scope
commercial
root# set extensions providers chef license-type customer deployment-scope
commercial
```

4. Configure the DHCP server to provide the necessary information to the device.

Configure IP address assignment.

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

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

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

```
option space NEW_OP; option;
option NEW_OP.config-file-name code 1 = text;
option NEW_OP.image-file-type code 2 = text;
option NEW_OP.transfer-mode code 3 = text;
option NEW_OP.alt-image-file-name code 4 = text;
option NEW_OP-encapsulation code 43 = encapsulate NEW_OP;
```

6. Configure the following DHCP option 43 suboptions:

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



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

```
option NEW_OP.image-file-name
"/dist/images/jinstall-ex-4300-13.2R1.1-domestic-signed.tgz";
```

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



**NOTE:** On EX4300 and QFX5100 devices running Enhanced Layer 2 Software, and QFX5100 devices running a Junos OS image that contains enhanced automation, you can specify the name of a script file or a configuration file. ZTP determines if the file is a script file based on the first line that is included in the file. If the first line contains `#!` characters followed by an interpreter path, ZTP determines that the file is a script file, and executes the script file with the specified interpreter path. In order for a script to execute, the script file must provide the ability to fetch and load a valid configuration file on the device during the ZTP process.

The following list provides the types of scripts and their associated interpreter paths:

- Shell script interpreter path: `#!/bin/sh`
- SLAX script interpreter path: `#!/usr/libexec/ui/cscript`
- Python script interpreter path: `#!/usr/bin/python`

Unsigned Python scripts are only supported on limited platforms, such as the QFX5100 device. If you try to execute unsigned Python scripts on devices that do not provide support, error messages will be issued.

If the file does not contain special characters (`#!`), ZTP determines that the file is a configuration file and loads the configuration file.

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

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

**option NEW\_OP.image-file-type "symlink";**



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

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

**option NEW\_OP.transfer-mode "ftp";**



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

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



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



**NOTE:** DHCP option 43 suboptions 05 through 255 are reserved.

**option NEW\_OP.alt-image-file-name  
"/dist/images/jinstall-ex-4300-13.2R1.1-domestic-signed.tgz";**

7.



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

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

**option option-150 code 150 "10.100.31.71";**

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

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

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

**option log-servers 10.100.31.72;**

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

**option ntp-servers 10.100.31.73;**

11. (Optional) Configure DHCP option 12 to specify the hostname of the device.

**option hostname "jn-switch35";**

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

```

host jn-switch35 {
  hardware ethernet ac:4b:c8:29:5d:02;
  fixed-address 10.100.31.36;
  option tftp-server-name "10.100.31.71";
  option host-name "jn-switch35";
  option log-servers 10.100.31.72;
  option ntp-servers 10.100.31.73;
  option NEW_OP.image-file-name
    "/dist/images/jinstall-ex-4300-13.2R1.1-domestic-signed.tgz";
  option NEW_OP.transfer-mode "ftp";
  option NEW_OP.config-file-name "/dist/config/jn-switch35.config";
}

```

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

```

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

```

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



**NOTE:** When SLAX (live operating system based on Linux) scripts are issued, the `op-script.log` and `event-script.log` files are produced.

- `/var/log/dhcp_logfile`
- `/var/log/event-script.log`
- `/var/log/image_load_log`
- `/var/log/messages`
- `/var/log/op-script.log`
- `/var/log/script_output`

You can also monitor the ZTP process by looking at error messages and issuing operational commands. See [“Monitoring Zero Touch Provisioning” on page 336](#) for more information.

**Related Documentation**

- [Understanding Zero Touch Provisioning on page 32](#)

- [Understanding NTP Time Servers on page 16](#)
- [Op Script Overview](#)
- [Monitoring Zero Touch Provisioning on page 336](#)
- [Understanding DHCP Services for Switches on page 21](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 188](#)

## Junos OS Package Names

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

A software package name is in the following format:



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

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

where:

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

A sample switch software package name is:

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

### Related Documentation

- [Upgrading Software on page 134](#)
- [Upgrading Software on a QFabric System](#)
- [Software Installation Overview on page 122](#)

## Launching a Guest Virtual Machine (VM) to Run a Third Party Application on Junos OS Release 13.2X51-D15

- [Understanding Guest VMs on page 93](#)
- [Prerequisites for Setting up a Virtual Build Environment in the JunosV App Engine on page 93](#)
- [Setting up the Virtual Build Environment for the JunosV App Engine on page 96](#)
- [Downloading and Installing the JunosV App Engine Software on page 96](#)
- [Launching the VNC Server on page 97](#)
- [Launching the FreeBSD Virtual Build Environment \(VBE\) Virtual Machine \(VM\) on page 98](#)
- [Installing the Junos SDK Packages on the Virtual Build Environment on page 98](#)
- [Prerequisites for Using the Virtual Build Environment on page 99](#)
- [Obtaining Junos SDK Certificate Request File and Certificate Key File for the Virtual Build Environment on page 99](#)
- [Processing and Obtaining the Certificate File on page 100](#)
- [Prerequisites for Packaging the Guest VM on page 100](#)
- [Launching the Guest VM on the CentOS Server on page 100](#)
- [Copying Required Application to Package with the Guest VM on page 101](#)
- [Editing Packaging Tool Scripts on page 102](#)
- [Executing Packaging Scripts on page 104](#)
- [Copying the Third Party Application to the Switch on page 104](#)
- [Configure the Provider Name, License Type, and Deployment Scope on page 104](#)
- [Configure the Guest VM Options on page 105](#)

### Understanding Guest VMs

You can use a guest virtual machine (VM) to run third party software applications. Guest VMs provide a native environment in which third party applications can be executed, and eliminate the need for porting or adapting third party applications to work on the host OS. You can use the Junos SDK Virtual Build Environment in the JunosV App Engine to package the guest VM images. Once the guest VMs are packaged, you can launch them from the Junos OS CLI.



**NOTE:** Only one guest VM is supported at this time.

### Prerequisites for Setting up a Virtual Build Environment in the JunosV App Engine

Make sure the following prerequisites are met before you set up a Virtual Build Environment in the JunosV App Engine:

- Dedicated server running CentOS 6.2 with a 64-bit processor capable of full hardware virtualization

To find out if the server running CentOS is capable of full hardware virtualization, issue the following command at the shell:

```
egrep '(vmx|svm)' --color=always /proc/cpuinfo
```

If you receive a result with `vmx` or `svm`, the server is capable of virtualization. If you receive a null result, then the server is not capable of virtualization.

The server must have access to the Internet to download and install various Linux, Junos SDK, and JVAE packages.

- CentOS packages installed on the server:

- `kvm`
- `vnc`
- `gcc`
- `make`
- `wget`
- `libvirt`
- `dhcp`
- `dnsmasq`
- `bridge-utils`
- `flex`
- `bison`
- `gcc-c++`
- `glib2-devel`
- `vnc-server`
- `which`
- `xterm`
- `xorg-x11-twm`
- `xorg-x11-server-utils`
- `libXfont`

You can install these packages using the **yum** tool in CentOS.



**NOTE:** For **yum** to work properly, the server must have Internet connectivity, and the DNS servers must be configured.

To see which packages are installed, issue the following command:

```
yum list installed
```



To install all of the packages in the list, issue the following command:

```
yum install kvm vnc gcc make wget libvirt dhcp dnsmasq bridge-utils flex bison gcc-c++ glib2-devel
vnc-server which xterm xorg-x11-twm xorg-x11-server-utils libXfont
```

- Quick Emulator (QEMU) installed for managing VMs.

Issue the following commands to download and install QEMU on your server:

```
wget http://wiki.qemu.org/download/qemu-1.0.1.tar.gz
tar xvf *.gz
cd qemu-1.0.1
./configure
make
make install
```

- Virtual bridge for VM network connectivity is created.

To ensure network connectivity for the Virtual Build Environment (VBE) VM, create a virtual bridge. The virtual bridge allows you to connect the VM to the physical Ethernet interface of the host machine. After performing the following steps, you will have a virtual bridge interface named `virbr0`, which links to the `eth0` interface of the system. When the VM is created, the `virbr0` interface is added to the bridge, and the IP address of the `eth0` is assigned to the `virbr0` interface.

1. Before you configure a virtual bridge, make sure there are no virtual bridges already configured. Issue the following commands to remove any virtual bridges that have been created previously:

```
rm /etc/sysconfig/network-scripts/ifcfg-virbr0
virsh net-destroy default
virsh net-undefine default
service libvirtd restart
```

2. To create a bridge interface named `virbr0` for the virtual bridge and then link it to the physical interface of the system (`eth0`), issue the following commands:

For example:

```
echo 10 > /var/tmp/tapno
chmod 644 /var/tmp/tapno
/etc/sysconfig/network-scripts/qifup
!/bin/sh
/sbin/ifconfig \${1} 0.0.0.0 promisc up
/usr/sbin/brctl addif virbr0 \${1}
/etc/sysconfig/network-scripts/qifdown
!/bin/sh
/sbin/ifconfig \${1} down
/usr/sbin/brctl delif virbr0 \${1}
chmod 755 /etc/sysconfig/network-scripts/qif*
/usr/sbin/brctl addbr virbr0
/usr/sbin/brctl addif virbr0 eth0
cd /etc/sysconfig/network-scripts
cp ifcfg-eth0 ifcfg-virbr0
```

3. Edit the `ifcfg-eth0` file located at the `/etc/sysconfig/network-scripts/` directory and add the following line to specify that the `virbr0` interface is being used as the bridge interface:

```
BRIDGE="virbr0"
```

4. Edit the `ifcfg-virbr0` file located at the `/etc/sysconfig/network-scripts/` directory and modify the following values:

```
DEVICE="virbr0"  
TYPE="Bridge"
```

5. Enable the VNC connection to the VM by adding the following lines to the `iptables` file located at the `/etc/sysconfig/` directory:

```
-A INPUT -i eth0 -j ACCEPT  
-A INPUT -i virbr0 -j ACCEPT
```

6. Issue the following command to restart the network and `iptables` service for the changes to take effect:

```
/sbin/service iptables restart
```

7. Issue the following command to verify that all the changes have taken effect:

```
/usr/sbin/brctl show
```

The output of this command should show that the `virbr0` interface is configured. Verify that there is still network connectivity to the server.

---

### Setting up the Virtual Build Environment for the JunosV App Engine

The JunosV App Engine (JVAE) enables third party applications—applications written in Linux—to run on a guest VM. JunosV App Engine also enables third party applications to run in their native environment without porting to Junos OS.

JVAE provides a virtualized environment with a Kernel-based Virtual Machine (KVM) hypervisor, which runs on the host OS. The host OS controls the creation of virtual machines (VMs) on top of the hypervisor. The hypervisor and host OS run within a compute node. The compute node is connected to a device running Junos OS. In this case, the compute node is connected to a QFX5100 switch.

---

### Downloading and Installing the JunosV App Engine Software

Download the following JunosV App Engine development tools, packaging tools, and sample guest OS packages located at <http://www.juniper.net/support/csc/swdist-junos-sdk/#sw> to your server:

- `junos-sdk-remote-devtools-13.1R1.6.tgz`
- `junos-sdk-remote-pkgtools-13.1R1.6.tgz`
- `junos-sdk-os-13.1R1.6.tgz`

1. Copy the `junos-sdk-remote-devtools-13.1R1.6.tgz` file to the `/usr/src/remote-devtools/` directory.

For example:

```
scp junos-sdk-remote-devtools-13.1R1.6.tgz /usr/src/remote-devtools/
```

2. Extract the `junos-sdk-remote-devtools-13.1R1.6.tgz` file.

For example:

```
tar -zxvf junos-sdk-remote-devtools-13.1R1.6.tgz
```

3. Install the `junos-sdk-remote-devtools-13.1R1.6`.

For example:

```
./setup
```

4. Copy the `junos-sdk-remote-pkgtools-13.1R1.6.tgz` file to the `/usr/src/remote-pkgtools` directory.

For example:

```
scp junos-sdk-remote-pkgtools-13.1R1.6.tgz /usr/src/remote-pkgtools
```

5. Extract the `junos-sdk-remote-pkgtools-13.1R1.6.tgz` file.

For example:

```
tar -zxvf junos-sdk-remote-pkgtools-13.1R1.6.tgz
```

6. Install the `junos-sdk-remote-pkgtools-13.1R1.6` software.

For example:

```
./setup
```

7. Copy the `junos-sdk-os-13.1R1.6.tgz` file to the `/usr/src/sdk-os` directory.

For example:

```
scp junos-sdk-os-13.1R1.6.tgz /usr/src/sdk-os
```

8. Extract the `junos-sdk-os-13.1R1.6.tgz` file.

For example:

```
tar -zxvf junos-sdk-os-13.1R1.6.tgz
```

9. Install the `junos-sdk-os-13.1R1.6` software.

For example:

```
./setup
```

---

## Launching the VNC Server

Launch a VNC server, so you can access a VM.

1. Before you can launch the VNC server, use **yum** to install the **Desktop** package.

For example:

```
yum groupinstall -y Desktop
```

2. Issue the following command to set the VNC password.

You need to set the password when you use VNC for the first time.

For example:

```
vncpasswd
```

3. Issue the following command to launch the VNC server.

For example:

```
vncserver &
```

The name of the desktop is displayed.

For example:

```
vnc-test.juniper.net:1
```

4. Issue the following command to verify that the VNC viewer is working correctly.

For example:

```
vncviewer vnc-test.juniper.net:1
```

A VNC session is created on the server.

### Launching the FreeBSD Virtual Build Environment (VBE) Virtual Machine (VM)

Use the Kernel-based Virtual Machine (KVM) hypervisor to launch the VBE VM.

1. Issue the following commands to launch the VBE VM.

For example:

```
cd /usr/src/remote-pkgtools/junos-sdk-remote-pkgtools  
./start_vm --img /usr/src/sdk-os/junos-sdk-20110408a1/junos-sdk-20110408.img
```

This command launches the VM and returns a port number. The port number is used to establish a VNC connection to the VM.

2. Issue the following command to access the VM.

For example:

```
vncviewer localhost::<port> &
```

This command generates a VNC session to the FreeBSD VBE VM and enables the root password of **letmein**.

### Installing the Junos SDK Packages on the Virtual Build Environment

Before you install the Junos SDK packages, configure an IP address and default gateway on the VBE VM, and ensure that the VBE has proper network connectivity. Also, use the **adduser** tool to create user profiles for the VBE.

1. Issue the following commands to configure the IP address and default gateway on the VBE VM:

For example:

```
ipconfig em0 inet 10.204.42.20 netmask 255.255.255.0  
route add default 10.204.42.20
```

2. Issue the **adduser** command to add user profiles for the VBE.

The **adduser** command provides an interactive guided procedure.

3. Download the following Junos SDK packages located at <http://www.juniper.net/support/csc/swdist-junos-sdk/> to the VBE:

- junos-sdk-ui-sim-13.1R1.6-signed.tgz
- junos-sdk-toolchain-13.1R1.6-signed.tgz
- junos-sdk-sb-13.1R1.6-signed.tgz

4. Issue the following commands to install the Junos SDK packages.

For example:

```
pkg_add junos-sdk-sb-13.1R1.6-signed.tgz  
pkg_add junos-sdk-toolchain-13.1R1.6-signed.tgz  
pkg_add junos-sdk-ui-sim-13.1R1.6-signed.tgz
```

### Prerequisites for Using the Virtual Build Environment

---

Before you can use the Virtual Build Environment to create sandboxes for development, you need to meet the following hardware and software requirements:

- Dedicated server running CentOS 6.2 with a 64-bit processor capable of full hardware virtualization
- Junos SDK Virtual Build Environment (VBE) Virtual Machine (VM) running with access to the Internet on a dedicated server
- Junos SDK packages installed on the VBE VM:
  - junos-sdk-ui-sim-13.1R1.6-signed.tgz
  - junos-sdk-toolchain-13.1R1.6-signed.tgz
  - junos-sdk-sb-13.1R1.6-signed.tgz

### Obtaining Junos SDK Certificate Request File and Certificate Key File for the Virtual Build Environment

---

1. Log into Virtual Build Environment (VBE) Virtual Machine (VM) as root.
2. Issue the following command to launch the **sdk-certificate-request** script.

For example:

```
/usr/local/junos-sdk/13.1R1.6/bin/sdk-certificate-request
```

3. Provide the following information when the script prompts you. Press **Enter** after you provide a response.

- City, state, and country
- Organization and unit
- Provider prefix

This is the unique provider name assigned by Juniper to each SDK partner

- User string

The user string can be a project name, product name, or any generic word.

- Deployment scope

Juniper assigns this string to differentiate multiple certificate for the same partner. If Juniper did not assign this string, you can leave this field empty

- Index number

This number is also referred to as a certificate generations number. The number 1 is used for the initial certificate. After the certificate expires and a new one is requested, this number is increased incrementally.

After you provide all of the information, the script will generate the following files in the **/usr/local/junos-sdk/cert** directory:

- Certificate Key *filename\_key.pem*

This file contains the Junos SDK package-signing key. Ensure that no one outside of the development organization has access to the certificate key. Do not send this file to Juniper

- Certificate Request File *filename\_req.pem*

This file contains the certificate request. Send this file to Juniper for processing

---

### Processing and Obtaining the Certificate File

1. Send the certificate request file to Juniper Junos SDK Certificate Processing Team at [sdk-cert@juniper.net](mailto:sdk-cert@juniper.net)

Once the processing is complete, the Junos SDK Certificate Processing Team will send you the certificate.

2. When you receive the certificate, rename the certificate file as ***filename.pem*** and copy it to the ***/usr/local/junos-sdk/certs*** directory in the VBE VM.
3. Delete the Certificate Request File from the directory.

There should only be one key and certificate pair in the ***/usr/local/junos-sdk/certs*** directory.

---

### Prerequisites for Packaging the Guest VM

Before you can package the guest VM, make sure you meet the following hardware and software requirements:

- Dedicated server running CentOS 6.2 with a 64-bit processor capable of full hardware virtualization.
- Junos SDK Virtual Build Environment (VBE) Virtual Machine (VM) running with access to the Internet on a dedicated server.
  - You must be able to issue a successful ping request to the VBE VM, and SSH and SCP must be enabled on the VBE VM.
  - The VBE VM must have a valid Junos SDK certificate-and-key pair in the ***/usr/local/junos-sdk/certs*** directory.
- Junos SDK packages installed on the VBE VM:
  - *junos-sdk-ui-sim-13.1R1.6-signed.tgz*
  - *junos-sdk-toolchain-13.1R1.6-signed.tgz*
  - *junos-sdk-sb-13.1R1.6-signed.tgz*

---

### Launching the Guest VM on the CentOS Server

1. Issue the ***start\_vm*** script from the Junos SDK Remote Packaging tools directory to launch the guest VM.

For example:

```
cd /usr/src/remote-pkgtools/junos-sdk-remote-pkgtools
./start_vm --img <path>/third-party-app.img -- tapno 1
```

This command launches the VM and returns a port number to which a VNC connection to the VM can be established.

2. Issue the following command to launch the guest VM

For example:

```
vncviewer localhost::port-number &
```

This command generates a VNC session to the Guest VM.

### Copying Required Application to Package with the Guest VM

Before you use `scp` to copy the required applications contained in the **junos-sdk-remote-devtools-13.1R1.6.tgz** file to the Guest VM, configure an IP address and default gateway on the VBE VM, and ensure that the VBE has proper network connectivity.

In this example, an IP address of 10.204.42.40 has been assigned to the Guest VM, and the required application is **remote-helloworld**.

1. Issue the following commands to configure the IP address and default gateway on the VBE VM:

For example:

```
ipconfig em0 inet 10.204.42.20 netmask 255.255.255.0  
route add default 10.204.42.20
```

2. Copy the **remote-helloworld** binary to the **/usr/local/bin** directory and the script to the **/etc/init.d** directory in the Guest VM.

For example:

```
scp /usr/src/remote-devtools/examples/remote-helloworld/remote-helloworld  
root@10.204.42.40:/usr/local/bin  
scp /usr/src/remote-devtools/examples/remote-helloworld/remote-helloworld.sh  
root@10.204.42.40:/etc/init.d
```

3. Rename the **remote-helloworld.sh** script in the **/etc/init.d** folder to **remote-helloworld**.

For example:

```
cd /etc/init.d  
mv remote-helloworld.sh remote-helloworld
```

4. You can now add the **remote-helloworld** application as a startup service.

For example:

```
chkconfig --add remote-helloworld  
chkconfig remote-helloworld on
```

5. Shutdown the VM before packaging the application.

For example:

```
shutdown -h now
```

## Editing Packaging Tool Scripts

---

Before you package the software, you need to modify the **sample.manifest** and **export-user-data.sh** files.

1. Modify the **sample.manifest** file to include the source and destination pairs of the application binaries and scripts to be packaged into the Guest VM.

For example:

```
cd /usr/src/remote-pkgtools/junos-sdk-remote-pkgtools
vi sample.manifest
```

Here is an example of a sample.manifest file:

```
# $Id: sample.manifest 131 2013-02-19 19:24:45Z tomwright $

# Copyright (c) 2012, Juniper Networks
# All rights reserved

# blank lines and lines beginning with # ignored

# dest (guest OS VM) user and host/IP are specified in export-user-data.sh
# dest ending with slash indicates directory

# src dest

/usr/src/remote-devtools/examples/remote-helloworld/remote-helloworld
/usr/local/bin/
/usr/src/remote-devtools/examples/remote-helloworld/remote-helloworld.sh
/etc/init.d/remote-helloworld
```

2. Modify the export-user-data.sh file to include details about the Guest VM .img file, DDL, ODL, and VBE VM details. In this example, an IP address of 10.204.42.20 has been assigned to the VBE VM

For example:

```
cd /usr/src/remote-pkgtools/junos-sdk-remote-pkgtools
vi export-user-data.sh
```

Modify the following values with the correct values:

- APP\_CMD

Name of the command DDL file, including the path (.cmd.dd). Leave this variable undefined if your application does not require a user interface.

In this example,

```
APP_CMD=/usr/src/remote-devtools/examples/remote-helloworld/extensions/libdl/input/remote-helloworldcmd
```

- APP\_CNF

Name of config DDL file, including the path (.cnf.dd). Leave this variable undefined if your application does not require a user interface.

In this example,

```
APP_CNF=/usr/src/remote-devtools/examples/remote-helloworld/extensions/libdl/input/remote-helloworldcnf
```

- APP\_MANIFEST



Name of application manifest file, including the path.

In this example,

`APP_MANIFEST=/usr/src/remote-pkgtools/junos-sdk-remote-pkgtools/sample.manifest`

- APP\_ODL

Name of ODL file, including the path (.odl). Leave this variable undefined if your application does not require formatted output.

In this example,

`APP_ODL=/usr/src/remote-devtools/examples/remote-helloworld/extensions/libodl/input/remote-helloworld`

- LOG\_DIR

Path on development system for log files.

In this example, `LOG_DIR=/var/tmp`

- ROUTER\_PLATFORM

router platform (aka machine) can be i386 (default), octeon, powerpc or xlr

In this example, `ROUTER_PLATFORM=i386`

- VBE\_BSB\_PATH

Path to the Junos SDK Backing Sandbox (BSB) on the VBE.

In this example, `VBE_BSB_PATH=/usr/local/junos-sdk/13.1R1.6`

- VBE\_DSB\_PATH

Path to your application code on the VBE.

In this example, `VBE_DSB_PATH=/usr/home/<user1>/sandboxes/hello-world`

- VBE\_IP

VBE IP address or hostname.

In this example, `VBE_IP=10.204.42.20`

- VBE\_USER

VBE user account.

In this example, `VBE_USER=<user1>`

- VE\_SDK\_DATA\_DEFINED

Set this to 1.

In this example, `VE_SDK_DATA_DEFINED=1`

- VM\_IMG

Guest OS VM file name, including the path.

In this example, `VM_IMG=/root/test/third-party-app.img`

- VM\_IP

Guest OS VM IP address or hostname.

In this example, VM\_IP=10.204.42.40

- VM\_USER

Guest OS VM user account

In this example, VM\_USER=root

---

### Executing Packaging Scripts

You can now execute the Junos SDK packaging scripts to package the guest VM with the required applications into a .tgz file. The Junos SDK VBE VM must be running on the server, with SCP and SSH enabled. The scripts will prompt for the VBE VM user password that you entered in the export-user-data.sh file.

1. Issue the following commands to execute the packaging scripts.

For example:

```
cd /usr/src/remote-pkgtools/junos-sdk-remote-pkgtools
./mksb-vbe-dsb.sh
./do-setup-re-ve-pkg.sh
./update-vbe-dsb.sh
./mk-vbe-dsb.sh
```

The guest VM application packages are now available in the ship directory in the VBE VM in the sandbox location you specified in the VBE\_DSB\_PATH in the export-user-data.sh file.

```
cd /usr/home/user1/sandboxes/hello-world/13.1R1.6-obj/ship
ls -lrt
-rw-r--r--  1 user1 wheel      21040 Aug 22 12:44
third-party-app-i386-13.1I20130822_1944.tgz

-rw-r--r--  1 user1 wheel  223695254 Aug 22 12:45
third-party-app-i386-13.1I20130822_1944.gz

-rw-r--r--  1 user1 wheel  223767502 Aug 22 12:46
third-party-app-bundle-i386-13.1I20130822_1944.gz
```

---

### Copying the Third Party Application to the Switch

1. Copy the third party application to the switch using either FTP or SCP:

For example:

```
root% scp //hostname/pathname/third-party-app.gz /var/tmp
```

---

### Configure the Provider Name, License Type, and Deployment Scope

1. Configure the provider name, the license type, and the deployment scope (describes the certificate associated with the third party application), at the **[edit system]** hierarchy. The certificate contains parameters regarding the provider's partnership with Juniper Networks. Configure these options to ensure that the third party application is installed.

For example:

```
{master:0}
root# set extensions providers [ Provider Name ] license-type customer deployment-scope
[ private commercial ]
```

2. Commit the configuration.

For example:

```
{master:0}
root# commit
```

Here are the results of your configuration:

```
system {
  extensions {
    providers {
      [Provider Name] {
        license-type customer deployment-scope [ private commercial ]
      }
    }
  }
}
```

### Configure the Guest VM Options

1. Configure the following options for guest VM support in the Junos OS CLI at the **[edit]** hierarchy.
  - Compute cluster name
  - Compute node name
  - Virtual machine instance name
  - Dedicated management interface for guest VM
  - Third party package name
  - Internal IP address of the guest VM
2. Configure the name of the compute cluster and compute node.

The name of the compute cluster must be default-cluster, and the name of the name of the compute node must be default-node, otherwise launching the guest VM fails.

For example:

```
{master:0}
root# set services app-engine compute-cluster default-cluster compute-node default-node
hypervisor
```

3. Configure the name of the virtual machine instance and the name of the third party application.



**NOTE:** The package names in the `show app-engine virtual-machine-package` and `show version` commands should match.

```
{master:0}
root# set services app-engine virtual-machines instance instance-name package
package-name
```

For example:

```
{master:0}
root# set services app-engine virtual-machines instance test package third-party-app ve
```

4. Associate the virtual machine instance with the configured compute cluster and compute node.



**NOTE:** The name of the compute cluster must be default-cluster, and the name of the compute node must be default-node, otherwise launching the guest VM fails.

```
{master:0}
root# set services app-engine virtual-machines instance instance-name compute-cluster
name compute-node name
```

For example:

```
{master:0}
root# set services app-engine virtual-machines instance test compute-cluster default-cluster
compute-node default-node
```

5. Configure the local management IP address.

This IP address is used for the internal bridging interface. The host uses this IP address to check the availability of the guest VM. The IP address must be 192.168.1.X, where X is from 100 to 200.

```
{master:0}
root# set services app-engine virtual-machines instance instance-name local-management
family inet address 192.168.1.X
```

For example:

```
{master:0}
root# set services app-engine virtual-machines instance test local-management family inet
address 192.168.1.100
```

6. Configure the management interface for the guest VM.

This management interface is separate from the one used for Junos OS.



**NOTE:** The management interface name must be either em0 or em1. The configuration will fail if you do not configure a management interface and then commit the configuration.

For example:

```
{master:0}
root # set services app-engine virtual-machines instance test management-interface em1
```

The new management interface is provisioned for the guest VM.

7. Commit the configuration.

For example:

```
{master:0}
root# commit
```

Here are the results of the configuration:

```
services {
  app-engine {
    compute-cluster default-cluster {
      compute-node default-node {
```

```

        hypervisor;
    }
}
virtual-machines {
    instance test {
        package third-party-app;
        local-management {
            family inet {
                address 192.168.1.100;
            }
        }
        compute-cluster default-cluster {
            compute-node default-node;
        }
        management-interface em1;
    }
}
}
}
}

```

8. Configure the internal IP address of the guest VM.

- Log into the host shell by specifying the internal management IP address:

For example:

```
shell% ssh -JU __juniper_private4__ 192.168.1.1
```

- Issue the **virsh list** command to see which VMs are running. From the output, you can see that the guest VM (named **test** in this example) is running:

```

{master:0}
shell# virsh list
Id      Name                                State
-----
3       vjunos1                             running
4       test                                running

```

- Log into the guest VM console (named **4** in this example).

For example:

```

shell# virsh console 4
Connected to domain test
Escape character is ^]
CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

[root@localhost ~]

```

From the output, you can see that you are connected to the guest VM console (root@localhost)

9. Configure the internal IP address of the guest VM on the Ethernet interface.

- Issue the **ifconfig -a** command to see the name of the management interface that is used to access the guest VM from outside of the network, and the name of the management interface that is used for internal use.

The interface names are either eth6 or eth7, or eth7 or eth8. You can associate one of the interfaces to the guest VM by issuing the **set services app-engine**

**virtual-machines instance *name* management-interface *interface-name*.** command.  
 Use the same IP address as the one you configured using the **set services app-engine virtual-machines instance test local-management family inet address 192.168.1.100**.  
 The MAC addresses associated with these interfaces are used for internal bridging.

For example:

```
root@localhost ifconfig -a
eth6      Link encap:Ethernet  HWaddr 52:54:00:5D:DB:01
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:982 errors:0 dropped:0 overruns:0 frame:0
          TX packets:970 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:92188 (90.0 KiB)  TX bytes:91468 (89.3 KiB)
eth7      Link encap:Ethernet  HWaddr 52:54:00:5D:DB:02
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:982 errors:0 dropped:0 overruns:0 frame:0
          TX packets:970 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:92188 (90.0 KiB)  TX bytes:91468 (89.3 KiB)
```

In this example, eth6 is a management interface that is used to access the guest VM, and eth7 is a management interface that is used for internal use.

- Issue the **ifconfig** command to configure the internal IP address on the Ethernet interface. This is the same IP address you configured in the Junos OS CLI.

For example:

```
root@localhost ifconfig eth7 192.168.1.100 netmask 255.255.255.0
```

10. Issue the following show commands to verify that everything is working correctly:

For example:

- **root# show app-engine status**

```
Compute cluster: default-cluster
                  Compute node   Status
                  default-node   Online
```

The status should be Online.

- **root# show app-engine virtual-machine instance**

```
VM name          Compute cluster   VM status
                test               default-cluster   ACTIVE
```

The VM status should be active.

- **root# show app-engine virtual-machine instanceshow app-engine virtual-machine package**

```
VM package: cust-vm-ve
              VM disk image:
third-party-app-ve/20140409_015447/third-party-app.img.gz
              Compute cluster   Package download status
              default-cluster   DOWNLOADED
```

The package downloaded status should be either download in progress or downloaded.

- `root# show interfaces terse management-interface`  
error: device em1 not found

This interface should be detached from Junos OS.

11. To remove the guest VM, delete the configuration statements and uninstall the third party software package.

For example, to remove the **app-engine** statement:

```
root # delete services app-engine
Commit the configuration.
```

For example:

```
root# commit
```

Issue the **show version** command to see what the name of the third party application package is.

Issue the **request system software delete <package-name>** command to uninstall the third party application:

For example:

```
root> request system software delete third-party-app.tgz
fpc0:
```

```
-----
Notifying sdk-vmmd ...
```

```
{master:0}
```

#### Related Documentation

- [QFX5100 Guest VM Data Monitoring Application \(ZIP - 2MB\)](#)
- [JunosV App Engine Quick Start Guide](#)
- [Junos SDK Packaging and Deploying Remote Applications Guide](#)
- [Junos SDK Installation Guide](#)

## Launching a Guest Virtual Machine (VM) to Run a Third Party Application on Junos OS Release 13.2X51-D20

- [Understanding Guest VMs on page 109](#)
- [Troubleshooting Tips on page 110](#)
- [Copying the Third Party Application to the Switch on page 110](#)
- [Install the Third Party Application on the Switch on page 110](#)
- [Configure the Guest VM Options to Launch the Guest VM on the Host on page 111](#)

### Understanding Guest VMs

You can use a guest virtual machine (VM) to run third party software applications. Guest VMs provide a native environment in which third party applications can be executed, and eliminate the need for porting or adapting third party applications to work on the host OS. You can use the Junos SDK Virtual Build Environment in the JunosV App Engine to

package the guest VM images. Once the guest VMs are packaged, you can launch them from the Junos OS CLI.



**NOTE:** Only one guest VM is supported at this time.

---

### Troubleshooting Tips

Configure traceoption and System Log options to troubleshoot issues that occur while you are launching a guest VM:

- **set system processes app-engine-virtual-machine-management-service traceoptions level all**
- **set system processes app-engine-virtual-machine-management-service traceoptions flag all**
- **set system syslog file messages any any**

---

### Copying the Third Party Application to the Switch

1. Copy the third party application to the switch using any file transfer protocol:

For example:

```
root% scp //hostname/pathname/third-party-app.img.gz /var/tmp
```

---

### Install the Third Party Application on the Switch

1. Install the third party application package on the switch.

This might take a few minutes.

For example:

```
{master:0}
root> request system software add virtual-machine-package /var/tmp/third-party-app.img.gz
Installing virtual-machine package..
Copying virtual-machine package..
Uncompressing virtual-machine package..
Finished virtual-machine package installation.
```

2. Issue the **show version** command to verify that the installation was successful.

For example:

```
{master:0}
root> show version
Apr 02 09:12:13
fpc0:
-----
Hostname: host
Model: qfx5100-96s-8q
JUNOS Base OS Software Suite [13.2-20140401_x_132_x51_vjunos.0]
JUNOS Base OS boot [13.2-20140401_x_132_x51_vjunos.0]
JUNOS Crypto Software Suite [13.2-20140401_x_132_x51_vjunos.0]
JUNOS Online Documentation [13.2-20140401_x_132_x51_vjunos.0]
JUNOS Kernel Software Suite [13.2-20140401_x_132_x51_vjunos.0]
JUNOS Packet Forwarding Engine Support (qfx-ex-x86-32)
```



```
[13.2-20140401_x_132_x51_vjunos.0]
JUNOS Routing Software Suite [13.2-20140401_x_132_x51_vjunos.0]
JUNOS Enterprise Software Suite [13.2-20140401_x_132_x51_vjunos.0]
JUNOS py-base-i386 [13.2-20140401_x_132_x51_vjunos.0]
third-party-app-ve Virtual Engine package [13.1I20130918_2234]
```

```
JUNOS Host Software [13.2-20140401_x_132_x51_vjunos.0]
```

The CLI output shows that the application named third-party-app was installed.

### Configure the Guest VM Options to Launch the Guest VM on the Host

1. Configure the following options for guest VM support in the Junos OS CLI at the **[edit]** hierarchy.
  - Compute cluster name
  - Compute node name
  - Virtual machine instance name
  - Dedicated management interface for guest VM
  - Third party package name
  - Internal IP address of the guest VM

2. Configure the name of the compute cluster and compute node.

The name of the compute cluster must be default-cluster, and the name of the name of the compute node must be default-node, otherwise launching the guest VM fails.

For example:

```
{master:0}
root# set services app-engine compute-cluster default-cluster compute-node default-node
hypervisor
```

3. Configure the name of the virtual machine instance and the name of the third party application.



**NOTE:** The package names in the `show app-engine virtual-machine-package` and `show version` commands should match.

```
{master:0}
root# set services app-engine virtual-machines instance instance-name package
package-name
```

For example:

```
{master:0}
root# set services app-engine virtual-machines instance test package third-party-app ve
```

4. Associate the virtual machine instance with the configured compute cluster and compute node.



**NOTE:** The name of the compute cluster must be default-cluster, and the name of the compute node must be default-node, otherwise launching the guest VM fails.

```
{master:0}
root# set services app-engine virtual-machines instance instance-name compute-cluster
name compute-node name
```

For example:

```
{master:0}
root# set services app-engine virtual-machines instance test compute-cluster default-cluster
compute-node default-node
```

5. Configure the local management IP address.

This IP address is used for the internal bridging interface. The host uses this IP address to check the availability of the guest VM. The IP address must be 192.168.1.X, where X is from 100 to 200.

```
{master:0}
root# set services app-engine virtual-machines instance instance-name local-management
family inet address 192.168.1.X
```

For example:

```
{master:0}
root# set services app-engine virtual-machines instance test local-management family inet
address 192.168.1.100
```

6. Configure the management interface for the guest VM.

This management interface is separate from the one used for Junos OS.



**NOTE:** The management interface name must be either em0 or em1. The configuration will fail if you do not configure a management interface and then commit the configuration.

For example:

```
{master:0}
root # set services app-engine virtual-machines instance test management-interface em1
```

The new management interface is provisioned for the guest VM.

7. Commit the configuration.

For example:

```
{master:0}
root# commit
```

Here are the results of the configuration:

```
services {
  app-engine {
    compute-cluster default-cluster {
      compute-node default-node {
        hypervisor;
      }
    }
  }
  virtual-machines {
    instance test {
      package third-party-app;
      local-management {
        family inet {
          address 192.168.1.100;
        }
      }
    }
  }
}
```

```

    }
  }
  compute-cluster default-cluster {
    compute-node default-node;
  }
  management-interface em1;
}
}
}
}

```

8. Configure the internal IP address of the guest VM.

- Log into the host shell by specifying the internal management IP address:

For example:

```
shell% ssh -JU _juniper_private4_ 192.168.1.1
```

- Issue the **virsh list** command to see which VMs are running. From the output, you can see that the guest VM (named **test** in this example) is running:

```
{master:0}
shell# virsh list
Id      Name                                State
-----
3       vjunos1                             running
4       test                                 running
```

- Log into the guest VM console (named **4** in this example).

For example:

```
shell# virsh console 4
Connected to domain test
Escape character is ^]
CentOS release 6.4 (Final)
Kernel 2.6.32-358.el6.x86_64 on an x86_64

[root@localhost ~]
```

From the output, you can see that you are connected to the guest VM console (root@localhost)

9. Configure the internal IP address of the guest VM on the Ethernet interface.

- Issue the **ifconfig -a** command to see the name of the management interface that is used to access the guest VM from outside of the network, and the name of the management interface that is used for internal use.

The interface names are either eth6 or eth7, or eth7 or eth8. You can associate one of the interfaces to the guest VM by issuing the **set services app-engine virtual-machines instance *name* management-interface *interface-name*** command. Use the same IP address as the one you configured using the **set services app-engine virtual-machines instance test local-management family inet address 192.168.1.100**. The MAC addresses associated with these interfaces are used for internal bridging.

For example:

```
root@localhost ifconfig -a
```

```

eth6      Link encap:Ethernet  HWaddr 52:54:00:5D:DB:01
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:982 errors:0 dropped:0 overruns:0 frame:0
          TX packets:970 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:92188 (90.0 KiB)  TX bytes:91468 (89.3 KiB)

eth7      Link encap:Ethernet  HWaddr 52:54:00:5D:DB:02
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:982 errors:0 dropped:0 overruns:0 frame:0
          TX packets:970 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:92188 (90.0 KiB)  TX bytes:91468 (89.3 KiB)

```

In this example, eth6 is a management interface that is used to access the guest VM, and eth7 is a management interface that is used for internal use.

- Issue the **ifconfig** command to configure the internal IP address on the Ethernet interface. This is the same IP address you configured in the Junos OS CLI.

For example:

```
root@localhost ifconfig eth7 192.168.1.100 netmask 255.255.255.0
```

10. Issue the following show commands to verify that everything is working correctly:

For example:

- root# **show app-engine status**

```

Compute cluster: default-cluster
                  Compute node      Status
                  default-node      Online

```

The status should be Online.

- root# **show app-engine virtual-machine instance**

```

VM name          Compute cluster      VM status
                test                  default-cluster      ACTIVE

```

The VM status should be active.

- root# **show app-engine virtual-machine instanceshow app-engine virtual-machine package**

```

VM package: cust-vm-ve
              VM disk image:
third-party-app-ve/20140409_015447/third-party-app.img.gz
                  Compute cluster      Package download status
                  default-cluster      DOWNLOADED

```

The package downloaded status should be either download in progress or downloaded.

- root# **show interfaces terse *management-interface***

```
error: device em1 not found
```

This interface should be detached from Junos OS.

11. To remove the guest VM, delete the configuration statements and uninstall the third party software package.

For example, to remove the **app-engine** statement:

```
root # delete services app-engine
```

Commit the configuration.

For example:

```
root# commit
```

Issue the **show version** command to see what the name of the third party application package is.

For example:

```
{master:0}
root> show version
fpc0:
-----
Hostname: st-96s-p2b-03
Model: qfx5100-96s-8q
JUNOS Base OS Software Suite [13.2-20140406_x_132_x51_vjunos.0]
JUNOS Base OS boot [13.2-20140406_x_132_x51_vjunos.0]
JUNOS Crypto Software Suite [13.2-20140406_x_132_x51_vjunos.0]
JUNOS Online Documentation [13.2-20140406_x_132_x51_vjunos.0]
JUNOS Kernel Software Suite [13.2-20140406_x_132_x51_vjunos.0]
JUNOS Packet Forwarding Engine Support (qfx-ex-x86-32)
[13.2-20140406_x_132_x51_vjunos.0]
JUNOS Routing Software Suite [13.2-20140406_x_132_x51_vjunos.0]
JUNOS Enterprise Software Suite [13.2-20140406_x_132_x51_vjunos.0]
JUNOS py-base-i386 [13.2-20140406_x_132_x51_vjunos.0]
  third-party-app-ve Virtual Engine package [20140409_015447]
JUNOS Host Software [13.2-20140406_x_132_x51_vjunos.0]
```

Issue the **request system software delete virtual-machine-package <package-name>** command to uninstall the third party application:

For example:

```
root> request system software delete virtual-machine-package third-party-app-ve
fpc0:
-----
Deleted virtual-machine package cust-vm-ve ...
```

#### Related Documentation

- [QFX5100 Guest VM Data Monitoring Application \(ZIP - 2MB\)](#)
- [JunosV App Engine Quick Start Guide](#)
- [Junos SDK Packaging and Deploying Remote Applications Guide](#)
- [Junos SDK Installation Guide](#)

## Performing a Recovery Installation

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

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

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device” on page 176](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



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

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

To reinstall Junos OS:

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



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

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

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

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

3. The software prompts you with the following options:

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

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

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

The following messages are displayed:

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

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

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

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

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

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

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

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

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

**Related Documentation**

- [Creating an Emergency Boot Device on page 176](#)

## Performing a Recovery Installation

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

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

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device” on page 176](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



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

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

To reinstall Junos OS:

1. Insert the emergency boot device into the device.
2. Power cycle the device.

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

3. The software prompts you with the following options:

```
Junos Snapshot Installer - (c) Juniper Networks 2013
Reboot
Install Junos Snapshot
[13.2-20131115_x_132_x51_vjunos.0Boot to host shell [debug]
```

4. Select **Install Junos** to format the internal flash storage and install the Junos OS on the emergency boot device onto the internal flash storage.
5. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

When the software is finished being copied from the emergency device to the device, the device reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the Junos OS login prompt:



```
root@switch#
```

6. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.
7. Remove the emergency boot device.

**Related Documentation**

- [Creating an Emergency Boot Device on page 176](#)

## Performing an In-Service Software Upgrade (ISSU)

You can use an in-service software upgrade to upgrade the software running on the switch with minimal traffic disruption during the upgrade.



**NOTE:** ISSU is supported in Junos OS Release 13.2X51-D15 and later on QFX5100 switches, and in Junos OS Release 13.2X51-D25 and later on EX4600 switches.

This topic covers:

1. [Preparing the Switch for Software Installation on page 119](#)
2. [Upgrading the Software Using ISSU on page 120](#)

### Preparing the Switch for Software Installation

Before you begin software installation using ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:



**NOTE:** If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see [“Configuring Nonstop Active Routing on Switches” on page 2277](#) for information on how to enable it.

- Enable nonstop bridging (NSB). See [“Configuring Nonstop Bridging on Switches \(CLI Procedure\)” on page 2272](#) for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the [request system snapshot](#) command.

## Upgrading the Software Using ISSU

---

This procedure describes how to upgrade the software running on a standalone switch:

To upgrade the switch using ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in [“Upgrading Software” on page 134](#).
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:

- On the switch, enter:

```
user@switch> request system software in-service-upgrade  
/var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, `jinstall-132_x51_vjunos.domestic.tgz`.



**NOTE:** During the upgrade, you will not be able to access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get  
lost!  
ISSU: Validating Image  
ISSU: Preparing Backup RE  
Prepare for ISSU  
ISSU: Backup RE Prepare Done  
Extracting jinstall-qfx-5-13.2X51-D15.4-domestic ...  
Install jinstall-qfx-5-13.2X51-D15.4-domestic completed  
Spawning the backup RE  
Spawn backup RE, index 0 successful  
GRES in progress  
GRES done in 0 seconds  
Waiting for backup RE switchover ready  
GRES operational  
Copying home directories  
Copying home directories successful  
Initiating Chassis In-Service-Upgrade  
Chassis ISSU Started  
ISSU: Preparing Daemons  
ISSU: Daemons Ready for ISSU  
ISSU: Starting Upgrade for FRUs  
ISSU: FPC Warm Booting  
ISSU: FPC Warm Booted  
ISSU: Preparing for Switchover  
ISSU: Ready for Switchover  
Checking In-Service-Upgrade status
```

Item	Status	Reason
FPC 0	Online (ISSU)	
Send ISSU done to chassisd on backup RE		
Chassis ISSU Completed		
ISSU: IDLE		
Initiate em0 device handoff		



**NOTE:** An ISSU might stop instead of abort if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).



**NOTE:** If the ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

- Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

#### Related Documentation

- [Understanding In-Service Software Upgrade \(ISSU\) on page 25](#)
- [request system software in-service-upgrade on page 436](#)

## Recovering from a Failed Software Installation

**Problem** **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution** If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

- Power on the switch. The loader script starts.
- After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

- Enter the following command:

```
loader> install [--format] [--external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Software Installation Overview

A device is delivered with the Junos OS preinstalled. As new features and software fixes become available, you can upgrade your software to use them.

When you power on the switch, it starts (boots) using the installed software.

You upgrade the Junos OS on a switch by copying a software package to a switch or other system on your local network and then using the CLI to install the new software on the switch. You then reboot the switch, which boots from the upgraded software. After a successful upgrade, you should back up the new current configuration to a secondary device.

During a successful upgrade, the installation package removes all files from the `/var/tmp` directory of the switch and completely reinstalls the existing software. It retains configuration files, and similar information, such as secure shell and host keys, from the previous version. The previous software package is preserved in a separate disk partition, and you can manually revert to it if necessary. If the software installation fails for any reason, such as loss of power during the installation process, the system returns to the originally active installation when you reboot.

If you encounter any difficulties during software installation or an upgrade, you can use the recovery installation procedure to install the Junos OS on the switch.

### Related Documentation

- [Upgrading Software on page 134](#)
- [Upgrading Software on a QFabric System](#)
- [Recovering from a Failed Software Installation on page 121](#)
- [Performing a Nonstop Software Upgrade on the QFabric System](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 7233](#)

- [Performing a Recovery Installation on page 116](#)

## Upgrading Jloader Software on QFX Series Devices

Jloader software contains a boot loader (Uboot), which is used to bring up QFX Series devices and load the Junos OS from the flash memory of these devices. You can upgrade Jloader software on QFX3500 switches, QFX3500 and QFX3600 Node devices, and QFX3600-I and QFX3008-I Interconnect devices.



**NOTE:** Before you upgrade the Jloader software, see [Table 16 on page 123](#), [Table 17 on page 123](#), and [Table 18 on page 124](#) to make sure that you are upgrading to the right version of Jloader software for the Junos OS software release running on your QFX3500 switches, or Node devices and Interconnect devices in your QFabric system.

See [Table 19 on page 124](#) to see which Uboot software versions are available and the filenames of the Jloader software packages.

**Table 16: Junos OS and Jloader Software Compatibility Matrix for the QFX3500 Switch and QFX3500 Node Device**

Junos OS Software Version	1.1.2	1.1.4	1.1.5	1.1.8
11.3R1 and later (QFX3500 switch)	Supported	Supported	Not supported	Supported and recommended
11.3X30.6 and later (QFX3500 Node device)	Supported	Supported	Not supported	Supported and recommended
12.1X49-D1 and later (QFX3500 switch)	Supported	Supported	Not supported	Supported and recommended
12.2X50-D1 and later (QFX3500 switch and QFX3500 Node device)	Supported	Supported	Not supported	Supported and recommended



**NOTE:** An en dash means that the item is not applicable.

**Table 17: Junos OS and Jloader Software Compatibility Matrix for the QFX3008-I Interconnect Device**

Junos OS Software Version	1.1.2	1.1.4	1.1.5	1.1.8
11.3X30.9 and later (QFX3008-I Interconnect device)	Supported	Supported	Not supported	Supported and recommended

**Table 17: Junos OS and Jloader Software Compatibility Matrix for the QFX3008-I Interconnect Device (continued)**

Junos OS Software Version	11.2	11.4	11.5	11.8
11.3X30.6 and later (QFX3008-I Interconnect device)	Supported	Supported	Not supported	Supported and recommended
12.2X50-D10.3 and later (QFX3008-I Interconnect device)	Supported	Supported	Not supported	Supported and recommended



**NOTE:** An en dash means that the item is not applicable.

**Table 18: Junos OS and Jloader Software Compatibility Matrix for the QFX3600-I Interconnect Device and QFX3600 Node Device**

Junos OS Software Version	11.2	11.4	11.5	11.8
12.2X50-D10.3 and later (QFX3600-I Interconnect Device and QFX3600 Node Device)	-	-	Supported	Supported and recommended
12.2X50-D20 and later (QFX3600 switch)	-	-	Supported	Supported and recommended

**Table 19: Uboot Software Release and Jloader Software Compatibility Matrix**

Uboot Software Release Number	Jloader Software Package Name
11.2	jloader-qfx-11.3X30.9-signed.tgz
11.4 (11.3R3 and 11.3R2 releases only. Not supported on 11.3R1)	jloader-qfx-11.3I20120127_0733_dc-builder-signed.tgz
11.4 (12.1R1 release and later)	jloader-qfx-12.1-20120125_pr.0-signed.tgz
11.5 (12.2X50-D10.3 and later)	jloader-qfx-12.2X50.D10.3-signed.tgz
11.8 (13.1X50-D15.1 and later)	jloader-qfx-13.3-20130831_pr_branch_qfd.0.tgz

#### Jloader Software Version 1.1.4 Guidelines

Jloader Release 1.1.4 is compatible with Junos OS Release 11.3R3 and 11.3R2, and Junos OS Release 12.1R1 and later. Jloader Release 1.1.4 is not compatible with Junos OS Release 11.3R1. The Jloader software package names are different for versions 1.1.4 (Junos OS 11.3R3 and 11.3R2) and 1.1.4 (Junos OS 12.2R1 release and later), but the binaries are the same. Because the binaries are the same, you can upgrade or downgrade to any Junos OS release.

- If you have Junos OS Release 11.3 installed and want to upgrade the Jloader software from version 1.1.2 to version 1.1.4, you need to upgrade using the **jloader-qfx-11.3I20120127\_0733\_dc-builder-signed.tgz** software package.
- If you have Junos OS Release 11.3R2 installed and want to upgrade to Junos OS Release 12.1, you do not need to upgrade the Jloader Release and can continue to use Jloader Release 1.1.2.
- If you have Junos OS Release 12.1 installed and want to upgrade the Jloader software from version 1.1.2 to version 1.1.4, you need to upgrade using the **jloader-qfx-12.1-20120125\_pr.0-signed.tgz** software package.
- If you upgrade to Junos OS Release 12.1, you can upgrade to Jloader Release 1.1.4 using the **jloader-qfx-12.1-20120125\_pr.0-signed.tgz** software package.

### Upgrading Jloader Software on a QFX3500 Switch

The Jloader software for a QFX3500 switch resides in two flash memory banks. At any time, one bank acts as the primary bank, and the QFX3500 switch boots from it. The other bank is the backup bank—if the QFX3500 switch cannot boot from the primary bank, it boots from the backup bank. When you upgrade the Jloader software, the upgraded software is installed in the backup bank, which then becomes the new primary bank. Thus the primary and backup banks alternate each time you upgrade the Jloader software, with the primary bank containing the most recently installed version of the software, and the backup bank containing the previous version. To upgrade the Jloader software on a QFX3500 switch, you must perform the upgrade twice: once for each bank. Each upgrade requires that you to reboot the QFX3500 switch.



**NOTE:** If you are running Junos OS Release 11.3R1 or Junos OS Release 11.3R2, you must use the **no-validate** option when you issue the **request system software add** command to upgrade the Jloader software. Otherwise, the installation will fail and you receive a configuration error. The **no-validate** option is not required for Junos OS Release 11.3R3 and later.



**NOTE:** After you upgrade the Jloader software on the first bank, the software package is deleted after you reboot. Make sure that you have either downloaded the Jloader software package to either a remote site or in a local directory on the switch, such as the **/var/tmp** directory on the QFX3500 device.

1. In a browser, go to <http://www.juniper.net/support/downloads/junos.html> .  
The Junos Platforms Download Software page appears.
2. In the QFX Series section of the Junos Platforms Download Software download page, select the QFX Series platform software you want to download.
3. Select the number of the software version that you want to download in the Release: pull-down window to the right of the tabs on the Download Software page.

4. Select the Software tab and then select the install package you want to download in the Install Package section.
5. In the pop-up Alert box, click the link to the Product Support Notification (PSN) document.
6. Enter your name and password and press **Enter**.
7. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
8. Open or save the **jloader-qfx-version-signed.tgz** file either to a local system or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.
9. Log in to the QFX3500 switch and enter the shell. We recommend using a console connection.
10. Determine the version of the Jloader software package installed on the switch.

For example:

```
root@switch% ls
gres-tp krt_gencfg_filter.txt
jloader-qfx-11.3-20110510.0-signed.tgz
```

11. Determine the version of the Uboot software that is running in the bank:

For example:

```
root@switch% kenv | grep boot.version
boot.version="1.0.7"
```

12. Enter the CLI and install the Jloader software package.

- To install a Jloader software package that is located in the **/var/tmp** directory, issue the **request system software add /var/tmp/jloader-qfx-version.tgz no-validate** command:

For example:

```
user@switch> request system software add
/var/tmp/jloader-qfx-11.3-20110510.0-signed.tgz no-validate
```

You see the following messages during the installation:

```
Verified jloader-qfx-11.3-20110510.0.tgz signed by PackageProduction_11_3_0
Adding jloader-qfx...
Installation in progress, please wait...
Mounted jloader-qfx package on /dev/md8...
Verified manifest signed by PackageProduction_11_3_0
Verified jloader-qfx-11.3-20110510.0 signed by PackageProduction_11_3_0
Registering jloader-qfx as unsupported
```

```
Installation finished successfully.
Please reboot to activate the package
Saving package file in /var/sw/pkg/jloader-qfx-11.3-20110510.0-signed.tgz
...
Saving state for rollback ...
```

```
juniper@qfx3500>
```

- To install a Jloader software package located on a remote server using FTP, issue the **request system software add**



**/ftp://hostname/pathname/jloader-qfx-version-signed.tgz no-validate**  
command.

For example:

```
user@switch> request system software add
/ftp://hostname/pathname/jloader-qfx-11.3-20110510.0-signed.tgz no-validate
```

- To install a Jloader software package located on a remote server using HTTP, issue the **request system software add /http://hostname/pathname/jloader-qfx-version-signed.tgz no-validate** command.

For example:

```
user@switch> request system software add
/http://hostname/pathname/jloader-qfx-11.3-20110510.0-signed.tgz no-validate
```

13. When prompted, reboot the Control Board by issuing the **request system reboot** command.

For example:

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

14. Enter the shell and verify that the version of the Uboot software in the primary bank is the version you just installed.

For example:

```
root@switch% kenv | grep boot.version
boot.version="1.1.1"
```

15. To install the Jloader software package on the current backup bank, repeat Step 10 through Step 14.

## Upgrading Jloader Software on a QFabric System

This procedure explains how to upgrade the Jloader software on your Node devices and Interconnect devices. The example shows how to upgrade the Jloader Release 1.1.1 to 1.1.2 on a Node device with the serial number BBAK1186.



**NOTE:** Before you upgrade the Jloader software, make sure you have the serial numbers of the Node devices, Interconnect devices, and Control Boards in the Interconnect devices you want to upgrade.

1. Issue the **show chassis hardware node-device ?** command to view the serial numbers of the Node devices.

For example:

```
user@qfabric> show chassis hardware node-device ?
<node-device>      Node device identifier
BBAK1186            Node device
BBAK3149            Node device
BBAK3177            Node device
BBAK8063            Node device
BBAK8799            Node device
P2443-C             Node device
P2515-C             Node device
P3708-C             Node device
P3885-C             Node device
P3916-C             Node device
node0               Node device
node1               Node device
node2               Node device
node3               Node device
node4               Node device
node5               Node device
node6               Node device
node7               Node device
node8               Node device
```

An example of a Node device serial number is BBAK1186.

2. Issue the **show chassis hardware interconnect-device ?** command to view the serial numbers of the Interconnect devices.

For example:

```
user@qfabric> show chassis hardware interconnect-device ?
Possible completions:
interconnect-device  Interconnect device identifier
IC-F1052             Interconnect device
IC-F3947             Interconnect device
```

The Interconnect device serial numbers are IC-F1052 and IC-F3947.

3. Issue the **show chassis hardware interconnect-device name** command to view the serial numbers of the Control Boards in the Interconnect device.

For example:

```
user@qfabric> show chassis hardware interconnect-device IC-F3947
```

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis	REV 10		F3947	QFXC08-3008
Midplane	REV 10	750-035835	F3947-C	QFX Midplane
CB 0 Board	REV 14	750-035855	ZJ9432	QFX Chassis Control
Routing Engine 0		BUILTIN	BUILTIN	QFX Routing Engine
CB 1 Board	REV 14	750-035855	ZJ9404	QFX Chassis Control

The Control Board serial numbers are ZJ9432 and ZJ9404.

4. Issue the **show chassis firmware node-device *name*** command to see which version of Uboot software you have installed on your Node device.

For example:

```
user@qfabric> show chassis firmware node-device BBAK1186
```

Part	Type	Version
node4	U-Boot	1.1.6 (May 10 2011 - 04:52:59) 1.1.1
	loader	FreeBSD/MIPS U-Boot bootstrap loader 0.1

The Uboot software version is 1.1.1. The loader software version appears after the timestamp for U-Boot 1.1.6.

5. Issue the **show chassis firmware interconnect-device *name*** command to see which version of Uboot software you have installed on the Routing Engines located on the Control Boards of the Interconnect device.

For example:

```
user@qfabric> show chassis firmware interconnect-device IC-F3947
```

Part	Type	Version
Routing Engine 0	U-Boot	U-Boot 1.1.6 (Jan 27 2012 - 03:24:34) 1.1.4
	loader	FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1	U-Boot	U-Boot 1.1.6 (Jan 27 2012 - 03:24:34) 1.1.4
	loader	FreeBSD/MIPS U-Boot bootstrap loader 0.1

The Uboot software version is 1.1.4. The loader software version appears after the timestamp for U-Boot 1.1.6.

6. In a browser, go to <http://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

7. In the QFX Series section of the Junos Platforms Download Software download page, select the QFX Series platform software you want to download.
8. Select the number of the software version that you want to download in the Release: pull-down window to the right of the tabs on the Download Software page.
9. Select the **Software** tab and then select the install package you want to download in the Install Package section.

10. In the pop-up Alert box, click the link to the Product Support Notification (PSN) document.
11. Enter your username and password, and press **Enter**.
12. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
13. Open or save the **jloader-qfx-version-signed.tgz** file either to a local system or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.
14. Retrieve the software from the location in which you downloaded it. To do this, issue the **request system software download /path/package-name** command.

For example:

```
user@qfabric> request system software download
ftp://server/files/jloader-qfx-11.3X30.9-signed.tgz
```

15. Log in to the Director device as root and enter the shell to verify that you have downloaded the Jloader software package. We recommend using a console connection. The software package is copied from where you downloaded it and is placed locally on the QFabric system in the **/pbdata/packages** directory.

For example:

```
[root@dg0] # pwd
/pbdata/packages

[root@dg0] # ls
jloader-qfx-11.3X30.9-signed.tgz
```

16. Before you copy over the Jloader software package to the Node device or Interconnect device, determine the directory that matches the serial number of the Node device or Interconnect device that you want to upgrade. View the remote logs and the Node device and Interconnect device serial numbers by issuing the **ls /pbdata/export/rlogs** command at the command line of the Director device before you copy the software package over to the device.



**NOTE:** The **/pbdata/export/rlogs/node-device-serial-ID** and **/pbdata/export/rlogs/interconnect-device-serial-ID** directories on the Director device are NFS mounted as the **/tftpboot/logfiles** directories on the Node device and Interconnect device. These directories are created for all Node devices and Interconnect devices in a QFabric system. The Jloader files are stored in the **/tftpboot/logfiles** directories for each Node device and Interconnect device.

For example:

```
[root@dg0 tmp] # ls /pbdata/export/rlogs
02de4930-828b-11e1-a319-00e081c57938 c9898afe-828b-11e1-956c-00e081c57938
04103b2a-29d5-e011-bf8a-0e6bdf3aa1e6 eeba4aac-828b-11e1-85e2-00e081c57938
1e2739e0-828b-11e1-bf74-00e081c57938 F1052
8d8a978c-828b-11e1-a833-00e081c57938 F3947
ad55b89e-828b-11e1-b70e-00e081c57938 P2443-C
BBAK1186 P2515-C
```

BBAK3149	P3708-C
BBAK3177	P3885-C
BBAK8063	P3916-C
BBAK8799	

BBAK1186 is the serial number of the Node device that needs to be upgraded.

17. Copy the Jloader software package from the `/var/tmp` directory to the `/pbdata/export/rlogs/BBAK1186` directory.

For example:

```
[root@dg0 tmp] # cp jloader-qfx-11.3X30.9-signed.tgz /pbdata/export/rlogs/BBAK1186
```

18. Confirm that the Jloader software package you copied over is in the `/pbdata/export/rlogs/BBAK1186` directory.

For example:

```
[root@dg0 tmp] # ls /pbdata/export/rlogs/BBAK1186
jloader-qfx-11.3X30.9-signed.tgz
```

19. Issue the `/root/dns.dump` command to find out the internal IP addresses of the Node device or Interconnect device.

```
[root@dg0 tmp] # /root/dns.dump
; <<>> DiG 9.3.6-P1-RedHat-9.3.6-4.P1.e15 <<>> -t axfr pkg.dcbg.juniper.net
@169.254.0.1
;; global options: printcmd
pkg.dcbg.juniper.net. 600 IN SOA ns.pkg.dcbg.juniper.net.
mail.pkg.dcbg.juniper.net. 152 3600 600 7200 3600
pkg.dcbg.juniper.net. 600 IN NS ns.pkg.dcbg.juniper.net.
pkg.dcbg.juniper.net. 600 IN A 169.254.0.1
pkg.dcbg.juniper.net. 600 IN MX 1 mail.pkg.dcbg.juniper.net.
dcfnnode---DCF-ROOT.pkg.dcbg.juniper.net. 45 IN A 169.254.192.17
dcfnnode---DRE-0.pkg.dcbg.juniper.net. 45 IN A 169.254.3.3
dcfnnode-8d8a978c-828b-11e1-a833-00e081c57938.pkg.dcbg.juniper.net. 45 IN A
169.254.128.19
dcfnnode-ad55b89e-828b-11e1-b70e-00e081c57938.pkg.dcbg.juniper.net. 45 IN A
169.254.128.20
dcfnnode-BBAK1186.pkg.dcbg.juniper.net. 45 IN A 169.254.128.14
```

The internal IP address for BBAK1186 is 169.254.128.14.

20. Upgrade the Jloader software on the Node device or Interconnect device.

Before you can upgrade the Jloader software, you need to use SSH to log in to the Node device or Interconnect device and verify that the software is in the `/tftpboot/logfiles` directory.

- a. Use SSH to log in to the Node device or Interconnect device.

For example:

```
[root@dg0 tmp] # ssh 160.254.128.14
root@169.254.128.14's password:
--- JUNOS 11.3X30.10 built 2012-03-11 22:55:43 UTC
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@sng3%
```

- b. Verify that the Jloader software package is in the `tftpboot/logfiles` directory of the Node device or Interconnect device.

For example:

```
root@sng3% ls /tftpboot/logfiles
.index                               jloader-qfx-11.3X30.9-signed.tgz
```

- c. Copy the Jloader software package from the **/tftpboot/logfiles** directory to the **/var/tmp** directory of the Node device or Interconnect device.

For example:

```
root@sng3% cp /tftpboot/logfiles/jloader-qfx-11.3X30.9-signed.tgz /var/tmp
```

- d. Verify that the Jloader software package is in the **/var/tmp** directory of the Node device or Interconnect device.

For example:

```
root@sng3% ls /var/tmp
.snap                               jloader-qfx-11.3X30.9-signed.tgz
    tmp
gres-tp                           krt_gencfg_filter.txt
    vc-autoupgrade
if-rtbdb                           rtsdb
```

- e. Enter CLI mode and issue the **request system software add /var/tmp/jloader-qfx-version-signed.tgz** command.

For example:

```
root@sng3% cli
root@sng3> request system software add /var/tmp/jloader-qfx-11.3X30.9-signed.tgz
Validating on fpc0
Checking compatibility with configuration
Initializing...
Using jbase-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jbase-11.3X30.10 signed by PackageProduction_11_3_0
Using /var/tmp/jloader-qfx-11.3X30.9-signed.tgz
Verified jloader-qfx-11.3X30.9.tgz signed by PackageProduction_11_3_0
Using jloader-qfx-11.3X30.9.tgz
Checking jloader-qfx requirements on /
Verified manifest signed by PackageProduction_11_3_0
Verified jloader-qfx-11.3X30.9 signed by PackageProduction_11_3_0
Using jkernel-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jkernel-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jroute-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jroute-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jcrypto-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jcrypto-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jweb-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jweb-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Using jswitch-qfx-11.3X30.10
Verified manifest signed by PackageProduction_11_3_0
Verified jswitch-qfx-11.3X30.10 signed by PackageProduction_11_3_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
```

Done with validate on all chassis

```
fpc0:
Verified jloader-qfx-11.3X30.9.tgz signed by PackageProduction_11_3_0
Adding jloader-qfx...
Installation in progress, please wait...
Mounted jloader-qfx package on /dev/md10...
Verified manifest signed by PackageProduction_11_3_0
Verified jloader-qfx-11.3X30.9 signed by PackageProduction_11_3_0
#####
#####
Installation finished successfully.
Please reboot to activate the package
Saving package file in /var/sw/pkg/jloader-qfx-11.3X30.9-signed.tgz ...
Saving state for rollback ...
```

Upgrade has completed successfully.  
Reboot is now required.

- f. Reboot both the Node device and Interconnect device twice, because they each contain two partitions.

For example:

```
root@sng3> request system reboot
Reboot the system ? [yes,no] (no) yes
Shutdown NOW!
[pid 37663]
```

```
root@sng3>
```

```
*** FINAL System shutdown message from root@sng3 ***
```

```
System going down IMMEDIATELY
```

- g. Verify that the Uboot software on the Node device or Interconnect device has been upgraded to the new Uboot software by logging in to the QFabric CLI and issuing either the **show chassis firmware node-device *name*** command or the **show chassis firmware interconnect-device *name*** command.

For example:

```
user@qfabric> show chassis firmware node-device BBAK1186
Part                Type      Version
node4               U-Boot   1.1.6 (Nov 19 2011 - 11:42:07) 1.1.2
                                loader   FreeBSD/MIPS U-Boot bootstrap loader
0.1
```

The Uboot software version is now 1.1.2. The loader software version appears after the timestamp for U-Boot 1.1.6.

## Upgrading Software

To upgrade Junos OS, you need to install the appropriate upgrade package on the device. Upgrading involves these tasks:

1. [Downloading Software Files with a Browser on page 134](#)
2. [Accessing Software Downloaded to a Remote Location on page 135](#)
3. [Connecting to the Console Port on page 135](#)
4. [Backing Up the Current Configuration Files on page 135](#)
5. [Installing a Standard Software Package on page 136](#)
6. [Upgrading to an ELS-Based Software Package on page 137](#)

---

### Downloading Software Files with a Browser

To download the software package from the Juniper Networks Support website, go to <http://www.juniper.net/support/>.



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

This procedure shows you how to upgrade software on a QFX Series device, but you can follow the same procedure for any device unless otherwise specified.

- 
1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
  2. Click **Download Software**.
  3. In the **Switching** box, click **Junos OS Platforms**.
  4. In the **QFX Series** section, click the name of the platform for which you want to download software.
  5. Click the **Software** tab and select the release number from the **Release** drop-down list.
  6. In the **Install Package** section of the **Software** tab, select the **Install Package** for the release.  
A login screen appears.
  7. Enter your name and password and press **Enter**.
  8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.



9. Save the **jinstall-qfx-<version>-domestic-signed.tgz** file on your computer.
10. Open or save the installation package either to the local system in the **var/tmp** directory or to a remote location. If you are saving the installation package to a remote system, make sure that you can access it using HTTP, TFTP, FTP, or scp.

### Accessing Software Downloaded to a Remote Location

To access the installation package if you downloaded it to a remote location (for example, any system other than the switch):

1. From the command line, make sure you are in the **/var/tmp** directory of the switch.
2. Start the shell interface:  
`user@switch> start shell`
3. Initiate an FTP, TFTP, or scp session.

In this example, FTP is used.

- `>ftp`
4. Use FTP to access the remote location where the installation package resides.  
`ftp ftp://<hostname>/<pathname>/<package-name-m.mZx-distribution>.tgz.`  
where `<package-name-m.mZx-distribution>.tgz` is  
`jinstall-qfx-11.1R1.5-domestic-signed.tgz`
5. When prompted, enter your username and password.
6. Use the **get** command to transfer the installation package from the remote location to your **/var/tmp** directory on your switch.  
`get <package-name-m.mZx-distribution>.tgz`
7. Close the FTP session:  
`bye`

### Connecting to the Console Port

We recommend that you connect to the console port while installing the installation package so you can respond to any required user input and detect any errors that may occur.

### Backing Up the Current Configuration Files

Before you install the new installation package, we strongly recommend that you back up your current configuration files because the upgrade process removes all of the stored files on the switch.

To back up your current configuration files, enter the **save** command:

```
user@switch# save filename
```

Executing this command saves a copy of your configuration files to a remote location such as an external USB device.

## Installing a Standard Software Package

---



**NOTE:** Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <http://www.juniper.net/support>.

---



**NOTE:** If you are upgrading from a standard software package to an ELS-based package, see the *Upgrading to an ELS-Based Software Package* section.

---

Install the software in one of three ways:

If the installation package resides locally on the switch, execute the **request system software add validate <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add validate
/var/tmp/jinstall-qfx-11.1R1.5-domestic-signed.tgz reboot
```

If the Install Package resides remotely, execute the **request system software add validate <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add validate
ftp://ftpsrvr/directory/jinstall-qfx-11.1R1.5-domestic-signed.tgz reboot
```



**NOTE:** On the QFX5100 switch, use the force-host option to force installing the latest version of the Host OS.

If the installation package resides locally on the QFX5100 switch, execute the **request system software add force-host validate <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add force-host validate
/var/tmp/jinstall-qfx-5.13.2X51-D10.6-domestic-signed.tgz reboot
```

If the install Package resides remotely from the QFX5100 switch, execute the **request system software add force-host validate <pathname> <source> reboot** command.

For example:

```
user@switch> request system software add force-host validate
ftp://ftpsrvr/directory/jinstall-qfx-5.13.2X51-D10.6-domestic-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

### Upgrading to an ELS-Based Software Package

To upgrade your switch from a version of Junos OS that does not support Enhanced Layer 2 Software (ELS) to a version of Junos OS that supports ELS, we recommend performing the following procedure.



**NOTE:** Because this procedure can cause service outages, we recommend that you avoid performing this procedure on switches carrying traffic in a production network.

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



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

2. Set your device to standalone mode by issuing the **request chassis device-mode standalone** command. Do not reboot your system at this time.



**NOTE:** This step is only required for new devices shipped from the factory or QFabric system Node devices that you plan to redeploy in a QFX Series Virtual Chassis.

3. Choose whether you wish to reuse your previous configuration or not.
  - To reuse your previous configuration as part of the software upgrade, you must convert the configuration from the original style Junos OS CLI to the ELS CLI format using the following steps:



**NOTE:** We recommend this procedure for customers currently using a QFX3500 or QFX3600 switch as a standalone device.

- a. Copy your entire existing configuration into a text file. Save the file to a remote location or USB drive.
- b. Retain the portion of your existing configuration related to management network connectivity (such as **[edit system]** and management interfaces). Delete all other configuration elements (such as the **[edit protocols]** and **[edit vlans]** hierarchy levels, non-management interfaces, and so on). Issue a **commit** operation to remove the deleted configuration.
- c. Perform the software upgrade with the **validate** option and reboot your device to complete the upgrade by issuing the **request system software add validate reboot** command. Maintain your console port connection during the reboot.
- d. Using a web browser, navigate to the [ELS Translator Tool](#). Follow the instructions on the page to convert your saved configuration file to the new ELS CLI format.
- e. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator Tool and load it in to your switch.
- f. Issue a **commit** operation to activate the translated configuration.
- To delete your current configuration and upgrade the software, follow these steps:



**NOTE:** We recommend this procedure for customers with new QFX3500 or QFX3600 devices shipped from the factory or QFabric system Node devices that will be redeployed in a QFX Series Virtual Chassis.

- a. Perform a software upgrade with the **no-validate** option by issuing the **request system software add no-validate** command.
- b. Delete the configuration and set the device to factory defaults by issuing the **request system zeroize** command. The device automatically reboots and reverts to a factory default configuration.
- c. Configure your device using the ELS CLI format.

**Related Documentation**

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 17](#)
- [Software Installation Overview on page 122](#)
- [Recovering from a Failed Software Installation on page 121](#)
- [Upgrading Jloader Software on QFX Series Devices on page 123](#)
- [request system software add on page 421](#)
- [Installation and Upgrade Guide](#)

## Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member switches on a Virtual Chassis Fabric with minimal network traffic disruption during the upgrade. A Virtual Chassis Fabric can contain 20 members—up to 2 members can be in the routing engine role, and up to 18 line cards can be configured in the line card role. You can upgrade software for a fixed configuration of switches (QFX3500/QFX3600 and QFX5100 switches, or EX4300 and QFX5100 switches) or for a mixed mode of switches (combination of QFX3500/QFX3600, QFX5100, and EX4300 switches) in a Virtual Chassis Fabric.

This topic covers:

- [Preparing the Switch for Software Installation on page 139](#)
- [Upgrading the Software Using NSSU on page 141](#)

### Preparing the Switch for Software Installation

Before you begin software installation using NSSU:

- Ensure that the Virtual Chassis Fabric is configured correctly to support NSSU. Verify that:
  - The Virtual Chassis Fabric members are connected in a spine and leaf topology. A spine and leaf topology prevents the Virtual Chassis from splitting during an NSSU. Each leaf device must be connected to both spine devices.
  - The Virtual Chassis Fabric must be preprovisioned so that the line card role has been explicitly assigned to member switches acting in a line card role, and that the routing engine role has been explicitly assigned to member switches acting in a routing engine role. During an NSSU, the Virtual Chassis Fabric members must maintain their roles—the master and backup must maintain their master and backup roles (although

mastership will change), the member switches must remain their routing engine roles, and the remaining switches must maintain their line card roles.

- Only two pre-provisioned members in the routing engine role are supported. If more than two routing engines are configured, a warning will be issued, and NSSU will stop.
- A two-member Virtual Chassis has **no-split-detection** configured so that the Virtual Chassis Fabric does not split when an NSSU upgrades a member.
- Verify that the members are running the same version of the software:

```
user@switch> show version
```

If you are going to perform an NSSU on a fixed configuration of switches (QFX3500/QFX3600, QFX5100, or EX4300 switches) or a mixed mode configuration of switches (QFX3500/QFX3600, QFX5100, or EX4300 switches) that are not running the same version of the software, use the **request system software nonstop-upgrade <set [package-name package-name]> reboot** command to upgrade the software on the inconsistent members.



**NOTE:** This command can require up to three software images depending on devices configured in the Virtual Chassis Fabric.

For example:

```
user@switch> request system software nonstop-upgrade add set [jinstall-qfx5100.tgz  
jinstall-qfx3500.tgz jinstall-ex-4300.tgz] reboot
```

- Ensure that graceful Routing Engine switchover (GRES) is enabled. To verify that they are enabled, you need to check only the state of nonstop active routing—if nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

To verify that nonstop active routing is enabled:

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

Protocol	Synchronization Status
OSPF	Complete
BGP	Complete
PIM	Complete

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see “[Example: Configuring Nonstop Active Routing on Switches](#)” on page 2274 for information on how to enable it.

- Enable nonstop bridging (NSB). Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on each member to an external storage device with the **request system snapshot** command.

## Upgrading the Software Using NSSU

This procedure describes how to upgrade the software running on all Virtual Chassis Fabric members using NSSU. When the upgrade completes, all members are running the new version of the software. Because a graceful Routing Engine switchover occurs during the upgrade, the original Virtual Chassis Fabric backup is the new master.

To upgrade all members using NSSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in [“Upgrading Software” on page 134](#) and *Downloading Software Packages from Juniper Networks*. If you are upgrading the software running on a mixed mode Virtual Chassis Fabric, download the software packages for each switch type.
2. Copy the software package or packages to the Virtual Chassis Fabric. We recommend that you copy the file to the `/var/tmp` directory on the master.
3. Log in to the Virtual Chassis Fabric using the console connection of the master or the virtual management Ethernet (VME) interface. Without the console connection, you will not be able to view any CLI output during an NSSU reboot. The console connection enables you to view CLI output during an NSSU reboot and monitor the progress of the master switch reboot.

4. Start the NSSU:

- To perform an NSSU on a fixed configuration of switches (QFX3500/QFX3600, QFX5100, or EX4300 switches), enter:

```
user@switch> request system software nonstop-upgrade
/var/tmp/package-name.tgz
```

where `package-name.tgz` is, for example, `jinstall-qfx5100.tgz`.

- To perform an NSSU on a mixed mode configuration of switches (QFX3500/QFX3600, QFX5100, or EX4300 switches), enter:

```
user@switch> request system software nonstop-upgrade set [package-name
package-name
package-name]
```

where `[package-name.tgz package-name.tgz package-name.tgz]` is, for example, `[jinstall-qfx5100.tgz jinstall-qfx3500.tgz jinstall-ex-4300.tgz]`.

The switch displays status messages similar to the following messages as the upgrade executes:

```
switch# request system software nonstop-upgrade [ jinstall-qfx-3-mixed_vc-1-domestic.tgz
jinstall-qfx-5-mixed_vc-1-domestic-img.tgz jinstall-ex-4300-mixed_vc-1-domestic.tgz ]
```

```
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Installing image on other FPC's along with the backup
```

```
Retrieving software images. This process can take several minutes. Please be
patient..
```

```
Retrieving version and model information from
/var/tmp/jinstall-qfx-3-mixed_vc-1-domestic.tgz

Retrieving version and model information from
/var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz

Retrieving version and model information from
/var/tmp/jinstall-ex-4300-mixed_vc-1-domestic.tgz
Starting with package /var/tmp/jinstall-qfx-3-mixed_vc-1-domestic.tgz

Download done for package /var/tmp/jinstall-qfx-3-mixed_vc-1-domestic.tgz
Pushing bundle to fpc5
WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

Setting up /var/shared/v for validation ...
Checking compatibility with configuration
Initializing...
Using jbase-qfx-mixed_vc
Using /var/tmp/jinstall-qfx-3-mixed_vc-1-domestic.tgz
Using jbundle-qfx-3-mixed_vc-1-domestic.tgz
Using jkernel-qfx-3-mixed_vc
Using jroute-qfx-3-mixed_vc
Using jcrypto-qfx-mixed_vc
Using jswitch-qfx-3-mixed_vc
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
fpc5
WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

Image name:    /var/tmp/jinstall-qfx-3-mixed_vc-1-domestic.tgz
Package name:  jinstall-qfx-3
Junos revision: mixed_vc-1

==> STEP #1 of 20: Creating temporary file system <==
*** Fri Apr  4 16:28:23 PDT 2014 ***

==> STEP #2 of 20: Determining installation source <==
*** Fri Apr  4 16:28:23 PDT 2014 ***
Interactive installation: yes
Boot media formatting option: disabled
Installing packages from internal drive da0
Packages will be installed to da0, media size: 8G

==> STEP #3 of 20: Processing format options <==
*** Fri Apr  4 16:28:24 PDT 2014 ***
Formatting of boot media is not supported in this form of installation

==> STEP #4 of 20: Determining installation slice <==
*** Fri Apr  4 16:28:24 PDT 2014 ***
Physmem correction sectors: 0
Dump partition sectors: 4161792
package: INFO: Dump partition da0s3b size sufficient for kernel dumps.
package: INFO: da0s3b will be retained as dump partition.

==> STEP #5 of 20: Creating and labeling new slices <==
*** Fri Apr  4 16:28:25 PDT 2014 ***
```



```

Primary partition - /: /dev/da0s1a, /var: /dev/da0s1f
Shared partition - /var/tmp: /dev/da0s3d, /config: /dev/da0s3e, /var/shared:
/dev/da0s3f, recovery: /dev/da0s3a, /var/rundb: /dev/da0s3g
Checking integrity on da0s3

```

```

==> STEP #6 of 20: Create and mount new file system <==
*** Fri Apr  4 16:28:25 PDT 2014 ***
/dev/da0s1a: 464.0MB (950192 sectors) block size 16384, fragment size 2048
        using 4 cylinder groups of 116.00MB, 7424 blks, 14848 inodes.
super-block backups (for fsck -b #) at:
32, 237600, 475168, 712736
tunefs: soft updates set
/dev/da0s1f: 249.8MB (511656 sectors) block size 16384, fragment size 2048
        using 4 cylinder groups of 62.47MB, 3998 blks, 8064 inodes.
super-block backups (for fsck -b #) at:
32, 127968, 255904, 383840
tunefs: soft updates remains unchanged as disabled
Size of staging area: 1006M
Staging area setup in /instrootmnt/var/tmp ...

```

```

==> STEP #7 of 20: Getting OS bundles <==
*** Fri Apr  4 16:28:35 PDT 2014 ***

```

```

==> STEP #8 of 20: Updating recovery media <==
*** Fri Apr  4 16:28:35 PDT 2014 ***
System recovery media update option is disabled - skipping
set update_recovery_on_install="YES" in /etc/rc.conf.platform

```

```

==> STEP #9 of 20: Extracting incoming image <==
*** Fri Apr  4 16:28:35 PDT 2014 ***
install bundle source /var/tmp:jinstall-qfx-3-mixed_vc-1-domestic.tgz
Enabling platform watchdog
Starting pkg_add on /var/tmp/tmp - Fri Apr  4 16:29:09 PDT 2014

```

```

WARNING:      The software that is being installed has limited support.
WARNING:      Run 'file show /etc/notices/unsupported.txt' for details.

```

```

Completed pkg_add - Fri Apr  4 16:29:37 PDT 2014
Enabling platform watchdog

```

```

==> STEP #10 of 20: Unpacking OS packages <==
*** Fri Apr  4 16:29:37 PDT 2014 ***
Enabling platform watchdog

```

```

==> STEP #11 of 20: Mounting jbase package <==
*** Fri Apr  4 16:29:39 PDT 2014 ***
Mounted jbase package on /dev/md17...

```

```

==> STEP #12 of 20: Creating base OS symbolic links <==
*** Fri Apr  4 16:29:54 PDT 2014 ***

```

```

==> STEP #13 of 20: Creating fstab <==
*** Fri Apr  4 16:30:49 PDT 2014 ***

```

```

==> STEP #14 of 20: Creating new system files <==
*** Fri Apr  4 16:30:49 PDT 2014 ***

```

```

==> STEP #15 of 20: Adding jbundle package <==
*** Fri Apr  4 16:30:49 PDT 2014 ***
Checking package integrity...
Verified SHA1 checksum of jbase-qfx-mixed_vc-1.tgz

```

```
Verified SHA1 checksum of jboot-qfx-mixed_vc-1.tgz
Verified SHA1 checksum of jcrypto-qfx-mixed_vc-1.tgz
Verified SHA1 checksum of jdocs-qfx-mixed_vc-1.tgz
Verified SHA1 checksum of jkernel-qfx-3-mixed_vc-1.tgz
Verified SHA1 checksum of jpfe-qfx-3-mixed_vc-1.tgz
Verified SHA1 checksum of jroute-qfx-3-mixed_vc-1.tgz
Verified SHA1 checksum of jswitch-qfx-3-mixed_vc-1.tgz
Verified SHA1 checksum of py-base-xlr-mixed_vc-1.tgz
Running requirements check first for jbundle-qfx-3-mixed_vc-1-domestic...
Setting up alternate devfs for installation
```

```
WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.
```

```
Running pre-install for jbundle-qfx-3-mixed_vc-1-domestic...
Installing jbundle-qfx-3-mixed_vc-1-domestic in
/var/tmp/tmp/pa8070.50/jbundle-qfx-3-mixed_vc-1-domestic.x8070...
Running post-install for jbundle-qfx-3-mixed_vc-1-domestic...
Adding jkernel-qfx-3...
Registering jkernel as unsupported
Adding jcrypto-qfx...
Registering jcrypto as unsupported
Adding jdocs-qfx...
Registering jdocs as unsupported
Adding jswitch-qfx-3...
Registering jswitch as unsupported
Adding jpfe-qfx-3...
Registering jpfe as unsupported
Adding jroute-qfx-3...
Registering jroute as unsupported
Adding py-base-xlr...
Registering py-base-xlr as unsupported
Setting up shared filesystem (/dev/da0s3f) ...
```

```
==> STEP #16 of 20: Backing up system data <==
*** Fri Apr  4 16:33:18 PDT 2014 ***
Backup configuration file: /config/install_backup.conf
Files to be backed up: /var/db/scripts /var/home /root
Backup file saved at /var/tmp/install_backup.tar
Restoring system backed up data on new partition - Fri Apr  4 16:33:18 PDT
2014
Done - Fri Apr  4 16:33:19 PDT 2014
```

```
==> STEP #17 of 20: Setting up shared partition data <==
*** Fri Apr  4 16:33:19 PDT 2014 ***
Creating directories in /dev/da0s3f
```

```
==> STEP #18 of 20: Checking package sanity in installation <==
*** Fri Apr  4 16:33:19 PDT 2014 ***
```

```
==> STEP #19 of 20: Unmounting and cleaning up temporary file systems <==
*** Fri Apr  4 16:33:19 PDT 2014 ***
tunefs: soft updates remains unchanged as disabled
tunefs: soft updates cleared
Enabling platform watchdog
```

```
==> STEP #20 of 20: Setting da0s1 as new active partition <==
*** Fri Apr  4 16:33:26 PDT 2014 ***
WARNING: NOTE: A reboot is required to start using the new software
WARNING:      Use the 'request system reboot' command when ready
```

```

Installation log saved at /var/sw/install-04042014162823.log
Starting with package /var/tmp/jinstall-ex-4300-mixed_vc-1-domestic.tgz

Download done for package /var/tmp/jinstall-ex-4300-mixed_vc-1-domestic.tgz
Pushing bundle to fpc2
WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.
fpc2
WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.

WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Starting with package /var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz

Download done for package /var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz
Pushing bundle to fpc1
Pushing bundle to fpc4
WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.
WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.
fpc1
WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.

Image name: /var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz
Package name: jinstall-qfx-5-img
Junos revision: mixed_vc-1
NOTE: A reboot is required to install the software
      Use the 'request system reboot' command immediately
fpc4
WARNING: The software that is being installed has limited support.
WARNING: Run 'file show /etc/notices/unsupported.txt' for details.

Image name: /var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz
Package name: jinstall-qfx-5-img
Junos revision: mixed_vc-1
NOTE: A reboot is required to install the software
      Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting fpc1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
.
.
.
.
.
.
FPC 2 is undergoing a software upgrade
.
.

```



```

.
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 4         Online (ISSU)
  FPC 5         Online (ISSU)
Going to install image on master

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

Image name:    /var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz
Package name:  jinstall-qfx-5-img
Junos revision: mixed_vc-1

```

```

Going to install image on master

WARNING:    The software that is being installed has limited support.
WARNING:    Run 'file show /etc/notices/unsupported.txt' for details.

Image name:    /var/tmp/jinstall-qfx-5-mixed_vc-1-domestic-img.tgz
Package name:  jinstall-qfx-5-img
Junos revision: mixed_vc-1
NOTE: A reboot is required to install the software
      Use the 'request system reboot' command immediately
failover links
Rebooting Old master
ISSU: IDLE

```

```

*** FINAL System shutdown message from root@vcf ***

```

```

System going down IMMEDIATELY

```

```

Shutdown NOW!
[pid 5290]

```

5. Log in after the reboot of the original master switch completes. To verify that the software on all Routing Engines in the Virtual Chassis Fabric members has been upgraded, enter the following command:

```

user@switch> show version

```

#### Related Documentation

- [request system software nonstop-upgrade on page 449](#)
- [show chassis nonstop-upgrade on page 887](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2270](#)
- [Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric on page 26](#)

- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2272](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2274](#)
- [Example: Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade on EX Series Switches](#)
- [Understanding Resilient Dual-Root Partitions on Switches](#)

## Upgrading Software by Using Automatic Software Download

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The switch reboots automatically to complete the upgrade.

**Related  
Documentation**

- [Verifying That Automatic Software Download Is Working Correctly on page 341](#)
- [Understanding Software Installation on EX Series Switches](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 155](#)
- [Configuring DHCP Services \(J-Web Procedure\)](#)
- [Understanding DHCP Services for Switches on page 21](#)





## CHAPTER 4

# Configuration

- [Initial Configuration on page 151](#)
- [Configuration Examples on page 196](#)
- [Configuration Statements on page 204](#)

### Initial Configuration

---

- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 152](#)
- [Configuring a DHCP Client \(CLI Procedure\) on page 154](#)
- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 155](#)
- [Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 158](#)
- [Reaching a Domain Name System Server on page 158](#)
- [Configuring the Hostname of the Router or Switch on page 160](#)
- [Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types on page 161](#)
- [Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 162](#)
- [Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 162](#)
- [Configuring the Junos OS to Display a System Login Announcement on page 163](#)
- [Configuring the Junos OS to Display a System Login Message on page 163](#)
- [Configuring Junos OS to Extend the Default Port Address Range on page 164](#)
- [Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 165](#)
- [Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 165](#)
- [Configuring NTP Authentication Keys on page 165](#)
- [Configuring the NTP Time Server and Time Services on page 166](#)
- [Specifying the Physical Location of the Switch on page 169](#)
- [Configuring the Root Password on page 170](#)

- [Configuring the Router or Switch to Listen for Broadcast Messages Using NTP on page 171](#)
- [Configuring the Router or Switch to Listen for Multicast Messages Using NTP on page 172](#)
- [Configuring System Alarms to Appear Automatically Upon Login on page 172](#)
- [Configuring Time-Based User Access on page 173](#)
- [Configuring the Timeout Value for Idle Login Sessions on page 174](#)
- [Configuring a QFX3500 Device as a Standalone Switch on page 175](#)
- [Creating an Emergency Boot Device on page 176](#)
- [Creating a Snapshot and Using It to Boot a QFX Series Switch on page 178](#)
- [Creating a Snapshot and Using It to Boot QFX5100 and EX4600 Devices on page 180](#)
- [Including the Year or Millisecond in Timestamps on page 181](#)
- [Mapping the Hostname of the Switch to IP Addresses on page 182](#)
- [Methods for Configuring Junos OS on page 183](#)
- [Modifying the Default Time Zone for a Router or Switch Running Junos OS on page 186](#)
- [Rebooting and Halting a Device on page 186](#)
- [Reverting to the Default Factory Configuration on page 188](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 188](#)
- [Reverting to the Rescue Configuration on page 189](#)
- [Saving Core Files Generated by Junos OS Processes on page 189](#)
- [Updating the IANA Time Zone Database on Junos Devices on page 190](#)
- [Setting the Date and Time on page 192](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 192](#)
- [Synchronizing and Coordinating Time Distribution Using NTP on page 194](#)
- [Viewing Core Files from Junos OS Processes on page 196](#)

## Configuring Autoinstallation of Configuration Files (CLI Procedure)

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

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

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

- Have a service available—typically Dynamic Host Configuration Protocol (DHCP)—to assign an IP address to the switch
- Configure a DHCP server on your network to meet your network requirements. You can configure a switch to operate as a DHCP server. For more information, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 155](#).
- Create one of the following configuration files, and store it on a TFTP server (or HTTP server or FTP server) in the network:
  - A host-specific file with the name **hostname.conf** for each switch undergoing autoinstallation. Replace **hostname** with the name of a switch. The **hostname.conf** file typically contains all the configuration information necessary for the switch with this hostname.
  - A default configuration file named **switch.conf** with the minimum configuration necessary to enable you to telnet into the new switch for further configuration.
- Physically attach the switch to the network using a Gigabit Ethernet port.
- If you configure the DHCP server to provide only the TFTP server hostname, add an IP address-to-hostname mapping entry for the TFTP server to the DNS database file on the Domain Name System (DNS) server in the network.
- If the switch is not on the same network segment as the DHCP server (or other device providing IP address resolution), configure an existing device as an intermediate device to receive TFTP and DNS requests and forward them to the TFTP server and the DNS server. You must configure the LAN or serial interface on the intermediate device with the IP addresses of the hosts providing TFTP and DNS services. Connect this interface to the switch.
- If you are using **hostname.conf** files for autoinstallation, you must also complete the following tasks:
  - Configure the DHCP server to provide a **hostname.conf** filename to each switch. Each switch uses its **hostname.conf** filename to request a configuration file from the TFTP server. Copy the necessary **hostname.conf** configuration files to the TFTP server.
  - Create a default configuration file named **network.conf**, and copy it to the TFTP server. This file contains IP-address-to-hostname mapping entries. If the DHCP server does not send a **hostname.conf** filename to a new switch, the switch uses **network.conf** to resolve its hostname based on its IP address.

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

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

To configure autoinstallation:

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

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



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

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

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

#### Related Documentation

- [Verifying Autoinstallation Status on page 340](#)
- [Understanding Autoinstallation of Configuration Files on page 19](#)
- [Understanding DHCP Services for Switches on page 21](#)

## Configuring a DHCP Client (CLI Procedure)

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

DHCP configuration consists of two components, configuration of DHCP clients and configuration of a DHCP server. Client configuration determines how clients send a message requesting an IP address, whereas a DHCP server configuration enables the server to send an IP address configuration back to the client. This topic describes configuring a DHCP client. For directions for configuring a DHCP server, see [“Configuring a DHCP Server on Switches \(CLI Procedure\)” on page 155](#).

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

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

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

The options that you can configure are listed in [Table 20 on page 154](#). Replace the variable *configuration-statement* with one or more of the statements listed in this table. If you do not explicitly configure these options, the switch uses default values for them.

**Table 20: DHCP Client Settings**

Configuration Statement	Description
<b>client-identifier</b>	Unique client ID—By default this consists of the hardware type (01 for Ethernet) and the MAC address (a.b.c.d). For this example, the value would be 01abcd.
<b>lease-time</b>	Ttime in seconds that a client holds the lease for an IP address assigned by a DHCP server. If a client does not request a specific lease time, then the server sends the default lease time. The default lease time on a Junos OS DHCP server is 1 day.

Table 20: DHCP Client Settings (*continued*)

Configuration Statement	Description
<code>retransmission-attempt</code>	Number of times the client attempts to retransmit a DHCP packet.
<code>retransmission-interval</code>	Time between transmission attempts.
<code>server-address</code>	IP address of the server that the client queries for an IP address.
<code>update-server</code>	TCP/IP settings learned from an external DHCP server to the DHCP server running on the switch are propagated.
<code>vendor-option</code>	Vendor class ID (CPU's manufacturer ID string) for the DHCP client.

**Related  
Documentation**

- [Configuring a DHCP Server on Switches \(CLI Procedure\) on page 155](#)
- [Understanding DHCP Services for Switches on page 21](#)

## Configuring a DHCP Server on Switches (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring a DHCP Server on Switches (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 43.

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

A DHCP configuration consists of two components—an optional reconfiguration of default settings on DHCP clients and the configuration of a DHCP server. This topic covers configuration of the DHCP server. For information about reconfiguring a DHCP client, see “[Configuring a DHCP Client \(CLI Procedure\)](#)” on page 154.

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

This topic includes the following tasks:

1. [Configuring an Extended DHCP Server on a Switch on page 156](#)
2. [Configuring a Legacy DHCP Server on a Switch \(CLI Procedure\) on page 156](#)

### Configuring an Extended DHCP Server on a Switch

---

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

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

1. Create an address pool for DHCP IP addresses:

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

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

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

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

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

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

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

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

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

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

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

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

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

---

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

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

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



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

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

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

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

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

4. (Optional) Change the DHCP options:

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

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

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

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

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

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

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

8. (Optional) Change the SIP server:

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

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

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

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

10. (Optional) Change the NetBIOS name server:

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

- Related Documentation**
- [Configuring a DHCP Client \(CLI Procedure\) on page 154](#)
  - [Configuring a DHCP SIP Server \(CLI Procedure\)](#)
  - [Understanding DHCP Services for Switches on page 21](#)

## Configuring a DNS Name Server for Resolving a Hostname into Addresses

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

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

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

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

- Related Documentation**
- [name-server on page 289](#)

## Reaching a Domain Name System Server

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

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

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



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

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

1. Reference the IP addresses of your DNS servers.

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

The following example shows how to reference two DNS servers:

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

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

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

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

```
[edit system]
domain-name domain-name;
```

The following example shows how to configure the domain name:

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

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

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

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

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

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

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

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

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

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

5. Commit the configuration.

```
user@host# commit
```

6. Verify the configuration.

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

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

For example:

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

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

**Related Documentation**

- [Understanding DNS](#)

## Configuring the Hostname of the Router or Switch

The hostname of the device provides its identification for many purposes. Junos OS uses the configured hostname as part of the command prompt, to prepend log files and other accounting information, as well as in other places where knowing the device identity is useful. We recommend that the hostname be descriptive and memorable.

Optionally, instead of configuring the hostname at the **[edit system]** hierarchy level, you can use a configuration group, as shown in this procedure. This is a recommended best practice for configuring the hostname, especially if the device has dual Routing Engines. This procedure uses groups called **re0** and **re1** as an example.

To set the hostname:

1. Include the **host-name** statement in the configuration.

The name value must be less than 256 characters.

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

For example:

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

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

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

For example:

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

3. Commit the changes.

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

The hostname subsequently appears in the device CLI prompt.

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

**Related Documentation**

- [Understanding Hostnames](#)

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

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



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

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

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

**alarm-name** is the name of an alarm.

**Related Documentation**

- [System-Wide Alarms and Alarms for Each Interface Type](#)
- [Chassis Conditions That Trigger Alarms](#)

- *Silencing External Devices Connected to Alarm Relay Contacts*

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

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

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

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

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

### Related Documentation

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

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

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

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

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

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

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

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

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

- Related Documentation**
- *Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets*

## Configuring the Junos OS to Display a System Login Announcement

By default, no login announcement is displayed. To configure a system login announcement, include the **announcement** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
announcement text;
```

If the announcement text contains any spaces, enclose the text in quotation marks.

A system login *announcement* appears after the user logs in. A system login *message* appears before the user logs in.



**TIP:** You can use the same special characters described to format your system login announcement.

- Related Documentation**
- *Defining Junos OS Login Classes*
  - *Configuring the Junos OS to Display a System Login Message*

## Configuring the Junos OS to Display a System Login Message

By default, no login message is displayed on the router or switch. To configure a system login message, include the **message** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
message text;
```

If the message text contains any spaces, enclose it in quotation marks.

You can format the message using the following special characters:

- \n—New line
- \t—Horizontal tab
- \'—Single quotation mark
- \"—Double quotation mark
- \\—Backslash

The following is a sample login message configuration:

```
[edit]
system {
  login {
    message "\n\n\n\tUNAUTHORIZED USE OF THIS SYSTEM\n\tIS STRICTLY PROHIBITED!\n\n\tPlease contact\n\t'company-noc@company.com\t' to gain\n\taccess\n\tto this equipment if you need authorization.\n\n\n";
```

```
    }  
  }  
}
```

The preceding login message configuration example produces a login message similar to the following:

```
server% telnet router1  
Trying 1.1.1.1...  
Connected to router1.  
Escape character is '^['.
```

```
UNAUTHORIZED USE OF THIS SYSTEM  
IS STRICTLY PROHIBITED!
```

```
Please contact 'company-noc@company.com' to gain  
access to this equipment if you need authorization.
```

```
router1 (tty0)
```

```
login:
```

A system login message appears before the user logs in. A system login announcement appears after the user logs in.

- Related Documentation**
- *Defining Junos OS Login Classes*
  - [message on page 283](#)

## Configuring Junos OS to Extend the Default Port Address Range

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

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

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

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

- Related Documentation**
- *Configuring Junos OS to Disable TCP RFC 1323 Extensions*
  - *Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses*
  - [source-port on page 309](#)

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

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

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

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

**Related Documentation**

- [icmpv4-rate-limit on page 275](#)

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

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

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

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

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

**Related Documentation**

- [Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 162](#)
- [default-address-selection on page 267](#)

## Configuring NTP Authentication Keys

Time synchronization can be authenticated to ensure that the switch obtains its time services only from known sources. By default, network time synchronization is unauthenticated. The switch will synchronize to whatever system appears to have the

most accurate time. We strongly encourage you to configure authentication of network time services.

To authenticate other time servers, include the **trusted-key** statement at the **[edit system ntp]** hierarchy level. Only time servers that transmit network time packets containing one of the specified key numbers are eligible to be synchronized. Additionally, the key needs to match the value configured for that key number. Other systems can synchronize to the local switch without being authenticated.

```
[edit system ntp]
trusted-key [ key-numbers ];
```

Each key can be any 32-bit unsigned integer except 0. Include the **key** option in the **peer**, **server**, or **broadcast** statements to transmit the specified authentication key when transmitting packets. The key is necessary if the remote system has authentication enabled so that it can synchronize to the local system.

To define the authentication keys, include the **authentication-key** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
authentication-key key-number type type value password;
```

**number** is the key number, **type** is the authentication type (only Message Digest 5 [MD5] is supported), and **password** is the password for this key. The key number, type, and password must match on all systems using that particular key for authentication.

#### Related Documentation

- [Understanding NTP Time Servers on page 16](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)
- [trusted-key on page 331](#)
- *authentication-key*

## Configuring the NTP Time Server and Time Services

When you use NTP, configure the router or switch to operate in one of the following modes:

- Client mode
- Symmetric active mode
- Broadcast mode
- Server mode

The following topics describe how to configure these modes of operation:

1. [Configuring the Router or Switch to Operate in Client Mode on page 167](#)
2. [Configuring the Router or Switch to Operate in Symmetric Active Mode on page 167](#)



3. [Configuring the Router or Switch to Operate in Broadcast Mode on page 168](#)
4. [Configuring the Router or Switch to Operate in Server Mode on page 168](#)

### Configuring the Router or Switch to Operate in Client Mode

To configure the local router or switch to operate in client mode, include the **server** statement and other optional statements at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
server address <key key-number> <version value> <prefer>;
authentication-key key-number type type value password;
boot-server address;
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement, as described in .

By default, the router or switch sends NTP version 4 packets to the time server. To set the NTP version level to 1, 2, or 3, include the **version** option.

If you configure more than one time server, you can mark one server preferred by including the **prefer** option.

The following example shows how to configure the router or switch to operate in client mode:

```
[edit system ntp]
authentication-key 1 type md5 value "$9$EgfcvX7VY4ZEcwgoHjkP5Q3CuREyv87";
boot-server 10.1.1.1;
server 10.1.1.1 key 1 prefer;
trusted-key 1;
```

### Configuring the Router or Switch to Operate in Symmetric Active Mode

To configure the local router or switch to operate in symmetric active mode, include the **peer** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
peer address <key key-number> <version value> <prefer>;
```

Specify the address of the remote system. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the remote system, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the remote system. To set the NTP version level to 1, 2 or 3, include the **version** option.

If you configure more than one remote system, you can mark one system preferred by including the **prefer** option:

```
peer address <key key-number> <version value> prefer;
```

### Configuring the Router or Switch to Operate in Broadcast Mode

---

To configure the local router or switch to operate in broadcast mode, include the **broadcast** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
broadcast address <key key-number> <version value> <tll value>;
```

Specify the broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be **224.0.1.1**.

To include an authentication key in all messages sent to the remote system, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the remote system. To set the NTP version level to 1, 2, or 3, include the **version** option.

### Configuring the Router or Switch to Operate in Server Mode

---

In server mode, the router or switch acts as an NTP server for clients when the clients are configured appropriately. The only prerequisite for "server mode" is that the router or switch must be receiving time from another NTP peer or server. No other configuration is necessary on the router or switch.

To configure the local router or switch to operate as an NTP server, include the following statements at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
authentication-key key-number type type value password;  
server address <key key-number> <version value> <prefer>;  
trusted-key [ key-numbers ];
```

Specify the address of the system acting as the time server. You must specify an address, not a hostname.

To include an authentication key in all messages sent to the time server, include the **key** option. The key corresponds to the key number you specify in the **authentication-key** statement.

By default, the router or switch sends NTP version 4 packets to the time server. To set the NTP version level to 1, or 2, or 3, include the **version** option.

If you configure more than one time server, you can mark one server preferred by including the **prefer** option.

The following example shows how to configure the router or switch to operate in server mode:

```
[edit system ntp]  
authentication-key 1 type md5 value "$9$tXERuBEreWx-wtuLNdboaUjH.T3AtOESe";  
server 172.17.27.46 prefer;
```

trusted-key 1;

**Related  
Documentation**

- [Understanding NTP Time Servers on page 16](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)

## Specifying the Physical Location of the Switch

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

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

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

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

**Related  
Documentation**

- [Example: Configuring the Name of the Switch, IP Address, and System ID on page 200](#)

## Configuring the Root Password

The Junos OS is preinstalled on the router or switch. When the router or switch is powered on, it is ready to be configured. Initially, you log in as the user **root** with no password. The root directory of a UNIX device is the entry point to all other folders and files on that device. As a result, access to the root directory is restricted by default to a predefined user account known as the *root user*. The root user (also referred to as *superuser*) has unrestricted access and full permissions within the system. The expression “log in as root” is commonly used when an action requires the user to log into the device as the root user.



**NOTE:** If you configure a blank password using the **encrypted-password** statement at the **[edit system root-authentication]** hierarchy level for root authentication, you can commit a configuration but you *cannot* log in as the root user and gain root level access to the router or switch.

After you log in, you should configure the root (superuser) password by including the **root-authentication** statement at the **[edit system]** hierarchy level and configuring one of the password options:

```
[edit system]
root-authentication {
  (encrypted-password "password"| plain-text-password);
  load-key-file URL filename;
  ssh-dsa "public-key" <from hostname>;
  ssh-ecdsa "public-key" <from hostname>;
  ssh-rsa "public-key" <from hostname>;
}
```

If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system]
user@host# set root-authentication plain-text-password
New password: type password here
Retype new password: retype password here
```

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
  - You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
  - Valid passwords must contain at least one change of case or character class.

You can use the **load-key-file** *URL filename* statement to load an SSH key file that was previously generated using **ssh-keygen**. The *URL filename* is the path to the file's location and name. When using this option, the contents of the key file are copied into the

configuration immediately after entering the **load-key-file** *URL* statement. This command loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

Optionally, you can use the **ssh-dsa**, **ssh-eccdsa**, or **ssh-rsa** statements to directly configure SSH RSA, DSA, or ECDSA keys to authenticate root logins. You can configure more than one public key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

To view the SSH keys entries, use the configuration mode **show** command. For example:

```
[edit system]
user@host# set root-authentication load-key-file my-host:.ssh/id_dsa.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@host# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
20740496252839038203869014158453496417001961060835872296
15634757491827360336127644187426594689320773910834481012
68312595772262546166799927831612350043866091586628382248
97467326056611921489539813965561563786211940327687806538
16960202749164163735913269396344008443 boojum@juniper.net"; #
  SECRET-DATA
}
```

Junos-FIPS software has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the router or switch, you cannot configure passwords unless they meet this standard. If you use the **encrypted-password** option, then a null-password (empty) is not permitted.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

#### Related Documentation

- *Configuring the Root Password*
- [Example: Configuring a Plain-Text Password for Root Logins on page 1382](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1386](#)
- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 196](#)
- *Recovering the Root Password*

## Configuring the Router or Switch to Listen for Broadcast Messages Using NTP

When you are using NTP, you can configure the local router or switch to listen for broadcast messages on the local network to discover other servers on the same subnet by including the **broadcast-client** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
broadcast-client;
```

When the router or switch detects a broadcast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *broadcast client* mode, in which it listens for, and synchronizes to, succeeding broadcast messages.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

**Related  
Documentation**

- [Configuring the Router or Switch to Listen for Multicast Messages Using NTP on page 172](#)
- [Configuring the NTP Time Server and Time Services on page 166](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)

## Configuring the Router or Switch to Listen for Multicast Messages Using NTP

When you are using NTP, you can configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet by including the **multicast-client** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]  
multicast-client <address>;
```

When the router or switch receives a multicast message for the first time, it measures the nominal network delay using a brief client-server exchange with the remote server. It then enters *multicast client* mode, in which it listens for, and synchronizes to, succeeding multicast messages.

You can specify one or more IP addresses. (You must specify an address, not a hostname.) If you do, the router or switch joins those multicast groups. If you do not specify any addresses, the software uses **224.0.1.1**.

To avoid accidental or malicious disruption in this mode, both the local and remote systems must use authentication and the same trusted key and key identifier.

**Related  
Documentation**

- [Configuring the Router or Switch to Listen for Broadcast Messages Using NTP on page 171](#)
- [Configuring the NTP Time Server and Time Services on page 166](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)

## Configuring System Alarms to Appear Automatically Upon Login

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

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

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

- Related Documentation**
- *System Alarms on J Series Routers*
  - [show system alarms on page 974](#)

## Configuring Time-Based User Access

The Junos OS enables you to configure time-based restrictions for user access to log in to a device. This is useful for restricting the time and duration of user logins for all users belonging to a login class. You can specify the days of the week when users can log in, the access start time, and the access end time.

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

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

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

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

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

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

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



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

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

#### Related Documentation

- *Examples: Configuring Time-Based User Access*
- *Defining Junos OS Login Classes*
- *access-end*
- *access-start*
- *allowed-days*
- [access-end on page 246](#)
- [access-start on page 247](#)
- [allowed-days on page 250](#)

## Configuring the Timeout Value for Idle Login Sessions

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

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

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

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

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

```
user@host# Session will be closed in 5 minutes if there is no activity.
```



Warning: session will be closed in 1 minute if there is no activity  
 Warning: session will be closed in 10 seconds if there is no activity  
 Idle timeout exceeded: closing session

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

**Related  
Documentation**

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

## Configuring a QFX3500 Device as a Standalone Switch

If you are using the QFX3500 device as a standalone switch, you must perform the initial configuration of the QFX3500 device through the console port using the command-line interface (CLI). If you are using the QFX3500 as a Node device in a QFX3000 QFabric system, you instead perform the initial setup of a QFabric system on a QFX3100 Director device (see *Performing the QFabric System Initial Setup on a QFX3100 Director Group*).

Before you begin connecting and configuring a QFX3500 device, set the following parameter values on the console server or PC:

- Baud Rate—9600
- Flow Control—None
- Data—8
- Parity—None
- Stop Bits—1
- DCD State—Disregard

To connect and configure the device from the console:

1. Connect the console port to a laptop or PC using the supplied RJ-45 cable and RJ-45 to DB-9 adapter. The console (**CON**) port is located on the front panel of the device.
2. Log in as **root**. There is no password. If the software booted before you connected to the console port, you might need to press the Enter key for the prompt to appear.

```
login: root
```

3. Start the CLI.

```
root@% cli
```

4. Enter configuration mode.

```
root> configure
```

5. Add a password to the root administration user account.

```
[edit]
```

```
root@# set system root-authentication plain-text-password
```

```
New password: password
```

```
Retype new password: password
```

6. (Optional) Configure the name of the device. If the name includes spaces, enclose the name in quotation marks (" ").

```
[edit]
root@# set system host-name host-name
```

7. Configure the default gateway.

```
[edit]
root@# set routing-options static route default next-hop address
```

8. Configure the IP address and prefix length for the device management interface.

```
[edit]
root@# set interfaces me0 unit 0 family inet address address/prefix-length
```



**CAUTION:** Configuring the two management Ethernet interfaces within the same subnet is not supported.



**NOTE:** The management ports are on the front panel of the QFX3500 device. They are labeled C0 and C1 on the front panel. In the CLI they are referred to as me0 and me1.

9. (Optional) Configure the static routes to remote prefixes with access to the management port.

```
[edit]
root@# set routing-options static route remote-prefix next-hop destination-ip retain
no-readvertise
```

10. Enable telnet service.

```
[edit]
root@# set system services telnet
```



**NOTE:** When Telnet is enabled, you cannot log in to a QFX3500 device through Telnet using root credentials. Root login is allowed only for SSH access.

11. Commit the configuration to activate it on the device.

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

#### Related Documentation

- *Installing and Connecting a QFX3500 Device*
- *QFX3000-G QFabric System Installation Overview*
- *Understanding QFX3000-G QFabric System Hardware Configurations*

## Creating an Emergency Boot Device

If Junos OS on the device is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



**NOTE:** In the following procedure, we assume that you are creating the emergency boot device on a QFX device or EX4600 device. You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device:

1. Use FTP to copy the installation media image into the `/var/tmp` directory on the device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the `su` command:

```
% su
Password: password
```



**NOTE:** The password is the root password for the device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command on the QFX3500, QFX3600, and QFX3600-I devices:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/install-media-qfx3500.junos_11.1 of=/dev/da1 bs=16k
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

6. Enter the following command on the QFX5100 and EX4600 devices:

```
root@device% dd if=/var/tmp/filename of=/dev/da0 bs=1048576
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/jinstall-vjunos-usb-13.2.img of=/dev/da0 bs=1048576
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

7. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

#### Related Documentation

- *USB Port Specifications for the QFX Series*
- [Performing a Recovery Installation on page 116](#)

- [Performing a QFabric System Recovery Installation on the Director Group on page 7233](#)
- [Performing a Recovery Installation on page 118](#)

## Creating a Snapshot and Using It to Boot a QFX Series Switch

The system snapshot feature takes a “snapshot” of the files currently used to run the QFX Series switch—the complete contents of the `/config` and `/var` directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration—and copies all of these files into an alternate (internal, meaning internal flash, or an external, meaning USB flash) memory source. You can then use these snapshots to boot the switch at the next bootup or as a backup boot option.

This topic includes the following tasks:

- [Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch on page 178](#)
- [Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch on page 179](#)
- [Creating a Snapshot on the Alternate Slice of the Boot Media on page 179](#)

### Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch

---

A snapshot can be created on USB flash memory after a switch is booted using files stored in internal memory.

Ensure that you have the following tools and parts available before creating a snapshot on a USB Flash drive:

- A USB flash drive that meets the QFX Series switch USB port specifications. See *USB Port Specifications for the QFX Series*.

To create a snapshot on USB flash memory and use it to boot the switch:

1. Place the snapshot into USB flash memory:

```
user@switch> request system snapshot partition
```



**NOTE:** This example uses the `partition` option. If you have already created a partition for the snapshot, you don't need to use the `partition` option.

2. (Optional) Perform this step if you want to boot the switch now using the snapshot stored on the external USB flash drive. If you created the snapshot as a backup, do not perform this step.

- To reboot the switch using the most recently created snapshot:

```
user@switch> request system reboot
```

- To reboot the switch using a snapshot in a specific partition on the USB flash drive:

```
user@switch> request system reboot slice 1
```

### Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch

A snapshot can be created on internal memory after a switch is booted using files stored in external memory.

To create a snapshot in internal memory and use it to boot the switch:

1. Place the snapshot files in internal memory:

```
user@switch> request system snapshot partition
```



**NOTE:** This example uses the `partition` option. If you have already created a partition for the snapshot, you don't need to use the `partition` option.

2. (Optional) Perform this step if you want to boot the switch now using the newly created snapshot. If you created the snapshot as a backup, do not perform this step.

- To reboot the switch using the most recently created snapshot:

```
user@switch> request system reboot
```

- To reboot the switch using a snapshot in a specific partition in internal memory:

```
user@switch> request system reboot slice 1
```

### Creating a Snapshot on the Alternate Slice of the Boot Media

The alternate slice of the boot media contains a backup software image that the switch can boot from if it is unable to boot from the primary slice. When you upgrade software, the new software image gets copied only to the primary slice of the boot media.

To create a snapshot of the currently booted software image on the backup slice of the boot media:

```
user@switch> request system reboot slice alternate
```

After the system boots up, you will see the following message before the login prompt:

**WARNING: THIS DEVICE HAS BOOTED FROM THE BACKUP JUNOS IMAGE**

It is possible that the primary copy of JUNOS failed to boot up properly, and so this device has booted up from the backup copy.

Please re-install JUNOS to recover the primary copy in case it has been corrupted.

The system will generate an alarm indicating that the switch has booted from the backup slice.

#### Related Documentation

- [Verifying That a System Snapshot Was Created on a QFX Series Switch](#)
- [Understanding System Snapshot on page 32](#)

## Creating a Snapshot and Using It to Boot QFX5100 and EX4600 Devices

The system snapshot feature takes a “snapshot” of the files currently used to run the device—the complete contents of the **/config** directories, which include the running Juniper Networks Junos OS, the active configuration, and the rescue configuration, as well as the host OS—and copies all of these files into an external USB flash drive. You can use the snapshot to boot the device at the next bootup or as a backup boot option.

This topic includes the following tasks:

- [Creating a Snapshot on an External USB Flash Drive and Using It to Boot the Device on page 180](#)

### Creating a Snapshot on an External USB Flash Drive and Using It to Boot the Device

---

A snapshot can be created on an external USB flash drive after a device is booted using files stored in internal memory.

Ensure that you have the following tools and parts available before creating a snapshot on an external USB flash drive:

- An external USB flash drive that meets the device USB port specifications. See *USB Port Specifications for the QFX Series*.

To create a snapshot on the external USB flash drive and use it to boot the device:

1. Insert the external USB flash drive.
2. Issue the **request system snapshot** command.  

```
user@device> request system snapshot
```
3. (Optional) Perform this step if you want to boot the device now using the snapshot stored on the external USB flash drive. If you created the snapshot as a backup, do not perform this step.
  - Insert the external USB flash drive.
  - Power cycle the device.

The external USB flash drive is detected.

- The software prompts you with the following options:

```
Junos Snapshot Installer - (c) Juniper Networks 2013
Reboot
Install Junos Snapshot [13.2-20131115_x_132_x51_vjunos.0
Boot to host shell [debug]
```

- Select **Install Junos Snapshot** to install the snapshot located on the external USB flash drive to the device.

The device copies the software from the external USB flash drive, occasionally displaying status messages. When the software is finished being copied from the external USB flash drive to the device, the device then reboots from the internal

flash storage on which the software was just installed. When the reboot is complete, the device displays the Junos OS login prompt:

```
root@device#
```

- Related Documentation**
- [Verifying That a System Snapshot Was Created on a QFX Series Switch](#)
  - [Understanding System Snapshot on page 32](#)

## Including the Year or Millisecond in Timestamps

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

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

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

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

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

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

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

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



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

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

- Related Documentation**
- [Single-Chassis System Logging Configuration Overview](#)
  - [Examples: Configuring System Logging](#)

## Mapping the Hostname of the Switch to IP Addresses

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

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

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

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

### Related Documentation

- [Reaching a Domain Name System Server on page 158](#)
- [Example: Configuring the Name of the Router, IP Address, and System ID](#)
- [static-host-mapping on page 311](#)



## Methods for Configuring Junos OS

You can use any of the methods shown in [Table 21 on page 183](#) to configure Junos OS.

**Table 21: Methods for Configuring Junos OS**

Method	Description
Command-line interface (CLI)	Create the configuration for the device using the CLI. You can enter commands from a single command line, and scroll through recently executed commands.
ASCII file	Load an ASCII file containing a configuration that you created earlier, either on this system or on another system. You can then activate and run the configuration file, or you can edit it using the CLI and then activate it.
J-Web graphical user interface (GUI)	Use the J-Web GUI to configure the device. J-Web enables you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser. The J-Web GUI is preinstalled on J Series Routers and is an optional software package that can be installed on M Series and T Series routers. J-Web is not available for the QFX Series.
Junos XML management protocol (API)	Use Junos XML protocol Perl client modules to develop custom applications for configuring information on devices that run Junos OS. Client applications use the Junos XML management protocol to request and change configuration information on Juniper Networks J Series, M Series, and T Series routers. The Junos XML management protocol is customized for Junos OS, and operations in the API are equivalent to those in the Junos OS CLI.
NETCONF application programming interface (API)	Use NETCONF Perl client modules to develop custom applications for configuring information on devices that run Junos OS. Client applications use the NETCONF XML management protocol to request and change configuration information on Juniper Networks J Series, M Series, and T Series routers. The NETCONF XML management protocol includes features that accommodate the configuration data models of multiple vendors.
Configuration commit scripts	Create scripts that run at commit time to enforce custom configuration rules. Commit scripts are written in Extensible Stylesheet Language Transformations (XSLT). Commit scripts are not available for the QFX Series.

The following sections contain complete descriptions of the methods you can use to configure Junos OS:

- [Junos OS Command-Line Interface on page 184](#)
- [ASCII File on page 184](#)
- [J-Web Package on page 184](#)
- [Junos XML Management Protocol Software on page 185](#)

- [NETCONF XML Management Protocol Software on page 185](#)
- [Configuration Commit Scripts on page 185](#)

---

## Junos OS Command-Line Interface

The Junos OS CLI is a straightforward command interface. You use Emacs-style keyboard sequences to move around on a command line and scroll through a buffer that contains recently executed commands. You type commands on a single line, and the commands are executed when you press the Enter key. The CLI also provides command help and command completion. For more information about the CLI, see the *CLI User Guide* and the [CLI Explorer](#).

---

## ASCII File

You can load an ASCII file containing a configuration that you created earlier, either on this system or another system. You can then activate and run the configuration file as is, or you can edit it using the CLI and then activate it.

---

## J-Web Package

As an alternative to entering CLI commands, Junos OS supports the J-Web GUI. The J-Web user interface enables you to monitor, configure, troubleshoot, and manage the router on a client by means of a Web browser with Hypertext Transfer Protocol (HTTP) or HTTP over Secure Sockets Layer (HTTPS) enabled.

The J-Web user interface is preinstalled on J Series Routers. It is provided as an optional, licensed software package (jweb package) on M Series and T Series routers. The jweb package is not included in jinstall and jbundle software bundles. It must be installed separately. To install the package on M Series and T Series routers, follow the procedure described in the *Installation and Upgrade Guide*.

J-Web supports weak (56-bit) encryption by default. This enables international customers to install J-Web and use HTTPS connections for J-Web access. Domestic customers can also install the jcrypto strong encryption package. This package automatically overrides the weak encryption. For more information about the J-Web GUI, see the *J-Web Interface User Guide*.



**NOTE:** Because the J-Web package is bundled separately from other packages, it is possible to have a version mismatch between J-Web and other Junos OS packages you have installed.

To check for a version mismatch, use the `show system alarms` CLI command. If the version number does not match exactly, a system alarm appears. For example, if you install the 7.4R1.2 jroute package and the 7.4R1.1 jweb package, an alarm is activated. For more information on the `show system alarms` command, see the [CLI Explorer](#).

---

---

### Junos XML Management Protocol Software

---

The Junos XML management protocol is an Extensible Markup Language (XML) application that client applications use to request and change configuration information on Juniper Networks J Series, M Series, MX Series, and T Series routers. This API is customized for Junos OS, and operations in the API are equivalent to Junos OS CLI configuration mode commands. The Junos XML management protocol includes a set of Perl modules that enable client applications to communicate with a Junos XML protocol server on the router. The Perl modules are used to develop custom applications for configuring and monitoring Junos OS.

For a complete description of how to use Junos XML and Junos XML management protocol software, see the *Junos XML Management Protocol Developer Guide*.

---

### NETCONF XML Management Protocol Software

---

The NETCONF XML management protocol is an Extensible Markup Language (XML) application that client applications can use to request and change configuration information on Juniper Networks J Series, M Series, MX Series, and T Series routers. This API is customized for Junos OS, and includes features that accommodate the configuration data models of multiple vendors. The NETCONF XML management protocol includes a set of Perl modules that enable client applications to communicate with a NETCONF server on the router. The Perl modules are used to develop custom applications for configuring and monitoring Junos OS.

For a complete description of how to use Junos XML and NETCONF XML management protocol software, see the *NETCONF XML Management Protocol Developer Guide*.

---

### Configuration Commit Scripts

---

You can create and use scripts that run at commit time to enforce custom configuration rules. If a configuration breaks the custom rules, the script can generate actions that the Junos OS performs. These actions include:

- Generating custom error messages
- Generating custom warning messages
- Generating custom system log messages
- Making changes to the configuration

Configuration commit scripts also enable you to create macros, which expand simplified custom aliases for frequently used configuration statements into standard Junos OS configuration statements. Commit scripts are written in Extensible Stylesheet Language Transformations (XSLT). For more information, see the *Junos OS Automation Library*.

#### Related Documentation

- *Configuring Junos OS from External Devices*

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

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

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

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

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

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



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

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

The following example shows how to change the current time zone to **America/New\_York**:

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

### Related Documentation

- [Understanding NTP Time Servers on page 16](#)
- [Updating the IANA Time Zone Database on Junos Devices on page 190](#)

## Rebooting and Halting a Device

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

```
user@switch> request system reboot ?
Possible completions:
<[Enter]>           Execute this command
```

```

all-members      Reboot all virtual chassis members
at               Time at which to perform the operation
both-routing-engines Reboot both the Routing Engines
fast-boot        Enable fast reboot
in               Number of minutes to delay before operation
local            Reboot local virtual chassis member
member           Reboot specific virtual chassis member (0..9)
message          Message to display to all users
other-routing-engine Reboot the other Routing Engine
|               Pipe through a command
{master:0}

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

```



**NOTE:** Not all options shown in the preceding command output are available on all QFX Series and EX4600 devices. For example, the `fast-boot` option is available only on QFX5100. See the documentation for the [request system reboot](#) command for details about options.

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



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

```

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

```



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

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

- Related Documentation**
- [clear system reboot on page 355](#)
  - [request system reboot on page 415](#)
  - [request system halt on page 400](#)
  - [request system power-off on page 410](#)
  - [Connecting a QFX Series Device to a Management Console](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

- Related Documentation**
- [Understanding Configuration Files on page 1242](#)
  - [Loading a Previous Configuration File on page 1252](#)
  - [Reverting to the Rescue Configuration on page 189](#)

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

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

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



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

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

1. `user@switch> request system zeroize`  
warning: System will be rebooted and may not boot without configuration  
Erase all data, including configuration and log files? [yes,no] (yes)
2. Type **yes** to remove configuration and log files and revert to the factory default configuration.
3. Complete the initial configuration of the switch. See or [“Configuring a QFX3500 Device as a Standalone Switch” on page 175](#)

**Related Documentation** • [request system zeroize on page 476](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.  
  
[edit]  
`user@switch# load override filename`
2. Commit your changes.  
  
[edit]  
`user@switch# commit filename`

**Related Documentation** • [Setting or Deleting the Rescue Configuration on page 1261](#)  
• [Reverting to the Default Factory Configuration on page 188](#)  
• [Configuration File Terms on page 11](#)

## Saving Core Files Generated by Junos OS Processes

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

- To disable the saving of core files and associated context information:  
  
[edit system]  
`no-saved-core-context;`
- To save the core files only:  
  
[edit system]  
`saved-core-files number;`

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

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

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

#### Related Documentation

- [Viewing Core Files from Junos OS Processes on page 196](#)

## Updating the IANA Time Zone Database on Junos Devices

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

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

1. [Importing and Installing Time Zone Files on page 190](#)
2. [Configuring a Custom Time Zone on page 191](#)

### Importing and Installing Time Zone Files

---

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

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

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

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

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



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

```
# ftp ftp.iana.org/tz  
# bin  
# get tzdata-latest.tar.gz
```





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

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

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

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

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

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

7. Set the time zone and commit the configuration:

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

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

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

### Configuring a Custom Time Zone

To use a custom time zone, follow these steps:

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

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

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

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

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

#### Related Documentation

- *Modifying the Default Time Zone for a Router or Switch Running Junos OS*
- *NTP Overview*
- [Understanding NTP Time Servers on page 16](#)

- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)
- *use-imported-time-zones*

## Setting the Date and Time

1. Enter operational mode in the CLI.

2. Enter the following command:

```
user@switch> set date YYYYMMDDHHMM.ss source-address
```

For example, the following command sets the date and time.

```
user@switch# set date 201102151010.55
```

3. To set the date and time from an NTP server, enter the following command:

```
user@switch# set date ntp servers
```

For example, the following command sets the date and time from an NTP server:

```
user@switch# set date ntp 200.40.40.1
```

4. To set the date and time from more than one NTP server, enter the same command:

```
user@switch# set date ntp servers
```

For example, the following command sets the date and time from more than one NTP server:

```
user@switch# set date ntp 200.40.40.1 200.40.40.2
```

**Related Documentation**

- *set date*

## Specifying Access Privileges for Junos OS Operational Mode Commands

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

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

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

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

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

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

allow-commands "show interfaces";



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

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

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

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

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

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

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

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

Anchors are required when specifying complex regular expressions with **allow-commands** or **deny-commands** statements. For example, when specifying multiple commands using

the pipe (|) symbol for **allow-commands**, the following syntax is incorrect:

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

#### Related Documentation

- [Example: Configuring Access Privileges for Operational Mode Commands on page 1381](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1339](#)
- *allow-commands*
- *deny-commands*

## Synchronizing and Coordinating Time Distribution Using NTP

Using NTP to synchronize and coordinate time distribution in a large network involves these tasks:

1. [Configuring NTP on page 194](#)
2. [Configuring the NTP Boot Server on page 194](#)
3. [Specifying a Source Address for an NTP Server on page 195](#)

### Configuring NTP

---

- To configure NTP on the router or switch, include the **ntp** statement at the **[edit system]** hierarchy level:

```
[edit system]
ntp {
  authentication-key number type type value password;
  boot-server (address | hostname);
  broadcast <address> <key key-number> <version value> <ttl value>;
  broadcast-client;
  multicast-client <address>;
  peer address <key key-number> <version value> <prefer>;
  server address <key key-number> <version value> <prefer>;
  source-address source-address;
  trusted-key [ key-numbers ];
}
```

### Configuring the NTP Boot Server

---

When you boot the router or switch, it issues an **ntpdate** request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. Otherwise, NTP will not be able to synchronize to a time server if the server's time appears to be very far off of the local router's or switch's time.

- To configure the NTP boot server, include the **boot-server** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
boot-server (address | hostname);
```

Specify either the IP address or the hostname of the network server.

### Specifying a Source Address for an NTP Server

For IP version 4 (IPv4), you can specify that if the NTP server configured at the **[edit system ntp]** hierarchy level is contacted on one of the loopback interface addresses, the reply always uses a specific source address. This is useful for controlling which source address NTP uses to access your network when it is either responding to or sending an NTP client request from your network.

To configure the specific source address that the reply will always use, and the source address that requests initiated by NTP server will use, include the **source-address** statement at the **[edit system ntp]** hierarchy level:

```
[edit system ntp]
source-address source-address;
```

**source-address** is a valid IP address configured on one of the router or switch interfaces.



**NOTE:** If a firewall filter is applied on the loopback interface, ensure that the source address specified for the NTP server at the **[edit system ntp]** hierarchy level is explicitly included as one of the match criteria in the firewall filter. This enables the Junos OS to accept traffic on the loopback interface from the specified source address.

The following example shows a firewall filter with the source address 10.0.10.100 specified in the **from** statement included at the **[edit firewall filter firewall-filter-name]** hierarchy:

```
[edit firewall filter Loopback-Interface-Firewall-Filter]
term Allow-NTP {
  from {
    source-address {
      172.17.27.46/32; // IP address of the NTP server
      10.0.10.100/32; // Source address specified for the NTP server
    }
  }
  then accept;
}
```

If no source address is configured for the NTP server, include the primary address of the loopback interface in the firewall filter.

#### Related Documentation

- [Understanding NTP Time Servers on page 16](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)

## Viewing Core Files from Junos OS Processes

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

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

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

### Related Documentation

- [Saving Core Files from Junos OS Processes](#)
- [Saving Core Files Generated by Junos OS Processes on page 189](#)

## Configuration Examples

---

- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 196](#)
- [Reaching a Domain Name System Server on page 198](#)
- [Example: Configuring the Name of the Switch, IP Address, and System ID on page 200](#)
- [Example: Configuring NTP on page 200](#)
- [Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization on page 203](#)

### Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 196](#)
- [Overview on page 196](#)
- [Configuration on page 197](#)

#### Requirements

---

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**s, **minimum-punctuations**, or **minimum-upper-cases**.

#### Overview

---

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login password minimum-length 12
set system login password maximum-length 22
set system login password minimum-numeric 1
set system login password minimum-upper-cases 1
set system login password minimum-lower-cases 1
set system login password minimum-punctuations 1
```

## Configuring Requirements for Plain-Text Passwords

**Step-by-Step Procedure** This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the [system login password] hierarchy level.
 

```
user@host> edit
[edit]
user@host# edit system login password
```
2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.
 

```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.
 

```
[edit system login password]
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```
4. Require users to set a password that has at least one punctuation-class character and at least one number.
 

```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

### Results

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
minimum-length 12;
maximum-length 22;
minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;
```

#### Related Documentation

- [Special Requirements for Junos OS Plain-Text Passwords on page 1339](#)
- *password (Login)*

## Reaching a Domain Name System Server

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

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

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

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

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

1. Reference the IP addresses of your DNS servers.

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



The following example shows how to reference two DNS servers:

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

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

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

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

```
[edit system]
domain-name domain-name;
```

The following example shows how to configure the domain name:

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

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

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

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

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

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

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

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

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

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

5. Commit the configuration.

```
user@host# commit
```

6. Verify the configuration.

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

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

For example:

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

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

**Related Documentation** • [Understanding DNS](#)

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

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

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

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

### Example: Configuring NTP

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. NTP uses a returnable-time design in which a distributed subnet of time servers operating in a self-organizing, hierarchical primary-secondary configuration synchronizes local clocks within the subnet and to

national time standards by means of wire or radio. The servers also can redistribute reference time using local routing algorithms and time daemons.

This example shows how to configure NTP:

- [Requirements on page 201](#)
- [Overview on page 201](#)
- [Configuration on page 201](#)
- [Verification on page 202](#)

### Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later
- A switch connected to a network on which an NTP boot server and NTP server reside

### Overview

Debugging and troubleshooting are much easier when the timestamps in the log files of all switches are synchronized, because events that span a network can be correlated with synchronous entries in multiple logs. We recommend using the Network Time Protocol (NTP) to synchronize the system clocks of your switch and other network equipment.

In this example, an administrator wants to synchronize the time in a switch to a single time source. We recommend using authentication to make sure that the NTP peer is trusted. The **boot-server** statement identifies the server from which the initial time of day and date are obtained when the router or switch boots. The **server** statement identifies the NTP server used for periodic time synchronization. The **authentication-key** statement specifies that an HMAC-Message Digest 5 (MD5) scheme is used to hash the key value for authentication, which prevents the switch from synchronizing with an attacker's host that is posing as the time server.

### Configuration

To configure NTP:

#### CLI Quick Configuration

To quickly configure NTP, copy the following commands and paste them into the switch's terminal window:

```
[edit system]
set ntp boot-server 10.1.4.1
set ntp server 10.1.4.2
set ntp authentication-key 2 type md5 value "$9$ah1j8"
```

#### Step-by-Step Procedure

To configure NTP :

1. Specify the boot server:
 

```
[edit system]
user@switch# set ntp boot-server 10.1.4.1
```
2. Specify the NTP server:

- ```
[edit system]
user@switch# set ntp server 10.1.4.2
```
3. Specify the key number, authentication type (MD5), and key for authentication:

```
[edit system]
user@switch# set ntp authentication-key 2 type md5 value "$9$aH1j8"
```

**Results** Check the results:

```
[edit system]
user@switch# show
ntp {
  boot-server 10.1.4.1;
  authentication-key 2 type md5 value "$9$aH1j8"; ## SECRET-DATA
  server 10.1.4.2;
}
```

---

### Verification

To confirm that the configuration is correct, perform these tasks:

- [Checking the Time on page 202](#)
- [Displaying the NTP Peers on page 202](#)
- [Displaying the NTP Status on page 203](#)

#### *Checking the Time*

**Purpose** Check the time that has been set on the switch.

**Action** Enter the **show system uptime** operational mode command to display the time.

```
user@switch> show system uptime
fpc0:
-----
Current time: 2009-06-12 12:49:03 PDT
System booted: 2009-05-15 06:24:43 PDT (4w0d 06:24 ago)
Protocols started: 2009-05-15 06:27:08 PDT (4w0d 06:21 ago)
Last configured: 2009-05-27 14:57:03 PDT (2w1d 21:52 ago) by admin1
12:49PM up 28 days, 6:24, 1 user, load averages: 0.05, 0.06, 0.01
```

**Meaning** The output shows that the current date and time are June 12, 2009 and 12:49:03 PDT. The switch booted 4 weeks, 6 hours, and 24 minutes ago, and its protocols were started approximately 3 minutes before it booted. The switch was last configured by user **admin1** on May 27, 2009, and there is currently one user logged in to the switch.

The output also shows that the load average is 0.05 seconds for the last minute, 0.06 seconds for the last 5 minutes, and 0.01 seconds for the last 15 minutes.

#### *Displaying the NTP Peers*

**Purpose** Verify that the time has been obtained from an NTP server.

**Action** Enter the **show ntp associations** operational mode command to display the NTP server from switch obtained its time.

```
user@switch> show ntp associations
      remote          refid      st t when poll reach  delay  offset  jitter
=====
*ntp5.domain1.ne .GPS.          1 u  414 1024  377   3.435   4.002   0.765
```

**Meaning** The asterisk (\*) in front of the NTP server name, or peer, indicates that the time is synchronized and obtained from this server. The delay, offset, and jitter are displayed in milliseconds.

### *Displaying the NTP Status*

**Purpose** View the configuration of the NTP server and the status of the system.

**Action** Enter the **show ntp status** operational mode command to view the status of the NTP.

```
user@switch> show ntp status
status=0644 leap_none, sync_ntp, 4 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Mon Apr 13 19:09:05 UTC 2009 (1)",
processor="powerpc", system="JUNOS9.5R1.8", leap=00, stratum=2,
precision=-18, rootdelay=2.805, rootdispersion=42.018, peer=48172,
refid=172.17.28.5,
reftime=cddd397a.60e6d7bf Fri, Jun 12 2009 13:30:50.378, poll=10,
clock=cddd3b1b.ec5a2bb4 Fri, Jun 12 2009 13:37:47.923, state=4,
offset=3.706, frequency=-23.018, jitter=1.818, stability=0.303
```

**Meaning** The output shows status information about the switch and the NTP.

- Related Documentation**
- [Understanding NTP Time Servers on page 16](#)
  - [ntp on page 294](#)
  - [Configuring the NTP Time Server and Time Services on page 166](#)
  - [CLI Explorer](#)
  - *Junos OS Baseline Network Operations Guide*

## **Example: Configuring NTP as a Single Time Source for Router and Switch Clock Synchronization**

Debugging and troubleshooting are much easier when the timestamps in the log files of all the routers or switches are synchronized, because events that span the network can be correlated with synchronous entries in multiple logs. We strongly recommend using the Network Time Protocol (NTP) to synchronize the system clocks of routers, switches, and other network equipment.

By default, NTP operates in an entirely unauthenticated manner. If a malicious attempt to influence the accuracy of a router or switch's clock succeeds, it could have negative effects on system logging, make troubleshooting and intrusion detection more difficult, and impede other management functions.

The following sample configuration synchronizes all the routers or switches in the network to a single time source. We recommend using authentication to make sure that the NTP

peer is trusted. The **boot-server** statement identifies the server from which the initial time of day and date is obtained when the router boots. The **server** statement identifies the NTP server used for periodic time synchronization. The **authentication-key** statement specifies that an HMAC-Message Digest 5 (MD5) scheme should be used to hash the key value for authentication, which prevents the router or switch from synchronizing with an attacker's host posing as the time server.

```
[edit]
system {
  ntp {
    authentication-key 2 type md5 value "$9$aHlj8gqQ1gijjghgjgiiii"; # SECRET-DATA
    boot-server 10.1.4.1;
    server 10.1.4.2;
  }
}
```

#### Related Documentation

- [NTP Overview](#)
- [Understanding NTP Time Servers on page 16](#)
- [authentication-key](#)
- [boot-server on page 258](#)
- [server on page 306](#)
- [show ntp associations on page 951](#)
- [show ntp status on page 953](#)

---

## Configuration Statements

- [QFX Series CLI Hierarchy on page 207](#)
- [access-end on page 246](#)
- [access-start on page 247](#)
- [accounting on page 248](#)
- [accounting-port on page 249](#)
- [allow-commands on page 249](#)
- [allow-configuration on page 250](#)
- [allowed-days on page 250](#)
- [allow-transients on page 251](#)
- [announcement on page 251](#)
- [archival on page 252](#)
- [arp \(System\) on page 253](#)
- [authentication \(Login\) on page 254](#)
- [authentication-key on page 255](#)
- [authentication-order on page 256](#)
- [auxiliary on page 257](#)

- [boot-server \(NTP\) on page 258](#)
- [broadcast on page 259](#)
- [broadcast-client on page 260](#)
- [change-type on page 260](#)
- [checksum on page 261](#)
- [class \(Defining Login Classes\) on page 262](#)
- [class \(Assigning a Class to an Individual User\) on page 263](#)
- [commit on page 264](#)
- [compress-configuration-files \(System\) on page 265](#)
- [console \(Physical Port\) on page 266](#)
- [default-address-selection on page 267](#)
- [deny-commands on page 268](#)
- [deny-configuration on page 269](#)
- [destination \(Accounting\) on page 270](#)
- [destination-override on page 271](#)
- [direct-access on page 271](#)
- [domain-name on page 272](#)
- [domain-search on page 272](#)
- [explicit-priority on page 273](#)
- [events on page 274](#)
- [format on page 274](#)
- [host-name on page 275](#)
- [icmpv4-rate-limit on page 275](#)
- [idle-timeout on page 276](#)
- [internet-options on page 276](#)
- [l2-learning on page 277](#)
- [load-key-file on page 278](#)
- [location on page 279](#)
- [login on page 280](#)
- [login-alarms on page 281](#)
- [login-tip on page 281](#)
- [max-configurations-on-flash on page 282](#)
- [maximum-length on page 282](#)
- [message on page 283](#)
- [minimum-changes on page 283](#)
- [minimum-length on page 284](#)
- [minimum-lower-cases on page 285](#)

- [minimum-numeric on page 286](#)
- [minimum-punctuations on page 287](#)
- [minimum-upper-cases on page 288](#)
- [multicast-client on page 288](#)
- [name-server on page 289](#)
- [no-multicast-echo on page 290](#)
- [no-ping-record-route on page 291](#)
- [no-ping-time-stamp on page 291](#)
- [no-redirects \(IPv4 Traffic\) on page 292](#)
- [no-split-detection on page 293](#)
- [ntp on page 294](#)
- [optional on page 294](#)
- [password \(Login\) on page 295](#)
- [peer on page 296](#)
- [permissions on page 297](#)
- [port \(TACACS+ Server\) on page 297](#)
- [ports on page 298](#)
- [radius \(System\) on page 299](#)
- [refresh \(Commit Scripts\) on page 300](#)
- [refresh-from \(Commit Scripts\) on page 300](#)
- [retry on page 301](#)
- [retry-options on page 302](#)
- [root-authentication on page 303](#)
- [saved-core-context on page 304](#)
- [saved-core-files on page 304](#)
- [secret on page 305](#)
- [server \(TACACS+ Accounting\) on page 305](#)
- [server \(NTP\) on page 306](#)
- [server \(RADIUS Accounting\) on page 307](#)
- [single-connection on page 307](#)
- [source \(Commit Scripts\) on page 308](#)
- [source-address \(NTP, RADIUS, System Logging, or TACACS+\) on page 308](#)
- [source-port \(Port Addresses\) on page 309](#)
- [ssh-dsa on page 309](#)
- [ssh-rsa on page 310](#)
- [static-host-mapping on page 311](#)
- [structured-data on page 312](#)



- [syslog \(System\) on page 313](#)
- [system on page 315](#)
- [tacplus on page 320](#)
- [tacplus-server on page 321](#)
- [timeout on page 322](#)
- [time-format on page 323](#)
- [time-zone on page 324](#)
- [traceoptions \(Commit Scripts\) on page 326](#)
- [traceoptions \(Layer 2 Learning\) on page 328](#)
- [tracing on page 330](#)
- [trusted-key on page 331](#)
- [uid on page 331](#)
- [use-imported-time-zones on page 332](#)
- [user \(Access\) on page 332](#)

## QFX Series CLI Hierarchy

This topic contains the full command-line interface (CLI) statement hierarchy for the QFX Series.

- [\[edit access\] Hierarchy on page 207](#)
- [\[edit accounting-options\] Hierarchy on page 208](#)
- [\[edit chassis\] Hierarchy on page 209](#)
- [\[edit class-of-service\] Hierarchy on page 211](#)
- [\[edit ethernet-switching-options\] Hierarchy on page 213](#)
- [\[edit fabric\] Hierarchy on page 215](#)
- [\[edit fc-fabrics\] Hierarchy on page 216](#)
- [\[edit fc-options\] Hierarchy on page 217](#)
- [\[edit firewall\] Hierarchy on page 217](#)
- [\[edit groups\] Hierarchy on page 218](#)
- [\[edit interfaces\] Hierarchy on page 218](#)
- [\[edit policy-options\] Hierarchy on page 224](#)
- [\[edit protocols\] Hierarchy on page 224](#)
- [\[edit security\] Hierarchy on page 237](#)
- [\[edit snmp\] Hierarchy on page 237](#)
- [\[edit system\] Hierarchy on page 241](#)
- [\[edit vlans\] Hierarchy on page 246](#)

### [\[edit access\] Hierarchy](#)

---

```
access {
```

```
address-assignment
pool pool-name
address-pool pool-name
profile profile-name {
  accounting (Access Profile) {
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
    (authentication-order (ldap radius | none);
    order (radius | none);
  }
  radius {
    accounting-server [server-addresses];
    authentication-server [server-addresses];
  }
}
```

---

### [\[edit accounting-options\] Hierarchy](#)

---

```
accounting-options {
  class-usage-profile profile-name {
    destination-classes {
      destination-class-name;
    }
    file filename;
    interval minutes;
    source-classes {
      source-class-name;
    }
  }
  file filename {
    archive-sites {
      site-name;
    }
    files number;
    nonpersistent;
    size bytes;
    start-time time;
    transfer-interval minutes;
  }
  filter-profile profile-name {
    counters {
      counter-name;
    }
    file filename;
    interval minutes;
  }
  interface-profile profile-name {
    fields {
      input-bytes;
      input-errors;
      input-multicast;
      input-packets;
      input-unicast;
      output-bytes;
      output-errors;
```

```

        output-multicast;
        output-packets;
        output-unicast;
        rpf-check-bytes;
        rpf-check-packets;
        rpf-check6-bytes;
        rpf-check6-packets;
        unsupported-protocol;
    }
    file filename;
    interval minutes;
}
mib-profile profile-name {
    file filename;
    interval minutes;
    object-names {
        mib-object-name;
    }
    operation (get | get-next | walk);
}
policy-decision-statistics-profile profile-name {
    application-aware-access-list-fields {
        address;
        application;
        application-group;
        input-bytes;
        input-interface;
        input-packets;
        mask;
        output-bytes;
        output-packets;
        subscriber-name;
        timestamp;
        vrf-name;
    }
    file filename;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}

```

#### [\[edit chassis\] Hierarchy](#)

```

interconnect-device {
    alarm {
        interface-type {
            link-down (red | yellow | ignore);
        }
    }
    container-devices {
        device-count number;
    }
}

```

```
}
craft-lockout {
  alarm {
    interface-type {
      link-down (red | yellow | ignore);
    }
  }
  container-devices {
    device-count number;
  }
  fpc slot {
    power (on | off);
  }
  routing-engine {
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
fpc slot {
  power (on | off);
}
routing-engine {
  on-disk-failure {
    disk-failure-action (halt | reboot);
  }
}
}
chassis {
  routing-engine {
    redundancy {
      failover {
        on-disk-failure {
          disk-failure-action (halt | reboot);
        }
        on-loss-of-keepalives;
      }
      graceful-switchover;
    }
  }
  aggregated-devices {
    ethernet {
      device-count number;
    }
    alarm {
      interface-type {
        alarm-name (red | yellow | ignore);
      }
    }
  }
}
forwarding-options profile-name {
  num-65-127-prefix value
}
fpc slot {
  auto-speed-detection disable
  pic pic-number{
    port port-number{
```

```

        tunnel-port port-number tunnel-services;
        channel-speed speed;
    }
    port-range port-range-low port-range-high {
        channel-speed speed;
    }
}
}
maximum-ecmp next-hops;
}

```

### [\[edit class-of-service\] Hierarchy](#)

```

class-of-service {
  classifiers {
    (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority level {
          code-points [aliases] [bit-patterns];
        }
      }
    }
  }
  code-point-aliases {
    (dscp | dscp-ipv6 | ieee-802.1) {
      alias-name bits;
    }
  }
  congestion-notification-profile profile-name {
    input {
      ieee-802.1 {
        code-point [code-point-bits] {
          pfc {
            mru mru-value;
          }
        }
      }
    }
    cable-length cable-length-value;
  }
  output {
    ieee-802.1 {
      code-point [code-point-bits] {
        flow-control-queue [queue | list-of-queues];
      }
    }
  }
  drop-profiles {
    profile-name {
      interpolate {
        fill-level low-value fill-level high-value drop-probability 0 drop-probability
          high-value;
      }
    }
  }
}

```

```
forwarding-class class-name {
  loss-priority level {
    code-points [ aliases ] [ bit-patterns ];
  }
}
forwarding-class class-name {
  scheduler scheduler-name;
}
forwarding-class-sets forwarding-class-set-name {
  class class-name;
}
forwarding-classes {
  class {
    class-name {
      queue-num queue-number <no-loss>;
    }
  }
}
host-outbound-traffic {
  forwarding-class class-name;
  dscp-code-point code-point;
}
interfaces {
  interface-name {
    congestion-notification-profile profile-name {
    }
    forwarding-class lossless-forwarding-class-name;
    forwarding-class-set forwarding-class-set-name {
      output-traffic-control-profile profile-name;
    }
    rewrite-value {
      input {
        ieee-802.1 {
          code-point code-point-bits;
        }
      }
    }
  }
  unit logical-unit-number {
    classifiers {
      (dscp | dscp-ipv6 | ieee-802.1 exp) (classifier-name | default);
    }
    forwarding-class class-name;
    rewrite-rules {
      (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
    }
  }
}
multi-destination {
  classifiers {
    (dscp | ieee-802.1) classifier-name;
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
    import (rewrite-name | default);
  }
}
```

```

        forwarding-class class-name {
            loss-priority priority code-point (alias | bits);
        }
    }
}
scheduler-maps {
    map-name {
        forwarding-class class-name scheduler scheduler-name;
    }
}
schedulers {
    scheduler-name {
        buffer-size (percent percentage | remainder);
        drop-profile-map loss-priority (low | medium-high | high) protocol protocol
            drop-profile drop-profile-name;
        explicit-congestion-notification;
        priority priority;
        shaping-rate (rate | percent percentage);
        transmit-rate (percent percentage);
    }
}
shared-buffer {
    egress {
        percent percent;
        buffer-partition (lossless | lossy | multicast) {
            percent percent
        }
    }
    ingress {
        percent percent;
        buffer-partition (lossless | lossless-headroom | lossy) {
            percent percent
        }
    }
}
system-defaults {
    classifiers exp classifier-name;
}
traffic-control-profiles profile-name {
    guaranteed-rate(rate| percent percentage);
    scheduler-map map-name;
    shaping-rate (rate| percent percentage);
}
}

```

#### [\[edit ethernet-switching-options\] Hierarchy](#)

```

ethernet-switching-options {
    analyzer {
        name {
            input {
                egress {
                    interface (all | interface-name);
                }
                ingress {
                    interface (all | interface-name);
                }
            }
        }
    }
}

```

```

        vlan (vlan-id | vlan-name);
    }
    output {
        interface interface-name;
        ip-address ip-address;
        vlan (vlan-id | vlan-name);
    }
}
}
bpdv-block {
    interface (all | [interface-name]);
    disable-timeout timeout;
}
dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100)
}
interfaces interface-name {
    no-mac-learning;
}
mac-table-aging-time seconds {
}
port-error-disable {
    disable-timeout timeout;
}
secure-access-port {
    dhcp-snooping-file {
        location local_pathname | remote_URL;
        timeout seconds;
        write-interval seconds;
    }
    interface (all | interface-name) {
        allowed-mac {
            mac-address-list;
        }
        (dhcp-trusted | no-dhcp-trusted);
        fcoe-trusted;
        mac-limit limit action action;
        no-allowed-mac-log;
    }
}
vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection) [
        forwarding-class (for DHCP Snooping or DAI Packets) class-name;
    ]
}
dhcp-option82 {
    circuit-id {
        prefix (Circuit ID for Option 82) hostname;
        use-interface-description;
        use-vlan-id;
    }
    remote-id {
        prefix (Remote ID for Option 82) hostname | mac | none;
        use-interface-description;
        use-string string;
    }
    vendor-id <string>;
}
}

```



```

(examine-dhcp | no-examine-dhcp) {
    forwarding-class (for DHCP Snooping or DAI Packets) class-name;
}
examine-fip {
    examine-vn2vn {
        beacon-period milliseconds;
    }
    fc-map fc-map-value;
}
mac-move-limit limit <fabric-limit limit action action;
}
}
static {
    vlan vlan-id {
        mac mac-address next-hop interface-name;
    }
}
storm-control {
    interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-multicast;
        no-unknown-unicast;
    }
}
traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
    flag flag <disable>;
}
}

```

#### [edit fabric] Hierarchy

```

fabric
  aliases {
    director-device director-device-name {
        assigned-director-device-name;
    }
    interconnect-device interconnect-device-name {
        assigned-interconnect-device-name;
    }
    node-device node-device-name {
        assigned-node-device-name;
    }
  }
  protocols {
    fabric-control {
        graceful-restart {
            restart-timesecs seconds;
            stale-routes-time seconds;
        }
    }
  }
  resources {
    node-group node-group-name {

```

```

        node-device node-device-name;
        network-domain;
    }
}

```

### [edit fc-fabrics] Hierarchy

```

fc-fabrics {
  fc-fabric-name {
    description
    fabric-id fc-fabric-id;
    fabric-type proxy;
    interface {
      interface-name {
        max-login-sessions max-login-sessions;
      }
      interface-name {
        max-login-sessions max-login-sessions;
      }
      <...>;
      max-login-sessions max-login-sessions;
    }
    vlan.interface-name;
  }
  fc2 {
    traceoptions {
      file filename <replace> <size size> <files number> <no-stamp>;
      <world-readable | no-world-readable>;
      flag flag <flag-modifier>;
    }
  }
  max-login-sessions max-login-sessions;
  protocols {
    fip {
      fcoe-trusted;
      fc-map fc-map-value;
      fka-adv-period milliseconds;
      interface {
        interface-name {
          fka-adv-period milliseconds;
          priority priority;
        }
      }
      max-sessions-per-enode max-sessions-per-enode;
      priority priority;
      traceoptions {
        file filename <replace> <size size> <files number> <no-stamp>;
        <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
      }
    }
  }
  proxy {
    auto-load-rebalance
    load-balance-algorithm (simple | enode-based | flogi-based);
  }
}

```

```

no-fabric-wwn-verify;
traceoptions {
  file filename <replace> <size size> <files number> <no-stamp>;
  <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
}
}
}

```

#### [\[edit fc-options\] Hierarchy](#)

```

fc-options
max-login-sessions-per-node max-login-sessions-per-node;
no-fip-snooping-scaling;
traceoptions {
  file filename <replace> <size size> <files number> <no-stamp>;
  <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}

```

#### [\[edit firewall\] Hierarchy](#)

```

firewall {
  family family-name {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
}

policer policer-name {
  filter-specific;
  if-exceeding {
    bandwidth-limit bps;
    burst-size-limit bytes;
  }
  then {
    policer-action;
  }
}

three-color-policer policer-name {
  action {
    loss-priority high then discard;
  }
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
  }
}

```

```
    excess-burst-size bytes;  
  }  
  two-rate {  
    (color-aware | color-blind);  
    committed-information-rate bps;  
    committed-burst-size bytes;  
    peak-information-rate bps;  
    peak-burst-size bytes;  
  }  
}  
}
```

---

### [edit groups] Hierarchy

```
groups {  
  group-name {  
    configuration-data;  
  }  
  global {  
    configuration-data  
  }  
  if-config {  
    configuration-data  
  }  
  rel {  
    configuration-data  
  }  
}
```

---

### [edit interfaces] Hierarchy

```
interfaces {  
  aex {  
    disable;  
    aggregated-ether-options {  
      configured-flow-control {  
        rx-buffers (on | off);  
        tx-buffers (on | off);  
      }  
      (fcoe-lag | no-fcoe-lag);  
      (flow-control | no-flow-control);  
      lacp mode {  
        admin-key key;  
        force-up;  
        periodic interval;  
        system-id mac-address;  
      }  
      link-speed speed;  
      local-bias;  
      loopback;  
      no-loopback;  
      minimum-links number;  
    }  
    mc-ae {  
      chassis-id chassis-id;  
      mc-ae-id mc-ae-id;  
    }  
  }  
}
```

```

        mode (active-active);
        status-control (active | standby);
    }
    description text;
    gratuitous-arp-reply | no-gratuitous-arp-reply
    hold-time down milliseconds up milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no traps);
    unit logical-unit-number {
        disable;
        description text;
        family {
            ethernet-switching {
                filter input filter-name;
                filter output filter-name;
                native-vlan-id vlan-id;
                port-mode mode;
                reflective-relay;
                vlan {
                    members [ (all | names | vlan-ids) ];
                }
            }
            inet {
                address address {
                    primary;
                }
                filter input filter-name;
                filter output filter-name;
                primary;
                targeted-broadcast;
            }
            (traps | no traps);
            vlan-id vlan-id-number;
        }
        vlan-tagging;
    }
    interface-range interface-range-name {
        disable;
        description text;
        ether-options {
            802.3ad aex {
                lacp {
                    force-up;
                }
            }
        }
        (auto-negotiation | no-auto-negotiation);
        configured-flow-control {
            rx-buffers (on | off);
            tx-buffers (on | off);
        }
        (flow-control | no-flow-control);
        link-mode mode;
        speed (auto-negotiation | speed);
    }
}

```

```
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
    disable;
    description text;
    family family-name {...}
    (traps | no traps);
    vlan-id vlan-id-number;
}
}
lo0 {
    disable;
    description text;
    hold-time milliseconds down milliseconds;
    traceoptions;
    (traps | no traps);
    unit logical-unit-number {
        disable;
        description text;
        family {
            inet {
                address address {
                    primary;
                }
                filter input filter-name;
                filter output filter-name;
                primary;
                targeted-broadcast;
            }
        }
        (traps | no traps);
    }
}
mex {
    disable;
    description text;
    hold-time milliseconds down milliseconds;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    no-gratuitous-arp-request;
    traceoptions;
    traps;
    unit logical-unit-number {
        disable;
        description text;
        family {
            ethernet-switching {
                filter input filter-name;
                filter output filter-name;
                native-vlan-id vlan-id;
                port-mode mode;
                reflective-relay;
                vlan {
                    members [ (all | names | vlan-ids) ];
                }
            }
        }
    }
}
```

```

    inet {
        address address {
            primary;
            filter input filter-name;
            filter output filter-name;
            primary;
            targeted-broadcast;
        }
    }
    traps;
    vlan-id vlan-id-number;
}
vlan-tagging;
vlan {
    disable;
    description text;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no traps);
    unit logical-unit-number {
        description text;
        disable;
        family {
            inet {
                address address {
                    primary;
                }
                filter input filter-name;
                filter output filter-name;
                primary;
                targeted-broadcast;
            }
        }
        (traps | no traps);
    }
}
fc-0/0/port {
    fibrechannel-options {
        bb-sc-n;
        (loopback | no-loopback);
        speed (auto-negotiation | 2g | 4g | 8g);
    }
    unit logical-unit-number {
        disable;
        description text;
        family {
            fibre-channel {
                port-mode np-port;
            }
        }
        (traps | no traps);
    }
}
ge-0/0/port {
    disable;
    description text;

```

```
ether-options {
  802.3ad aex {
    lacp {
      force-up;
      primary;
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  loopback;
  no-loopback;
  speed (auto-negotiation | speed);
}
gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
  description text;
  disable;
  family {
    ethernet-switching {
      filter input filter-name;
      filter output filter-name;
      native-vlan-id vlan-id;
      port-mode mode;
      reflective-relay;
      vlan {
        members [ (all | names | vlan-ids) ];
      }
    }
  }
  inet {
    address address {
      primary;
    }
    filter input filter-name;
    filter output filter-name;
    primary;
    targeted-broadcast;
  }
  (traps | no traps);
  vlan-id vlan-id-number;
}
vlan-tagging;
}
vrrp-group group-id {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
```



```

authentication-type authentication;
fast-interval milliseconds;
(preempt | no-preempt) {
    hold-time seconds;
}
priority number;
track {
    interface interface-name {
        bandwidth-threshold bits-per-second priority-cost priority;
        priority-cost priority;
    }
    priority-hold-time seconds;
    route prefix/prefix-length routing-instance instance-name priority-cost priority;
}
}
virtual-address [ addresses ];
}
xe-0/0/port {
    disable;
    description text;
    ether-options {
        802.3ad aex {
            lacp {
                force-up;
                (primary | backup);
            }
        }
    }
    configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
    }
    (flow-control | no-flow-control);
    loopback;
    no-loopback;
}
(gratuitous-arp-reply | no-gratuitous-arp-reply)
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
    disable;
    description text;
    family {
        ethernet-switching {
            filter input filter-name;
            filter output filter-name;
            native-vlan-id vlan-id;
            port-mode mode;
            reflective-relay;
            vlan {
                members [ (all | names | vlan-ids) ];
            }
        }
    }
    fibre-channel {

```

```
        port-mode (f-port | np-port);
    }
    inet {
        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
    }
    (traps | no traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
```

---

#### [edit policy-options] Hierarchy

```
policy-options
  application-maps application-map-name {
    application application-name {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name match-type <actions>;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
    }
  }
```

---

#### [edit protocols] Hierarchy

```
protocols {
  bgp {
    disable;
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
```

```

authentication-key-chain key-chain;
bfd-liveness-detection {
  authentication {
    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
      meticulous-keyed-sha-1 | simple-password);
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  hold-down-interval milliseconds;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  session-mode (automatic | multihop | single-hop);
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family family-name {
  ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
}
graceful-restart {
  disable;
  restart-time seconds;
  stale-routes-time seconds;
}
group group-name {
  ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
}
hold-time seconds;
import [ policy-names ];
include-mp-next-hop;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> < alias> <private>;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
  no-nexthop-change;
  ttl ttl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {

```

```
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
```

```

    backup-peer-ip ip-address;
}
liveness-detection {
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (Liveness Detection) (1 | automatic);
}
local-ip-addr ipv4-address;
session-establishment-hold-time seconds;
}
session-establishment-hold-time seconds;
traceoptions {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable;
}
interface interface-name {
    group-limit limit;
    multicast-router-interface;
    static {
        group ip-address;
    }
}
robust-count number;
}
}
isis {
    disable;
    export [ policy-names ];
    ignore-attached-bit;
    interface interface-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
            }
        }
    }
}

```

```
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
checksum;
csnp-interval (seconds | disable);
disable;
hello-padding (adaptive | loose | strict);
level (1 | 2) {
    disable;
    hello-authentication-key key;
    hello-authentication-type authentication;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    metric metric;
    passive;
    priority number;
}
lsp-interval milliseconds;
mesh-group (value | blocked);
no-ipv4-multicast;
no-unicast-topology;
passive;
point-to-point;
}
level (1 | 2) {
    disable;
    authentication-key key;
    authentication-type authentication;
    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    prefix-export-limit number;
    wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
```

```

        advertise-high-metrics;
        timeout seconds;
    }
    reference-bandwidth reference-bandwidth;
    rib-group {
        inet group-name;
    }
    topologies {
        ipv4-multicast;
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    traffic-engineering {
        disable;
        family inet {
            shortcuts {
                multicast-rpf-routes:
            }
        }
    }
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
        flag flag (detail | disable | receive | send);
    }
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name (MSTP) name;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            block;
            alarm;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
}

```

```
max-age seconds;
max-hops hops;
msti msti-id {
  vlan (vlan-id | vlan-name);
  interface interface-name {
    disable;
    cost cost;
    edge;
    mode mode;
    priority priority;
  }
}
revision-level revision-level;
traceoptions {
  file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
  flag flag;
}
}
ospf {
  disable;
  area area-id {
    area-range ip-prefix </prefix-length > <exact> <override-metric metric > <restrict>;
    context-identifier identifier
    interface interface-name {
      disable;
      authentication {
        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
      }
      bandwidth-based-metrics {
        bandwidth value metric number;
      }
      bfd-liveness-detection {
        authentication {
          algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
          key-chain key-chain-name;
          loose-check;
        }
        detection-time {
          threshold milliseconds;
        }
      }
      full-neighbors-only;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  dead-interval seconds;
  dynamic-neighbors;
```



```

flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix </prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication {
        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
    }
    dead-interval seconds;
    demand-circuit;
    flood-reduction;
    hello-interval seconds;
    ipsec-sa sa-name;
    no-neighbor-down-notification;
    retransmit-interval seconds;
    topology (name | default | ipv4-multicast) {
        disable;
        metric metric;
    }
}

```

```
    }
    transit-delay seconds;
  }
}
database-protection {
  ignore-count number;
  ignore-time seconds;
  maximum-lsa number;
  reset-time seconds;
  warning-only;
  warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
  disable;
  helper-disable <both | restart-signaling | standard>;
  no-strict-lsa-checking;
  notify-duration seconds;
  restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
  overload;
  prefix-export-limit number;
  topology-id number;
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
traffic-engineering {
  advertise-unnumbered-interfaces;
  credibility-protocol-preference;
  ignore-lsp-metrics;
  multicast-rpf-routes;
  no-topology;
  shortcuts <lsp-metric-into-summary>;
}
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
```

```

family (inet | inet6) {
    disable;
}
graceful-restart {
    disable;
    restart-duration seconds;
}
import [ policy-names ];
interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
        disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
}
join-load-balance;
join-prune-timeout;
nonstop-routing;
override-interval milliseconds;
propagation-delay milliseconds;
reset-tracking-bit;
rib-group group-name;
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-import [ policy-names ];
    bootstrap-export [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
    }
    local {
        family (inet | inet6) {
            address address;
            anycast-pim {

```

```
        disable;
        rp-set {
            address address <forward-msdp-sa>;
        }
        local-address address;
    }
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    hold-time seconds;
    priority number;
}
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix </prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
    tunnel-devices [ mt-fpc/pic/port ];
}
rip {
    authentication-key password;
    authentication-type type;
    (check-zero | no-check-zero);
    group group-name {
        bfd-liveness-detection {
            authentication {
```

```

        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
            meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
export [ policy-names ];
import [ policy-names ];
metric-out metric;
neighbor neighbor-name {
    any-sender;
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
        ... same statements as at the [edit protocols rip group group-name
            bfd-liveness-detection] hierarchy level ...
    }
    (check-zero | no-check-zero);
    import [ policy-names ];
    message-size number;
    metric-in metric;
    receive (both | none | version-1 | version-2);
    route-timeout seconds;
    send (broadcast | multicast | none | version-1);
    update-interval seconds;
}
preference preference;
route-timeout seconds;
update-interval seconds;
}
holddown seconds;
import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;

```

```
}
rstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
}
traceoptions {
  file filename <files number > <size size > <no-stamp > <world-readable |
    no-world-readable>;
  flag flag;
}
}
stp {
  disable;
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
}
traceoptions {
  file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
  flag flag;
}
uplink-failure-detection {
  group group-name {
    link-to-monitor interface-name;
    link-to-disable interface-name;
  }
}
```

```

}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number > <size size > <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
}
}

```

---

#### [edit security] Hierarchy

```

security {
  certificates
  pki
  ssh-known-hosts
  traceoptions
}

```

---

#### [edit snmp] Hierarchy

```

snmp {
  client-list client-list-name {
    ip-addresses;
  }
  community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
      address restrict;
    }
    logical-system logical-system-name {
      routing-instance routing-instance-name {
        clients {
          addresses;
        }
      }
    }
  }
}

```

```

    }
  }
}
routing-instance routing-instance-name {
  clients {
    addresses;
  }
}
view view-name;
}
contact contact;
description description;
filter-duplicates;
filter-interfaces;
health-monitor {
  falling-threshold integer;
  interval seconds;
  rising-threshold integer;
}
interface [ interface-names ];
location location;
name name;
nonvolatile {
  commit-delay seconds;
}
rmon {
  alarm index {
    description description;
    falling-event-index index;
    falling-threshold integer;
    falling-threshold-interval seconds;
    interval seconds;
    request-type;
    rising-event-index index;
    rising-threshold integer;
    sample-type (absolute-value | delta-value);
    startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
    syslog-subtag syslog-subtag;
    variable oid-variable;
  }
  event index {
    community community-name;
    description description;
    type type;
  }
  history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regular-expression>;
  flag flag;

```



```

}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance routing-instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
  snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance routing-instance-name;
    tag-list tag-list;
    target-parameters target-parameters-name;
    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
}
usm {
  local-engine {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {

```

```
        authentication-password authentication-password;
    }
    authentication-none;
    privacy-aes128 {
        privacy-password privacy-password;
    }
    privacy-des {
        privacy-password privacy-password;
    }
    privacy-3des {
        privacy-password privacy-password;
    }
    privacy-none;
}
}
remote-engine engine-id {
    user username {
        authentication-sha {
            authentication-password authentication-password;
        }
        authentication-md5 {
            authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
            privacy-password privacy-password;
        }
        privacy-des {
            privacy-password privacy-password;
        }
        privacy-3des {
            privacy-password privacy-password;
        }
        privacy-none {
            privacy-password privacy-password;
        }
    }
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
```

```

        group group-name;
    }
}
}
}
view view-name {
    oid object-identifier (include | exclude);
}
}

```

### [edit system] Hierarchy

```

system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
}

```

```
default-address-selection;
domain-name domain-name;
domain-search [ domain-list ];
host-name hostname;
internet-options {
    icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
    source-port upper-limit <upper-limit>;
}
location {
    altitude feet;
    building name;
    country-code code;
    floor number;
    hcoord horizontal-coordinate;
    lata service-area;
    latitude degrees;
    longitude degrees;
    npa-nxx number;
    postal-code postal-code;
    rack number;
    vcoord vertical-coordinate;
}
login {
    announcement text;
    class class-name {
        access-end;
        access-start;
        allow-configuration "regular-expression";
        allowed-days "regular-expression";
        deny-commands "regular-expression";
        deny-configuration "regular-expression";
        idle-timeout minutes;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    password {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    retry-options {
        backoff-factor seconds;
        backoff-threshold number;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        authentication {
            (encrypted-password "password" | plain-text-password);
            load-key-file URL;
            remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
            ssh-rsa "public-key";
            ssh-dsa "public-key";
        }
    }
}
```

```

    }
    uid uid-value;
    class class-name;
    full-name complete-name;
  }
}
name-server {
  address;
}
no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
  authentication-key number type type value password;
  serveraddress <key key-number> <version value> <prefer>;
}
ports {
  auxiliary {
    disable;
    insecure;
    type terminal-type;
  }
  console {
    disable;
    insecure;
    log-out-on-disconnect;
    type terminal-type;
  }
}
radius-server server-address {
  accounting-port port-number;
  port number;
  retry number;
  secret password;
  source-address source-address;
  timeout seconds;
}
radius-options {
  password-protocol mschap-v2;
}
attributes {
  nas-ip-address ip-address;
}
root-authentication {
  (encrypted-password "password" | plain-text-password);
  ssh-rsa "public-key";
  ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
  finger {
    connection-limit limit;
    rate-limit limit;
  }
}

```

```
flow-tap-dtcp {
  ssh {
    connection-limit limit;
    rate-limit limit;
  }
}
ftp {
  connection-limit limit;
  rate-limit limit;
}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
```

```

    }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
  console {
    facility severity;
  }
  file filename {
    archive {
      files number;
      size maximum-file-size;
      start-time "YYYY-MM-DD.hh:mm";
      transfer-interval minutes;
      (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
      brief;
    }
  }
}
host (hostname | other-routing-engine | scc-master) {
  explicit-priority;
  facility-override facility;
  facility severity;
  log-prefix string;
  match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
  facility severity;
  match "regular-expression";
}
}
tacplus-options {
  service-name service-name;
  (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
  port
  secret password;
  single-connection;
  source-address source-address;
  timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {

```

```
    destination-override {  
      syslog host;  
    }  
  }  
  use-imported-time-zones;  
}
```

#### [edit vlans] Hierarchy

---

```
vlans {  
  vlan-name {  
    description text-description;  
    dot1q-tunneling {  
      customer-vlans (id | range);  
    }  
    filter input filter-name;  
    filter output filter-name;  
    interface interface-name {  
      isolated;  
      mapping (policy | tag push | native push);  
      promiscuous;  
    }  
    isolation-vlan-id;  
    l3-interface vlan.logical-interface-number;  
    mac-limit number;  
    mac-table-aging-time seconds;  
    no-local-switching;  
    no-mac-learning;  
    primary-vlan vlan-name;  
    pvlan extend-secondary-vlan-id vlan-id;  
    vlan-id number;  
    vlan-range vlan-id-low-vlan-id-high;  
  }  
}
```

## access-end

---

|                          |                                                                                                                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------|
| Syntax                   | access-end <i>HH:MM</i> ;                                                                                       |
| Hierarchy Level          | [edit system <a href="#">login</a> class]                                                                       |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                               |
| Description              | Configure the end time for login access.                                                                        |
| Required Privilege Level | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration. |
| Related Documentation    | <ul style="list-style-type: none"><li><a href="#">Configuring Time-Based User Access on page 173</a></li></ul>  |



---

## access-start

---

|                                 |                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>access-start <i>HH:MM</i>;</code>                                                                          |
| <b>Hierarchy Level</b>          | [edit system <a href="#">login</a> class]                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                |
| <b>Description</b>              | Configure the start time for login access.                                                                       |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Time-Based User Access on page 173</a></li></ul> |

## accounting

---

```
Syntax  accounting {  
        destination {  
            radius {  
                server {  
                    server-address {  
                        accounting-port port-number;  
                        secret password;  
                        source-address address;  
                        retry number;  
                        timeout seconds;  
                    }  
                }  
            }  
        }  
        tacplus {  
            server {  
                server-address {  
                    port port-number;  
                    secret password;  
                    single-connection;  
                    timeout seconds;  
                }  
            }  
        }  
    }
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure audit of TACACS+ or RADIUS authentication events, configuration changes, and interactive commands.

**Options** The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring RADIUS Accounting*
- [Configuring TACACS+ System Accounting on page 1366](#)

## accounting-port

---

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>accounting-port <i>port-number</i>;</code>                                                                                 |
| <b>Hierarchy Level</b>          | [edit system accounting destination radius server <i>server-address</i> ],<br>[edit system radius server <i>server-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                |
| <b>Description</b>              | Configure the accounting port number on which to contact the RADIUS server.                                                      |
| <b>Options</b>                  | <i>number</i> —Port number on which to contact the RADIUS server.<br><b>Default:</b> 1813                                        |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS Accounting</a></li> </ul>                                |

## allow-commands

---

|                                 |                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>allow-commands "<i>regular-expression</i>";</code>                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit system login class <i>class-name</i> ]                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                    |
| <b>Description</b>              | Specify the operational mode commands that members of a login class can use.                                                                                                                                                                         |
| <b>Default</b>                  | If you omit this statement and the <b>deny-commands</b> statement, users can issue only those commands for which they have access privileges through the <b>permissions</b> statement.                                                               |
| <b>Options</b>                  | <i>regular-expression</i> —Extended (modern) regular expression as defined in POSIX 1003.2.<br>If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Specifying Access Privileges for Junos OS Operational Mode Commands on page 192</a></li> <li>• <a href="#">deny-commands on page 268</a></li> <li>• <a href="#">user on page 332</a></li> </ul> |

## allow-configuration

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>allow-configuration "regular-expression";</code>                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system login class <i>class-name</i> ]                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Explicitly allow configuration access to the specified levels in the hierarchy even if the permissions set with the <b>permissions</b> statement do not grant such access by default.                                                                                                                                                                                                   |
| <b>Default</b>                  | If you omit this statement and the <b>deny-configuration</b> statement, users can edit only those commands for which they have access privileges through the <b>permissions</b> statement.                                                                                                                                                                                              |
| <b>Options</b>                  | <b>regular-expression</b> —Extended (modern) regular expression as defined in POSIX 1003.2.<br>If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.                                                                                                                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1371</a></li><li>• <a href="#">Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1338</a></li><li>• <a href="#">deny-configuration on page 269</a></li><li>• <a href="#">user on page 332</a></li></ul> |

## allowed-days

---

|                                 |                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>allowed-days [ <i>days-of-the-week</i> ];</code>                                                           |
| <b>Hierarchy Level</b>          | [edit system <b>login</b> class <i>class-name</i> ]                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                |
| <b>Description</b>              | Specify the days of the week when users can log in.                                                              |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Time-Based User Access on page 173</a></li></ul> |



## allow-transients

|                                 |                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | allow-transients;                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit systems scripts commit]                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>              | For Junos OS commit scripts, enable transient configuration changes to be committed.                                                                                                                                                                   |
| <b>Default</b>                  | Transient changes are disabled by default. If you do not include the <b>allow-transients</b> statement, and an enabled script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails. |
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Generating a Persistent or Transient Change</i></li> <li>• <i>Creating a Macro to Read the Custom Syntax and Generate Related Configuration Statements</i></li> </ul>                                      |

## announcement

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | announcement <i>text</i> ;                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system login]                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | Configure a system login announcement. This announcement appears after a user logs in.                                                                                                                                                                                      |
| <b>Options</b>                  | <i>text</i> —Text of the announcement. If the text contains any spaces, enclose it in quotation marks.                                                                                                                                                                      |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Junos OS to Display a System Login Announcement</i></li> <li>• <a href="#">Configuring the Junos OS to Display a System Login Message on page 163</a></li> <li>• <a href="#">message on page 283</a></li> </ul> |

## archival

|                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                        | <pre> archival {   configuration {     archive-sites {       file://&lt;path&gt;/&lt;filename&gt;;       ftp://username@host:&lt;port&gt;url-path password password;       http://username@host:&lt;port&gt;url-path password password;       pasvftp://username@host:&lt;port&gt;url-path password password;       scp://username@host:&lt;port&gt;url-path password password;     }     transfer-interval interval;     transfer-on-commit;   } } </pre> |
| <b>Hierarchy Level</b>                                                                                                                                                                               | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                                                                                                                                                                           | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                              |
| <b>Description</b>                                                                                                                                                                                   | Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.                                                                                                                                                                                                                                                                                                                      |
| <div>  <b>NOTE:</b> The <code>edit system archival</code> hierarchy is not available on QFabric systems. </div>   |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                                                                                                                                                                                       | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                         |
| <div>  <b>NOTE:</b> The <code>[edit system archival]</code> hierarchy is not available on QFabric systems. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                      | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                         | <ul style="list-style-type: none"> <li>Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1263</li> </ul>                                                                                                                                                                                                                                                                                             |

## arp (System)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>arp {     aging-timer <i>minutes</i>;     gratuitous-arp-delay <i>seconds</i>;     gratuitous-arp-on-ifup;     interfaces {         <i>interface-name</i> {             aging-timer <i>minutes</i>;         }     }     passive-learning;     purging; }</pre> <p>For EX-Series switches:</p> <pre>arp {     aging-timer <i>minutes</i>; }</pre>                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Specify ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.</p> <p>For EX-Series switches, set only the time interval between ARP updates.</p>                                                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>aging-timer</b>—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.</p> <p><b>passive-learning</b> (QFX-Series only)—Configure switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests.</p> <p><b>Default:</b> 20 minutes</p> <p><b>Range:</b> 1 to 240 minutes</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                   |

- For more information about ARP updates, see the [Junos OS System Basics Configuration Guide](#).

## authentication (Login)

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>authentication {<br/>  encrypted-password <i>password</i>;<br/>  load-key-file <i>URL</i>;<br/>  plain-text-password <i>password</i>;<br/>  remote-debug-permission (qfabric-admin   qfabric-operator   qfabric-user);<br/>  ssh-dsa "<i>public-key</i>";<br/>  ssh-rsa "<i>public-key</i>";<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Hierarchy Level          | [edit system login user <i>username</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Description              | Authentication methods that a user can use to log in to the switch. You can assign multiple authentication methods to a single user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Options                  | <p><b>encrypted-password "<i>password</i>"</b>—Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.</p> <p>You cannot configure a blank password for <b>encrypted-password</b> using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p> <p><b>load-key-file</b>—Load RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys from a file. The file is a URL containing one or more SSH keys.</p> <p><b>plain-text-password</b>—Plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it.</p> <p><b>remote-debug-permission</b> (QFabric systems only)—QFabric component authentication. Specifies permission levels for users to access individual components in a QFabric system.</p> <p><b>ssh-dsa "<i>public-key</i>"</b>—SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p> <p><b>ssh-rsa "<i>public-key</i>"</b>—SSH version 1 and SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.</p> |
| Required Privilege Level | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Junos OS User Accounts on page 1344</a></li><li>• <a href="#">root-authentication on page 303</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |



## authentication-key

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-key <i>key-number</i> type <i>type</i> value <i>password</i>;</code>                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit system ntp]                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure Network Time Protocol (NTP) authentication keys so that the router or switch can send authenticated packets. If you configure the router or switch to operate in authenticated mode, you must configure a key.</p> <p>Both the keys and the authentication scheme (MD5) must be identical between a set of peers sharing the same key number.</p> |
| <b>Options</b>                  | <p><b><i>key-number</i></b>—An integer in the range of 1 to 65533.</p> <p><b><i>type type</i></b>—Authentication type. It can only be <b>md5</b>.</p> <p><b><i>value password</i></b>—Key itself, consisting of 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.</p>                                                   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring NTP Authentication Keys (QFabric System)</i></li> <li>• <i>NTP Time Server and Time Services Overview (QFabric System)</i></li> </ul>                                                                                                                                                                  |

## authentication-order

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>authentication-order [ <i>authentication-methods</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.                                                                                                                                                                                                                                            |
| <b>Default</b>                  | If you do not include the <b>authentication-order</b> statement, users are verified based on their configured passwords.                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b><i>authentication-methods</i></b>—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:</p> <ul style="list-style-type: none"><li>• <b>password</b>—Use the password configured for the user with the <b>authentication</b> statement at the [edit system login user] hierarchy level.</li><li>• <b>radius</b>—Use RADIUS authentication services.</li><li>• <b>tacplus</b>—Use TACACS+ authentication services.</li></ul> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1346</a></li><li>• <a href="#">authentication on page 254</a></li></ul>                                                                                                                                                                                                                                                                          |

## auxiliary

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> auxiliary {   disable;   insecure;   type <i>terminal-type</i>; }</pre>                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system ports]                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure the characteristics of the auxiliary port.                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | The auxiliary port is disabled.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>disable</b>—Disable the port.</p> <p><b>insecure</b>—Disable superuser access or root logins to establish a terminal connection.</p> <p><b>type <i>terminal-type</i></b>—Type of terminal that is connected to the port.</p> <p><b>Range:</b> ansi, vt100, small-xterm, xterm</p> <p><b>Default:</b> The terminal type is unknown, and the user is prompted for the terminal type.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Console and Auxiliary Port Properties on page 6583</a></li> </ul>                                                                                                                                                                                                                                                           |

## boot-server (NTP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>boot-server (address   hostname);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit system <a href="#">ntp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure the server that NTP queries when the router or switch boots to determine the local date and time.</p> <p>When you boot the router or switch, it issues an <b>ntpdate</b> request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. Otherwise, NTP cannot synchronize to a time server if the server time significantly differs from the local router's or switch's time. You can configure either an IP address or a hostname for the boot server. If you configure a hostname instead of an IP address, the <b>ntpdate</b> request resolves the hostname to an IP address when the router or switch boots up.</p> |
| <b>Options</b>                  | <ul style="list-style-type: none"><li>• <b>address</b>—IP address of an NTP boot server.</li><li>• <b>hostname</b>—Hostname of an NTP boot server.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Synchronizing and Coordinating Time Distribution Using NTP on page 194</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## broadcast

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>broadcast address &lt;key key-number&gt; &lt;version value&gt; &lt;tll value&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit system <a href="#">ntp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the local router or switch to operate in broadcast mode with the remote system at the specified address to send periodic broadcast messages to a client population. Normally, you include this statement only when the local router or switch is operating as a transmitter.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>address</b>—Broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be <b>224.0.1.1</b>.</p> <p><b>key key-number</b>—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number (any unsigned 32-bit integer).</p> <p><b>tll value</b>—(Optional) Time-to-live (TTL) value to use.<br/> <b>Range:</b> 1 through 255<br/> <b>Default:</b> 1</p> <p><b>version value</b>—(Optional) Specify the version number to be used in outgoing NTP packets.<br/> <b>Range:</b> 1 through 4<br/> <b>Default:</b> 4</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the NTP Time Server and Time Services on page 166</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## broadcast-client

---

|                                 |                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>broadcast-client;</code>                                                                                                                            |
| <b>Hierarchy Level</b>          | <code>[edit system ntp]</code>                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                         |
| <b>Description</b>              | Configure the local router or switch to listen for broadcast messages on the local network to discover other servers on the same subnet.                  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Router or Switch to Listen for Broadcast Messages Using NTP on page 171</a></li></ul> |

## change-type

---

|                                 |                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>change-type (character-sets   set-transitions);</code>                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit system login password]</code>                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Set requirements for using character sets in plain-text passwords. When you combine this statement with the <b>minimum-changes</b> statement, you can check for the total number of character sets included in the password or for the total number of character-set changes in the password. Newly created passwords must meet these requirements.   |
| <b>Options</b>                  | Specify one of the following: <ul style="list-style-type: none"><li>• <b>character-sets</b>—Number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters.</li><li>• <b>set-transitions</b>—Number of transitions between character sets.</li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li><li>• <a href="#">minimum-changes on page 283</a></li></ul>                                                                                                                                                           |

## checksum

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>checksum (md5   sha-256   sha1) <i>hash</i>;</code>                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit event-options event-script file <i>filename</i> ],<br>[edit system scripts commit file <i>filename</i> ],                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | For Junos commit scripts and op scripts, specify the MD5, SHA-1, or SHA-256 checksum hash. When it executes a local event or commit script, the Junos OS verifies the authenticity of the script by using the configured checksum hash.                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>md5 <i>hash</i></b>—MD5 checksum of this script.</p> <p><b>sha-256 <i>hash</i></b>—SHA-256 checksum of this script.</p> <p><b>sha1 <i>hash</i></b>—SHA-1 checksum of this script.</p>                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i></li> <li>• <a href="#">file checksum md5 on page 363</a></li> <li>• <a href="#">file checksum sha-256 on page 365</a></li> <li>• <a href="#">file checksum sha1 on page 364</a></li> </ul> |

## class (Defining Login Classes)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>class <i>class-name</i> {<br/>    access-end;<br/>    access-start;<br/>    allow-commands "<i>regular-expression</i>";<br/>    allow-configuration "<i>regular-expression</i>";<br/>    deny-commands "<i>regular-expression</i>";<br/>    deny-configuration "<i>regular-expression</i>";<br/>    idle-timeout <i>minutes</i>;<br/>    login-tip;<br/>    permissions [ <i>permissions</i> ];<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system login]                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Define a login class.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><i>class-name</i>—A name you choose for the login class.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining Junos OS Login Classes on page 1368</a></li><li>• <a href="#">user on page 332</a></li></ul>                                                                                                                                                                                                                                                        |



---

## class (Assigning a Class to an Individual User)

---

|                                 |                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>class <i>class-name</i> {<br/>    operator;<br/>    read-only;<br/>    super-user;<br/>    unauthorized;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system login user <i>username</i> ]                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                         |
| <b>Description</b>              | Configure a user's login class. You must configure one class for each user.                                               |
| <b>Options</b>                  | <i>class-name</i> —One of the classes defined at the [edit system login class] hierarchy level.                           |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Junos OS User Accounts on page 1344</a></li></ul>         |

## commit

---

**Syntax**    `commit {  
          allow-transients;  
          direct-access;  
          file filename {  
              checksum (md5 | sha-256 | sha1) hash;  
              optional;  
              refresh;  
              refresh-from url;  
              sourceurl;  
          }  
          refresh;  
          refresh-from url;  
          traceoptions {  
              file <filename> <files number> <size size> <world-readable | no-world-readable>;  
              flag flag;  
              no-remote-trace;  
          }  
          }`

**Hierarchy Level**    [edit system scripts]


**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    For Junos OS commit scripts, configure the commit-time scripting mechanism.

**Options**    The statements are explained separately.

**Required Privilege Level**    maintenance—To view this statement in the configuration.  
                                  maintenance-control—To add this statement to the configuration.

## compress-configuration-files (System)

|                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                  | (compress-configuration-files   no-compress-configuration-files);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                         | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                                                                                                                                                                                                                     | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                                                                                                                                                                                                                             | Compress the current operational configuration file. The file is stored in the file <b>juniper.conf</b> , in the <b>/config</b> file system, along with the last three committed versions of the configuration. However, with large networks, the current configuration file might exceed the available space in the <b>/config</b> file system. Compressing the current configuration file allows the file to fit in the file system, typically reducing the size of the file by 90 percent. The current configuration file is compressed on the second commit of the configuration after the first commit is made to include the <b>compress-configuration-files</b> statement. |
| <div>  <p><b>NOTE:</b> We recommend that you enable compression of the configuration files to minimize the amount of disk space that they require.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Default</b>                                                                                                                                                                                                                                 | The current operational configuration file is uncompressed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>                                                                                                                                                                                                                   | <ul style="list-style-type: none"> <li>• <a href="#">Compressing the Current Configuration File on page 1247</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## console (Physical Port)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>console {<br/>  disable;<br/>  insecure;<br/>  log-out-on-disconnect;<br/>  type <i>terminal-type</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit system ports]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Configure the characteristics of the console port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>                  | The console port is enabled and its speed is 9600 baud.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b>disable</b>—Disable console login connections.</p> <p><b>insecure</b>—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode. This option can be used to prevent a user from attempting password recovery by booting into single-user mode, if the user does not know the root password.</p> <p><b>log-out-on-disconnect</b>—Log out the session when the data carrier on the console port is lost.</p> <p><b>type <i>terminal-type</i></b>—Type of terminal that is connected to the port: <b>ansi</b>, <b>vt100</b>, <b>small-xterm</b>, or <b>xterm</b>.</p> |
| <b>Required Privilege Level</b> | <p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Console and Auxiliary Port Properties on page 6583</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## default-address-selection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | default-address-selection;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Use the loopback interface, <b>lo0</b>, as the source address for all locally generated IP packets when the packet is sent through a routed interface, but not when the packet is sent through a local interface such as <b>fxp0</b>. The <b>lo0</b> interface is the interface to the switch's Routing Engine.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Default</b>                  | <p>The default address is used as the source address for all locally generated IP packets on outgoing interfaces that are unnumbered. If an outgoing interface is numbered, the default address is chosen using the following sequence:</p> <ul style="list-style-type: none"> <li>• The primary address on the loopback interface <b>lo0</b> that is <i>not</i> <b>127.0.0.1</b> is used.</li> <li>• The primary address for the primary interface or the preferred address (if configured) for the primary interface is used.</li> </ul> <p>By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.</p> <p>An interface's <i>primary address</i> is used by default as the local address for broadcast and multicast packets sourced locally and sent out through the interface. An interface's <i>preferred address</i> is the default local address used for packets sourced by the local switch to destinations on the subnet. By default, the numerically lowest local address configured for the interface is chosen as the preferred address on the subnet.</p> <p>To configure a different primary address or preferred address, include the <b>primary</b> or <b>preferred</b> statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>] or [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>] hierarchy levels.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets on page 165</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## deny-commands

---

|                                 |                                                                                                                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>deny-commands "regular-expression";</code>                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system login class]                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                 |
| <b>Description</b>              | Specify the operational mode commands that the user is denied permission to issue even though the permissions set with the <b>permissions</b> statement would allow their use.                                                                    |
| <b>Default</b>                  | If you omit this statement and the <b>allow-commands</b> statement, users can issue only those commands for which they have access privileges through the <b>permissions</b> statement.                                                           |
| <b>Options</b>                  | <b>regular-expression</b> —Extended (modern) regular expression as defined in POSIX 1003.2.<br>If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.                                   |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Specifying Access Privileges for Junos OS Operational Mode Commands on page 192</a></li><li>• <a href="#">allow-commands on page 249</a></li><li>• <a href="#">user on page 332</a></li></ul> |

---

## deny-configuration

---

|                                 |                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>deny-configuration "regular-expression";</code>                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system login class]                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                    |
| <b>Description</b>              | Explicitly deny configuration access to the specified levels in the hierarchy even if the permissions set with the <b>permissions</b> statement grant such access by default.                                                        |
| <b>Default</b>                  | If you omit this statement and the <b>allow-configuration</b> statement, users can edit those levels in the configuration hierarchy for which they have access privileges through the <b>permissions</b> statement.                  |
| <b>Options</b>                  | <b>regular-expression</b> —Extended (modern) regular expression as defined in POSIX 1003.2.<br>If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.                      |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Specifying Access Privileges Using allow/deny-configuration Statements</i></li><li>• <a href="#">allow-configuration on page 250</a></li><li>• <a href="#">user on page 332</a></li></ul> |

## destination (Accounting)

---

**Syntax**

```
destination {  
  radius {  
    server {  
      server-address {  
        accounting-port port-number;  
        secret password;  
        source-address address;  
        retry number;  
        timeout seconds;  
      }  
    }  
  }  
  tacplus {  
    server {  
      server-address {  
        port port-number;  
        secret password;  
        single-connection;  
        timeout seconds;  
      }  
    }  
  }  
}
```

**Hierarchy Level** [edit system accounting]

**Release Information** Statement introduced before Junos OS Release 7.4.  
**radius** statement added in Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the authentication server.

**Options** The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring RADIUS System Accounting on page 1349](#)
- [Configuring TACACS+ System Accounting](#)



## destination-override

---

|                                 |                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>destination-override {<br/>  syslog host <i>ip-address</i>;<br/>}</code>                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit system tracing]                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                          |
| <b>Description</b>              | Override the system-wide configuration of the switch at the <b>[edit system tracing]</b> hierarchy level. This statement has no effect if system tracing is not configured.                                                                                                                                |
| <b>Options</b>                  | <p><b>syslog</b>—System process log files to send to the remote tracing host.</p> <ul style="list-style-type: none"> <li>• <b>syslog</b>—System process log files to send to the remote tracing host.</li> <li>• <b>host <i>ip-address</i></b>—IP address to which to send tracing information.</li> </ul> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Tracing and Logging Operations on page 6468</a></li> <li>• <a href="#">tracing on page 330</a></li> </ul>                                                                                                                               |

## direct-access

---

|                                 |                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>direct-access;</code>                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit system scripts commit]                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                           |
| <b>Description</b>              | Specify that commit scripts read input configurations directly from the database when inspecting these scripts for errors.                                                                                                                                                  |
| <b>Required Privilege Level</b> | <p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <code>commit</code></li> <li>• <a href="#">scripts on page 6663</a></li> <li>• <a href="#">How Commit Scripts Work on page 6474</a></li> <li>• <a href="#">Controlling the Execution of Commit Scripts on page 6588</a></li> </ul> |

## domain-name

---

|                                 |                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>domain-name <i>domain-name</i>;</code>                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                            |
| <b>Description</b>              | Configure the name of the domain in which the switch is located. This is the default domain name that is appended to hostnames that are not fully qualified. |
| <b>Options</b>                  | <i>domain-name</i> —Name of the domain.                                                                                                                      |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Reaching a Domain Name System Server on page 158</a></li></ul>                                           |

## domain-search

---

|                                 |                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>domain-search <i>domain-list</i>;</code>                                                                                       |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                    |
| <b>Description</b>              | Configure a list of domains to be searched.                                                                                          |
| <b>Options</b>                  | <i>domain-list</i> —List of domain names to search. The list can contain up to 6 domain names, with a total of up to 256 characters. |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Reaching a Domain Name System Server on page 158</a></li></ul>                   |

## explicit-priority

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>explicit-priority {   archive &lt;files <i>number</i>&gt; &lt;size <i>size</i> &lt;start-time<i>time</i>&gt; &lt;transfer-interval     <i>interval</i>&gt;&lt;world-readable   no-world-readable&gt;;   archive-sites {     file://&lt;path&gt;/&lt;filename&gt;;     ftp://username@host:&lt;port&gt;url-path password password;     scp://username@host:&lt;port&gt;url-path password password;   }   structured-data {     brief;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit system syslog file <i>filename</i> ],<br>[edit system syslog host]                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination.</p> <p>When the <b>structured-data</b> statement is also included at the <b>[edit system syslog file <i>filename</i>]</b> hierarchy level, this statement is ignored for the file.</p>                                                                                                                      |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Including Priority Information in System Log Messages on page 6622</a></li> </ul>                                                                                                                                                                                                                                                                                                                |

## events

---

|                                 |                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>events [ <i>events</i> ];</code>                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit system accounting]                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the types of events to track and log.                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b><i>events</i></b> —Event types; can be one or more of the following: <ul style="list-style-type: none"><li>• <b>change-log</b>—Audit configuration changes.</li><li>• <b>interactive-commands</b>—Audit interactive commands (any command-line input).</li><li>• <b>login</b>—Audit logins.</li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring TACACS+ System Accounting on page 1366</a></li></ul>                                                                                                                                                                                     |

## format

---

|                                 |                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>format (des   md5   sha1);</code>                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit system login password]                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure the authentication algorithm for plain-text passwords.                                                                                                                                                                                                                                            |
| <b>Default</b>                  | For Junos OS, the default encryption format is <b>md5</b> . For Junos OS-FIPS software, the default encryption format is <b>sha1</b> .                                                                                                                                                                      |
| <b>Options</b>                  | The hash algorithm that authenticates the password can be one of three algorithms: <ul style="list-style-type: none"><li>• <b>des</b>—Has a block size of 8 bytes; its key size is 48 bits long.</li><li>• <b>md5</b>—Produces a 128-bit digest.</li><li>• <b>sha1</b>—Produces a 160-bit digest.</li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li></ul>                                                                                                                                                                       |

## host-name

---

|                                 |                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>host-name <i>hostname</i>;</code>                                                                                          |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                |
| <b>Description</b>              | Set the hostname of the switch.                                                                                                  |
| <b>Options</b>                  | <i>hostname</i> —Name of the switch.                                                                                             |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Hostname of the Router or Switch on page 160</a></li> </ul> |

## icmpv4-rate-limit

---

|                                 |                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>icmpv4-rate-limit {     bucket-size <i>seconds</i>;     packet-rate <i>pps</i>; }</pre>                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit system internet-options]                                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure rate-limiting parameters for ICMPv4 messages sent.                                                                                                                                                                                                                                                       |
| <b>Options</b>                  | <p><b>bucket-size <i>seconds</i></b>—Number of seconds in the rate-limiting bucket.<br/> <b>Range:</b> 0 through 4294967295 seconds<br/> <b>Default:</b> 5</p> <p><b>packet-rate <i>pps</i></b>—Rate-limiting packets earned per second.<br/> <b>Range:</b> 0 through 4294967295 pps<br/> <b>Default:</b> 1000</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>ping</i></li> <li>• <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i></li> <li>• <i>Configuring Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages</i></li> </ul>                                                            |

## idle-timeout

---

|                                 |                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>idle-timeout <i>minutes</i>;</code>                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | <code>[edit system login class <i>class-name</i>]</code>                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                   |
| <b>Description</b>              | For a login class, configure the maximum time that a session can be idle before the user is logged off the switch. The session times out after remaining at the CLI operational mode prompt for the specified time. |
| <b>Default</b>                  | If you omit this statement, a user is never forced off the system after extended idle times.                                                                                                                        |
| <b>Options</b>                  | <i>minutes</i> —Maximum idle time.<br><b>Range:</b> 0 through 4294967295 minutes                                                                                                                                    |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Timeout Value for Idle Login Sessions on page 174</a></li></ul>                                                                                 |

## internet-options

---

|                                 |                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>internet-options {<br/>    <i>icmpv4-rate-limit</i> bucket-size <i>bucket-size</i> packet-rate <i>packet-rate</i>;<br/>    <i>source-port</i> upper-limit <i>upper-limit</i>;<br/>}</code>                                                                |
| <b>Hierarchy Level</b>          | <code>[edit system]</code>                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                               |
| <b>Description</b>              | Configure system IP options to protect against certain types of denial-of-service (DoS) attacks.<br><br>The remaining statements are explained separately.                                                                                                      |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages on page 165</a></li><li>• <a href="#">Configuring Junos OS to Extend the Default Port Address Range on page 164</a></li></ul> |


## l2-learning

---

|                                 |                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>l2-learning {<br/>    global-mac-limit <i>limit</i>;<br/>    global-mac-statistics;<br/>    global-mac-table-aging-time <i>seconds</i>;<br/>    global-no-mac-learning;<br/>}</pre>                |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.4.                                                                                                                                                           |
| <b>Description</b>              | <p>(MX Series routers, and EX Series switches, and QFX Series switches only) Configure Layer 2 address learning and forwarding properties globally.</p> <p>The statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Layer 2 Learning and Forwarding Overview</i></li></ul>                                                                                                       |

## load-key-file

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | load-key-file <i>URL filename</i> ;                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system root-authentication],<br>[edit system login user <i>username</i> authentication]                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <div> <b>NOTE:</b> ECDSA is not supported on the QFabric system.</div> <p>Load RSA (SSH version 1 and SSH version 2) and DSA or ECDSA (SSH version 2) public keys from a previously-generated named file at a specified URL location or local path. The file contains one or more SSH keys that are copied into the configuration when the command is issued.</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Root Password on page 170</a></li><li>• <a href="#">Configuring the Root Password on page 1354</a></li><li>• <a href="#">Configuring Junos OS User Accounts</a></li><li>• <a href="#">Configuring Junos OS User Accounts on page 1344</a></li></ul>                                                                                                                            |



## location

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>location {   altitude <i>feet</i>;   building <i>name</i>;   country-code <i>code</i>;   floor <i>number</i>;   hcoord <i>horizontal-coordinate</i>;   lata <i>service-area</i>;   latitude <i>degrees</i>;   longitude <i>degrees</i>;   npa-nxx <i>number</i>;   postal-code <i>postal-code</i>;   rack <i>number</i>;   vcoord <i>vertical-coordinate</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure the system location.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | <p><b>altitude <i>feet</i></b>—Number of feet above sea level.</p> <p><b>building <i>name</i></b>—Name of the building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").</p> <p><b>country-code <i>code</i></b>—Two-letter country code.</p> <p><b>floor <i>number</i></b>—Floor in the building.</p> <p><b>hcoord <i>horizontal-coordinate</i></b>—Bellcore Horizontal Coordinate.</p> <p><b>lata <i>service-area</i></b>—Long-distance service area.</p> <p><b>latitude <i>degrees</i></b>—Latitude in degree format.</p> <p><b>longitude <i>degrees</i></b>—Longitude in degree format.</p> <p><b>npa-nxx <i>number</i></b>—First six digits of the phone number (area code and exchange).</p> <p><b>postal-code <i>postal-code</i></b>—Postal code.</p> <p><b>rack <i>number</i></b>—Rack number.</p> <p><b>vcoord <i>vertical-coordinate</i></b>—Bellcore Vertical Coordinate.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Specifying the Physical Location of the Switch on page 169</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## login

```
Syntax login {
    announcement text;
    class class-name {
        access-end "regular-expression";
        access-start "regular-expression";
        allow-commands "regular-expression";
        allow-configuration "regular-expression";
        deny-commands "regular-expression";
        deny-configuration "regular-expression";
        idle-timeout minutes;
        login-tip;
        permissions [ permissions ];
    }
    message text;
    password {
        change-type (set-transitions | character-set);
        format (md5 | sha1 | des);
        maximum-length length;
        minimum-changes number;
        minimum-length length;
    }
    retry-options {
        backoff-factor seconds;
        backoff-threshold number;
        minimum-time seconds;
        tries-before-disconnect number;
    }
    user username {
        authentication authentication;
        (encrypted-password "password" | plain-text-password);
        load-key-file URL;
        remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);
        ssh-dsa "public-key";
        ssh-rsa "public-key";
    }
    class class-name;
    full-name complete-name;
    uid uid-value;
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure user access to the switch.

**Options** The remaining statements are explained separately.

**Required Privilege Level** admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

- Related Documentation**
- [Defining Junos OS Login Classes on page 1368](#)

## login-alarms

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | login-alarms;                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system login class <i>class-name</i> ]                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Show system alarms automatically when an <b>admin</b> user logs in to the router or switch.                                                                                                    |
| <b>Options</b>                  | <i>class-name</i> —Login class name.                                                                                                                                                           |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring System Alarms to Appear Automatically Upon Login on page 172</a></li> </ul>                                                   |

## login-tip

---

|                                 |                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | login-tip;                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system login class <i>class-name</i> ]                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                 |
| <b>Description</b>              | Enable CLI tips at login.                                                                                         |
| <b>Default</b>                  | Disabled.                                                                                                         |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Login Tips on page 41</a></li> </ul>             |

## max-configurations-on-flash

---

|                                 |                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>max-configurations-on-flash</code> <i>number</i> ;                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Specify the number of configurations stored on the internal fixed media storage (for example, USB device).                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>number</i> —The number of configurations stored on the CompactFlash card.<br><b>Range:</b> 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.                                                                                                                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Saving a Configuration to a File on page 1257</a></li><li>• <a href="#">Setting or Deleting the Rescue Configuration on page 1261</a></li><li>• <a href="#">Uploading a Configuration File on page 1261</a></li><li>• <a href="#">Uploading a Configuration File</a></li></ul> |

## maximum-length

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>maximum-length</code> <i>length</i> ;                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system login passwords]                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                          |
| <b>Description</b>              | Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.                                                              |
| <b>Default</b>                  | For Junos OS-FIPS software, the maximum number of characters for plain-text passwords is 20. For Junos OS, no maximum is set.                                                              |
| <b>Options</b>                  | <i>length</i> —Maximum number of characters the password can include.<br><b>Range:</b> 1 to 64 characters                                                                                  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li><li>• <a href="#">minimum-length on page 284</a></li></ul> |

## message

---

|                                 |                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>message text;</code>                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit system login]                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                              |
| <b>Description</b>              | Configure a system login message. This message appears before a user logs in.                                                                                                                  |
| <b>Options</b>                  | <i>text</i> —Text of the message.                                                                                                                                                              |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Junos OS to Display a System Login Message on page 163</a></li> <li>• <a href="#">announcement on page 251</a></li> </ul> |

## minimum-changes

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-changes number;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system login passwords]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Specify the minimum number of character sets (or character set changes) required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement is used in combination with the <b>change-type</b> statement. If the change type is <b>character-sets</b>, then the number of character sets included in the password is checked against the specified minimum. If the change type is <b>set-transitions</b>, then the number of character set changes in the password is checked against the specified minimum.</p> |
| <b>Default</b>                  | For Junos OS, the minimum number of changes is 1. For Junos-FIPS software, the minimum number of changes is 3.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>number</i> —Minimum number of character sets (or character set changes) required for the password.                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li> <li>• <a href="#">change-type on page 260</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                  |

## minimum-length

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | minimum-length <i>length</i> ;                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit system login password]                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                          |
| <b>Description</b>              | Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.                                                             |
| <b>Default</b>                  | For Junos OS, the minimum number of characters for plain-text passwords is six. For Junos-FIPS software, the minimum number of characters for plain-text passwords is 10.                  |
| <b>Options</b>                  | <b>length</b> —Minimum number of characters the password must include.<br><b>Range:</b> 6 to 20 characters                                                                                 |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li><li>• <a href="#">maximum-length on page 282</a></li></ul> |

## minimum-lower-cases

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-lower-cases <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit system login password]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Specify the minimum number of lower-case letters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as <b>minimum-length</b>, <b>minimum-punctuations</b>, <b>minimum-upper-cases</b>, and so on.</p> <p>Using several password minimum requirement options will cause the <b>minimum-length</b> to be reset if the total sum of the required minimums exceeds the <b>minimum-length</b> setting.</p> |
| <b>Options</b>                  | <i>number</i> —The minimum number of lower-case letters required for the password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li> <li>• <a href="#">Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 196</a></li> <li>• <i>password (Login)</i></li> </ul>                                                                                                                                                                                                                                                                                |

## minimum-numeric

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-numeric <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit system login password]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Specify the minimum number of numeric class characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as <b>minimum-length</b>, <b>minimum-punctuations</b>, <b>minimum-lower-cases</b>, and so on.</p> <p>Using several password minimum requirement options will cause the <b>minimum-length</b> to be reset if the total sum of the required minimums exceeds the <b>minimum-length</b> setting.</p> |
| <b>Options</b>                  | <i>number</i> —The minimum number of numeric class characters required for the password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <code>system</code> —To view this statement in the configuration.<br><code>system-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li><li>• <a href="#">Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 196</a></li><li>• <i>password (Login)</i></li></ul>                                                                                                                                                                                                                                                                                          |



## minimum-punctuations

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-punctuations <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit system login password]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Specify the minimum number of punctuation class characters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as <b>minimum-length</b>, <b>minimum-upper-cases</b>, <b>minimum-lower-cases</b>, and so on.</p> <p>Using several password minimum requirement options will cause the <b>minimum-length</b> to be reset if the total sum of the required minimums exceeds the <b>minimum-length</b> setting.</p> |
| <b>Options</b>                  | <i>number</i> —The minimum number of punctuation class characters required for the password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li> <li>• <a href="#">Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 196</a></li> <li>• <i>password (Login)</i></li> </ul>                                                                                                                                                                                                                                                                                         |

## minimum-upper-cases

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>minimum-upper-cases <i>number</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit system login password]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | <p>Specify the minimum number of upper-case letters required in plain-text passwords. Newly created passwords must meet this requirement.</p> <p>This statement can be used in combination with all of the other requirement options for plain-text passwords, such as <b>minimum-length</b>, <b>minimum-punctuations</b>, <b>minimum-lower-cases</b>, and so on.</p> <p>Using several password minimum requirement options will cause the <b>minimum-length</b> to be reset if the total sum of the required minimums exceeds the <b>minimum-length</b> setting.</p> |
| <b>Options</b>                  | <i>number</i> —The minimum number of upper-case letters required for the password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li><li>• <a href="#">Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 196</a></li><li>• <i>password (Login)</i></li></ul>                                                                                                                                                                                                                                                                                    |

## multicast-client

---

|                                 |                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>multicast-client &lt;<i>address</i>&gt;;</code>                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit system <a href="#">ntp</a> ]                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                         |
| <b>Description</b>              | For Network Time Protocol (NTP), configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet. |
| <b>Options</b>                  | <i>address</i> —(Optional) One or more IP addresses. If you specify addresses, the router or switch joins those multicast groups.<br><b>Default:</b> 224.0.1.1.           |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Router or Switch to Listen for Multicast Messages Using NTP on page 172</a></li></ul>                 |

---

## name-server

---

|                                 |                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>name-server {<br/>    <i>address</i>;<br/>}</code>                                                                                            |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                   |
| <b>Description</b>              | Configure one or more Domain Name System (DNS) name servers.                                                                                        |
| <b>Options</b>                  | <i>address</i> —Address of the name server. To configure multiple name servers, include multiple <i>address</i> options.                            |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring a DNS Name Server for Resolving a Hostname into Addresses on page 158</a></li></ul> |

## no-multicast-echo

---

**Syntax**   no-multicast-echo {  
            arp {  
                aging-timer *minutes*;  
                gratuitous-arp-delay*seconds*;  
                gratuitous-arp-on-ifup;  
                interfaces {  
                    *interface-name* {  
                        aging-timer *minutes*;  
                    }  
                }  
                passive-learning;  
                purging;  
            }  
            host-name *hostname*;  
            location {  
                altitude *feet*;  
                building *name*;  
                country-code *code*;  
                floor *number*;  
                hcoord *horizontal-coordinate*;  
                lata *service-area*;  
                latitude *degrees*;  
                longitude *degrees*;  
                npa-nxx *number*;  
                postal-code *postal-code*;  
                rack *number*;  
                vcoord *vertical-coordinate*;  
            }  
            license {  
                autoupdate *URL*;  
                }  
                renew before-expiration (*number* | interval *number*)  
            }  
            }  
            }

**Hierarchy Level**   [edit system]

**Release Information**   Statement introduced in Junos OS Release 8.1.  
                          Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                          Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**       Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

**Default**           The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

**Required Privilege Level**   system—To view this statement in the configuration.  
                              system-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets](#)

## no-ping-record-route

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ping-record-route;                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.4 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the Junos OS to disable the reporting of the IP address in ping responses.                                                                                                       |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 162</a></li> </ul>               |

## no-ping-time-stamp

---

|                                 |                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-ping-time-stamp;                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4.<br>Statement introduced in Junos OS Release 9.4 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Configure the Junos OS to disable the recording of timestamps in ping responses.                                                                                                           |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses on page 162</a></li> </ul>               |

## no-redirects (IPv4 Traffic)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-redirects;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit system],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | <p>Stop protocol redirect messages for IPv4 traffic from being sent on the entire switch or on an interface on the router or switch.</p> <p>To disable the sending of protocol redirect messages for the entire router or switch, include the <b>no-redirects</b> statement at the [edit system] hierarchy level.</p> <p>To disable the sending of protocol redirect messages on a specific interface, include the <b>no-redirects</b> statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] hierarchy level.</p> |
| <b>Default</b>                  | The router or switch sends redirect messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Junos OS to Disable Protocol Redirect Messages on the Router or Switch on page 162</a></li><li>• <i>Understanding the Protocol Redirect Mechanism on EX Series Switches</i></li><li>• <i>Configuring Junos OS to Disable Sending Protocol Redirect Messages on EX Series Switches (CLI Procedure)</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>                                                                                                                  |

## no-split-detection

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | no-split-detection;                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <a href="#">virtual-chassis</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | <p>Disable the split and merge feature in a Virtual Chassis or VCF configuration.</p> <p>We recommend using this statement to disable the split and merge feature when configuring a two-member Virtual Chassis. Enabling this statement on a two-member Virtual Chassis ensures that both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split.</p>                                                                                                                          |
| <b>Default</b>                  | The split and merge feature is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li> <li>• <a href="#">Disabling Split and Merge in a Virtual Chassis (CLI Procedure) on page 6944</a></li> <li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 6946</a></li> <li>• <a href="#">Understanding Split and Merge in a Virtual Chassis on page 6922</a></li> </ul> |

## ntp

---

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>ntp {<br/>  authentication-key <i>number</i> type <i>type</i> value <i>password</i>;<br/>  boot-server <i>address</i>;<br/>  broadcast &lt;<i>address</i>&gt; &lt;<i>key key-number</i>&gt; &lt;<i>version value</i>&gt; &lt;<i>ttl value</i>&gt;;<br/>  broadcast-client;<br/>  multicast-client &lt;<i>address</i>&gt;;<br/>  peer <i>address</i> &lt;<i>key key-number</i>&gt; &lt;<i>version value</i>&gt; &lt;<i>prefer</i>&gt;;<br/>  server <i>address</i> &lt;<i>key key-number</i>&gt; &lt;<i>version value</i>&gt; &lt;<i>prefer</i>&gt;;<br/>  source-address <i>source-address</i>;<br/>  trusted-key [ <i>key-numbers</i> ];<br/>}</pre> |
| Hierarchy Level          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Description              | <p>Configure Network Time Protocol (NTP) on the switch.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Required Privilege Level | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Synchronizing and Coordinating Time Distribution Using NTP on page 194</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## optional

---

|                          |                                                                                                                                                                                                       |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>optional;</pre>                                                                                                                                                                                  |
| Hierarchy Level          | [edit system scripts commit file <i>filename</i> ]                                                                                                                                                    |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                     |
| Description              | For Junos OS commit scripts, allow a commit operation to succeed even if the script specified in the <b>file</b> statement is missing from the <b>/var/db/scripts/commit</b> directory on the router. |
| Required Privilege Level | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration.                                                                           |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Controlling Execution of Commit Scripts During Commit Operations</a></li></ul>                                                                    |



---

## password (Login)

---

|                                 |                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>password {<br/>  <b>change-type</b> (set-transitions   character-set);<br/>  <b>format</b> (md5   sha1   des);<br/>  <b>maximum-length</b> <i>length</i>;<br/>  <b>minimum-changes</b> <i>number</i>;<br/>  <b>minimum-length</b> <i>length</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system login]                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                 |
| <b>Description</b>              | <p>Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirements.</p> <p>The remaining statements are explained separately.</p>                                    |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Special Requirements for Junos OS Plain-Text Passwords on page 1339</a></li></ul>                                                                                                                             |

## peer

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>peer address &lt;key key-number&gt; &lt;version value&gt; &lt;prefer&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit system <a href="#">ntp</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | For NTP, configure the local router or switch to operate in symmetric active mode with the remote system at the specified address. In this mode, the local router or switch and the remote system can synchronize with each other. This configuration is useful in a network in which either the local router or switch or the remote system might be a better source of time.                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>address</b>—Address of the remote system. You must specify an address, not a hostname.</p> <p><b>key key-number</b>—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.</p> <p><b>Range:</b> Any unsigned 32-bit integer</p> <p><b>prefer</b>—(Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p><b>version value</b>—(Optional) Specify the NTP version number to be used in outgoing NTP packets.</p> <p><b>Range:</b> 1 through 4</p> <p><b>Default:</b> 4</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the NTP Time Server and Time Services on page 166</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## permissions

---

|                                 |                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>permissions {<br/>    storage;<br/>    storage-control;<br/>}</code>                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system login class]                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                 |
| <b>Description</b>              | Configure the login access privileges to be provided on the switch.                                                                                                                                               |
| <b>Options</b>                  | <i>permissions</i> —Privilege type.                                                                                                                                                                               |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Access Privilege Levels on page 1344</a></li> <li>• <a href="#">Table 77 on page 1326</a></li> <li>• <a href="#">user on page 332</a></li> </ul> |

## port (TACACS+ Server)

---

|                                 |                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>port <i>port-number</i>;</code>                                                                                  |
| <b>Hierarchy Level</b>          | [edit system accounting destination tacplus server <i>server-address</i> ]                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                      |
| <b>Description</b>              | Configure the port number on which to contact the TACACS+ server.                                                      |
| <b>Options</b>                  | <i>number</i> —Port number on which to contact the TACACS+ server.<br><b>Default:</b> 49                               |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring TACACS+ System Accounting on page 1366</a></li> </ul> |

## ports

---

|                                 |                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ports {<br/>  auxiliary {<br/>    disable;<br/>    insecure;<br/>    type <i>terminal-type</i>;<br/>  }<br/>  console {<br/>    disable;<br/>    insecure;<br/>    log-out-on-disconnect;<br/>    type <i>terminal-type</i>;<br/>  }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure the properties of the console and auxiliary ports. The ports are located on the craft interface.</p> <p>See the switch hardware documentation for port locations.</p> <p>The remaining statements are explained separately.</p>          |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Console and Auxiliary Port Properties on page 6583</a></li></ul>                                                                                                                      |

## radius (System)

|                                 |                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>radius {   server {     server-address {       accounting-port <i>port-number</i>;       secret <i>password</i>;       source-address <i>address</i>;       retry <i>number</i>;       timeout <i>seconds</i>;     }   } }</pre> |
| <b>Hierarchy Level</b>          | [edit system accounting destination]                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                             |
| <b>Description</b>              | Configure the RADIUS accounting server.                                                                                                                                                                                               |
| <b>Options</b>                  | <p><b><i>server-address</i></b>—Address of the RADIUS accounting server.</p> <p>The remaining statements are explained separately.</p>                                                                                                |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS System Accounting on page 1349</a></li> </ul>                                                                                                                 |

## refresh (Commit Scripts)

---

|                                 |                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>refresh;</code>                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | <code>[edit system scripts commit],</code><br><code>[edit system scripts file <i>filename</i>]</code>                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                             |
| <b>Description</b>              | For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level. |
| <b>Required Privilege Level</b> | <code>maintenance</code> —To view this statement in the configuration.<br><code>maintenance-control</code> —To add this statement to the configuration.                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">refresh-from on page 300</a></li><li>• <a href="#">source on page 308</a></li></ul>                                                                                                                                                       |

## refresh-from (Commit Scripts)

---

|                                 |                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>refresh-from url;</code>                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | <code>[edit system scripts commit],</code><br><code>[edit system scripts commit file <i>filename</i>]</code>                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                       |
| <b>Description</b>              | For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at a URL other than the URL specified in the <b>source</b> statement. |
| <b>Options</b>                  | <b>url</b> —The source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.                                                                                                                            |
| <b>Required Privilege Level</b> | <code>maintenance</code> —To view this statement in the configuration.<br><code>maintenance-control</code> —To add this statement to the configuration.                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">refresh on page 300</a></li><li>• <a href="#">source on page 308</a></li></ul>                                                                                                                                      |

## retry

|                            |                                                                                                                                  |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>retry number;</code>                                                                                                       |
| <b>Hierarchy Level</b>     | [edit system radius server <i>server-address</i> ],<br>[edit system accounting destination radius server <i>server-address</i> ] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                |
| <b>Description</b>         | Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.                  |
| <b>Options</b>             | <i>number</i> —Number of retries allowed for contacting a RADIUS server.<br><b>Range:</b> 1 through 10<br><b>Default:</b> 3      |



**NOTE:** The [edit system accounting] hierarchy is not available on QFabric systems.

|                                 |                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li> <li>• <a href="#">Configuring RADIUS Accounting</a></li> <li>• <a href="#">timeout on page 322</a></li> </ul> |

## retry-options

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>retry-options {<br/>    backoff-threshold <i>number</i>;<br/>    backoff-factor <i>seconds</i>;<br/>    maximum-time <i>seconds</i>;<br/>    minimum-time <i>seconds</i>;<br/>    tries-before-disconnect <i>number</i>;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit system login]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Maximum number of times a user can attempt to enter a password while logging in through SSH or Telnet before being disconnected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>backoff-threshold <i>number</i></b>—Threshold for the number of failed login attempts before the user experiences a delay when attempting to reenter a password. Use the <b>backoff-factor</b> option to specify the length of delay, in seconds.</p> <p><b>Range:</b> 1 through 3</p> <p><b>Default:</b> 2</p> <p><b>backoff-factor <i>seconds</i></b>—Length of delay after each failed login attempt. The length of delay increases by this value for each subsequent login attempt after the value specified in the <b>backoff-threshold</b> option.</p> <p><b>Range:</b> 5 through 10</p> <p><b>Default:</b> 5</p> <p><b>maximum-time <i>seconds</i></b>—Maximum length of time that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the configured <b>maximum-time</b>, the connection is closed.</p> <p><b>Range:</b> 20 through 300</p> <p><b>Default:</b> 120</p> <p><b>minimum-time <i>seconds</i></b>—Minimum length of time that the connection remains open while the user is attempting to enter a password to log in.</p> <p><b>Range:</b> 20 through 60</p> <p><b>Default:</b> 20</p> <p><b>tries-before-disconnect <i>number</i></b>—Maximum number of times a user is allowed to attempt to enter a password to log in through SSH or Telnet.</p> <p><b>Range:</b> 1 through 10</p> <p><b>Default:</b> 10</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



- Related Documentation
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1369](#)

## root-authentication

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax                   | <pre>root-authentication {     (encrypted-password "password"   load-key-password URL   plain-text-password);     ssh-dsa "public-key";     ssh-rsa "public-key"; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Hierarchy Level          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Description              | Configure the authentication methods for the root-level user, whose username is <b>root</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Options                  | <p><b>encrypted-password "password"</b>— Specify the MD5 or other encrypted authentication password. You can specify only one encrypted password.</p> <p>You cannot configure a blank password for the <b>encrypted-password</b> option using blank quotation marks (" "). You must configure a password of 1 through 128 characters and enclose the password in quotation marks.</p> <p><b>plain-text-password</b>—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.</p> <p><b>ssh-dsa "public-key"</b>—SSH version 2 authentication. Specify the DSA (SSH version 2) public key. You can specify one or more public keys.</p> <p><b>ssh-rsa "public-key"</b>—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.</p> |
| Required Privilege Level | <p><b>admin</b>—To view this statement in the configuration.</p> <p><b>admin-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Root Password on page 1354</a></li> <li>• <a href="#">Recovering the Root Password</a></li> <li>• <a href="#">authentication on page 254</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



**CAUTION:** Do not use the **encrypted-password** option unless the password is *already* encrypted, and you are entering the encrypted version of the password. If you commit the **encrypted-password** option with a plain-text password or with blank quotation marks (" "), you will not be able to log in to the device as root, and you will need to use the password recovery process.

## saved-core-context

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (saved-core-context   no-saved-core-context);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure whether the switch saves core files generated by internal Junos OS processes, along with contextual information (system log files and a copy of the current configuration):</p> <ul style="list-style-type: none"><li>• <b>saved-core-context</b>—The switch saves each core file and its associated context in a compressed tar file named <code>/var/tmp/process-name.core.core-number.tgz</code>.</li><li>• <b>no-saved-core-context</b>—The switch does not save core files and their associated context.</li></ul> |
| <b>Default</b>                  | The switch saves core files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Saving Core Files from Junos OS Processes</i></li><li>• <a href="#">saved-core-files on page 304</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                            |

## saved-core-files

---

|                                 |                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | saved-core-files <i>number</i> ;                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                           |
| <b>Description</b>              | Save core files generated by internal Junos OS processes, but not the associated contextual information (configuration and system log files).               |
| <b>Options</b>                  | <p><i>number</i>—Maximum number of core files to save.</p> <p><b>Range:</b> 1 through 10</p>                                                                |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Saving Core Files from Junos OS Processes</i></li><li>• <a href="#">saved-core-context on page 304</a></li></ul> |

## secret

|                                 |                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>secret password;</code>                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit system accounting destination radius server <i>server-address</i> ],<br>[edit system accounting destination tacplus server <i>server-address</i> ],<br>[edit system radius-server <i>server-address</i> ],<br>[edit system tacplus-server <i>server-address</i> ]                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local switch must match that used by the server.                                                                                                                                                                                                              |
| <b>Options</b>                  | <i>password</i> —Password to use; can include spaces included in quotation marks.                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS Accounting</a></li> <li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li> <li>• <a href="#">Configuring TACACS+ Authentication (QFX Series) on page 1364</a></li> <li>• <a href="#">Configuring TACACS+ System Accounting on page 1366</a></li> </ul> |

## server (TACACS+ Accounting)

|                                 |                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>server {   server-address {     port <i>port-number</i>;     secret <i>password</i>;     single-connection;     timeout <i>seconds</i>;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit system accounting destination tacplus]                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                        |
| <b>Description</b>              | <p>Configure TACACS+ logging.</p> <p>The remaining statements are explained separately.</p>                                                              |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring TACACS+ System Accounting on page 1366</a></li> </ul>                                   |

## server (NTP)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>server address &lt;key key-number&gt; &lt;version value&gt; &lt;prefer&gt;;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit system ntp]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | For NTP, configure the switch to operate in client mode with the remote system at the specified server address. In this mode, the local switch can be synchronized with the remote system, but the remote system can never be synchronized with the local switch.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>address</b>—Address of the remote system. You must specify an address, not a hostname.</p> <p><b>key key-number</b>—(Optional) Use the specified key number to encrypt authentication fields in all packets sent to the specified address.</p> <p><b>Range:</b> Any unsigned 32-bit integer</p> <p><b>prefer</b>—(Optional) Mark the remote system as preferred host, which means that if all other things are equal, this remote system is chosen for synchronization among a set of correctly operating systems.</p> <p><b>version value</b>—(Optional) Specify the version number to be used in outgoing NTP packets.</p> <p><b>Range:</b> 1 through 4</p> <p><b>Default:</b> 4</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>ntp</i></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## server (RADIUS Accounting)

|                                 |                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>server {   server-address {     accounting-port port-number;     retry number     secret password;     source-address address;     timeout seconds;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit system accounting destination radius]                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                    |
| <b>Description</b>              | <p>Configure RADIUS logging.</p> <p>The remaining statements are explained separately.</p>                                                                           |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS System Accounting on page 1349</a></li> </ul>                                                |

## single-connection

|                                 |                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | single-connection;                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | <p>[edit system accounting destination tacplus server <i>server-address</i>],</p> <p>[edit system tacplus <i>server-address</i>]</p>                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                              |
| <b>Description</b>              | Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests rather than opening a connection for each connection attempt.             |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring TACACS+ Authentication (QFX Series) on page 1364</a></li> <li>• <a href="#">Configuring TACACS+ System Accounting on page 1366</a></li> </ul> |

## source (Commit Scripts)

---

|                                 |                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source url;</code>                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit system scripts commit file <i>filename</i> ]                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | For Junos OS commit scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/commit</code> directory. When you include the <b>refresh</b> statement at the same hierarchy level and commit the configuration, the local copy is overwritten by the version stored at the specified URL. |
| <b>Options</b>                  | <i>url</i> —The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <b>maintenance</b> —To view this statement in the configuration.<br><b>maintenance-control</b> —To add this statement to the configuration.                                                                                                                                                                                                |

## source-address (NTP, RADIUS, System Logging, or TACACS+)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-address source-address;</code>                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system accounting destination radius server <i>server-address</i> ],<br>[edit system accounting destination tacplus server <i>server-address</i> ],<br>[edit system ntp],<br>[edit system radius-server <i>server-address</i> ],<br>[edit system syslog],<br>[edit system tacplus-server <i>server-address</i> ]                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify a source address for each configured TACACS+ server, RADIUS server, NTP server, or the source address to record in system log messages that are directed to a remote machine.                                                                                                                                                                      |
| <b>Options</b>                  | <i>source-address</i> —Valid IP address configured on one of the switch interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all <b>host hostname</b> statements at the [edit system syslog] hierarchy level.                                                                |
| <b>Required Privilege Level</b> | <b>system</b> —To view this statement in the configuration.<br><b>system-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li><li>• <a href="#">Synchronizing and Coordinating Time Distribution Using NTP on page 194</a></li><li>• <a href="#">Specifying an Alternative Source Address for System Log Messages Directed to a Remote Destination</a></li></ul> |

## source-port (Port Addresses)

---

|                                 |                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>source-port upper-limit &lt;upper-limit&gt;;</code>                                                                                     |
| <b>Hierarchy Level</b>          | [edit system internet-options]                                                                                                                |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                        |
| <b>Description</b>              | Configure the range of port addresses.                                                                                                        |
| <b>Options</b>                  | <b>upper-limit <i>upper-limit</i></b> —(Optional) The range of port addresses and can be a value from 5000 through 65,355.                    |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Junos OS to Extend the Default Port Address Range on page 164</a></li> </ul> |

## ssh-dsa

---

|                                 |                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>ssh-dsa "public-key";</code>                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit system root-authentication]<br>[edit system login user authentication]                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                         |
| <b>Description</b>              | Specify the DSA (SSH version 2) public key. You can specify one or more public keys.                                                                                                                      |
| <b>Options</b>                  | <b>ssh-dsa "public-key"</b> —SSH version 2 authentication.                                                                                                                                                |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Root Password on page 170</a></li> <li>• <a href="#">authentication on page 254</a></li> <li>• <i>root-authentication</i></li> </ul> |

## ssh-rsa

---

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | ssh-dsa " <i>public-key</i> ";                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit system root-authentication]<br>[edit system login user authentication]                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | Specify the RSA (SSH version 1) public key. You can specify one or more public keys.                                                                                                                   |
| <b>Options</b>                  | ssh-rsa " <i>public-key</i> "—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.                                     |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Root Password on page 1354</a></li><li>• <a href="#">authentication on page 254</a></li><li>• <i>root-authentication</i></li></ul> |




## static-host-mapping

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>static-host-mapping {     hostname {         alias [ <i>alias</i> ];         inet [ <i>address</i> ];         sysid <i>system-identifier</i>;     } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Map a hostname to one or more IP addresses and aliases, and configure an International Organization for Standardization (ISO) system identifier (system ID).                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><b>alias <i>alias</i></b>—Alias for the hostname.</p> <p><b>hostname</b>—Fully qualified hostname.</p> <p><b>inet <i>address</i></b>—IP address. You can specify one or more IP addresses for the host.</p> <p><b>sysid <i>system-identifier</i></b>—ISO system identifier (system ID). This is the 6-byte portion of the Intermediate System-to-Intermediate System (IS-IS) network service access point (NSAP). We recommend that you use the host's IP address represented in binary-coded decimal (BCD) format. For example, the IP address 208.197.169.18 is 2081.9716.9018 in BCD.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Hostname of the Router or Switch on page 160</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## structured-data

---

|                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                  | structured-data {<br>brief;<br>}                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                         | [edit system syslog file <i>filename</i> ]                                                                                                                                                                                                                                                             |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                     | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                             | Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, <i>The syslog Protocol</i> ( <a href="http://tools.ietf.org/html/draft-ietf-syslog-protocol-23">http://tools.ietf.org/html/draft-ietf-syslog-protocol-23</a> ). |
| <div> <b>NOTE:</b> When this statement is included, other statements that specify the format for messages written to the file are ignored (the explicit-priority statement at the [edit system syslog file <i>filename</i>] hierarchy level and the time-format statement at the [edit system syslog] hierarchy level).</div> |                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                   | <ul style="list-style-type: none"><li>• <i>Logging Messages in Structured-Data Format</i></li><li>• <a href="#">explicit-priority on page 273</a></li><li>• <a href="#">time-format on page 323</a></li></ul>                                                                                          |

## syslog (System)

```

Syntax  syslog {
        allow-duplicates;
        archive {
            (binary-data | no-binary-data);
            files number;
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
        }
        console {
            facility severity;
        }
        file filename {
            facility severity;
            explicit-priority;
            match "regular-expression";
            archive {
                (binary-data | no-binary-data);
                files number;
                size maximum-file-size;
                start-time "YYYY-MM-DD.hh:mm";
                transfer-interval minutes;
                (world-readable | no-world-readable);
            }
            structured-data {
                brief;
            }
        }
        host (hostname | other-routing-engine | scc-master) {
            facility severity;
            explicit-priority;
            facility-override facility;
            log-prefix string;
            match "regular-expression";
            source-address source-address;
            structured-data {
                brief;
            }
            port port number;
        }
        log-rotate-frequency frequency;
        server server name;
        source-address source-address;
        time-format (millisecond | year | year millisecond);
        user (username | *) {
            facility severity;
            match "regular-expression";
        }
    }

```

Hierarchy Level [edit system]

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure the types of system log messages to send to files, to a remote destination, to user terminals, or to the system console.<br><br>The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <b>archive</b> —Define parameters for archiving log messages.<br><br><b>console</b> —Send log messages of a specified class and severity to the console.<br><br><b>file</b> —Send log messages to a named file.<br><br><b>host</b> —Remote location to be notified of specific log messages.<br><br><b>log-rotate-frequency</b> —Configure the interval for checking logfile size and archiving messages.<br><br><b>server</b> —Name of the system log server in the inet.0 routing instance.<br><br><b>source-address</b> —Include a specified address as the source address for log messages.<br><br><b>time-format</b> —Additional information to include in the system log time stamp.<br><br><b>user</b> —Notify a specific user of the log event. |
| <b>Required Privilege Level</b> | <b>system</b> —To view this statement in the configuration.<br><b>system-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Junos OS System Log Overview</i></li><li>• <i>Junos OS System Log Messages Reference</i></li><li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6561</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## system

```
Syntax  system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
    default-address-selection;
    domain-name domain-name;
    domain-search [ domain-list ];
    host-name hostname;
    internet-options {
        icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
        source-port upper-limit <upper-limit>;
    }
    location {
```

```
altitude feet;  
building name;  
country-code code;  
floor number;  
hcoord horizontal-coordinate;  
lata service-area;  
latitude degrees;  
longitude degrees;  
npa-nxx number;  
postal-code postal-code;  
rack number;  
vcoord vertical-coordinate;  
}  
login {  
  announcement text;  
  class class-name {  
    access-end;  
    access-start;  
    allow-configuration "regular-expression";  
    allowed-days "regular-expression";  
    deny-commands "regular-expression";  
    deny-configuration "regular-expression";  
    idle-timeout minutes;  
    login-tip;  
    permissions [ permissions ];  
  }  
  message text;  
  password {  
    change-type (set-transitions | character-set);  
    format (md5 | sha1 | des);  
    maximum-length length;  
    minimum-changes number;  
    minimum-length length;  
  }  
  retry-options {  
    backoff-factor seconds;  
    backoff-threshold number;  
    minimum-time seconds;  
    tries-before-disconnect number;  
  }  
  user username {  
    authentication {  
      (encrypted-password "password" | plain-text-password);  
      load-key-file URL;  
      remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);  
      ssh-rsa "public-key";  
      ssh-dsa "public-key";  
    }  
    uid uid-value;  
    class class-name;  
    full-name complete-name;  
  }  
}  
name-server {  
  address;  
}
```

```

no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
    authentication-key number type type value password;
    serveraddress <key key-number> <version value> <prefer>;
}
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
radius-server server-address {
    accounting-port port-number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
radius-options {
    password-protocol mschap-v2;
}
attributes {
    nas-ip-address ip-address;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
}

```

```
}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
}
```



```

}
console {
  facility severity;
}
file filename {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
  explicit-priority;
  facility severity;
  match "regular-expression";
  structured-data {
    brief;
  }
}
host (hostname | other-routing-engine | scc-master) {
  explicit-priority;
  facility-override facility;
  facility severity;
  log-prefix string;
  match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
  facility severity;
  match "regular-expression";
}
}
tacplus-options {
  service-name service-name;
  (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
  port
  secret password;
  single-connection;
  source-address source-address;
  timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
  destination-override {
    syslog host;
  }
}
use-imported-time-zones;
}

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure system management properties.



**NOTE:** The `radius-server source-address` and `radius-options` statements are not available on the QFabric system.

---

**Required Privilege** system—To view this statement in the configuration.

**Level** system-control—To add this statement to the configuration.

---

## tacplus

---

**Syntax**

```
tacplus {  
  server {  
    server-address {  
      port port-number;  
      secret password;  
      single-connection;  
      timeout seconds;  
    }  
  }  
}
```

**Hierarchy Level** [edit system accounting destination]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure TACACS+.

**Options** *server-address*—Address of the TACACS+ authentication server.

The remaining statements are explained separately.

**Required Privilege** system—To view this statement in the configuration.

**Level** system-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring TACACS+ System Accounting on page 1366](#)

---

## tacplus-server

---


|                                 |                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>tacplus-server server-address {<br/>    port<br/>    secret password;<br/>    single-connection;<br/>    source-address source-address;<br/>    timeout seconds;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                 |
| <b>Description</b>              | Configure the TACACS+ server.                                                                                                                                                     |
| <b>Options</b>                  | <p><b>server-address</b>—Address of the TACACS+ authentication server.</p> <p>The remaining statements are explained separately.</p>                                              |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring TACACS+ Authentication (QFX Series) on page 1364</a></li></ul>                                                    |

## timeout

---

|                                 |                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>timeout seconds;</code>                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit system radius-server <i>server-address</i> ],<br>[edit system tacplus-server <i>server-address</i> ],<br>[edit system accounting destination radius server <i>server-address</i> ],<br>[edit system accounting destination tacplus server <i>server-address</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                       |
| <b>Description</b>              | Configure the length of time that the local router or switch waits to receive a response from a RADIUS or TACACS+ server.                                                                                                                                               |
| <b>Options</b>                  | <b>seconds</b> —Length of time to wait.<br><b>Range:</b> 1 through 90 seconds<br><b>Default:</b> 3 seconds                                                                                                                                                              |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring RADIUS Accounting</a></li><li>• <a href="#">Configuring TACACS+ System Accounting on page 1366</a></li><li>• <a href="#">retry on page 301</a></li></ul>                                                |

## time-format

|                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                      | time-format (year   millisecond   year millisecond);                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                             | [edit system syslog]                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>                                                                                                                                                                                                                                                         | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                                                                                                                                                                                                                                                 | Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a <b>file</b> , <b>console</b> , or <b>user</b> statement at the [edit system syslog] hierarchy level, but not to destinations configured by a <b>host</b> statement. |
| <b>Default</b>                                                                                                                                                                                                                                                                     | The timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, <b>Aug 21 12:36:30</b> .                                                                                                                                                                                                                                            |
| <div>  <p><b>NOTE:</b> When the <b>structured-data</b> statement is included at the [edit system syslog file <i>filename</i>] hierarchy level, this statement is ignored for the file.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                                                                                                                                                                                                                                                                     | <p><b>millisecond</b>—Include the millisecond in the timestamp.</p> <p><b>year</b>—Include the year in the timestamp.</p>                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                    | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                       | <ul style="list-style-type: none"> <li>• <a href="#">Including the Year or Millisecond in Timestamps on page 181</a></li> <li>• <a href="#">structured-data on page 312</a></li> </ul>                                                                                                                                                                                         |

## time-zone

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>time-zone (GMT <i>hour-offset</i>   <i>time-zone</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>     | [edit system]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | Set the local time zone. To have the time zone change take effect for all processes running on the switch, you must reboot the switch.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Default</b>             | UTC                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>             | <p><b>GMT <i>hour-offset</i></b>—Set the time zone relative to UTC time.</p> <p><b>Range:</b> -14 through +12</p> <p><b>Default:</b> 0</p> <p><b><i>time-zone</i></b>—Specify the time zone as <b>UTC</b>, which is the default time zone, or as a string such as PDT (Pacific Daylight Time), or use one of the following continents and major cities:</p> <p>Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako, Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali, Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi, Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu, Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis, Africa/Windhoek</p> <p>America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires, America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit, America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung, America/Paramaribo, America/Phoenix, America/Port-au-Prince, America/Port_of_Spain, America/Porto_Acre, America/Puerto_Rico, America/Rainy_River, America/Rankin_Inlet, America/Regina, America/Rosario, America/Santiago,</p> |

America/Santo\_Domingo, America/Sao\_Paulo, America/Scoresbysund,  
 America/Shiprock, America/St\_Johns, America/St\_Kitts, America/St\_Lucia,  
 America/St\_Thomas, America/St\_Vincent, America/Swift\_Current, America/Tegucigalpa,  
 America/Thule, America/Thunder\_Bay, America/Tijuana, America/Tortola,  
 America/Vancouver, America/Whitehorse, America/Winnipeg, America/Yakutat,  
 America/Yellowknife  
 Antarctica/Casey, Antarctica/DumontDURville, Antarctica/Mawson, Antarctica/McMurdo,  
 Antarctica/Palmer, Antarctica/South\_Pole  
 Arctic/Longyearbyen  
 Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqtau, Asia/Aqtobe,  
 Asia/Ashkhabad, Asia/Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut,  
 Asia/Bishkek, Asia/Brunei, Asia/Calcutta, Asia/Chungking, Asia/Colombo, Asia/Dacca,  
 Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/Harbin, Asia/Hong\_Kong,  
 Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/Jerusalem, Asia/Kabul,  
 Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Krasnoyarsk,  
 Asia/Kuala\_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/Manila,  
 Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom\_Penh,  
 Asia/Pyongyang, Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul,  
 Asia/Shanghai, Asia/Singapore, Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran,  
 Asia/Thimbu, Asia/Tokyo, Asia/Ujung\_Pandang, Asia/Ulan\_Bator, Asia/Urumqi,  
 Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk, Asia/Yekaterinburg, Asia/Yerevan  
 Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape\_Verde, Atlantic/Faeroe,  
 Atlantic/Jan\_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South\_Georgia,  
 Atlantic/St\_Helena, Atlantic/Stanley  
 Australia/Adelaide, Australia/Brisbane, Australia/Broken\_Hill, Australia/Darwin,  
 Australia/Hobart, Australia/Lindeman, Australia/Lord\_Howe, Australia/Melbourne,  
 Australia/Perth, Australia/Sydney  
 Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade,  
 Europe/Berlin, Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest,  
 Europe/Chisinau, Europe/Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki,  
 Europe/Istanbul, Europe/Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana,  
 Europe/London, Europe/Luxembourg, Europe/Madrid, Europe/Malta, Europe/Minsk,  
 Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/Paris, Europe/Prague,  
 Europe/Riga, Europe/Rome, Europe/Samara, Europe/San\_Marino, Europe/Sarajevo,  
 Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn,  
 Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius,  
 Europe/Warsaw, Europe/Zagreb, Europe/Zurich  
 Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro,  
 Indian/Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte,  
 Indian/Reunion  
 Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate,  
 Pacific/Enderbury, Pacific/Fakaofo, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos,  
 Pacific/Gambier, Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston,  
 Pacific/Kiritimati, Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas,  
 Pacific/Midway, Pacific/Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea,  
 Pacific/Pago\_Pago, Pacific/Palau, Pacific/Pitcairn, Pacific/Ponape, Pacific/Port\_Moresby,  
 Pacific/Rarotonga, Pacific/Saipan, Pacific/Tahiti, Pacific/Tarawa, Pacific/Tongatapu,  
 Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/Yap

**Required Privilege Level** system—To view this statement in the configuration.  
 system-control—To add this statement to the configuration.

**Related Documentation** • [Modifying the Default Time Zone for a Router or Switch Running Junos OS on page 186](#)

## traceoptions (Commit Scripts)

---

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntax              | <pre>traceoptions {<br/>    file &lt;filename&gt; &lt;files number&gt; &lt;size size&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag flag;<br/>    no-remote-trace;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Hierarchy Level     | [edit system scripts commit],                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Description         | Define tracing operations for commit or op scripts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Default             | If you do not include this statement, no script-specific tracing operations are performed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Options             | <p><b>filename</b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <code>/var/log</code>. By default, commit script process tracing output is placed in the file <code>cscript.log</code> and op script process tracing is placed in the file <code>op-script.log</code>. If you include the <b>file</b> statement, you must specify a filename. To retain the default, you can specify <code>cscript.log</code> or <code>op-script.log</code> as the filename.</p> <p><b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed and compressed to <i>trace-file.0.gz</i>. When <i>trace-file</i> again reaches its maximum size, <i>trace-file.0.gz</i> is renamed <i>trace-file.1.gz</i> and <i>trace-file</i> is renamed and compressed to <i>trace-file.0.gz</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 10 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—Log all operations</li><li>• <b>events</b>—Log important events</li><li>• <b>input</b>—Log script input data</li><li>• <b>offline</b>—Generate data for offline development</li><li>• <b>output</b>—Log script output data</li><li>• <b>rpc</b>—Log script RPCs</li><li>• <b>xslt</b>—Log the XSLT library</li></ul> <p><b>no-world-readable</b>—Restrict file access to owner. This is the default.</p> |



**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed and compressed to **trace-file.0.gz**. When **trace-file** again reaches its maximum size, **trace-file.0.gz** is renamed **trace-file.1.gz** and **trace-file** is renamed and compressed to **trace-file.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—Enable unrestricted file access.

|                                 |                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | maintenance—To view this statement in the configuration.<br>maintenance-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Tracing Commit Script Processing</i></li> </ul>                                 |

## traceoptions (Layer 2 Learning)

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre>traceoptions {<br/>    file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;<br/>    flag <i>flag</i> (detail   disable   receive   send);<br/>    in-memory-debug;<br/>    level;<br/>    no-remote-trace;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>     | [edit protocols l2-learning]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | Define tracing operations for Layer 2 learning.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>             | The <b>traceoptions</b> feature is disabled by default.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p>You can specify the following options:</p> <ul style="list-style-type: none"><li>• <b>no-world-readable</b>—(Optional) Restrict file access to the user who created the file.</li><li>• <b>size <i>size</i></b> —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>files</b> option. Use <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes.</li><li>• <b>world-readable</b>—(Optional) Enable unrestricted file access.</li></ul> <p><b>flag <i>flag</i></b> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li><li>• <b>bmac-next-hop</b>—Trace backbone MAC next hop operations.</li><li>• <b>bridge-bmac-next-hop</b>—Trace backbone MAC next hop bridge operations.</li><li>• <b>bridging-interface</b>—Trace interface bridge operations.</li><li>• <b>bridging-domain</b>—Trace bridging domain operations.</li><li>• <b>configuration</b>—Trace configuration operations.</li><li>• <b>flood-next-hop</b>—Trace flood next hop operations.</li><li>• <b>initialization</b>—Trace initialization operations.</li><li>• <b>interface-device</b>—Trace interface device operations.</li><li>• <b>interface-family</b>—Trace interface family operations.</li></ul> |

- **interface-logical**—Trace logical interface operations.
- **ipc**—Trace inter-process communications operations.
- **irb**—Trace integrated routing and bridging operations.
- **isid**—Trace i-tagged service ID operations.
- **kack**—Trace kernel-acknowledgment.
- **learning-domain**—Trace learning domain operations.
- **logical-system**—Trace logical system operations.
- **mac-learning**—Trace MAC address learning.
- **mc-ae**—Trace multichassis aggregated Ethernet interface operations.
- **redundant-trunk-group**—Trace redundant trunk group operations.
- **routing-instance**—Trace routing instance operations.
- **routing-socket**—Trace routing socket operations.
- **storm-control**—Trace storm control operations.
- **unknown-unicast-forwarding**—Trace unknown unicast forwarding events.
- **vpls-ping**—Trace Virtual Private VLAN Service (VPLS) ping operations.

**in-memory-debug**—Enable trace parameters in the memory.


**level**—Specify level of debugging output.

**no-remote-trace**—Disable remote tracing.

|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

## tracing

---

|                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                             | <pre>tracing {<br/>    destination-override syslog host <i>ip-address</i>;<br/>}</pre>                                                                                                 |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                    | [edit system]                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                      |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                        | Configure the switch to enable remote tracing to a specified host IP address.                                                                                                          |
| <hr/>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                        |
| <div> <b>NOTE:</b> The <code>tracing</code> statement is not supported on the QFX3000 QFabric system.</div> <hr/>                                                                                                                                                                                                                                                        |                                                                                                                                                                                        |
| <p>The following processes are supported:</p> <ul style="list-style-type: none"><li>• <b>chassisd</b>—Chassis-control process</li><li>• <b>eventd</b>—Event-processing process</li><li>• <b>cosd</b>—Class-of-service process</li></ul> <p>If you enabled remote tracing but wish to disable it for specific processes on the switch, use the <b>no-remote-trace</b> statement at the [edit system <i>process-name</i> traceoptions] hierarchy level.</p> |                                                                                                                                                                                        |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                            | Remote tracing is disabled by default.                                                                                                                                                 |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>destination-override syslog host <i>ip-address</i></b> —Overrides the global configuration for system tracing and has no effect if the <b>tracing</b> statement is not configured.  |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                           | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                              | <ul style="list-style-type: none"><li>• <a href="#">Understanding Tracing and Logging Operations on page 6468</a></li><li>• <a href="#">destination-override on page 271</a></li></ul> |

## trusted-key

---

|                                 |                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>trusted-key [ <i>key-numbers</i> ];</code>                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit system <a href="#">ntp</a> ]                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                      |
| <b>Description</b>              | For NTP, configure the keys to use when you configure the switch to synchronize its time with other systems on the network.                                                                            |
| <b>Options</b>                  | <i>key-numbers</i> —One or more key numbers. Each key can be any 32-bit unsigned integer except 0.                                                                                                     |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring NTP Authentication Keys on page 165</a></li> <li>• <i>authentication-key</i></li> <li>• <a href="#">server on page 306</a></li> </ul> |

## uid

---

|                                 |                                                                                                                                                   |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>uid <i>uid-value</i>;</code>                                                                                                                |
| <b>Hierarchy Level</b>          | [edit system login user]                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                 |
| <b>Description</b>              | Configure a user identifier for a login account.                                                                                                  |
| <b>Options</b>                  | <i>uid-value</i> —Number associated with the login account. This value must be unique on the router or switch.<br><b>Range:</b> 100 through 64000 |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Junos OS User Accounts on page 1344</a></li> </ul>                               |

## use-imported-time-zones

---

|                                 |                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | use-imported-time-zones;                                                                                                            |
| <b>Hierarchy Level</b>          | [edit system]                                                                                                                       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                   |
| <b>Description</b>              | Configure a custom time zone from a locally generated time zone database.                                                           |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Updating the IANA Time Zone Database on Junos Devices on page 190</a></li></ul> |

## user (Access)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>user username {<br/>  authentication {<br/>    (encrypted-password "password"   plain-text-password);<br/>    load-key-file URL;<br/>    remote-debug-permission (qfabric-admin   qfabric-operator   qfabric-user);<br/>    ssh-dsa "public-key" &lt;from hostname&gt;;<br/>    ssh-rsa "public-key" &lt;from hostname&gt;;<br/>  }<br/>  class class-name;<br/>  full-name "complete-name";<br/>  uid uid-value;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit system login]                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | Configure access permission for individual users.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Junos OS User Accounts on page 1344</a></li><li>• <a href="#">class on page 262</a></li></ul>                                                                                                                                                                                                                                                                      |

## CHAPTER 5

# Administration

- [Routine Monitoring on page 333](#)
- [Operational Commands on page 342](#)

### Routine Monitoring

---

- [Monitoring System Process Information on page 333](#)
- [Monitoring System Properties on page 334](#)
- [Monitoring Interface Status and Traffic on page 335](#)
- [Monitoring Zero Touch Provisioning on page 336](#)
- [Other Tools to Configure and Monitor Devices Running Junos OS on page 339](#)
- [Verifying a Unified In-Service Software Upgrade on page 339](#)
- [Verifying Autoinstallation Status on page 340](#)
- [Verifying That Automatic Software Download Is Working Correctly on page 341](#)

### Monitoring System Process Information

**Purpose** View the processes running on the device.

**Action** To view the software processes running on the device:  
[edit system]  
  
user@switch> [show system processes](#)

**Meaning** [Table 22 on page 333](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

**Table 22: Summary of System Process Information Output Fields**

| Field | Values                     |
|-------|----------------------------|
| PID   | Identifier of the process. |
| Name  | Owner of the process.      |

Table 22: Summary of System Process Information Output Fields (*continued*)

| Field              | Values                                                   |
|--------------------|----------------------------------------------------------|
| State              | Current state of the process.                            |
| CPU Load           | Percentage of the CPU that is being used by the process. |
| Memory Utilization | Amount of memory that is being used by the process.      |
| Start Time         | Time of day when the process started.                    |

- Related Documentation**
- [Monitoring System Properties on page 334](#)
  - [show system uptime on page 1137](#)

## Monitoring System Properties

**Purpose** View system properties such as the name, IP address, and resource usage.

**Action** To monitor system properties in the CLI, enter the following commands:

- [show system uptime](#)
- [show system users](#)
- [show system storage](#)

**Meaning** [Table 23 on page 334](#) summarizes key output fields in the system properties display.

Table 23: Summary of Key System Properties Output Fields

| Field                      | Values                                                                                                  | Additional Information                                 |
|----------------------------|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------|
| <b>General Information</b> |                                                                                                         |                                                        |
| Serial Number              | Serial number of device.                                                                                |                                                        |
| Junos OS Version           | Version of Junos OS active on the switch, including whether the software is for domestic or export use. | Export software is for use outside the USA and Canada. |
| Hostname                   | Name of the device.                                                                                     |                                                        |
| IP Address                 | IP address of the device.                                                                               |                                                        |
| Loopback Address           | Loopback address.                                                                                       |                                                        |
| Domain Name Server         | Address of the domain name server.                                                                      |                                                        |
| Time Zone                  | Time zone on the device.                                                                                |                                                        |



Table 23: Summary of Key System Properties Output Fields (*continued*)

| Field                          | Values                                                                                                                                       | Additional Information                                                           |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| <b>Time</b>                    |                                                                                                                                              |                                                                                  |
| Current Time                   | Current system time, in Coordinated Universal Time (UTC).                                                                                    |                                                                                  |
| System Booted Time             | Date and time when the device was last booted and how long it has been running.                                                              |                                                                                  |
| Protocol Started Time          | Date and time when the protocols were last started and how long they have been running.                                                      |                                                                                  |
| Last Configured Time           | Date and time when a configuration was last committed. This field also shows the name of the user who issued the last <b>commit</b> command. |                                                                                  |
| Load Average                   | CPU load average for 1, 5, and 15 minutes.                                                                                                   |                                                                                  |
| <b>Storage Media</b>           |                                                                                                                                              |                                                                                  |
| Internal Flash Memory          | Usage details of internal flash memory.                                                                                                      |                                                                                  |
| External Flash Memory          | Usage details of external USB flash memory.                                                                                                  |                                                                                  |
| <b>Logged in Users Details</b> |                                                                                                                                              |                                                                                  |
| User                           | Username of any user logged in to the switch.                                                                                                |                                                                                  |
| Terminal                       | Terminal through which the user is logged in.                                                                                                |                                                                                  |
| From                           | System from which the user has logged in. A hyphen indicates that the user is logged in through the console.                                 |                                                                                  |
| Login Time                     | Time when the user logged in.                                                                                                                | This is the <b>user@switch</b> field in <b>show system users</b> command output. |
| Idle Time                      | How long the user has been idle.                                                                                                             |                                                                                  |

- Related Documentation**
- [Monitoring System Process Information on page 333](#)
  - [show system processes on page 1051](#)

## Monitoring Interface Status and Traffic

**Purpose** View interface status to monitor interface bandwidth utilization and traffic statistics.

**Action** • To view interface status for all the interfaces, enter [show interfaces xe](#).

- To view status and statistics for a specific interface, enter **show interfaces xe interface-name**.
- To view status and traffic statistics for all interfaces, enter either **show interfaces xe detail** or **show interfaces xe extensive**.

**Meaning** For details about output from the CLI commands, see **show interfaces xe**.

## Monitoring Zero Touch Provisioning

You can use the console and operational commands to monitor Zero Touch Provisioning.

1. [Using the Console to Monitor Zero Touch Provisioning on page 336](#)
2. [Using System Log Alerts to Monitor Zero Touch Provisioning on page 336](#)
3. [Using Error Messages to Monitor Zero Touch Provisioning on page 337](#)
4. [Using System Log Files to Monitor Zero Touch Provisioning on page 337](#)
5. [Using the show dhcp client binding Command on page 338](#)
6. [Using the show dhcp client statistics Command on page 338](#)

### Using the Console to Monitor Zero Touch Provisioning

---

The following Zero Touch Provisioning (ZTP) activities are displayed on the console during the ZTP process:

- Starting and ending times of ZTP process.
- Lists of bound and unbound DHCP client interfaces.
- DHCP options that DHCP servers send to DHCP clients.
- Logs indicating which interfaces are used for ZTP.
- ZTP parameters that DHCP clients obtain from DHCP servers.
- File names of configuration and image files, names of file servers, protocols used to fetch files, and times when DHCP servers fetch configuration and image files.
- Failure states caused by files not being on servers, or unreachable servers, and time outs.
- Number of attempts made, and number of attempts remaining, for retry in current ZTP cycle.
- Completion of file transfers.
- Installation, reboot, and state of ZTP process.
- Internal state errors and termination of ZTP process.
- Logs for when default routes were added or deleted.

### Using System Log Alerts to Monitor Zero Touch Provisioning

---

**Purpose** In this example, the system log alert alerts you that the auto-image upgrade will start.

**Action** Use the following system log alert to monitor the auto-image upgrade process.

"ALERT:Auto-image upgrade will start. This can terminate config CLI session(s). Modified configuration will be lost. To stop Auto-image, in CLI do the following: 'edit; delete chassis auto-image-upgrade; commit'."

"Checking whether image upgrade is already invoked"

**Meaning** This system log alert indicates that the auto-image upgrade will start, and provides information on how to stop the auto-image upgrade process.

### Using Error Messages to Monitor Zero Touch Provisioning

**Purpose** Error messages provide information on which DHCP options are not configured.

**Action** Use the information in the following error message to find out which DHCP options are not configured.

"DHCP Log Server Option"  
"DHCP Host Name Option"  
"DHCP NTP Server Option"

**Meaning** The error message indicates that the DHCP log server, hostname, and NTP server options are not configured.

### Using System Log Files to Monitor Zero Touch Provisioning

**Purpose** System log files provide information on the state of the auto-upgrade process, lists of bound and unbound DHCP client interfaces, IP addresses of file servers, names and locations of image and configuration files, and successful and failed attempts at fetching configuration and image files.

**Action** Use the information in the following system log files to monitor the auto-upgrade process.

Auto Image Upgrade: Start fetching config-file file from server 1.1.1.1 through irb using ftp

Auto Image Upgrade: Tried [2] attempts to fetch config-file file from server 1.1.1.1 through irb. Summary: "Retrieving /config-file  
:: Failed to open file.". To retry [4] times.

Auto Image Upgrade: Tried [4] attempts to fetch config-file file from server 1.1.1.1 through irb. Summary: "Retrieving /config-fileconfig-file  
:: Failed to open file.". To retry [2] times.

Auto Image Upgrade: Tried [6] attempts to fetch config-file file from server 1.1.1.1 through irb. Summary: "Retrieving /config-file  
:: Failed to open file.". To retry [0] times.

Auto Image Upgrade: All [6] attempts to fetch config-file file from server 1.1.1.1 through irb FAILED. Start retry again in few minutes.

**Meaning** These system log files indicate that there were six failed attempts to fetch the configuration file from the file server, the IP address of the file server, the DHCP client interface name, and the number of times the retry process occurred.

### Using the show dhcp client binding Command

**Purpose** Issue the **show dhcp client binding** command to display DHCP client binding information

**Action** Issue the **show dhcp client binding** command to display the IP address of the DHCP client, the hardware address of the DHCP client, number of seconds in which the DHCP client's IP address lease expires, state of the DHCP client IP address in the binding table, and the name of the interface that has active client bindings.

#### **show dhcp client binding**

```
user@switch# show dhcp client binding
IP address      Hardware address Expires    State      Interface
0.0.0.0         00:22:83:2a:db:dc 0          SELECTING  irb.0
6.6.6.13        00:22:83:2a:db:dd 49201      BOUND      vme.0
0.0.0.0         00:22:83:2a:db:df 0          SELECTING  xe-0/0/0.0
0.0.0.0         00:22:83:2a:db:e0 0          SELECTING  xe-0/0/1.0
```

**Meaning** The output of this command shows that there is one client interface that is bound, and that there are three interfaces that are receiving DHCP offers from the DHCP server.

### Using the show dhcp client statistics Command

**Purpose** Issue the **show dhcp client statistics** command to display DHCP client statistics.

**Action** Issue the **show dhcp client statistics** command to display DHCP client statistics, such as the number of packets dropped, and the number DHCP and BOOTP messages sent and received.

#### **show dhcp client statistics**

```
user@switch# show dhcp client statistics
Packets dropped:
  Total          14
  Send error     14
Messages received:
  BOOTREPLY      5
  DHCPOFFER      1
  DHCPACK        4
  DHCPNAK        0
  DHCPFORCERENEW 0
Messages sent:
  BOOTREQUEST    6751
  DHCPDECLINE    0
  DHCPDISCOVER   6747
  DHCPREQUEST    4
  DHCPINFORM     0
  DHCPRELEASE    0
  DHCPRENEW      0
  DHCPREBIND     0
```

**Meaning** The output of this command displays how many packets were dropped with errors, the number of BOOTREPLY and DHCP OFFER messages that were received, and the number of BOOTREQUEST and DHCPREQUEST messages that were sent.

**Related Documentation**

- [Understanding Zero Touch Provisioning on page 32](#)
- [Configuring Zero Touch Provisioning on page 87](#)

## Other Tools to Configure and Monitor Devices Running Junos OS

Apart from the command-line interface, Junos OS also supports the following applications, scripts, and utilities that enable you to configure and monitor devices running Junos OS:

- Junos XML Management Protocol Application Programming Interface (API)—Application programmers can use the Junos XML Management Protocol API to monitor and configure Juniper Networks devices. Juniper Networks provides a Perl module with the API to help you more quickly and easily develop custom Perl scripts for configuring and monitoring the devices.
- NETCONF Application Programming Interface (API)—Application programmers can also use the NETCONF API to monitor and configure Juniper Networks devices.
- Junos OS commit scripts—You can define scripts to enforce custom configuration tasks, enforce consistency, prevent common mistakes, and more. Every time you commit a new candidate configuration, the active commit scripts are called to inspect the new candidate configuration. If a configuration violates your custom rules, the script can instruct the Junos OS to perform various actions, including making changes to the configuration and generating custom, warning, and system log messages.
- Junos OS Op scripts—You can add your own commands to the operation-mode CLI. You can use these scripts to automate troubleshooting of known network problems and correct them.
- Junos OS event scripts—You can use event scripts to diagnose and fix issues, monitor the overall status of the system, and examine errors periodically. Event scripts are similar to op scripts except that certain events on the switch will trigger these scripts.
- Junos Space—The Junos Space application design allows multiple users concurrent access to its user interface. It also includes applications for network infrastructure automation.

**Related Documentation**

- [CLI User Interface Overview on page 39](#)
- [QFX Series Software Features Overview](#)
- [NETCONF XML Management Protocol Developer Guide](#)
- [Understanding Device and Network Management Features on page 6463](#)

## Verifying a Unified In-Service Software Upgrade

**Purpose** Verify the status of FPCs and their corresponding PICs after the most recent unified ISSU.

**Action** Issue the **show chassis in-service-upgrade** command on the master Routing Engine:

```
user@host> show chassis in-service-upgrade
Item           Status           Reason
FPC 0          Online
FPC 1          Online
FPC 2          Online
  PIC 0        Online
  PIC 1        Online
FPC 3          Offline        Offlined by CLI command
FPC 4          Online
  PIC 1        Online
FPC 5          Online
  PIC 0        Online
FPC 6          Online
  PIC 3        Online
FPC 7          Online
```

**Meaning** See [show chassis in-service-upgrade](#) for more information.

- Related Documentation**
- *Performing a Unified ISSU*
  - *Troubleshooting Unified ISSU Problems*
  - *Managing and Tracing BFD Sessions During Unified ISSU Procedures*

## Verifying Autoinstallation Status

**Purpose** Display the status of the autoinstallation feature.

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

## Sample Output

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

```

Acquired address: None
Protocol: RARP Client
Acquired address: None

```

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

**Related Documentation**

- [Configuring Autoinstallation of Configuration Files \(CLI Procedure\) on page 152](#)

## Verifying That Automatic Software Download Is Working Correctly

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

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

```

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

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

```

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

**Related Documentation**

- [Upgrading Software by Using Automatic Software Download on page 148](#)
- [Understanding DHCP Services for Switches on page 21](#)

## Operational Commands

---

- `commit`
- `clear log`
- `clear chassis display message`
- `clear system commit`
- `clear system reboot`
- `file`
- `file archive`
- `file checksum md5`
- `file checksum sha1`
- `file checksum sha-256`
- `file compare`
- `file delete`
- `file list`
- `file rename`
- `file show`
- `load`
- `ping`
- `request chassis beacon`
- `request chassis fpc`
- `request chassis pic`
- `request chassis routing-engine master`
- `request message`
- `request system configuration rescue delete`
- `request system configuration rescue save`
- `request system halt`
- `request system license add`
- `request system license delete`
- `request system license save`
- `request system logout`
- `request system power-off`
- `request system reboot`
- `request system snapshot`
- `request system software add`
- `request system software delete`
- `request system software download`



- request system software in-service-upgrade
- request system software nonstop-upgrade
- request system software rollback
- request system software validate
- request system storage cleanup
- request system zeroize
- restart
- rollback
- save
- show chassis alarms
- show chassis beacon
- show chassis environment
- show chassis environment fpc
- show chassis environment pem
- show chassis environment routing-engine
- show chassis fan
- show chassis firmware
- show chassis fpc
- show chassis hardware
- show chassis in-service-upgrade
- show chassis lcd
- show chassis led
- show chassis location
- show chassis mac-addresses
- show chassis nonstop-upgrade
- show chassis pic
- show chassis routing-engine
- show chassis zones
- show cli
- show cli authorization
- show cli directory
- show cli history
- show host
- show interfaces diagnostics optics
- show log
- show ntp associations
- show ntp status

- [show subscribers](#)
- [show system alarms](#)
- [show system audit](#)
- [show system boot-messages](#)
- [show system buffers](#)
- [show system certificate](#)
- [show system commit](#)
- [show system configuration archival](#)
- [show system configuration rescue](#)
- [show system connections](#)
- [show system core-dumps](#)
- [show system directory-usage](#)
- [show system license](#)
- [show system processes](#)
- [show system reboot](#)
- [show system resource-cleanup processes](#)
- [show system rollback](#)
- [show system services service-deployment](#)
- [show system software](#)
- [show system statistics](#)
- [show system storage](#)
- [show system uptime](#)
- [show system users](#)
- [show system virtual-memory](#)
- [show version](#)
- [start shell](#)
- [test configuration](#)
- [traceroute](#)
- [traceroute monitor](#)

## commit

**Syntax** `commit <<at <"string">> <and-quit> <check> <comment <"comment-string">>  
<confirmed> <display detail> <fast-synchronize> <minutes>  
<synchronize <force> <scripts>>`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 11.1 for the QFX Series.  
Option **fast-synchronize** added in Junos OS Release 12.2.  
Option **synchronize scripts** introduced in Junos OS Release 13.2.

**Description** Commit the set of changes to the database and cause the changes to take operational effect.



**NOTE:** The **fast-synchronize** option is not supported in a QFX Series Virtual Chassis.



**NOTE:** Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot fpc-slot offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

**Options** `at <"string">`—(Optional) Save software configuration changes and activate the configuration at a future time, or upon reboot.

**string** is **reboot** or the future time to activate the configuration changes. Enclose the **string** value (including **reboot**) in quotation marks (" "). You can specify time in two formats:

- A time value in the form **hh:mm[:ss]** (hours, minutes, and optionally seconds)—Commit the configuration at the specified time, which must be in the future but before 11:59:59 PM on the day the **commit at** configuration command is issued. Use 24-hour time for the **hh** value; for example, **04:30:00** is 4:30:00 AM, and **20:00** is 8:00 PM. The time is interpreted with respect to the clock and time zone settings on the router.
- A date and time value in the form **yyyy-mm-dd hh:mm[:ss]** (year, month, date, hours, minutes, and, optionally, seconds)—Commit the configuration at the specified day and time, which must be after the **commit at** command is issued. Use 24-hour time for the **hh** value. For example, **2003-08-21 12:30:00** is 12:30 PM on August 21, 2003. The time is interpreted with respect to the clock and time zone settings on the router.

For example, **commit at "18:00:00"**. For date and time, include both values in the same set of quotation marks. For example, **commit at "2005-03-10 14:00:00"**.

A *commit check* is performed when you issue the **commit at** configuration mode command. If the result of the check is successful, then the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit can be performed until the scheduled commit is completed.



**NOTE:** If Junos OS fails before the configuration changes become active, all configuration changes are lost.

You cannot enter the **commit at** configuration command when there is a pending reboot.

You cannot enter the **request system reboot** command once you schedule a commit operation for a specific time in the future.

You cannot commit a configuration when a scheduled commit is pending. For information about how to use the **clear** command to cancel a scheduled configuration, see the [CLI Explorer](#).

**and-quit**—(Optional) Commit the configuration and, if the configuration contains no errors and the commit succeeds, exit from configuration mode.

**check**—(Optional) Verify the syntax of the configuration, but do not activate it.

**comment** <"*comment-string*">—(Optional) Add a comment that describes the committed configuration. The comment can be as long as 512 bytes and must be typed on a single line. You cannot include a comment with the **commit check** command. Enclose *comment-string* in quotation marks (" "). For example, **commit comment "Includes changes recommended by SW Lab"**.

**confirmed** <*minutes*>—(Optional) Require that the commit be confirmed within the specified amount of time. To confirm a commit, enter either a **commit** or **commit check** command. If the commit is not confirmed within the time limit, the configuration rolls back automatically to the precommit configuration and a broadcast message is sent to all logged-in users. To show when a rollback is scheduled, enter the **show system commit** command. The allowed range is 1 through 65,535 minutes, and the default is 10 minutes.

In Junos OS Release 11.4 and later, you can also use the **commit confirmed** command in the **[edit private]** configuration mode.

**display detail**—(Optional) Monitors the commit process.



**NOTE:** In Junos OS Release 10.4 and later, if the number of commit details or messages exceeds a page when used with the **| display detail** pipe option, the more pagination option on the screen is no longer available. Instead, the messages roll up on the screen by default, just like using the **commit** command with the **| no more** pipe option.

**fast-synchronize**—(Optional) Configure the commits to run in parallel on both the master and backup Routing Engines to reduce the time taken for commit synchronization.



**NOTE:** The **fast-synchronize** statement is not supported on QFX Series devices when used in a Virtual Chassis.

**synchronize <force> <scripts>**—(Optional) If your router has two Routing Engines, you can manually direct one Routing Engine to synchronize its configuration with the other by issuing the **commit synchronize** command. The Routing Engine on which you execute this command (request Routing Engine) copies and loads its candidate configuration to the other (responding Routing Engine). Both Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines. The **commit synchronize** command does not work if the responding Routing Engine has uncommitted configuration changes. However, you can enforce commit synchronization on the Routing Engines by using the **force** option. When you issue the **commit synchronize** command with the **force** option from one Routing Engine, the configuration sessions on the other Routing Engine are terminated and its configuration synchronized with that on the Routing Engine from which you issued the command.

When you issue the **commit synchronize** command with the **scripts** option, the device synchronizes all commit, event, lib, and op scripts from the requesting Routing Engine to the responding Routing Engine and also commits and synchronizes the configuration. If the commit check operation fails for the requesting Routing Engine, the process stops, and the scripts are not copied to the responding Routing Engine. If the commit check or commit operation fails for the responding Routing Engine, the scripts are still synchronized, since the synchronization occurs prior to the commit check operation on the responding Routing Engine.

If the **load-scripts-from-flash** statement is configured for the requesting Routing Engine, the device synchronizes the scripts from flash memory on the requesting Routing Engine to flash memory on the responding Routing Engine. Otherwise, the device synchronizes the scripts from the hard disk on the requesting Routing Engine to the hard disk on the responding Routing Engine. The device synchronizes all scripts regardless of whether they are enabled in the configuration or have been updated since the last synchronization.



**NOTE:** When you issue the **commit synchronize** command, you must use the **apply-groups re0** and **re1** commands. For information about how to use groups, see *Disabling Inheritance of a Junos OS Configuration Group*.

The responding Routing Engine must use Junos OS Release 5.0 or later.

**Required Privilege Level**

**configure**—To enter configuration mode.



**NOTE:** If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*

#### Related Documentation

- *Verifying a Junos OS Configuration, Committing a Junos OS Configuration*
- *Scheduling a Junos OS Commit Operation*
- *Deactivating and Reactivating Statements and Identifiers in a Junos OS Configuration*
- *Monitoring the Junos OS Commit Process*
- *Adding a Comment to Describe the Committed Configuration*

## Sample Output

### commit | display detail

```
user@host> commit | display detail
-----
2011-08-24 01:08:08.00691 PDT: begin creating snapshots
2011-08-24 01:08:09.00210 PDT: end creating snapshots
2011-08-24 01:08:09.00211 PDT: begin preparing metadata
2011-08-24 01:08:09.00228 PDT: end preparing metadata
2011-08-24 01:08:09.00229 PDT: begin computing dcf root changes
2011-08-24 01:08:09.00236 PDT: end computing dcf root changes
2011-08-24 01:08:09.00244 PDT: begin computing additions
2011-08-24 01:08:09.00251 PDT: end computing additions
2011-08-24 01:08:09.00251 PDT: begin local object validation
2011-08-24 01:08:09.00251 PDT: end local object validation
2011-08-24 01:08:09.00252 PDT: begin update instances
2011-08-24 01:08:09.00252 PDT: end update instances
2011-08-24 01:08:09.00252 PDT: begin adjust metadata
2011-08-24 01:08:09.00252 PDT: end adjust metadata
2011-08-24 01:08:09.00253 PDT: begin validate metadata
2011-08-24 01:08:09.00253 PDT: end validate metadata
2011-08-24 01:08:09.00253 PDT: begin adjust allocations
2011-08-24 01:08:09.00254 PDT: end adjust allocations
2011-08-24 01:08:09.00254 PDT: begin adjust dependencies
2011-08-24 01:08:09.00254 PDT: end adjust dependencies
2011-08-24 01:08:09.00255 PDT: begin instance validation
2011-08-24 01:08:09.00255 PDT: end instance validation
2011-08-24 01:08:09.00255 PDT: begin opening all sessions eagerly
2011-08-24 01:08:09.00277 PDT: begin request #1 [login]
2011-08-24 01:08:09.00278 PDT: end request #1 [login]
2011-08-24 01:08:09.00325 PDT: begin processing globals
2011-08-24 01:08:09.00330 PDT: begin waiting for stamp check
```

```
(qfabric-default---node0)
2011-08-24 01:08:09.00334 PDT: end reply #1 [login]
2011-08-24 01:08:09.00351 PDT: end reply #1 [login]
2011-08-24 01:08:09.00451 PDT: begin request #2 [open]
2011-08-24 01:08:09.00451 PDT: end request #2 [open]
2011-08-24 01:08:09.00451 PDT: begin request #3 [get commit history]
2011-08-24 01:08:09.00452 PDT: end request #3 [get commit history]
2011-08-24 01:08:09.00452 PDT: begin request #4 [load]
2011-08-24 01:08:09.00453 PDT: end request #4 [load]
2011-08-24 01:08:09.00453 PDT: begin request #5 [load]
2011-08-24 01:08:09.00454 PDT: begin reply #2 [open]
2011-08-24 01:08:09.00456 PDT: end reply #2 [open]
2011-08-24 01:08:09.00457 PDT: begin reply #3 [get commit history]
2011-08-24 01:08:09.00475 PDT: end reply #3 [get commit history]
2011-08-24 01:08:09.00476 PDT: begin reply #4 [load]
2011-08-24 01:08:09.00499 PDT: begin reply #5 [load]
2011-08-24 01:08:09.00501 PDT: end waiting for stamp check
(qfabric-default---node0)
2011-08-24 01:08:09.00501 PDT: begin waiting for open (qfabric-default---node0)
2011-08-24 01:08:09.00502 PDT: end waiting for open (qfabric-default---node0)
2011-08-24 01:08:09.00504 PDT: end processing globals
2011-08-24 01:08:09.00617 PDT: end request #5 [load]
2011-08-24 01:08:09.00617 PDT: begin request #6 [check]
2011-08-24 01:08:09.00617 PDT: end request #6 [check]
2011-08-24 01:08:09.00619 PDT: end reply #5 [load]
2011-08-24 01:08:09.00619 PDT: begin reply #6 [check]
2011-08-24 01:08:09.00730 PDT: end session
2011-08-24 01:08:09.00752 PDT: end request #5 [load]
2011-08-24 01:08:09.00754 PDT: begin request #6 [check]
2011-08-24 01:08:09.00755 PDT: end request #6 [check]
2011-08-24 01:08:09.00881 PDT: end request #5 [load]
2011-08-24 01:08:09.00961 PDT: begin commit to devices
2011-08-24 01:08:10.00668 PDT: begin request #8 [get commit history]
2011-08-24 01:08:10.00669 PDT: end request #8 [get commit history]
2011-08-24 01:08:10.00721 PDT: end session
2011-08-24 01:08:10.00727 PDT: end commit to devices
2011-08-24 01:08:10.00733 PDT: begin committing metadata
2011-08-24 01:08:10.00772 PDT: end committing metadata
2011-08-24 01:08:10.00772 PDT: begin calling commit callbacks
2011-08-24 01:08:10.00773 PDT: end calling commit callbacks
commit complete
```

## clear log

---

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>clear log <i>filename</i></code><br><code>&lt;all&gt;</code>                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Remove contents of a log file.                                                                                                                                                           |
| <b>Options</b>                  | <i>filename</i> —Name of the specific log file to delete.<br><br><code>all</code> —(Optional) Delete the specified log file and all archived versions of it.                             |
| <b>Required Privilege Level</b> | clear                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show log on page 948</a></li></ul>                                                                                                   |
| <b>List of Sample Output</b>    | <a href="#">clear log on page 350</a>                                                                                                                                                    |
| <b>Output Fields</b>            | See <a href="#">file list</a> for an explanation of output fields.                                                                                                                       |

## Sample Output

### clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```



## clear chassis display message

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 351</a><br><a href="#">Syntax (TX Matrix Router) on page 351</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 351</a><br><a href="#">Syntax (QFabric Systems) on page 351</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax</b>                         | clear chassis display message                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Syntax (TX Matrix Router)</b>      | clear chassis display message<br><lcc <i>number</i>   scc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Syntax (TX Matrix Plus Router)</b> | clear chassis display message<br><lcc <i>number</i>   sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (QFabric Systems)</b>       | clear chassis display message<br><node-device <i>name</i>   interconnect-device <i>name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Release Information</b>            | <p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option for the TX Matrix Plus routers introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                    | <p>(M40e, M160, M320, T Series routers, EX Series, and QFabric systems only) Clear or stop a text message on the craft interface display, which is on the front of the router or switch or on the LCD panel display on the router or switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.</p>                                                                                               |
| <b>Options</b>                        | <p><b>none</b>—Clear or stop a text message on the craft interface display.</p> <p><b>interconnect-device <i>name</i></b>—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Interconnect device.</p> <p><b>lcc <i>number</i></b>—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> </ul> |

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**node-device *name***—(QFabric systems only) (Optional) On a QFabric system, clear or stop a text message on the LCD panel display on the specified Node device in a Node group.

**scc**—(TX Matrix routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Clear or stop a text message on the craft interface on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**Required Privilege Level**

clear

**Related Documentation**

- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- *set chassis display message*
- *show chassis craft-interface*

**List of Sample Output** [clear chassis display message on page 352](#)

**Output Fields** See *show chassis craft-interface* for an explanation of output fields.

## Sample Output

### clear chassis display message

The following example displays and then clears the text message on the craft interface display:

```
user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
Host OK LED:    On
Host fail LED:  Off
FPCs           0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
      +-----+
      |NOC contact Dusty |
      |(888) 526-1234    |
      +-----+

user@host> clear chassis display message

user@host> show chassis craft-interface
Red alarm:      LED off, relay off
Yellow alarm:   LED off, relay off
```

```
Host OK LED:  On
Host fail LED: Off
FPCs      0  1  2  3  4  5  6  7
-----
Green  ..  *..  *  *.
Red    .....
LCD screen:
+-----+
|host    |
|Up: 0+17:05:47|
|        |
|Temperature OK|
+-----+
```

## clear system commit

---

|                                 |                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | clear system commit                                                                                                                                                                                                |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                           |
| <b>Description</b>              | Clear any pending commit operation.                                                                                                                                                                                |
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | maintenance (or the actual user who scheduled the commit)                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show system commit on page 1001</a></li></ul>                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">clear system commit on page 354</a><br><a href="#">clear system commit (None Pending) on page 354</a><br><a href="#">clear system commit (User Does Not Have Required Privilege Level) on page 354</a> |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                              |

### Sample Output

#### clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

#### clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

#### clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

## clear system reboot

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 355</a><br><a href="#">Syntax (EX Series Switches) on page 355</a><br><a href="#">Syntax (TX Matrix Router) on page 355</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 355</a><br><a href="#">Syntax (QFX Series) on page 355</a>                                                                                                                                                                                                                                                                                                          |
| <b>Syntax</b>                         | clear system reboot<br><both-routing-engines>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Syntax (EX Series Switches)</b>    | clear system reboot<br><all-members><br><both-routing-engines><br><local><br><member <i>member-id</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (TX Matrix Router)</b>      | clear system reboot<br><both-routing-engines><br><all-chassis   all-lcc   lcc <i>number</i>   scc>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Syntax (TX Matrix Plus Router)</b> | clear system reboot<br><both-routing-engines><br><all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (QFX Series)</b>            | clear system reboot<br><infrastructure <i>name</i> ><br><interconnect-device <i>name</i> ><br><node-group <i>name</i> >                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>            | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                               |
| <b>Description</b>                    | Clear any pending system software reboots or halts. When issued on a TX Matrix router without any options, the default behavior clears all pending system software reboots or halts on all T640 routers connected to the TX Matrix router. When issued on a TX Matrix Plus router without any options, the default behavior clears all pending system software reboots or halts on all T1600 or T4000 routers connected to the TX Matrix Plus router.                                                                                                                            |
| <b>Options</b>                        | <p><b>none</b>—Clear all pending system software reboots or halts.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for all the Routing Engines in the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests on the l connected T1600 or T4000 LCCs.</p> |

**all-members**—(EX4200 switches only) (Optional) Clear all halt or reboot requests on all members of the Virtual Chassis configuration.

**both-routing-engines**—(Systems with multiple Routing Engines) (Optional) Clear all halt or reboot requests on both Routing Engines. On a TX Matrix router, clear both Routing Engines on all chassis connected to the TX Matrix router. Likewise, on a TX Matrix Plus router, clear both Routing Engines on all chassis connected to the TX Matrix Plus router.

**infrastructure *name***—(QFabric systems) (Optional) Clear all halt or reboot requests on the fabric control Routing Engines or fabric manager Routing Engines.

**interconnect-device *name***—(QFabric systems) (Optional) Clear all halt or reboot requests on the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, clear all halt or reboot requests for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, clear all halt or reboot requests for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches only) (Optional) Clear all halt or reboot requests on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**node-group *name***—(QFabric systems) (Optional) Clear all halt or reboot requests on the Node group.

**scc**—(TX Matrix routers only) (Optional) Clear all halt or reboot requests for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Clear all halt or reboot requests for the TX Matrix Plus router. Replace *number* with 0.

**Required Privilege Level**      maintenance

|                              |                                                                                                                                                                                                                                                                                                      |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>request system reboot</i></li><li>• <a href="#">request system reboot on page 415</a></li><li>• <a href="#">Rebooting and Halting a Device on page 186</a></li><li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li></ul> |
| <b>List of Sample Output</b> | <a href="#">clear system reboot on page 358</a><br><a href="#">clear system reboot (TX Matrix Router) on page 358</a><br><a href="#">clear system reboot (QFX Series) on page 358</a>                                                                                                                |
| <b>Output Fields</b>         | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                |

## Sample Output

### clear system reboot

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

### clear system reboot (TX Matrix Router)

```
user@host> clear system reboot
scc-re0:
-----
No shutdown/reboot scheduled.
lcc0-re0:
-----
No shutdown/reboot scheduled.
lcc2-re0:
-----
No shutdown/reboot scheduled.
```

### clear system reboot (QFX Series)

```
user@switch> clear system reboot node-group node1
No shutdown/reboot scheduled.
```



## file

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | file <archive  change-owner   change-permission   checksum  compare   compress   copy   delete   delete-directory   link   list   make-directory   rename   show   source address>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>change-owner, change-permission, compress, delete-directory, link, and make-directory</b> options added in Junos OS Release 14.1.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Archive files from the device, copy files to and from the router or switch, calculate the file checksum, compare files, delete a file from the device, list files on the device, rename a file, show file contents, show the local address to initiate a connection, change owner of a file, change permission of a file, compress a file, delete a directory, create a link between files, or create a new directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>archive (Optional)</b> —Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.</p> <p><b>change-owner (Optional)</b> —Change owner of a file.</p> <p><b>change-permission (Optional)</b> —Change permission of a file.</p> <p><b>checksum (Optional)</b> —Calculate the Message Digest 5 (MD5) checksum of a file.</p> <p><b>compare (Optional)</b> —Compare two local files and describe the differences between them in default, context, or unified output styles.</p> <p><b>compress (Optional)</b> —Compress a file.</p> <p><b>copy (Optional)</b> —Copy files from one place to another on the local switch or between the local switch and a remote system.</p> <p><b>delete (Optional)</b> —Delete a file on the local switch.</p> <p><b>delete-directory (Optional)</b> —Delete a directory.</p> <p><b>link (Optional)</b> —Create a link between files.</p> <p><b>list (Optional)</b> —Display a list of files on the local switch.</p> <p><b>make-directory (Optional)</b> —Create a new directory.</p> <p><b>rename (Optional)</b> —Rename a file on the local switch.</p> <p><b>show (Optional)</b> —Display the contents of a file.</p> <p><b>source address (Optional)</b> —Specify the source address of the local file.</p> |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

- Related Documentation**
- *Viewing Files and Directories on a Device Running Junos OS*
  - [CLI Explorer](#)

## file archive

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file archive destination <i>destination</i> source <i>source</i> &lt;compress&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b>destination <i>destination</i></b>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none"> <li>• For archived files—The suffix <b>.tar</b></li> <li>• For archived and compressed files—The suffix <b>.tgz</b></li> </ul> <p><b>source <i>source</i></b>—Source of the original file or files. Specify the source as a URL or filename.</p> <p><b>compress</b>—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix <b>.tgz</b>.</p> |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 42</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>List of Sample Output</b>    | <p><a href="#">file archive (Multiple Files) on page 361</a></p> <p><a href="#">file archive (Single File) on page 361</a></p> <p><a href="#">file archive (with Compression) on page 362</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## Sample Output

### file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

### file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

### file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

## file checksum md5

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file checksum md5 &lt;pathname&gt; filename</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Calculate the Message Digest 5 (MD5) checksum of a file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <p><b>pathname</b>—(Optional) Path to a filename.</p> <p><b>filename</b>—Name of a local file for which to calculate the MD5 checksum.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <a href="#">file checksum sha-256 on page 365</a></li> <li>• <a href="#">file checksum sha1 on page 364</a></li> <li>• <i>op</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">file checksum md5 on page 363</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

### Sample Output

#### file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

## file checksum sha1

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file checksum sha1 &lt;pathname&gt; filename</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options</b>                  | <b>pathname</b> —(Optional) Path to a filename.<br><b>filename</b> —Name of a local file for which to calculate the SHA-1 checksum.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum md5 on page 363</a></li><li>• <a href="#">file checksum sha-256 on page 365</a></li><li>• <i>op</i></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">file checksum sha1 on page 364</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

### Sample Output

#### file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

## file checksum sha-256

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file checksum sha-256 &lt;pathname&gt; filename</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Options</b>                  | <p><b>pathname</b>—(Optional) Path to a filename.</p> <p><b>filename</b>—Name of a local file for which to calculate the SHA-256 checksum.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>maintenance</p> <p>view</p> <p>view-configuration</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <a href="#">file checksum md5 on page 363</a></li> <li>• <a href="#">file checksum sha1 on page 364</a></li> <li>• <i>op</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">file checksum sha-256 on page 365</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Sample Output

### file checksum sha-256

```

user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71

```

## file compare

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file compare (files <i>filename filename</i>)</code><br><code>&lt;context   unified&gt;</code><br><code>&lt;ignore-white-space&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"><li>• <b>Default</b>—In the first line of output, <b>c</b> means lines were changed between the two files, <b>d</b> means lines were deleted between the two files, and <b>a</b> means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (&lt;) in front of output lines refers to the first file. A right angle bracket (&gt;) in front of output lines refers to the second file.</li><li>• <b>Context</b>—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-).</li><li>• <b>Unified</b>—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.</li></ul> |
| <b>Options</b>                  | <p><b>files <i>filename</i></b>—Names of two local files to compare.</p> <p><b>context</b>—(Optional) Display output in context format.</p> <p><b>ignore-white-space</b>—(Optional) Ignore changes in the amount of white space.</p> <p><b>unified</b>—(Optional) Display output in unified format.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 42</a></li><li>• <a href="#">Viewing Core Files from Junos OS Processes on page 196</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>List of Sample Output</b>    | <a href="#">file compare files on page 367</a><br><a href="#">file compare files context on page 367</a><br><a href="#">file compare files unified on page 367</a><br><a href="#">file compare files unified ignore-white-space on page 367</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



## Sample Output

### file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

### file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

### file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

### file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

## file delete

|                                 |                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file delete <i>filename</i></code><br><code>&lt;purge&gt;</code>                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                  |
| <b>Description</b>              | Delete a file on the local router or switch.                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <b><i>filename</i></b> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.<br><br><b><i>purge</i></b> —(Optional) Overwrite regular files before deleting them. |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">file delete on page 369</a><br><a href="#">file delete (Routing Matrix) on page 369</a>                                                                                                                                                                                                       |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                     |

## Sample Output

### file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

### file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

## file list

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file list</code><br><code>&lt;detail   recursive&gt;</code><br><code>&lt;filename&gt;</code>                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                             |
| <b>Description</b>              | Display a list of files on the local router or switch.                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <b>none</b> —Display a list of all files for the current directory.<br><br><b>detail   recursive</b> —(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.<br><br><b>filename</b> —(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.                  |
| <b>Additional Information</b>   | The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the <b>file list</b> command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the <b>file list</b> command. |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                          |
| <b>List of Sample Output</b>    | <a href="#">file list on page 370</a><br><a href="#">file list (Routing Matrix) on page 370</a>                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                |

## Sample Output

### file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

### file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```

## file rename

---

|                                 |                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file rename <i>source destination</i></code>                                                                                                                                       |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Rename a file on the local router or switch.                                                                                                                                             |
| <b>Options</b>                  | <b><i>destination</i></b> —New name for the file.<br><br><b><i>source</i></b> —Original name of the file. For a routing matrix, the filename must include the chassis information.       |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">file rename on page 372</a><br><a href="#">file rename (Routing Matrix) on page 372</a>                                                                                      |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                    |

### Sample Output

#### file rename

The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

#### file rename (Routing Matrix)

The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
```

```
-----

/var/tmp:
.pccardd
sartre.conf
snmpd.rr
syslogd.core-tarball.0.tgz
```

## file show

|                                 |                                                                                                                                                                                                                                             |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>file show <i>filename</i></code><br><code>&lt;encoding (base64   raw)&gt;</code>                                                                                                                                                      |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                    |
| <b>Description</b>              | Display the contents of a file.                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>filename</i></b> —Name of a file. For a routing matrix, the filename must include the chassis information.<br><br><b><code>encoding (base64   raw)</code></b> —(Optional) Encode file contents with base64 encoding or show raw text. |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">file show on page 374</a><br><a href="#">file show (Routing Matrix) on page 374</a>                                                                                                                                             |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                       |

## Sample Output

### file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

### file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/.gdbinit
lcc0-re0:
-----
#####
# Settings
#####

set print pretty

#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbufp->msg_ptr
end
```



```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```

## load

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>load (factory-default   merge   override   patch   replace   set   update)<br/>load (<i>filename</i>   terminal) &lt;relative&gt;</code>                                                                                                                                                                                                                                      |
| <b>QFX Series</b>          | <code>load (dhcp-snooping <i>filename</i>)</code>                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b> | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                  |
| <b>Description</b>         | Load a configuration from an ASCII configuration file, from terminal input, or from the factory default. Your current location in the configuration hierarchy is ignored when the load operation occurs.                                                                                                                                                                            |
| <b>Options</b>             | <b>dhcp-snooping</b> —(QFX Series switches) Loads DHCP snooping entries.<br><br><b>factory-default</b> —Loads the factory configuration. The factory configuration contains the manufacturer's suggested configuration settings. The factory configuration is the router or switch's first configuration and is loaded when the router or switch is first installed and powered on. |



**NOTE:** To load the factory default configuration, you must first *unprotect* any protected hierarchies in the configuration.

On J Series Services Routers, pressing and holding down the Config button on the router for 15 seconds causes the factory configuration to be loaded and committed. However, this operation deletes all other configurations on the router; using the **load factory-default** command does not.

**filename**—Name of the file to load. For information about specifying the filename, see *Viewing Files and Directories on a Device Running Junos OS*.

**merge**—Combine the configuration that is currently shown in the CLI with the configuration.

**override**—Discard the entire configuration that is currently shown in the CLI and load the entire configuration. Marks every object as changed.

**patch**—Change part of the configuration and mark only those parts as changed.

**replace**—Look for a **replace** tag in *filename*, delete the existing statement of the same name, and replace it with the configuration.

**set**—Merge a set of commands with an existing configuration. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as **set**, **edit**, **exit**, and **top**.

**relative**—(Optional) Use the **merge** or **replace** option without specifying the full hierarchy level.

**terminal**—Use the text you type at the terminal as input to the configuration. Type Ctrl+d to end terminal input.

**update**—Discard the entire configuration that is currently shown in the CLI, and load the entire configuration. Marks changed objects only.



**NOTE:** If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

**Required Privilege Level** configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.

**Related Documentation**

- [Loading a Configuration from a File on page 1249](#)

## ping


---

**List of Syntax**   [Syntax on page 378](#)  
                          [Syntax \(QFX Series\) on page 378](#)

**Syntax**   `ping host`  
              `<bypass-routing>`  
              `<count requests>`  
              `<detail>`  
              `<do-not-fragment>`  
              `<inet | inet6>`  
              `<interface source-interface>`  
              `<interval seconds>`  
              `<logical-system logical-system-name>`  
              `<loose-source value>`  
              `<mac-address mac-address>`  
              `<no-resolve>`  
              `<pattern string>`  
              `<rapid>`  
              `<record-route>`  
              `<routing-instance routing-instance-name>`  
              `<size bytes>`  
              `<source source-address>`  
              `<strict >`  
              `<strict-source value.>`  
              `<tos type-of-service>`  
              `<ttl value>`  
              `<verbose>`  
              `<vpls instance-name>`  
              `<wait seconds>`

**Syntax (QFX Series)**   `ping host`  
                          `<bypass-routing>`  
                          `<count requests>`  
                          `<detail>`  
                          `<do-not-fragment>`  
                          `<inet>`  
                          `<interface source-interface>`  
                          `<interval seconds>`  
                          `<logical-system logical-system-name>`  
                          `<loose-source value>`  
                          `<mac-address mac-address>`  
                          `<no-resolve>`  
                          `<pattern string>`  
                          `<rapid>`  
                          `<record-route>`  
                          `<routing-instance routing-instance-name>`  
                          `<size bytes>`  
                          `<source source-address>`  
                          `<strict>`  
                          `< strict-source value>`  
                          `<tos type-of-service>`  
                          `<ttl value>`  
                          `<verbose>`

<wait *seconds*>

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Release Information</b> | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>         | <p>Check host reachability and network connectivity. The <b>ping</b> command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>             | <p><b>host</b>—IP address or hostname of the remote system to ping.</p> <p><b>bypass-routing</b>—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p><b>count requests</b>—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p><b>detail</b>—(Optional) Include in the output the interface on which the ping reply was received.</p> <p><b>do-not-fragment</b>—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> <b>NOTE:</b> In Junos OS Release 11.1 and later, when issuing the <b>ping</b> command for an IPv6 route with the <b>do-not-fragment</b> option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p> </div> <p><b>inet</b>—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p><b>inet6</b>—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p><b>interface source-interface</b>—(Optional) Interface to use to send the ping requests.</p> <p><b>interval seconds</b>—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p><b>logical-system logical-system-name</b>—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the <b>set cli logical-system logical-system-name</b> command and then run the <b>ping</b> command. To return to the main router or switch, enter the <b>clear cli logical-system</b> command.</p> |

**loose-source value**—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

**mac-address mac-address**—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**pattern string**—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

**rapid**—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

**record-route**—(Optional) Record and report the packet's path (IPv4).

**routing-instance routing-instance-name**—(Optional) Name of the routing instance for the ping attempt.

**size bytes**—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

**source source-address**—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

**strict**—(Optional) Use the strict source route option (IPv4).

**strict-source value**—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

**tos type-of-service**—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

If the device configuration includes the **dscp-code-point value** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

**ttl value**—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

**verbose**—(Optional) Display detailed output.

**vpls instance-name**—(Optional) Ping the instance to which this VPLS belongs.

**wait seconds**—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | network                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i></li> </ul>                                                                                                                                                                                                                                                                                      |
| <b>List of Sample Output</b>    | <a href="#">ping hostname on page 381</a><br><a href="#">ping hostname rapid on page 381</a><br><a href="#">ping hostname size count on page 381</a>                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | <p>When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.</p> |

## Sample Output

### ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

### ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

### ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

## request chassis beacon

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax (QFX Series)</b>      | <code>request chassis beacon</code><br><code>&lt;all (off   on)&gt;</code><br><code>&lt;fpc slot-number (off   on)&gt;</code><br><code>&lt;interconnect-device name (cb slot-number   fpc slot-number   (off   on)&gt;</code><br><code>&lt;node-device name (off   on)&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | (QFX Series only) Enable or disable the beacon LED on a QFX Series device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b>all</b>—Turn the beacon LED either <b>on</b> or <b>off</b> on all QFabric system Interconnect and Node devices.</p> <p><b>cb slot-number</b>—Turn the beacon LED either <b>on</b> or <b>off</b> on the Control Board of the QFX3008-I Interconnect device.</p> <p><b>fpc slot-number</b>—Turn the beacon LED either <b>on</b> or <b>off</b> on the Flexible PIC Concentrator on the standalone QFX3500 switch or the Interconnect device.</p> <p><b>interconnect-device name</b>—Turn the beacon LED either <b>on</b> or <b>off</b> on the Interconnect device.</p> <p><b>node-device name</b>—Turn the beacon LED either <b>on</b> or <b>off</b> on the Node device.</p> <p><b>off</b>—Turn the beacon LED <b>off</b>.</p> <p><b>on</b>—Turn the beacon LED <b>on</b>.</p> |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show chassis beacon on page 509</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>List of Sample Output</b>    | <a href="#">request chassis beacon fpc 0 on (QFX Series) on page 382</a><br><a href="#">request chassis beacon node-device (QFabric System) on page 382</a><br><a href="#">request chassis beacon on interconnect-device fpc (QFabric System) on page 383</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## Sample Output

### request chassis beacon fpc 0 on (QFX Series)

```
user@switch> request chassis beacon fpc 0 on

Beacon set to ON
```

### request chassis beacon node-device (QFabric System)

```
user@switch> request chassis beacon node-device node1 on
```



node1 ON

request chassis beacon on interconnect-device fpc (QFabric System)

user@switch> request chassis beacon on interconnect-device fpc 2

FPC 2 ON

## request chassis fpc

---

|                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 384</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 384</a><br><a href="#">Syntax (MX Series Routers) on page 384</a><br><a href="#">Syntax (MX2020 3D Universal Edge Routers) on page 384</a><br><a href="#">Syntax (MX2010 3D Universal Edge Routers) on page 384</a><br><a href="#">Syntax (QFabric System) on page 384</a><br><a href="#">Syntax (PTX Series Packet Transport Routers) on page 384</a> |
| <b>Syntax</b>                                        | <code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <code>request chassis fpc (offline   online   restart) slot <i>slot-number</i> &lt;lcc <i>number</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                |
| <b>Syntax (MX Series Routers)</b>                    | <code>request chassis fpc (offline   online   restart) slot <i>slot-number</i> &lt;all-members&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;member <i>member-id</i>&gt;</code>                                                                                                                                                                                                                                                                        |
| <b>Syntax (MX2020 3D Universal Edge Routers)</b>     | <code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (MX2010 3D Universal Edge Routers)</b>     | <code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Syntax (QFabric System)</b>                       | <code>request chassis fpc</code><br><code>&lt;interconnect-device <i>name</i> slot <i>slot-number</i> (offline   online)&gt;</code><br><code>&lt;(offline   online) interconnect-device <i>name</i> slot <i>slot-number</i>&gt;</code><br><code>&lt;slot <i>slot-number</i> interconnect-device <i>name</i> (offline   online)&gt;</code>                                                                                                                      |
| <b>Syntax (PTX Series Packet Transport Routers)</b>  | <code>request chassis fpc (offline   online   restart) slot <i>slot-number</i></code>                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                           | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS 11.3 for QFX Series.<br>Command introduced in Junos OS 12.1x48 for PTX Series Packet Transport Routers.<br>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.<br>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.                      |
| <b>Description</b>                                   | (M20, M40, M40e, M120, M160, M320, MX Series, and T Series routers, QFabric systems, EX Series switches, and PTX Series Packet Transport Routers only) Control the operation of the Flexible PIC Concentrator (FPC). For information about the meaning of “FPCs” on the switches, see <i>EX Series Switches Hardware and CLI Terminology Mapping</i> .                                                                                                         |



**NOTE:** Beginning in Junos OS 12.3, it is possible that FPCs brought offline using the `request chassis fpc slot fpc-slot offline` operational-mode CLI command can come online during a configuration commit or power-supply replacement procedure. As an alternative, use the `set fpc fpc-slot power off` configuration-mode command at the `[edit chassis]` hierarchy level to ensure that the FPCs remain offline.

**Options**    **offline**—Take the FPC offline.

**online**—Bring the FPC online.

**interconnect-device *name***—(QFabric systems only) Bring the Flexible Port Concentrator (FPC) on the QFX3008-I Interconnect device either offline or online:

- (QFabric System) On a QFabric system, specify the name of the QFX3008-I Interconnect device containing the Flexible Port Concentrator (FPC) you want to bring either offline or online.

**restart**—Restart the FPC.

**slot *slot-number***—FPC slot number:

- M20 router—0 through 3.
- M120 router—0 through 5.
- MX240 router—0 through 2. On the MX240 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX480 router—0 through 5. On the MX480 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX960 router—0 through 11. On the MX960 router, slot-number corresponds to the Dense Port Concentrator (DPC) slot number. If an MPC is installed, slot-number corresponds to the MPC slot number.
- MX2020 router—0 through 19.
- MX2010 router—0 through 9.
- TX Matrix and TX Matrix Plus routers only—On the TX Matrix router, if you specify the number of the T640 router by using the ***lcc number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 or T4000 router by using the ***lcc number*** option (the recommended method), replace ***slot-number*** with a value from 0 through 7. Otherwise, replace ***slot-number*** with a value from 0 through 31. In case of TX Matrix Plus router with 3D SIBs, replace

*slot-number* with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis fpc lcc 1 slot 1 offline
user@host> request chassis fpc slot 9 offline
```

- Other routers—0 through 7.
- QFabric System—Replace *slot-number* with a value from 0 through 2.
- EX Series switches:
  - EX4200 switches in a Virtual Chassis configuration—Replace *slot-number* with a value from 0 through 9.
  - EX6210 switches—Replace *slot-number* with a value from 0 through 9.



**NOTE:** These commands are not supported for slots 4 and 5 when a Switch Fabric and Routing Engine (SRE) module is installed in those slots. These commands are supported for slots 4 and 5 only if a line card is installed in them.

- EX8208 switches—Replace *slot-number* with a value from 0 through 7.
- EX8216 switches—Replace *slot-number* with a value from 0 through 15.
- PTX5000 Packet Transport Router—Replace *slot-number* with a value from 0 through 7.

**all-members**—(MX Series routers only) (Optional) Change FPC status of all members of the Virtual Chassis configuration.

**local**—(MX Series routers only) (Optional) Change FPC status of the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Change FPC status of the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis fpc on page 639](#)
- *show chassis fpc-feb-connectivity*
- *show chassis fabric fpcs*
- *Configuring the Junos OS to Make a Flexible PIC Concentrator Stay Offline*
- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*

**List of Sample Output**

- [request chassis fpc on page 387](#)
- [request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\) on page 387](#)
- [request chassis fpc \(MX2020 Router\) on page 387](#)
- [request chassis fpc \(MX2010 Router\) on page 387](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request chassis fpc](#)

```
user@host> request chassis fpc online slot 0
FPC 0 already online
```

### [request chassis fpc \(MX Series Routers with Media Services Blade \[MSB\]\)](#)

```
user@host> request chassis fpc slot 0
Possible completions:
  offline          Take FPC offline
  online           Bring FPC online
  restart          Restart FPC
```

### [request chassis fpc \(MX2020 Router\)](#)

```
user@host >request chassis fpc online slot 2
FPC 2 already online
```

### [request chassis fpc \(MX2010 Router\)](#)

```
user@host >request chassis fpc offline slot 5
Offline initiated, use "show chassis fpc" to verify
```

## request chassis pic

---

|                                                      |                                                                                                                                                                                                                                                              |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 388</a><br><a href="#">Syntax (ACX4000 Series Routers) on page 388</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 388</a>                                                                               |
| <b>Syntax</b>                                        | <code>request chassis pic (offline   online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>                                                                                                                                                  |
| <b>Syntax (ACX4000 Series Routers)</b>               | <code>request chassis pic (offline   online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>                                                                                                                                                  |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <code>request chassis pic (offline   online) fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> &lt;lcc <i>number</i>&gt;</code>                                                                                                                        |
| <b>Release Information</b>                           | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.3 for ACX4000 Routers.<br>Command introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>                                   | Control the operation of the PIC.                                                                                                                                                                                                                            |



**NOTE:** The `request chassis pic (offline | online) fpc-slot slot number pic-slot slot-number` command is not supported for built-in PICs on MX Series routers.

To view a list of built-in PICs on the router or switch chassis, use the `show chassis hardware` command.



**NOTE:** This command is not supported on MX960 and MX2020 routers with MPC5EQ.



**NOTE:** T1600 routers and TX Matrix Plus routers with 100-Gigabit Ethernet PICs require two adjacent PIC slots, 0 and 1, for each PIC. Therefore, only online and offline command options to PIC slot 0 are allowed. Use of the online and offline command options for PIC slot 1 with the described router and PIC combination is not allowed.



**NOTE:** In T Series routers, when the PIC state is set from offline to online or vice-versa before the processing is complete for the previous command, you are provided feedback on the status of your request. The following sample messages are displayed if you try to set a PIC offline or online:

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 online
fpc 1 pic 0 online initiated, use "show chassis fpc pic-status" to verify
```

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 online
FPC 1 PIC 0 already transitioning to online
```

When the same PIC is set to a different state while the transition is in progress, you are provided feedback on the status of your request.

```
user@switch> request chassis pic fpc-slot 1 pic-slot 0 offline
FPC 1, PIC 0 already transitioning to online. Please retry later.
```

**Options**    **offline**—Take the PIC offline.

**online**—Bring the PIC online.

**fpc-slot *slot-number***—Flexible PIC Concentrator (FPC) slot number. Replace *slot-number* with a value appropriate for your router or switch:

- ACX4000 routers—1 or 2.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—0.
  - EX4200 switches in a Virtual Chassis configuration—0 through 9 (switch's member ID).
  - EX8208 switches—0 through 7 (line card).
  - EX8216 switches—0 through 15 (line card).
- M5, M7i, M10, and M10i routers—0 or 1.
- M20 routers—0 through 3.
- M40 and M40e routers—0 through 7.
- M120 routers—0 through 5.
- M160 routers—0 through 7.
- M320 routers—0 through 7.
- MX 5, MX10, and MX40 routers—0 or 1.
- MX80 routers—0 or 1.
- MX240 routers—0 through 2
- MX480 routers—0 through 5
- MX2020 routers—0 through 19.

- MX2010 routers—0 through 9.
- MX960 routers—0 through 11.
- PTX5000 routers—0 or 1.
- T Series routers—0 through 7.
- TX Matrix and TX Matrix Plus routers only—On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the **number** of the T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, for the FPC slot number, replace **slot-number** with a value from 0 through 31. On a TX Matrix Plus router with 3D SIBs to assign the FPC slot number, replace **slot-number** with a value from 0 through 63. For example, the following commands have the same result:

```
user@host> request chassis pic fpc-slot 1 lcc 1 pic-slot 0 offline
user@host> request chassis pic fpc-slot 9 pic-slot 0 offline
```

- QFX5100 standalone switches—0.

**pic-slot slot-number**—PIC slot number.

- EX3200 and EX4200 switches—0 for built-in network interfaces and 1 for interfaces on uplink modules.
- EX8208 and EX8216 switches—0.
- M Series routers—0, 1, 2, or 3
- MX960 router—**slot-number** corresponds to the slot number of the Packet Forwarding Engine.
- PTX5000 routers—0 or 1.
- T320 router—0 or 1.
- T640 router—0, 1, 2, or 3.
- T1600 router —0, 1, 2, or 3.
- T4000 router—0, 1, 2, or 3.
- QFX5100 standalone switches—0, 1, or 2. PIC 0 is used for all interfaces that are not configured on expansion modules, and PIC 1 and PIC 2 are used for interfaces configured on expansion modules.

**lcc number**—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.



Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis hardware on page 676](#)
- [show chassis pic on page 889](#)
- *Configuring the PIC Type*

**List of Sample Output** [request chassis pic on page 391](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request chassis pic](#)

```
user@host> request chassis pic pic-slot 0 online fpc-slot 0
FPC 0, PIC 0 is already online
```

## request chassis routing-engine master

---

|                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                 | <a href="#">Syntax on page 392</a><br><a href="#">Syntax (M Series, MX Series, T Series Routers) on page 392</a><br><a href="#">Syntax (TX Matrix Routers) on page 392</a><br><a href="#">Syntax (TX Matrix Plus Routers) on page 392</a><br><a href="#">Syntax (MX Series Virtual Chassis) on page 392</a><br><a href="#">Syntax (QFX Series) on page 392</a>                                                                                                                                                                                                                                        |
| <b>Syntax</b>                                         | request chassis routing-engine master (acquire   release   switch)<br><force><br><no-confirm>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Syntax (M Series, MX Series, T Series Routers)</b> | request chassis routing-engine master (acquire   release   switch <check>)<br><no-confirm>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Syntax (TX Matrix Routers)</b>                     | request chassis routing-engine master (acquire   release   switch) (lcc <i>number</i>   scc   all-chassis)<br><force><br><no-confirm>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (TX Matrix Plus Routers)</b>                | request chassis routing-engine master (acquire   release   switch) (lcc <i>number</i>   sfc   all-chassis   all-lcc)<br><force><br><no-confirm>                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Syntax (MX Series Virtual Chassis)</b>             | request chassis routing-engine master (acquire   release   switch <check>)<br><all-members><br><local><br><member <i>member-id</i> ><br><no-confirm>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Syntax (QFX Series)</b>                            | request chassis routing-engine master (release   switch)<br><check><br><interconnect-device <i>name</i> ><br><node-group <i>name</i> ><br><no-confirm>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Release Information</b>                            | Command introduced before Junos OS Release 7.4.<br><b>all-chassis</b> option added in Junos OS Release 8.0.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.<br>Command introduced in Junos OS Release 11.3 for QFX Series.<br>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.<br>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.<br>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers. |
| <b>Description</b>                                    | For routers or switches with multiple Routing Engines, control which Routing Engine is the master.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



**CAUTION:** (Routing matrix based on the TX Matrix or TX Matrix Plus routers only) Within the routing matrix, we recommend that all Routing Engines run the same Junos OS Release. If you run different releases on the Routing Engines and a change in mastership occurs on any backup Routing Engine in the routing matrix, one or all routers (in a routing matrix based on the TX Matrix router or in a routing matrix based on a TX Matrix Plus router) might become logically disconnected from the TX Matrix router and cause data loss. For more information, see the [TX Matrix Router Hardware Guide](#) or the *Junos OS High Availability Library for Routing Devices*.



**NOTE:** Successive graceful Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to “Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset,” do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the Flexible PIC concentrators (FPCs) should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

You will receive an error message stating “Command aborted. Not ready for mastership switch, try after n seconds” when this command is re-entered before 240 seconds have elapsed on EX Series switches.



**NOTE:** On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the routing engine to the backup routing engine, and then reboot.

**Options** **acquire**—Attempt to become the master Routing Engine.

**release**—Request that the other Routing Engine become the master.

**switch**—Toggle mastership between Routing Engines.



**NOTE:** The **acquire** option should be used with caution because acquiring a Routing Engine may result in a corrupted database. If possible, use the **switch** option instead.

The **acquire**, **release**, and **switch** options have the following suboptions:

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) On a routing matrix composed of a TX Matrix router and the attached T640 routers, switch mastership on all the Routing Engines in the routing matrix. Likewise, on a routing matrix composed of a TX Matrix Plus router and the attached T1600 or T4000 routers, switch mastership on all the Routing Engines in the routing matrix.

**all-lcc**—(TX Matrix Plus routers only) Request to acquire mastership for all line-card chassis (LCC).

**all-members**—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in all member routers of the Virtual Chassis configuration.

**check**—(QFabric systems, MX104, MX480, MX960, MX2010, and MX2020 routers, and PTX5000 routers only) (Optional) Available only with the **switch** option. Check graceful switchover status of the standby Routing Engine before toggling mastership between Routing Engines.

**interconnect-device *name***—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on an Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Control Routing Engine mastership on the Routing Engines of the specified member in the Virtual Chassis Configuration. Replace *member-id* with a value of 0 or 1.

**no-confirm**—(Optional) Do not request confirmation for the switch.

**node-group *name***—(QFabric systems only) (Optional) Control Routing Engine mastership on the Routing Engines on a Node group.

**scc**—(TX Matrix routers only) TX Matrix (switch-card chassis).

**sfc**—(TX Matrix Plus routers only) TX Matrix Plus router (or switch-fabric chassis).

**force**—(Optional) Available only with the **acquire** option. Force the change to a new master Routing Engine.



**NOTE:** The **force** option is not supported on the M Series, MX Series, or T Series routers.

**Additional Information** Because both Routing Engines are always running, the transition from one to the other as the master Routing Engine is immediate. However, the changeover interrupts communication to the System and Switch Board (SSB). The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. Interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.

By default, the Routing Engine in slot 0 (**RE0**) is the master and the Routing Engine in slot 1 (**RE1**) is the backup. To change the default master Routing Engine, include the **routing-engine** statement at the **[edit chassis redundancy]** hierarchy level in the configuration. For more information, see the *Junos OS Administration Library for Routing Devices*

To have the backup Routing Engine become the master Routing Engine, use the **request chassis routing-engine master switch** command. If you use this command to change the master and then restart the chassis software for any reason, the master reverts to the default setting.



**NOTE:** Although the configurations on the two Routing Engines do not have to be the same and are not automatically synchronized, we recommend making both configurations the same.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis routing-engine on page 905](#)
- *Configuring Routing Engine Redundancy*
- *Switching the Global Master and Backup Roles in a Virtual Chassis Configuration*

**List of Sample Output**

- [request chassis routing-engine master acquire on page 396](#)
- [request chassis routing-engine master switch on page 396](#)
- [request chassis routing-engine master switch check on page 396](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request chassis routing-engine master acquire

```
user@host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted

Reset other routing engine and become master ? [yes,no] (no)
```

### request chassis routing-engine master switch

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between Routing Engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The other Routing Engine becomes the master.
```

Switch mastership back to the local Routing Engine:

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between routing engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.
```

### request chassis routing-engine master switch check

Usage shown for M Series, MX Series, and T Series routers.

```
{master}[edit]
user@host> request chassis routing-engine master switch check

warning: Standby Routing Engine is not ready for graceful switchover.

{master}[edit]
user@host> request chassis routing-engine master switch check
Switchover Ready

You can similarly check the backup Routing Engine.
```

## request message

---

|                                 |                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | request message all message " <i>text</i> "<br>request message message " <i>text</i> " (terminal <i>terminal-name</i>   user <i>user-name</i> )                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                 |
| <b>Description</b>              | Display a message on the screens of all users who are logged in to the router or switch or on specific screens.                                                                                                                                                                                                                          |
| <b>Options</b>                  | <b>all</b> —Display a message on the terminal of all users who are currently logged in.<br><br><b>message "<i>text</i>"</b> —Message to display.<br><br><b>terminal <i>terminal-name</i></b> —Name of the terminal on which to display the message.<br><br><b>user <i>user-name</i></b> —Name of the user to whom to direct the message. |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                              |
| <b>List of Sample Output</b>    | <a href="#">request message message on page 397</a>                                                                                                                                                                                                                                                                                      |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                    |

## Sample Output

### request message message

```
user@host> request message message "Maintenance window in 10 minutes" user maria
Message from user@host on tty0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```

## request system configuration rescue delete

---

|                            |                                                                                                                                                                                          |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | request system configuration rescue delete                                                                                                                                               |
| <b>Release Information</b> | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | Delete an existing rescue configuration.                                                                                                                                                 |



**NOTE:** The [edit system configuration] hierarchy is not available on QFabric systems.

---

|                                 |                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | This command has no options.                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">request system configuration rescue save on page 399</a></li><li>• <a href="#">request system software rollback on page 459</a></li><li>• <a href="#">show system commit on page 1001</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">request system configuration rescue delete on page 398</a>                                                                                                                                                                                  |
| <b>Output Fields</b>            | This command produces no output.                                                                                                                                                                                                                        |


### Sample Output

#### request system configuration rescue delete

```
user@host> request system configuration rescue delete
```



## request system configuration rescue save

|                                                                                                                                                                                            |                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                              | request system configuration rescue save                                                                                                                                                                                                          |
| <b>Release Information</b>                                                                                                                                                                 | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                          |
| <b>Description</b>                                                                                                                                                                         | Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the <b>rollback</b> command.                                                                                         |
| <div>  <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems. </div> |                                                                                                                                                                                                                                                   |
| <b>Options</b>                                                                                                                                                                             | This command has no options.                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b>                                                                                                                                                            | maintenance                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">request system software delete on page 430</a></li> <li>• <a href="#">request system software rollback on page 459</a></li> <li>• <a href="#">show system commit on page 1001</a></li> </ul> |
| <b>List of Sample Output</b>                                                                                                                                                               | <a href="#">request system configuration rescue save on page 399</a>                                                                                                                                                                              |
| <b>Output Fields</b>                                                                                                                                                                       | This command produces no output.                                                                                                                                                                                                                  |

### Sample Output

#### request system configuration rescue save

```
user@host> request system configuration rescue save
```

## request system halt

---

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 400</a><br><a href="#">Syntax (EX Series Switches) on page 400</a><br><a href="#">Syntax (PTX Series) on page 400</a><br><a href="#">Syntax (TX Matrix Router) on page 400</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 400</a><br><a href="#">Syntax (MX Series Router) on page 401</a><br><a href="#">Syntax (QFX Series) on page 401</a>                                                                                                                                      |
| <b>Syntax</b>                         | <code>request system halt</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;backup-routing-engine&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk   removable-compact-flash   usb)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code>                                                                                                                             |
| <b>Syntax (EX Series Switches)</b>    | <code>request system halt</code><br><code>&lt;all-members&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;backup-routing-engine&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;media (external   internal)&gt;</code><br><code>&lt;member <i>member-id</i>&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;slice <i>slice</i>&gt;</code> |
| <b>Syntax (PTX Series)</b>            | <code>request system halt</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;backup-routing-engine&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code>                                                                                                                                                             |
| <b>Syntax (TX Matrix Router)</b>      | <code>request system halt</code><br><code>&lt;all-lcc   lcc <i>number</i>   scc&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;backup-routing-engine&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code>                                                                                                   |
| <b>Syntax (TX Matrix Plus Router)</b> | <code>request system halt</code><br><code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                                   |

|                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                           | <pre> &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;other-routing-engine&gt; &lt;in <i>minutes</i>&gt; &lt;media (compact-flash   disk)&gt; &lt;message "text"&gt; </pre>                                                                                                                                                                                                                                                                                                                                              |
| Syntax (MX Series Router) | <pre> request system halt &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;backup-routing-engine&gt; &lt;both-routing-engines&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "text"&gt; &lt;other-routing-engine&gt; </pre>                                                                                                                                                                                                                                                         |
| Syntax (QFX Series)       | <pre> request system halt &lt;all-members&gt; &lt;at <i>time</i>&gt; &lt;both-routing-engines&gt; &lt;director-device <i>director-device-id</i>&gt; &lt;in <i>minutes</i>&gt; &lt;local&gt; &lt;media &gt; &lt;member <i>member-id</i>&gt; &lt;message "text"&gt; &lt;other-routing-engine&gt; &lt;slice <i>slice</i>&gt; </pre>                                                                                                                                                                                                                               |
| Release Information       | <p>Command introduced before Junos OS Release 7.4.</p> <p><b>other-routing-engine</b> option introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>director-device</b> option introduced for QFabric systems in Junos OS Release 12.2.</p> <p><b>backup-routing-engine</b> option introduced in Junos OS Release 13.1.</p> |
| Description               | Stop the router or switch software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |



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

When you issue this command on a QFX5100 switch, you are not prompted to reboot. You must power cycle the switch to reboot.

**Options** **none**—Stop the router or switch software immediately.

**all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Halt all chassis.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, halt all T1600 or T4000 routers connected to the TX Matrix Plus router.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Halt all members of the Virtual Chassis configuration.

**at time** —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- **+minutes**—Number of minutes from now to stop the software.
- **yymmddhhmm**—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to stop the software.

**backup-routing-engine**—(Optional) Halt the backup Routing Engine. This command halts the backup Routing Engine, regardless from which Routing Engine the command is executed. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. If you issue the command from the backup Routing Engine, the backup Routing Engine is halted.

**both-routing-engines**—(Optional) Halt both Routing Engines at the same time.

**director-device *director-device-id***—(QFabric systems only) Halt a specific Director device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, halt a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, halt a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Halt the local Virtual Chassis member.

**in *minutes***—(Optional) Number of minutes from now to stop the software. This option is an alias for the **at +*minutes*** option.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to J Series routers only.)

**media (external | internal)**—(EX Series and QFX Series switches and MX Series routers only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Halt the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping the software.

**other-routing-engine**—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

**scc**—(TX Matrix routers only) (Optional) Halt the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Halt the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with 0.

**slice *slice***—(EX Series and QFX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- 1—Halt partition 1.
- 2—Halt partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** On the M7i router, the **request system halt** command does not immediately power down the Packet Forwarding Engine. The power-down process can take as long as 5 minutes.

On a TX Matrix router and TX Matrix Plus router if you issue the **request system halt** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are halted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are halted.



**NOTE:** If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), and then halt the master Routing Engine. To halt a Routing Engine, issue the `request system halt` command. You can also halt both Routing Engines at the same time by issuing the `request system halt both-routing-engines` command.

|                                 |                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">clear system reboot on page 355</a></li><li>• <a href="#">request system power-off on page 410</a></li><li>• <a href="#">Rebooting and Halting a Device on page 186</a></li><li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">request system halt on page 405</a><br><a href="#">request system halt (In 2 Hours) on page 405</a><br><a href="#">request system halt (Immediately) on page 405</a><br><a href="#">request system halt (At 1:20 AM) on page 405</a>                                                                           |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                      |

## Sample Output

### request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

### request system halt (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

### request system halt (Immediately)

```
user@host> request system halt at now
```

### request system halt (At 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

## request system license add

---

|                                 |                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request system license add (<i>filename</i>   terminal)</code>                                                                                                                                                                                           |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 9.5 for SRX Series devices.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Add a license key.                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <b><i>filename</i></b> —License key from a file or URL. Specify the filename or the URL where the key is located.<br><br><b><i>terminal</i></b> —License key from the terminal.                                                                                |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Adding New Licenses (CLI Procedure) on page 76</a></li></ul>                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">request system license add on page 406</a>                                                                                                                                                                                                         |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                          |

## Sample Output

### request system license add

```
user@host> request system license add terminal
```



## request system license delete

---

|                                 |                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request system license delete ( <i>license-identifier</i>   license-identifier-list [ <i>licenseid001</i> <i>licenseid002</i> <i>licenseid003</i> ]   all )</code>                                                                                                                                                |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>license-identifier-list</b> introduced in Junos OS Release 13.1.</p>                               |
| <b>Description</b>              | Delete a license key. You can choose to delete one license at a time, all licenses at once, or a list of license identifiers enclosed in brackets.                                                                                                                                                                      |
| <b>Options</b>                  | <p><b>license-identifier</b>—Text string that uniquely identifies a license key.</p> <p><b>license-identifier-list [ <i>licenseid001</i> <i>licenseid002</i> <i>licenseid003</i>.... ]</b>—Delete multiple license identifiers as a list enclosed in brackets.</p> <p><b>all</b>—Delete all licenses on the device.</p> |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Deleting a License (CLI Procedure) on page 77</a></li> </ul>                                                                                                                                                                                                       |

## request system license save

---

|                                 |                                                                                                                                                                                                                                                                |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request system license save (<i>filename</i>   terminal)</code>                                                                                                                                                                                          |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 9.5 for SRX Series devices. |
| <b>Description</b>              | Save installed license keys to a file or URL.                                                                                                                                                                                                                  |
| <b>Options</b>                  | <b><i>filename</i></b> —License key from a file or URL. Specify the filename or the URL where the key is located.<br><br><b><i>terminal</i></b> —License key from the terminal.                                                                                |
| <b>Required Privilege Level</b> | maintenance                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Saving License Keys on page 78</a></li></ul>                                                                                                                                                               |
| <b>List of Sample Output</b>    | <a href="#">request system license save on page 408</a>                                                                                                                                                                                                        |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                          |

## Sample Output

### request system license save

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

## request system logout

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>request system logout (pid <i>pid</i>   terminal <i>terminal</i>   user <i>username</i>) &lt;all&gt;</code>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Log out users from the router or switch and the configuration database. If a user held the <b>configure exclusive</b> lock, this command clears the exclusive lock.                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b>all</b>—(Optional) Log out all sessions owned by a particular PID, terminal session, or user. (On a TX Matrix or TX Matrix Plus router, this command is broadcast to all chassis.)</p> <p><b>pid <i>pid</i></b>—Log out the user session using the specified management process identifier (PID). The PID type must be management process.</p> <p><b>terminal <i>terminal</i></b>—Log out the user for the specified terminal session.</p> <p><b>user <i>username</i></b>—Log out the specified user.</p> |
| <b>Required Privilege Level</b> | configure                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>List of Sample Output</b>    | <a href="#">request system logout on page 409</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.                                                                                                                                                                                                                                                                                                                                                                                                                           |

## Sample Output



### request system logout

```
user@host> request system logout user tammy all
Connection closed by foreign host.
```

## request system power-off

---

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 410</a><br><a href="#">Syntax (EX Series Switches) on page 410</a><br><a href="#">Syntax (TX Matrix Router) on page 410</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 410</a><br><a href="#">Syntax (MX Series Router) on page 410</a><br><a href="#">Syntax (QFX Series) on page 411</a>                                                                                                                                                |
| <b>Syntax</b>                         | <code>request system power-off</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk   removable-compact-flash   usb)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code>                                                                                                                             |
| <b>Syntax (EX Series Switches)</b>    | <code>request system power-off</code><br><code>&lt;all-members&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;local&gt;</code><br><code>&lt;media (external   internal)&gt;</code><br><code>&lt;member <i>member-id</i>&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;slice <i>slice</i>&gt;</code> |
| <b>Syntax (TX Matrix Router)</b>      | <code>request system power-off</code><br><code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code>                                                                                     |
| <b>Syntax (TX Matrix Plus Router)</b> | <code>request system power-off</code><br><code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;other-routing-engine&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;media (compact-flash   disk)&gt;</code><br><code>&lt;message "<i>text</i>"&gt;</code>                                                                       |
| <b>Syntax (MX Series Router)</b>      | <code>request system power-off</code><br><code>&lt;all-members&gt;</code><br><code>&lt;at <i>time</i>&gt;</code><br><code>&lt;both-routing-engines&gt;</code><br><code>&lt;in <i>minutes</i>&gt;</code><br><code>&lt;local&gt;</code>                                                                                                                                                                                                                                           |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <pre> &lt;media (external   internal)&gt; &lt;member <i>member-id</i>&gt; &lt;message "<i>text</i>"&gt; &lt;other-routing-engine&gt; </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax (QFX Series)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <pre> request system power-off &lt;at <i>time</i>&gt; &lt;in <i>minutes</i>&gt; &lt;media (external   internal)&gt; &lt;message "<i>text</i>"&gt; &lt;slice <i>slice</i>&gt; </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p>Command introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Power off the software.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <hr/> <div>  <p><b>NOTE:</b> When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member, use the <code>member</code> option. You cannot issue this command from the QFabric CLI.</p> </div> <hr/>                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <div>  <p><b>NOTE:</b> For a standalone chassis (such as MX Series, PTX Series, and T Series routers), the request to power off the system is applicable only to the Routing Engines. When you request to power off both Routing Engines, all the FPCs in the chassis shut down after approximately 10 minutes and the chassis fans run at full speed. The FPCs shut down because they no longer have communication with the Routing Engines and an Inter-Integrated Circuit (I2C) timeout occurred.</p> </div> <hr/> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p><b>none</b>—Power off the router or switch software immediately.</p> <p><b>all-chassis</b>—(Optional) (TX Matrix and TX Matrix Plus router only) Power off all Routing Engines in the chassis.</p> <p><b>all-lcc</b>—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, power off all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Power off all members of the Virtual Chassis configuration.</p> <p><b>at <i>time</i></b>—(Optional) Time at which to power off the software, specified in one of the following ways:</p> |

- **now**—Power off the software immediately. This is the default.
- **+minutes**—Number of minutes from now to power off the software.
- **yymmddhhmm**—Absolute time at which to power off the software, specified as year, month, day, hour, and minute.
- **hh:mm**—Absolute time on the current day at which to power off the software.

**both-routing-engines**—(Optional) Power off both Routing Engines at the same time.

**in minutes**—(Optional) Number of minutes from now to power off the software. This option is an alias for the **at +minutes** option.

**lcc number**—(Optional) (TX Matrix and TX Matrix Plus router only) On a TX Matrix router, power off a T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, power off a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Power off the local Virtual Chassis member.

**media (compact-flash | disk | removable-compact-flash | usb)**—(Optional) Boot medium for the next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

**media (external | internal)**—(EX Series and QFX Series switches and MX Series routers only) (Optional) Power off the boot media:

- **external**—Power off the external mass storage device.
- **internal**—Power off the internal flash device.

**member member-id**—(EX4200 switches and MX Series routers only) (Optional) Power off the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**message "text"**—(Optional) Message to display to all system users before powering off the software.

**other-routing-engine**—(Optional) Power off the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

**scc**—(Optional) (TX Matrix router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix router (or switch-card chassis). If you issue the command from the master Routing Engine, the master SCC is powered off. If you issue the command from the backup Routing Engine, the backup SCC is powered off.

**sfc number**—(Optional) (TX Matrix Plus router only) Power off only the master Routing Engine or the backup Routing Engine on the TX Matrix Plus router (or switch-fabric chassis). If you issue the command from the master Routing Engine, the master SFC is powered off. If you issue the command from the backup Routing Engine, the backup SFC is powered off. Replace *number* with zero.

**slice slice**—(EX Series and QFX Series switches only) (Optional) Power off a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** On a routing matrix composed of a TX Matrix router and T640 routers, if you issue the **request system power-off** command on the TX Matrix master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

Likewise, on a routing matrix composed of a TX Matrix Plus router and T1600 routers, if you issue the **request system power-off** command on the TX Matrix Plus master Routing Engine, all the master Routing Engines connected to the routing matrix are powered off. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are powered off.

If you issue the **request system power-off both-routing-engines** command on the TX Matrix or TX Matrix Plus router, all the Routing Engines on the routing matrix are powered off.

**Required Privilege Level** maintenance

**List of Sample Output** [request system power-off on page 414](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system power-off

```
user@host> request system power-off message "This router will be powered off in 30 minutes.  
Please save your data and log out immediately."  
warning: This command will not halt the other routing-engine.  
If planning to switch off power, use the both-routing-engines option.  
Power Off the system ? [yes,no] (no) yes  
  
*** FINAL System shutdown message from remote@nutmeg ***  
System going down IMMEDIATELY  
  
This router will be powered off in 30 minutes. Please save your data and log out  
immediately.  
  
Shutdown NOW!  
[pid 5177]
```



## request system reboot

**Syntax (QFX Series and EX4600)** request system reboot  
 <all <graceful>>  
 <all-members | local | member *member-id*>  
 <at time>  
 <both-routing-engines>  
 <director-device *name*>  
 <director-group <graceful>>  
 <fabric <graceful>>  
 <fast-boot>  
 <in minutes>  
 <media >  
 <message “text”>  
 <node-group *name*>  
 <other-routing-engine>  
 <slice (1 | 2 | alternate)>

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 13.2X51-D25 for EX4600 switches.

**Description** Reboot the Junos OS.



**NOTE:** On a QFabric system, to avoid traffic loss on the network Node group, switch mastership of the Routing Engine to the backup Routing Engine, and then reboot.

Reboot requests are recorded in the system log files, which you can view with the **show log messages** command. You can view the process names with the **show system processes** command.

**Options** **none**—Reboots the software immediately.

**all**—(QFabric systems only) (Optional) Reboots the software on the Director group, fabric control Routing Engines, fabric manager Routing Engines, Interconnect devices, and network and server Node groups.

**all-members | local | member *member-id***—(Optional) Specify which member of the Virtual Chassis to reboot:

- **all-members**—Reboots each switch that is a member of the Virtual Chassis.
- **local**—Reboots the local switch, meaning the switch you are logged into, only.
- **member *member-id***—Reboots the specified member switch of the Virtual Chassis.

**at time**—(Optional) Time at which to reboot the software, specified in one of the following ways:

- **+minutes**—Number of minutes from now to reboot the software.
- **hh:mm**—Absolute time on the current day at which to reboot the software, specified in 24-hour time.
- **now**—Stop or reboot the software immediately. This is the default.
- **yymmddhhmm**—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.

**both-routing-engines**—(Optional) Reboot both Routing Engines at the same time.

**director-device *name***—(QFabric systems only) (Optional) Reboots the software on the Director device and the default partition (QFabric CLI).

**director-group**—(QFabric systems only) (Optional) Reboots the software on the Director group and the default partition (QFabric CLI).

**fabric**—(QFabric systems only) (Optional) Reboots the fabric control Routing Engines and the Interconnect devices.

**fast-boot**—(QFX5100 only) (Optional) Enhances the reboot time. The switch reboots in such a way as to minimize downtime of network ports by not bringing the network ports down immediately as in the normal reboot option. There is minimal traffic loss while the forwarding device is reprogrammed.

**graceful**—(QFabric systems only) (Optional) Allows the QFabric component to reboot with minimal impact to network traffic. This option is only available for the **all**, **fabric**, and **director-group** options.

**in minutes**—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at +minutes** option.

**media (external | internal)**—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.

**message "text"**—(Optional) Message to display to all system users before rebooting the software.

**node-group *name***—(QFabric systems only) (Optional) Reboots the software on a server Node group or a network Node group.

**other-routing-engine**—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

**routing-engine**—(Optional) Reboot the Routing Engine.

**slice (1 | 2 | alternate)**—(Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:



**NOTE:** The slice option is not supported on the QFX5100 switch or the EX4600 switch, because there is no alternate slice when Junos OS boots as a Virtual Machine (VM). To switch to a previous version of Junos OS, issue the `request system software rollback` command.

- 1—Reboot from partition 1.
- 2—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition, which is the partition that did not boot the switch at the last bootup.

**Required Privilege Level** maintenance

**Related Documentation**

- [clear system reboot on page 355](#)
- [Rebooting and Halting a Device on page 186](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system reboot

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

### request system reboot (At 2300)

```
user@switch> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

### request system reboot (In 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request the system to reboot in 2 hours:

```
user@switch> request system reboot at +120
user@switch> request system reboot in 120
user@switch> request system reboot at 19:00
```

### request system reboot (Immediately)

```
user@switch> request system reboot at now
```

### request system reboot (At 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@switch> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

#### request system reboot director-device

```
user@switch> request system reboot director-device Node1
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

#### request system reboot director-group

```
user@switch> request system reboot director-group
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

#### request system reboot director-group graceful

```
user@switch> request system reboot director-group graceful
Issuing this command may interrupt traffic forwarding.
Continue? [yes,no] (no)
```

## request system snapshot

**Syntax** request system snapshot  
 <config-partition>  
 <media>  
 <partition>  
 <root-partition>  
 slice alternate

**Release Information** Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Copy the currently running Junos OS and configuration to alternate media. This command takes a snapshot of the contents of the / (root), and **/var** partitions on the media used to boot the switch and then copies the snapshot to alternate media. If the switch was booted from internal flash memory, the snapshot is copied to an external USB flash drive. If the switch was booted from an external USB flash drive, the snapshot is copied to internal flash memory.

**Options** **none**—Create a snapshot on the alternate media—that is, the external media if you booted the switch using software stored on internal media or internal media if you booted the switch using software stored on external media.

**config-partition**—(Optional) Create a snapshot of the configuration partition only and store it onto the default /altconfig on the hard disk device or an /altconfig on a USB device.

**media type**—(Optional) Specify the boot device the software is copied to:

- compact-flash—Copy software to the primary compact flash drive.
- external—Copy software to an external mass storage device, such as a USB flash drive. If a USB drive is not connected, the switch displays an error message.
- internal—Copy software to an internal flash drive.
- removable-compact-flash—Copy software to the removable compact flash drive.

**partition**—(Optional) Partition the destination media before copying over the snapshot.

**root-partition**—(Optional) Create a snapshot of the root partition only and store it onto the default /altroot on the hard disk device or an /altroot on a USB device.

**slice alternate**—(Optional) Take a snapshot of the active root partition and copy it to the alternate slice on the boot media.

**Required Privilege Level** view

**Related Documentation**

- *show system snapshot*
- [Creating a Snapshot and Using It to Boot a QFX Series Switch on page 178](#)
- *Verifying That a System Snapshot Was Created on a QFX Series Switch*

List of Sample Output [request system snapshot partition on page 420](#)

## Sample Output

[request system snapshot partition](#)

```
user@switch> request system snapshot partition
Clearing current label...
Partitioning external media (da1) ...
Verifying compatibility of destination media partitions...
Running newfs (334MB) on external media / partition ...
Running newfs (404MB) on external media /config partition ...
Running newfs (222MB) on external media /var partition ...
Copying '/dev/da0s2a' to '/dev/da1s1a' .. (this may take a few minutes)
Copying '/dev/da0s3e' to '/dev/da1s3e' .. (this may take a few minutes)
Copying '/dev/da0s2f' to '/dev/da1s1f' .. (this may take a few minutes)
The following filesystems were archived: / /config /var
```

## request system software add

|                                       |                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                 | <a href="#">Syntax on page 421</a><br><a href="#">Syntax (EX Series Switches) on page 421</a><br><a href="#">Syntax (TX Matrix Router) on page 421</a><br><a href="#">Syntax (TX Matrix Plus Router) on page 421</a><br><a href="#">Syntax (MX Series Router) on page 422</a><br><a href="#">Syntax (QFX Series) on page 422</a>                                                         |
| <b>Syntax</b>                         | <pre>request system software add <i>package-name</i> &lt;best-effort-load&gt; &lt;delay-restart&gt; &lt;force&gt; &lt;no-copy&gt; &lt;no-validate&gt; &lt;re0   re1&gt; &lt;reboot&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;unlink&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt; &lt;validate&gt;</pre>                                 |
| <b>Syntax (EX Series Switches)</b>    | <pre>request system software add <i>package-name</i> &lt;best-effort-load&gt; &lt;delay-restart&gt; &lt;force&gt; &lt;no-copy&gt; &lt;no-validate&gt; &lt;re0   re1&gt; &lt;reboot&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt; &lt;validate&gt;</pre>                                                |
| <b>Syntax (TX Matrix Router)</b>      | <pre>request system software add <i>package-name</i> &lt;best-effort-load&gt; &lt;delay-restart&gt; &lt;force&gt; &lt;lcc <i>number</i>   scc&gt; &lt;no-copy&gt; &lt;no-validate&gt; &lt;re0   re1&gt; &lt;reboot&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;unlink&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt; &lt;validate&gt;</pre> |
| <b>Syntax (TX Matrix Plus Router)</b> | <pre>request system software add <i>package-name</i> &lt;best-effort-load&gt;</pre>                                                                                                                                                                                                                                                                                                      |

```
<delay-restart>
<force>
<lcc number | sfc number>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name package-name]>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>
```

**Syntax (MX Series Router)**

```
request system software add package-name
<best-effort-load>
<delay-restart>
<force>
<member member-id>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<set [package-name package-name]>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>
```

**Syntax (QFX Series)**

```
request system software add package-name
<best-effort-load>
<component all>
<delay-restart>
<force>
<force-host>
<no-copy>
<no-validate>
<partition>
<reboot>
<unlink>
<upgrade-with-config>
<upgrade-with-config-format format>
<validate>
```

**Release Information**

Command introduced before Junos OS Release 7.4.

**best-effort-load** and **unlink** options added in Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

**sfc** option introduced for the TX Matrix Plus router in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.1 for the QFX Series.

**set [*package-name package-name*]** option added in Junos OS Release 11.1 for EX Series switches.

**set [*package-name package-name*]** option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.





**NOTE:** On EX Series switches, the set `[package-name package-name]` option allows you to install only two software packages on a mixed EX4200 and EX4500 Virtual Chassis, whereas, on M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways, the set `[package-name package-name]` option allows you to install multiple software packages and software add-on packages at the same time.

`upgrade-with-config` and `upgrade-with-config-format` *format* options added in Junos OS Release 12.3 for M Series routers, MX Series routers, T Series routers, EX Series Ethernet switches, and QFX Series devices.

#### Description



**NOTE:** We recommend that you always download the software image to `/var/tmp` only. On EX Series and QFX Series switches, you must use the `/var/tmp` directory. Other directories are not supported.

Install a software package or bundle on the router or switch.



**WARNING:** Any configuration changes performed after inputting the `request system software add` command will be lost when the system reboots with an upgraded version of JUNOS.

**Options** `package-name`—Location from which the software package or bundle is to be installed.

For example:

- `/var/tmp/package-name`—For a software package or bundle that is being installed from a local directory on the router or switch.
- `protocol://hostname/pathname/package-name`—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
  - **ftp**—File Transfer Protocol.  
Use `ftp://hostname/pathname/package-name`. To specify authentication credentials, use `ftp://<username>:<password>@hostname/pathname/package-name`. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.  
Use `http://hostname/pathname/package-name`. To specify authentication credentials, use `http://<username>:<password>@hostname/pathname/package-name`. If a password is required and you omit it, you are prompted for it.

- **scp**—Secure copy (available only for Canada and U.S. version).  
Use **scp://hostname/pathname/package-name**. To specify authentication credentials, use  
**scp://<username>:<password>@hostname/pathname/package-name**.

**NOTE:**

- The **pathname** in the protocol is the relative path to the user's home directory on the remote system and not the root directory.
- Do not use the **scp** protocol in the **request system software add** command to download and install a software package or bundle from a remote location. The previous statement does not apply to the QFabric switch. The software upgrade is handled by the MGD process which does not support **scp**.  
Use the **file copy** command to copy the software package or bundle from the remote location to the **/var/tmp** directory on the hard disk:  
**file copy scp://source/package-name /var/tmp**  
Then install the software package or bundle using the **request system software add** command:  
**request system software add /var/tmp/package-name**
- On a J Series Services Router, when you install the software from a remote location, the package is removed at the earliest opportunity in order to make room for the installation to be completed. If you copy the software to a local directory on the router and then install the new package, use the **unlink** option to achieve the same effect and allow the installation to be completed.

---

**best-effort-load**—(Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

**component all**—(QFabric systems only) (Optional) Install software package on all of the QFabric components.

**delay-restart**—(Optional) Install a software package or bundle, but do not restart software processes.

**force**—(Optional) Force the addition of the software package or bundle (ignore warnings).

**force-host**—(Optional) Force the addition of host software package or bundle (ignore warnings) on the QFX5100 device.

**lcc number** —(TX Matrix routers and TX Matrix Plus routers only) (Optional) In a routing matrix based on the TX Matrix router, install a software package or bundle on a T640 router that is connected to the TX Matrix router. In a routing matrix based on the TX Matrix Plus router, install a software package or bundle on a router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**member *member-id***—(MX Series routers only) (Optional) Install a software package on the specified Virtual Chassis member. Replace *member-id* with a value of 0 or 1.

**partition**—(QFX3500 switches only) (Optional) Format and repartition the media before installation.

**scc**—(TX Matrix routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Install a software package or bundle on a Routing Engine on a TX Matrix Plus router. Replace *number* with 0.

**no-copy**—(Optional) Install a software package or bundle, but do not save copies of the package or bundle files.

**no-validate**—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the **validate** option.

**re0 | re1**—(Optional) On routers or switches that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (re0) or the Routing Engine in slot 1 (re1).

**reboot**—(Optional) After adding the software package or bundle, reboot the system. On a QFabric switch, the software installation is not complete until you reboot the component for which you have installed the software.

**set [*package-name package-name*]**—(Mixed EX4200 and EX4500 Virtual Chassis only) (Optional) Install two software packages—a package for an EX4200 switch and the same release of the package for an EX4500 switch—to upgrade all member switches in a mixed EX4200 and EX4500 Virtual Chassis.

**set [*package-name package-name*]**—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages and software add-on packages at the same time.

**unlink**—(Optional) On J Series Services Routers, this option ensures that the software package is removed at the earliest opportunity in order to make room for the installation to be completed. On M Series, T Series, and MX Series routers, use the

**unlink** option to remove the software package from this directory after a successful upgrade is completed.

**upgrade-with-config**—(Optional) Install one or more configuration files.

**upgrade-with-config-format *format***—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



**NOTE:** The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

**validate**—(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.



**NOTE:** The **validate** option only works on systems that do not have **graceful-switchover (GRES)** enabled. To use the **validate** option on a system with GRES, either disable GRES for the duration of the installation, or install using the command **request system software in-service-upgrade**, which requires nonstop active routing (NSR) to be enabled when using GRES.

#### Additional Information

Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new package or bundle is successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.



**NOTE:** The **request system snapshot** command is currently not supported on the QFabric system. Also, you cannot add or install multiple packages on a QFabric system.

After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, jkernel, last. Add the operating system package, jkernel, first and the routing software package, jroute, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernel
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

By default, when you issue the **request system software add *package-name*** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, when you issue the **request system software add *package-name*** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are upgraded to the same version of software. If you issue the same command on the TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">request system software delete on page 430</a></li> <li>• <a href="#">request system software rollback on page 459</a></li> <li>• <a href="#">request system storage cleanup on page 466</a></li> <li>• <a href="#">Upgrading Software on page 134</a></li> <li>• <a href="#">Upgrading Software on a QFabric System</a></li> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
List of Sample Output	<a href="#">request system software add validate on page 428</a> <a href="#">request system software add (Mixed EX4200 and EX4500 Virtual Chassis) on page 428</a> <a href="#">request system software add component all (QFabric Systems) on page 429</a>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software add validate

```
user@host> request system software add validate /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING: This package will load JUNOS 7.2R1.7 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...
```

## Sample Output

### request system software add (Mixed EX4200 and EX4500 Virtual Chassis)

```
user@switch> request system software add set
[/var/tmp/jinstall-ex-4200-11.1R1.1-domestic-signed.tgz
/var/tmp/jinstall-ex-4500-11.1R1.1-domestic-signed.tgz]
...
```

**request system software add component all (QFabric Systems)**

```
user@switch> request system software add /pbdata/packages/jinstall-qfabric-12.2X50-D1.3.rpm  
component all  
...
```

## request system software delete

---

<b>List of Syntax</b>	<a href="#">Syntax on page 430</a> <a href="#">Syntax (TX Matrix Router) on page 430</a> <a href="#">Syntax (TX Matrix Plus Router) on page 430</a>
<b>Syntax</b>	<code>request system software delete <i>software-package</i></code> <code>&lt;force&gt;</code> <code>&lt;reboot&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>request system software delete <i>software-package</i></code> <code>&lt;force&gt;</code> <code>&lt;lcc <i>number</i>   scc&gt;</code> <code>&lt;reboot&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request system software delete <i>software-package</i></code> <code>&lt;force&gt;</code> <code>&lt;lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;reboot&gt;</code> <code>&lt;set [<i>package-name package-name</i>]&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <b>set [<i>package-name package-name</i>]</b> added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Services Gateways. Option <b>reboot</b> introduced in Junos OS Release 12.3.
<b>Description</b>	Remove a software package or bundle from the router or switch.



**CAUTION:** Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the router or switch.

- |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b> | <p><b><i>software-package</i></b>—Software package or bundle name. You can delete any or all of the following software bundles or packages:</p> <ul style="list-style-type: none"><li>• <b>jbase</b>—(Optional) Junos base software suite</li><li>• <b>jcrypto</b>—(Optional, in domestic version only) Junos security software</li><li>• <b>jdocs</b>—(Optional) Junos online documentation file</li><li>• <b>jkernel</b>—(Optional) Junos kernel software suite</li><li>• <b>jpfe</b>—(Optional) Junos Packet Forwarding Engine support</li></ul> |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



- **jroute**—(Optional) Junos routing software suite
- **junos**—(Optional) Junos base software



**NOTE:** On EX Series switches, some of the package names are different than those listed. To see the list of packages that you can delete on an EX Series switch, enter the command **show system software**.

**force**—(Optional) Ignore warnings and force removal of the software.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, remove an extension or upgrade package from a specific T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, remove an extension or upgrade package from a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**reboot**—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software delete** command.

**scc**—(TX Matrix routers only) (Optional) Remove an extension or upgrade package from the TX Matrix router (or switch-card chassis).

**set [package-name package-name]**—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

**sfc number**—(TX Matrix Plus routers only) (Optional) Remove an extension or upgrade package from the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the /altroot and /altconfig file systems (on routers) or the /, /altroot, /config, /var, and /var/tmp file systems (on switches). After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the /altroot and /altconfig file systems (on routers) or the /, /altroot,

/config, /var, and /var/tmp file systems (on switches). After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

**Required Privilege Level** maintenance

**Related Documentation**

- [request system software add on page 421](#)
- [request system software rollback on page 459](#)
- [request system software validate on page 463](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [request system software delete jdocs on page 432](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request system software delete jdocs](#)

The following example displays the system software packages before and after the **jdocs** package is deleted through the **request system software delete** command:

```
user@host> show system software
Information for jbase:
```

```
Comment:
JUNOS Base OS Software Suite [7.2R1.7]
```

```
Information for jcrypto:
```

```
Comment:
JUNOS Crypto Software Suite [7.2R1.7]
```

```
Information for jdocs:
```

```
Comment:
JUNOS Online Documentation [7.2R1.7]
```

```
Information for jkernel:
```

```
Comment:
JUNOS Kernel Software Suite [7.2R1.7]
```

```
...
```

```
user@host> request system software delete jdocs
Removing package 'jdocs' ...
```

```
user@host> show system software
```

Information for jbase:

Comment:

JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [7.2R1.7]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [7.2R1.7]

...

## request system software download

---

Syntax (QFabric System)	request system software download <i>path package-name</i>
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Download a software package from a location on the Director device, mounted external USB flash drive, remote FTP or SCP location, or other location.
Options	<p><b>path</b>—Location where the software package is located. For example:</p> <ul style="list-style-type: none"><li>• <b>/pbdata/packages/package-name</b>—For a software package that is being installed from a local directory on the switch.</li><li>• <b>protocol://hostname/pathname/package-name</b>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <b>protocol</b> with one of the following:<ul style="list-style-type: none"><li>• <b>ftp</b>—File Transfer Protocol. Use <b>ftp://hostname/pathname/package-name</b>. To specify authentication credentials, use <b>ftp://&lt;username&gt;:&lt;password&gt;@hostname/pathname/package-name</b>. To have the system prompt you for the password, specify <b>prompt</b> in place of the password. If a password is required, and you do not specify the password or <b>prompt</b>, an error message is displayed.</li><li>• <b>scp</b>—Secure copy (available only for Canada and U.S. version). Use <b>scp://hostname/pathname/package-name</b>. To specify authentication credentials, use <b>scp://&lt;username&gt;:&lt;password&gt;@hostname/pathname/package-name</b>.</li></ul></li></ul>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">request system software add on page 421</a></li><li>• <a href="#">request system software delete on page 430</a></li><li>• <a href="#">request system software rollback on page 459</a></li><li>• <a href="#">request system storage cleanup on page 466</a></li><li>• <a href="#">Upgrading Software on page 134</a></li><li>• <a href="#">Upgrading Software on a QFabric System</a></li></ul>
List of Sample Output	<a href="#">request system software download on page 435</a>
Output Fields	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software download

```
user@switch> request system software download
ftp://ftp.install-directory/jinstall-qfabric-11.3X30.6.rpm
% Total      % Received % Xferd  Average Speed   Time    Time     Time  Current
                               Dload  Upload   Total   Spent    Left   Speed
100  186M  100  186M    0     0  18.4M      0  0:00:10  0:00:10 --:--:-- 18.6M
```

## request system software in-service-upgrade

---

<b>Syntax</b>	<code>request system software in-service-upgrade <i>package-name</i></code> <code>&lt;no-old-master-upgrade&gt;</code> <code>&lt;reboot&gt;</code>
<b>Syntax (QFX5100 Switches)</b>	<code>request system software in-service-upgrade <i>package-name</i></code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 12.3R2, 13.1R2, and 13.2R1 for TX Matrix Plus routers.</p> <p>Command introduced in Junos OS Release 13.2 for PTX5000 routers.</p> <p>Command introduced in Junos OS Release 13.2 X51-D15 for the QFX Series.</p>
<b>Description</b>	<p>Perform a unified in-service software upgrade (ISSU). A unified ISSU enables you to upgrade from one Junos OS Release to another with no disruption on the control plane and with minimal disruption of traffic. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. On QFX5100 switches, nonstop bridging (NSB) must be enabled if you are using the Layer 2 Control Protocol process (l2cpd) to transmit Layer 2 spanning tree protocols in a Layer 2 bridge environment.</p>
<b>Options</b>	<p><b><i>package-name</i></b>—Location from which the software package or bundle is to be installed. For example:</p> <ul style="list-style-type: none"><li>• <b><i>/var/tmp/package-name</i></b>— For a software package or bundle that is being installed from a local directory on the router.</li><li>• <b><i>protocol://hostname/pathname/package-name</i></b>—For a software package or bundle that is to be downloaded and installed from a remote location. Replace <b><i>protocol</i></b> with one of the following:<ul style="list-style-type: none"><li>• <b>ftp</b>—File Transfer Protocol</li><li>• <b>http</b>—Hypertext Transfer Protocol</li><li>• <b>scp</b>—Secure copy (available only for Canada and U.S. version)</li></ul></li></ul> <p><b>no-old-master-upgrade</b>—(Optional) When the <b>no-old-master-upgrade</b> option is included, after the backup Routing Engine is rebooted with the new software package and a switchover occurs to make it the new master Routing Engine, the former master (new backup) Routing Engine will not be upgraded to the new software. In this case, you must manually upgrade the former master (new backup) Routing Engine. If you do not include the <b>no-old-master-upgrade</b> option, the system will automatically upgrade the former master Routing Engine.</p> <p><b>reboot</b>—(Optional) When the <b>reboot</b> option is included, the former master (new backup) Routing Engine is automatically rebooted after being upgraded to the new software. When the <b>reboot</b> option is not included, you must manually reboot the former master (new backup) Routing Engine using the <b>request system reboot</b> command.</p>



**NOTE:** The reboot option is not available on the QFX5100 switch.

<b>Additional Information</b>	<p>The following conditions apply to unified ISSUs:</p> <ul style="list-style-type: none"> <li>Unified ISSU is not supported on every platform. For a list of supported platforms, see <i>Unified ISSU System Requirements</i>.</li> <li>Unsupported PICs are restarted during a unified ISSU on certain routing devices. For information about supported PICs, see the <i>Junos OS High Availability Library for Routing Devices</i>.</li> <li>Unsupported protocols will experience packet loss during a unified ISSU. For information about supported protocols, see the <i>Junos OS High Availability Library for Routing Devices</i>.</li> <li>During a unified ISSU, you cannot bring any PICs online or offline on certain routing devices.</li> </ul> <p>For more information, see the <i>Junos OS High Availability Library for Routing Devices</i>.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>request system software abort</i></li> <li><a href="#">show chassis in-service-upgrade on page 851</a></li> <li><i>Unified ISSU Concepts</i></li> <li><i>Performing a Unified ISSU</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system software-in-service upgrade reboot on page 437</a> <a href="#">request system software-in-service upgrade reboot (TX Matrix Plus Router) on page 439</a> <a href="#">request system software-in-service upgrade (QFX5100 Switch) on page 447</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software-in-service upgrade reboot

```
{master}

user@host> request system software in-service-upgrade
/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz reboot
ISSU: Validating Image
PIC 0/3 will be offlined (In-Service-Upgrade not supported)
Do you want to continue with these actions being taken ? [yes,no] (no) yes

ISSU: Preparing Backup RE
Pushing bundle to re1
Checking compatibility with configuration
Initializing...
Using jbase-9.0-20080114.2
```

```
Verified manifest signed by PackageProduction_9_0_0
Using /var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Using jinstall-9.0-20080114.2-domestic.tgz
Using jbundle-9.0-20080114.2-domestic.tgz
Checking jbundle requirements on /
Using jbase-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jkernel-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jcrypto-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jpfe-9.0-20080114.2.tgz
Using jdocs-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Using jroute-9.0-20080114.2.tgz
Verified manifest signed by PackageProduction_9_0_0
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-9.0-20080114.2-domestic-signed.tgz' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING: This package will load JUNOS 9.0-20080114.2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz
...
Saving state for rollback ...
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU started
ISSU: Backup RE Prepare Done
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
```



```

ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
Installing package '/var/tmp/paKEuy' ...
Verified jinstall-9.0-20080114.2-domestic.tgz signed by PackageProduction_9_0_0
Adding jinstall...
Verified manifest signed by PackageProduction_9_0_0

WARNING:   This package will load JUNOS 9.0-20080114.2 software.
WARNING:   It will save JUNOS configuration files, and SSH keys
WARNING:   (if configured), but erase all other files and information
WARNING:   stored on this machine. It will attempt to preserve dumps
WARNING:   and log files, but this can not be guaranteed. This is the
WARNING:   pre-installation stage and all the software is loaded when
WARNING:   you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:   A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:   'request system reboot' command when software installation is
WARNING:   complete. To abort the installation, do not reboot your system,
WARNING:   instead use the 'request system software delete jinstall'
WARNING:   command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-9.0-20080114.2-domestic-signed.tgz
...
cp: /var/tmp/paKEuy is a directory (not copied).
Saving state for rollback ...
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
Reboot consistency check bypassed - jinstall 9.0-20080114.2 will complete
installation upon reboot
[pid 30227]

*** FINAL System shutdown message from root@host ***

System going down IMMEDIATELY

Connection to host closed.

```

### request system software-in-service upgrade reboot (TX Matrix Plus Router)

```

{master}

user@host> request system software in-service upgrade
/var/tmp/jinstall-12.3R2-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image

```

PIC 8/1 will be offlined (In-Service-Upgrade not supported)  
PIC 19/2 will be offlined (In-Service-Upgrade not supported)  
PIC 15/3 will be offlined (In-Service-Upgrade not supported)  
Do you want to continue with these actions being taken ? [yes,no] (no) yes

Checking compatibility with configuration  
Initializing...  
Using jbase-12.3R2  
Verified manifest signed by PackageProduction\_12\_3\_0  
Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz  
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0  
Using jinstall-12.3R2-domestic.tgz  
Using jbundle-12.3R2-domestic.tgz  
Checking jbundle requirements on /  
Using jbase-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jbase-12.3R2 signed by PackageProduction\_12\_3\_0  
Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz  
Using jcrypto-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jcrypto-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jdocs-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jdocs-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jkernel-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jkernel-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jpfe-12.3R2.tgz  
WARNING: jpfe-12.3R2.tgz: not a signed package  
WARNING: jpfe-common-12.3R2.tgz: not a signed package  
Verified jpfe-common-12.3R2 signed by PackageProduction\_12\_3\_0  
WARNING: jpfe-T-12.3R2.tgz: not a signed package  
Verified jpfe-T-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jplatform-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jplatform-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jroute-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jroute-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jruntime-12.3R2.tgz  
Verified manifest signed by PackageProduction\_12\_3\_0  
Verified jruntime-12.3R2 signed by PackageProduction\_12\_3\_0  
Using jservices-12.3R2.tgz  
Using jservices-crypto-12.3R2.tgz  
Hardware Database regeneration succeeded  
Validating against /config/juniper.conf.gz  
mgd: commit complete  
Validation succeeded  
ISSU: Preparing LCC Backup REs  
Pushing bundle to lcc0-re1  
Pushing bundle to lcc1-re1  
Pushing bundle to lcc2-re1  
Pushing bundle to lcc3-re1  
Pushing bundle to sfc0-re1  
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...  
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0  
Adding jinstall...  
Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 software.  
WARNING: It will save JUNOS configuration files, and SSH keys

```
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:      This package will load JUNOS 12.3R2 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software delete jinstall'
WARNING:      command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:      This package will load JUNOS 12.3R2 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
```

Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the  
WARNING: 'request system reboot' command when software installation is  
WARNING: complete. To abort the installation, do not reboot your system,  
WARNING: instead use the 'request system software delete jinstall'  
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...

Saving state for rollback ...

Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...

Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0

Adding jinstall...

Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 software.

WARNING: It will save JUNOS configuration files, and SSH keys

WARNING: (if configured), but erase all other files and information

WARNING: stored on this machine. It will attempt to preserve dumps

WARNING: and log files, but this can not be guaranteed. This is the

WARNING: pre-installation stage and all the software is loaded when

WARNING: you reboot the system.

Saving the config files ...

NOTICE: uncommitted changes have been saved in

/var/db/config/juniper.conf.pre-install

Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the

WARNING: 'request system reboot' command when software installation is

WARNING: complete. To abort the installation, do not reboot your system,

WARNING: instead use the 'request system software delete jinstall'

WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...

Saving state for rollback ...

ISSU: Preparing SFC Backup RE

NOTICE: Validating configuration against jinstall-12.3R2-domestic-signed.tgz.

NOTICE: Use the 'no-validate' option to skip this if desired.

Checking compatibility with configuration

Initializing...

Using jbase-12.3R2

Verified manifest signed by PackageProduction\_12\_3\_0

Using /var/tmp/jinstall-12.3R2-domestic-signed.tgz

Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0

Using jinstall-12.3R2-domestic.tgz

Using jbundle-12.3R2-domestic.tgz

Checking jbundle requirements on /

Using jbase-12.3R2.tgz

Verified manifest signed by PackageProduction\_12\_3\_0

Verified jbase-12.3R2 signed by PackageProduction\_12\_3\_0

Using /var/validate/chroot/tmp/jbundle/jboot-12.3R2.tgz

Using jcrypto-12.3R2.tgz

Verified manifest signed by PackageProduction\_12\_3\_0

Verified jcrypto-12.3R2 signed by PackageProduction\_12\_3\_0

Using jdocs-12.3R2.tgz

Verified manifest signed by PackageProduction\_12\_3\_0

Verified jdocs-12.3R2 signed by PackageProduction\_12\_3\_0

Using jkernel-12.3R2.tgz

Verified manifest signed by PackageProduction\_12\_3\_0

Verified jkernel-12.3R2 signed by PackageProduction\_12\_3\_0

```

Using jpfe-12.3R2.tgz
WARNING: jpfe-12.3R2.tgz: not a signed package
WARNING: jpfe-common-12.3R2.tgz: not a signed package
Verified jpfe-common-12.3R2 signed by PackageProduction_12_3_0
WARNING: jpfe-T-12.3R2.tgz: not a signed package
Verified jpfe-T-12.3R2 signed by PackageProduction_12_3_0
Using jplatform-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jplatform-12.3R2 signed by PackageProduction_12_3_0
Using jroute-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jroute-12.3R2 signed by PackageProduction_12_3_0
Using jruntime-12.3R2.tgz
Verified manifest signed by PackageProduction_12_3_0
Verified jruntime-12.3R2 signed by PackageProduction_12_3_0
Using jservices-12.3R2.tgz
Using jservices-crypto-12.3R2.tgz
Hardware Database regeneration succeeded
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
SFC Backup upgrade done
Rebooting SFC Backup RE

Rebooting sfc0-re1
ISSU: SFC Backup RE Prepare Done
Waiting for SFC Backup RE reboot

Rebooting lcc0-re1
Rebooting LCC [lcc0-re1]

Rebooting lcc1-re1
Rebooting LCC [lcc1-re1]

Rebooting lcc2-re1

```

Rebooting LCC [lcc2-re1]

Rebooting lcc3-re1

Rebooting LCC [lcc3-re1]

LCC Backup REs have rebooted

Waiting for LCC Backup REs come back online

ISSU: LCC Backup REs Prepare Done

GRES operational

Initiating Chassis In-Service-Upgrade

Chassis ISSU Started

ISSU: Preparing Daemons

ISSU: Daemons Ready for ISSU

ISSU: Starting Upgrade for FRUs

ISSU: Preparing for Switchover

ISSU: Ready for Switchover

Checking In-Service-Upgrade status

lcc0-re0:

Item	Status	Reason
FPC 1	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 1	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

lcc1-re0:

Item	Status	Reason
FPC 0	Online (ISSU)	
PIC 3	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	

lcc2-re0:

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	
PIC 1	Online (ISSU)	

lcc3-re0:

Item	Status	Reason
FPC 0	Online (ISSU)	
PIC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
PIC 2	Online (ISSU)	
FPC 4	Online (ISSU)	

```

FPC 5      Online (ISSU)
FPC 6      Online (ISSU)
FPC 7      Online (ISSU)
PIC 1      Online (ISSU)

lcc0-re0:
-----
Resolving mastership...
Complete. The other routing engine becomes the master.

lcc1-re0:
-----
Resolving mastership...
Complete. The other routing engine becomes the master.

lcc2-re0:
-----
Resolving mastership...
Complete. The other routing engine becomes the master.

lcc3-re0:
-----
Resolving mastership...
Complete. The other routing engine becomes the master.
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading SFC Old Master RE

lcc0-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING: This package will load JUNOS 12.3R2 software.
WARNING: It will save JUNOS configuration files, and SSH keys
WARNING: (if configured), but erase all other files and information
WARNING: stored on this machine. It will attempt to preserve dumps
WARNING: and log files, but this can not be guaranteed. This is the
WARNING: pre-installation stage and all the software is loaded when
WARNING: you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING: 'request system reboot' command when software installation is
WARNING: complete. To abort the installation, do not reboot your system,
WARNING: instead use the 'request system software delete jinstall'
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...

lcc1-re0:
Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...

```

Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 software.  
WARNING: It will save JUNOS configuration files, and SSH keys  
WARNING: (if configured), but erase all other files and information  
WARNING: stored on this machine. It will attempt to preserve dumps  
WARNING: and log files, but this can not be guaranteed. This is the  
WARNING: pre-installation stage and all the software is loaded when  
WARNING: you reboot the system.

Saving the config files ...

NOTICE: uncommitted changes have been saved in

/var/db/config/juniper.conf.pre-install

Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the  
WARNING: 'request system reboot' command when software installation is  
WARNING: complete. To abort the installation, do not reboot your system,  
WARNING: instead use the 'request system software delete jinstall'  
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...

Saving state for rollback ...

lcc2-re0:

Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...

Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0

Adding jinstall...

Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 software.  
WARNING: It will save JUNOS configuration files, and SSH keys  
WARNING: (if configured), but erase all other files and information  
WARNING: stored on this machine. It will attempt to preserve dumps  
WARNING: and log files, but this can not be guaranteed. This is the  
WARNING: pre-installation stage and all the software is loaded when  
WARNING: you reboot the system.

Saving the config files ...

NOTICE: uncommitted changes have been saved in

/var/db/config/juniper.conf.pre-install

Installing the bootstrap installer ...

WARNING: A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the  
WARNING: 'request system reboot' command when software installation is  
WARNING: complete. To abort the installation, do not reboot your system,  
WARNING: instead use the 'request system software delete jinstall'  
WARNING: command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...

Saving state for rollback ...

lcc3-re0:

Installing package '/var/tmp/jinstall-12.3R2-domestic-signed.tgz' ...

Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction\_12\_3\_0

Adding jinstall...

Verified manifest signed by PackageProduction\_12\_3\_0

WARNING: This package will load JUNOS 12.3R2 software.  
WARNING: It will save JUNOS configuration files, and SSH keys  
WARNING: (if configured), but erase all other files and information



```

WARNING:    stored on this machine. It will attempt to preserve dumps
WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:    'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed.tgz ...
Saving state for rollback ...
Installing package '/var/tmp/paBWTg' ...
Verified jinstall-12.3R2-domestic.tgz signed by PackageProduction_12_3_0
Adding jinstall...
Verified manifest signed by PackageProduction_12_3_0

WARNING:    This package will load JUNOS 12.3R2 software.
WARNING:    It will save JUNOS configuration files, and SSH keys
WARNING:    (if configured), but erase all other files and information
WARNING:    stored on this machine. It will attempt to preserve dumps
WARNING:    and log files, but this can not be guaranteed. This is the
WARNING:    pre-installation stage and all the software is loaded when
WARNING:    you reboot the system.

Saving the config files ...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Installing the bootstrap installer ...

WARNING:    A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:    'request system reboot' command when software installation is
WARNING:    complete. To abort the installation, do not reboot your system,
WARNING:    instead use the 'request system software delete jinstall'
WARNING:    command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-12.3R2-domestic-signed ...
cp: /var/tmp/paBWTg is a directory (not copied).
Saving state for rollback ...
ISSU: SFC Old Master Upgrade Done
ISSU: IDLE

```

#### request system software-in-service upgrade (QFX5100 Switch)

```

{master}

user@switch> request system software in-service-upgrade
/var/tmp/jinstall-qfx-132_x51_vjunos.0-domestic.tgz
ISSU: Validating Image
Prepare for ISSU
spawn the backup VM
ISSU: Preparing Backup RE
Backup upgrade done
ISSU: Backup RE Prepare Done
waiting for backup RE switchover ready
GRES operational

```

```
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
send ISSU done to chassisd on backup VM
Chassis ISSU Completed
ISSU: IDLE
mgd_package_opus_issu: Initiate em0 device handoff
```

## request system software nonstop-upgrade

<b>Syntax</b>	<pre>request system software nonstop-upgrade (<i>package-name</i>   set [<i>package-name</i> <i>package-name</i>]) &lt;no-copy&gt; &lt;no-old-master-upgrade&gt; &lt;reboot &gt; &lt;unlink&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Option <b>set [<i>package-name package-name</i>]</b> added in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D20 for the QFX Series.</p>
<b>Description</b>	<p>Perform a nonstop software upgrade (NSSU) on a switch with redundant Routing Engines or on a Virtual Chassis. The behavior of this command depends on which switch or Virtual Chassis it is executed on:</p> <ul style="list-style-type: none"> <li>When you execute this command on an EX3300, EX4200, EX4300, EX4500, or EX4550 Virtual Chassis or QFX3500 and QFX3600 Virtual Chassis, a fixed configuration of switches in a Virtual Chassis Fabric (QFX3500/QFX3600 and QFX5100 switches) or for a mixed Virtual Chassis Fabric composed of any combination of QFX3500/QFX3600, QFX5100, and EX4300 switches, or a mixed Virtual Chassis composed of any combination of EX4200, EX4500, and EX4550 switches, all members are upgraded. The original Virtual Chassis backup becomes the master. The original master is automatically upgraded and rebooted and rejoins the Virtual Chassis as the backup after the upgrade completes.</li> <li>When you execute this command on an EX6200 or EX8200 switch, both the backup and master Routing Engines are upgraded, with the original backup Routing Engine becoming the new master at the end of the upgrade. <p>The original master Routing Engine is automatically rebooted on an EX6200 switch.</p> <p>The original master Routing Engine is not automatically rebooted on an EX8200 switch unless you specify the <b>reboot</b> option.</p> </li> <li>When you execute this command on an EX8200 Virtual Chassis, all master and backup Routing Engines are upgraded in the Virtual Chassis, including the external Routing Engines. The original backup Routing Engines become the new master Routing Engines. The original master Routing Engines are not automatically rebooted, unless you specify the <b>reboot</b> option.</li> </ul> <p>This command has the following requirements:</p> <ul style="list-style-type: none"> <li>All Virtual Chassis members and all Routing Engines must be running the same Junos OS release.</li> <li>Graceful Routing Engine switchover (GRES) must be enabled.</li> <li>Nonstop active routing (NSR) must be enabled.</li> </ul>



**NOTE:** Although nonstop bridging (NSB) does not have to be enabled for you to use this command, we recommend that you enable NSB. Enabling NSB ensures that all NSB-supported Layer 2 protocols operate seamlessly during the Routing Engine switchover that is part of the NSSU. See *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*.

- The command must be executed from the master Routing Engine on a standalone switch or from the master on a Virtual Chassis.
- For minimal traffic disruption, you must define link aggregation groups (LAGs) such that the member links reside on different Virtual Chassis members (for EX3300, EX4200, EX4300, EX4500, EX4550, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis, and Virtual Chassis Fabric) or on different line cards (for EX6200 and EX8200 switches, and for EX8200 Virtual Chassis).
- For EX3300, EX4200, EX4300, EX4500, EX4550, QFX3500 and QFX3600 Virtual Chassis, and mixed Virtual Chassis:
  - The Virtual Chassis members must be connected in a ring topology. A ring topology prevents the Virtual Chassis from splitting during an NSSU.
  - The Virtual Chassis master and backup must be adjacent to each other in the ring topology. Adjacency permits the master and backup to always be in sync, even when the switches in linecard roles are rebooting.
  - The Virtual Chassis must be preprovisioned so that the linecard role has been explicitly assigned to member switches acting in a linecard role. During an NSSU, the Virtual Chassis members must maintain their roles—the master and backup must maintain their Routing Engine roles (although mastership will change), and the remaining switches must maintain their linecard roles.
  - A two-member Virtual Chassis must have **no-split-detection** configured so that the Virtual Chassis does not split when an NSSU upgrades a member.
- For Virtual Chassis Fabric:
  - Only two pre-provisioned members in the routing engine role are supported. If more than two routing engines are configured, a warning will be issued, and NSSU will stop.
  - The Virtual Chassis Fabric members are connected in a spine and leaf topology. A spine and leaf topology prevents the Virtual Chassis Fabric from splitting during an NSSU. Each leaf device must be connected to both spine devices.
  - The Virtual Chassis Fabric must be preprovisioned so that the line card role has been explicitly assigned to member switches acting in a line card role, and that the routing engine role has been explicitly assigned to member switches acting in a routing engine role. During an NSSU, the Virtual Chassis Fabric members must maintain their roles—the master and backup must maintain their master and backup roles (although

mastership will change), the member switches must remain their routing engine roles, and the remaining switches must maintain their linecard roles.

- A two-member Virtual Chassis Fabric must have **no-split-detection** configured so that the Virtual Chassis Fabric does not split when an NSSU upgrades a member.

**Options** *package-name*—Location from which the software package or bundle is to be installed. For example:

- */var/tmp/package-name*—For a software package or bundle that is being installed from a local directory on the switch.
- *protocol://hostname/pathname/package-name*—For a software package or bundle that is to be downloaded and installed from a remote location. Replace *protocol* with one of the following:
  - **ftp**—File Transfer Protocol.  
Use *ftp://hostname/pathname/package-name*. To specify authentication credentials, use *ftp://<username>:<password>@hostname/pathname/package-name*. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.  
Use *http://hostname/pathname/package-name*. To specify authentication credentials, use *http://<username>:<password>@hostname/pathname/package-name*. If a password is required and you omit it, you are prompted for it.
  - **scp**—Secure copy (available only for Canada and U.S. version).  
Use *scp://hostname/pathname/package-name*. To specify authentication credentials, use *scp://<username>:<password>@hostname/pathname/package-name*.



**NOTE:** The *pathname* in the protocol is the relative path to the user home directory on the remote system and not the root directory.

**set [package-name package-name]**—(Mixed Virtual Chassis only) Locations of the EX4200 and the EX4500 installation packages. These packages must be for the same Junos OS release. See the description of the *package-name* option for information about how to specify the location of the installation packages.

**no-copy**—(Optional) Install a software package or bundle, but do not save copies of package or bundle files.

**no-old-master-upgrade**—(Optional) (EX8200 switches only) Upgrade the backup Routing Engine only. After the upgrade completes, the original master Routing Engine becomes the backup Routing Engine and continues running the previous software version.

**reboot**—(Optional) (EX8200 switches and EX8200 Virtual Chassis only) When the **reboot** option is included, the original master (new backup) Routing Engines are automatically rebooted after being upgraded to the new software. When the **reboot** option is not included, you must manually reboot the original master (new backup) Routing Engines using the **request system reboot** command.



**NOTE:** If you do not use the **reboot** option on an EX8200 Virtual Chassis, you must establish a connection to the console port on the Switch Fabric and Routing Engine (SRE) module or Routing Engine (RE) module to perform the manual reboot of the backup Routing Engines.

**unlink**—(Optional) Remove the software package after a successful upgrade is completed.

**Required Privilege Level** maintenance

**Related Documentation**

- [show chassis nonstop-upgrade on page 887](#)
- *Upgrading Software on an EX3300, EX4200, EX4300, EX4500 and EX4550 Virtual Chassis, and Mixed Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)*
- *Upgrading Software on QFX3500, QFX3600, and QFX5100 Virtual Chassis Using Nonstop Software Upgrade*
- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade on page 139](#)

**List of Sample Output**

[request system software nonstop-upgrade \(EX4200 Virtual Chassis\) on page 452](#)  
[request system software nonstop-upgrade \(EX6200 Switch\) on page 454](#)  
[request system software nonstop-upgrade reboot \(EX8200 Switch\) on page 455](#)  
[request system software nonstop-upgrade no-old-master-upgrade \(EX8200 Switch\) on page 456](#)  
[request system software nonstop-upgrade reboot \(EX8200 Virtual Chassis\) on page 456](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### [request system software nonstop-upgrade \(EX4200 Virtual Chassis\)](#)

```
user@switch> request system software nonstop-upgrade
/var/tmp/install-ex-4200-12.1R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
```

Installing image on other FPC's along with the backup

Checking pending install on fpc1

Pushing bundle to fpc1

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc1

Checking pending install on fpc2

Pushing bundle to fpc2

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc2

Checking pending install on fpc3

Pushing bundle to fpc3

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc3

Checking pending install on fpc4

Pushing bundle to fpc4

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc4

Checking pending install on fpc5

Pushing bundle to fpc5

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc5

Checking pending install on fpc6

Pushing bundle to fpc6

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc6

Checking pending install on fpc7

Pushing bundle to fpc7

WARNING: A reboot is required to install the software

WARNING: Use the 'request system reboot' command immediately

Completed install on fpc7

Backup upgrade done

Rebooting Backup RE

Rebooting fpc1

ISSU: Backup RE Prepare Done

Waiting for Backup RE reboot

GRES operational

Initiating Chassis In-Service-Upgrade

Chassis ISSU Started

ISSU: Preparing Daemons

ISSU: Daemons Ready for ISSU

ISSU: Starting Upgrade for FRUs

ISSU: Preparing for Switchover

ISSU: Ready for Switchover

Checking In-Service-Upgrade status

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	

```
FPC 2      Online (ISSU)
FPC 3      Online (ISSU)
FPC 4      Online (ISSU)
FPC 5      Online (ISSU)
FPC 6      Online (ISSU)
FPC 7      Online (ISSU)
Going to install image on master
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE

*** FINAL System shutdown message from root@switch ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 9336]
```

#### request system software nonstop-upgrade (EX6200 Switch)

```
{master}
user@switch> request system software nonstop-upgrade
/var/tmp/jinstall-ex-6200-12.2R5.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re0
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re0
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
FPC 4	Online	
FPC 5	Online	
FPC 6	Online (ISSU)	
FPC 7	Online (ISSU)	
FPC 8	Online (ISSU)	
FPC 9	Online (ISSU)	



```

Going to install image on master
NOTICE: Validating configuration against
jinstall-ex-6200-12.2R5.5-domestic-signed.tgz.
NOTICE: Use the 'no-validate' option to skip this if desired.
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
relinquish mastership
ISSU: IDLE
Trying to relinquish mastership before rebooting...
Resolving mastership...
Complete. The other routing engine becomes the master.

*** FINAL System shutdown message from user@switch ***

System going down IMMEDIATELY

```

### request system software nonstop-upgrade reboot (EX8200 Switch)

```

{master}
user@switch> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status

```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 2	Offline	Offlined by CLI command
FPC 3	Online (ISSU)	

```

Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
ISSU: Upgrading Old Master RE
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
ISSU: Old Master Upgrade Done
ISSU: IDLE
Shutdown NOW!
[pid 2635]

*** FINAL System shutdown message from user@switch ***
System going down IMMEDIATELY

```

**request system software nonstop-upgrade no-old-master-upgrade (EX8200 Switch)**

```
{master}
user@switch> request system software nonstop-upgrade no-old-master-upgrade
/var/tmp/jinstall-ex-8200-10.4R1.5-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing Backup RE
Pushing bundle to re1
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
Backup upgrade done
Rebooting Backup RE

Rebooting re1
ISSU: Backup RE Prepare Done
Waiting for Backup RE reboot
GRES operational
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
  FPC 1         Online (ISSU)
  FPC 2         Online (ISSU)
  FPC 3         Offline          Offlined by CLI command
  FPC 4         Online (ISSU)
  FPC 5         Online (ISSU)
  FPC 6         Online (ISSU)
  FPC 7         Online (ISSU)
Resolving mastership...
Complete. The other routing engine becomes the master.
ISSU: RE switchover Done
Skipping Old Master Upgrade
ISSU: IDLE
```

**request system software nonstop-upgrade reboot (EX8200 Virtual Chassis)**

```
{master:9}
user@external-routing-engine> request system software nonstop-upgrade reboot
/var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
Chassis ISSU Check Done
ISSU: Validating Image
ISSU: Preparing LCC Backup REs
ISSU: Preparing Backup RE
Pushing bundle /var/tmp/jinstall-ex-xre200-11.1-20101130.0-domestic-signed.tgz
to member8
-----
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
VC Backup upgrade done
Rebooting VC Backup RE

Rebooting member8
ISSU: Backup RE Prepare Done
Waiting for VC Backup RE reboot
```

```

Pushing bundle to member0-backup
Pushing bundle to member1-backup
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately
WARNING: A reboot is required to install the software
WARNING: Use the 'request system reboot' command immediately

```

```

Rebooting member0-backup
Rebooting LCC [member0-backup]

```

```

Rebooting member1-backup
Rebooting LCC [member1-backup]
ISSU: LCC Backup REs Prepare Done
GRES operational
Initiating Chassis Nonstop-Software-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking Nonstop-Upgrade status
member0:

```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 5	Online (ISSU)	

```
member1:
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Offline	Offlined due to config
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 5	Online (ISSU)	
FPC 7	Online (ISSU)	

```
member0:
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Online (ISSU)	
FPC 2	Online (ISSU)	
FPC 5	Online (ISSU)	

```
member1:
```

Item	Status	Reason
FPC 0	Online (ISSU)	
FPC 1	Offline	Offlined due to config
FPC 2	Online (ISSU)	
FPC 3	Online (ISSU)	
FPC 4	Online (ISSU)	
FPC 5	Online (ISSU)	
FPC 7	Online (ISSU)	

```

ISSU: Upgrading Old Master RE
Pushing bundle /var/tmp/incoming-package-8200.tgz to member0-master

```

```
Pushing bundle /var/tmp/incoming-package-8200.tgz to member1-master
```

```
ISSU: RE switchover Done
```

```
WARNING: A reboot is required to install the software
```

```
WARNING: Use the 'request system reboot' command immediately
```

```
Rebooting ...
```

```
shutdown: [pid 2188]
```

```
Shutdown NOW!
```

```
ISSU: Old Master Upgrade Done
```

```
ISSU: IDLE
```

```
Shutdown NOW!
```

```
*** FINAL System shutdown message from root@ ***
```

```
System going down IMMEDIATELY
```

## request system software rollback

<b>List of Syntax</b>	<a href="#">Syntax on page 459</a> <a href="#">Syntax (EX Series Switches) on page 459</a> <a href="#">Syntax (TX Matrix Router) on page 459</a> <a href="#">Syntax (TX Matrix Plus Router) on page 459</a> <a href="#">Syntax (MX Series Router) on page 459</a>
<b>Syntax</b>	request system software rollback
<b>Syntax (EX Series Switches)</b>	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
<b>Syntax (TX Matrix Router)</b>	request system software rollback <lcc <i>number</i>   scc> <reboot>
<b>Syntax (TX Matrix Plus Router)</b>	request system software rollback <lcc <i>number</i>   sfc <i>number</i> > <reboot>
<b>Syntax (MX Series Router)</b>	request system software rollback <all-members> <local> <member <i>member-id</i> > <reboot>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command behavior changed in Junos OS Release 12.1. Option <b>reboot</b> introduced in Junos OS Release 12.3.
<b>Description</b>	<p>For all versions of Junos OS up to and including Junos OS 11.4, revert to the software that was loaded at the last successful <b>request system software add</b> command.</p> <p>As of Junos OS 12.1 and greater, revert to the last known good state before the most recent <b>request system software (add   delete)</b> command. For example, using rollback in Junos OS 12.1 after using <b>request system software add</b> restores the system to a known good state prior to using the <b>add</b> command. Similarly, using rollback in Junos OS 12.1 after using <b>request system software delete</b> restores the system to a known good state prior to using the <b>delete</b> command.</p> <p>A software rollback fails if any required package (or a <b>bundle</b> package containing the required package) cannot be found in <code>/var/sw/pkg</code>.</p> <p><i>Additional Information</i></p>

- On M Series and T Series routers, if **request system software add <jinstall> reboot** was used for the previous installation, then **request system software rollback** has no effect. In this case, use **jinstall** to reinstall the required package.
- On M Series and T Series routers, if **request system software add <sdk1>** was used for the previous installation, then **request system software rollback** removes the last installed SDK package (**sdk1** in this example).
- On SRX Series devices with dual root systems, when **request system software rollback** is run, the system switches to the alternate root. Each root can have a different version of Junos OS. Rollback takes each root back to the previously installed image.
- On QFX3500 and QFX3600 devices in a mixed Virtual Chassis, when the **request system software rollback** command is issued, the system does not rollback to the image stored in the alternate partition.
- On QFX5100 switches, the **reboot** option has been removed. To reboot the switch after a software rollback, issue the **request system reboot** command as a separate, secondary command.

**Options** **all-members**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, attempt to roll back to the previous set of packages on a T640 router connected to the TX Matrix router. On a TX Matrix Plus router, attempt to roll back to the previous set of packages on a connected router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

**member member-id**—(EX4200 switches and MX Series routers only) (Optional) Attempt to roll back to the previous set of packages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**none**—For all versions of Junos OS up to and including Junos OS 11.4, revert to the set of software as of the last successful **request system software add**. As of Junos OS 12.1 and greater, revert to the last known good state before the most recent **request system software (add | delete)** command.

**reboot**—As of Junos OS 12.3 and greater, automatically reboot upon completing the **request system software rollback** command.

**scc**—(TX Matrix routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix router (or switch-card chassis).

**sfc number**—(TX Matrix Plus routers only) (Optional) Attempt to roll back to the previous set of packages on the TX Matrix Plus router. Replace *number* with 0.

**Required Privilege Level**

maintenance

**Related Documentation**

- *request system software abort*
- [request system software add on page 421](#)
- [request system software delete on page 430](#)
- [request system software validate on page 463](#)
- [request system configuration rescue delete on page 398](#)
- [request system configuration rescue save on page 399](#)
- *Routing Matrix with a TX Matrix Plus Router Solutions Page*

**List of Sample Output** [request system software rollback on page 462](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software rollback

```
user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host
```



## request system software validate

<b>List of Syntax</b>	<a href="#">Syntax on page 463</a> <a href="#">Syntax (TX Matrix Router) on page 463</a> <a href="#">Syntax (TX Matrix Plus Router) on page 463</a> <a href="#">Syntax (MX Series Router) on page 463</a>
<b>Syntax</b>	<pre>request system software validate <i>package-name</i> &lt;set [<i>package-name package-name</i>]&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>request system software validate <i>package-name</i> &lt;lcc <i>number</i>   scc&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>request system software validate <i>package-name</i> &lt;lcc <i>number</i>   sfc <i>number</i>&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>request system software validate <i>package-name</i> &lt;member <i>member-id</i>&gt; &lt;set [<i>package-name package-name</i>]&gt; &lt;upgrade-with-config&gt; &lt;upgrade-with-config-format <i>format</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>set [<i>package-name package-name</i>]</b> option added in Junos OS Release 12.2 for M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways.</p> <p><b>upgrade-with-config</b> and <b>upgrade-with-config-format <i>format</i></b> options added in Junos OS Release 12.3 for M Series routers, MX Series routers, and T Series routers.</p>
<b>Description</b>	Validate candidate software against the current configuration of the router.
<b>Options</b>	<p><b>lcc <i>number</i></b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, validate the software bundle or package on a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, validate the software bundle or package for a specific router that is connected to the TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> </ul>

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**member *member-id***—(MX Series routers only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

***package-name***—Name of the software bundle or package to test.

**scc**—(TX Matrix routers only) (Optional) Validate the software bundle or package for the TX Matrix router (or switch-card chassis).

**set [*package-name package-name*]**—(M Series, MX Series, T Series routers, and Branch SRX Series Services Gateways only) (Optional) Install multiple software packages or software add-on packages at the same time.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Validate the software bundle or package for the TX Matrix Plus router.

**upgrade-with-config**—(Optional) Install one or more configuration files.

**upgrade-with-config-format *format***—(Optional) Specify the configuration file format, **text** or **xml**. The default format is **text**.



**NOTE:** The **upgrade-with-config** and **upgrade-with-config-format** options are only available locally on the router or switch. In a routing matrix, the configuration is applied only to the local router and is not propagated to other routers.

The options are validated during the validation process and applied to the router or switch during the upgrade process. If the upgrade process is successful, the options are removed from the configuration. If the upgrade process fails, the configuration file is renamed with the **.failed** suffix.

---

**Additional Information** By default, when you issue the **request system software validate** command on a TX Matrix master Routing Engine, all the T640 master Routing Engines that are connected to it are validated. If you issue the same command on the TX Matrix backup Routing Engine, all the T640 backup Routing Engines that are connected to it are upgraded to the same version of software.

Likewise, if you issue the **request system software validate** command on a TX Matrix Plus master Routing Engine, all the T1600 or T4000 master Routing Engines that are connected to it are validated. If you issue the same command on a TX Matrix Plus backup Routing Engine, all the T1600 or T4000 backup Routing Engines that are connected to it are upgraded to the same version of software.

<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>request system software abort</i></li> <li>• <a href="#">request system software add on page 421</a></li> <li>• <a href="#">request system software delete on page 430</a></li> <li>• <a href="#">request system software rollback on page 459</a></li> <li>• <i>Routing Matrix with a TX Matrix Plus Router Solutions Page</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system software validate (Successful Case) on page 465</a> <a href="#">request system software validate (Failure Case) on page 465</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request system software validate (Successful Case)

```

user@host> request system software validate /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)

```

### request system software validate (Failure Case)

```

user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0

user@host> request system software validate test
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test

```

## request system storage cleanup

---

<b>List of Syntax</b>	<a href="#">Syntax on page 466</a> <a href="#">Syntax (EX Series Switches) on page 466</a> <a href="#">Syntax (MX Series Router) on page 466</a> <a href="#">Syntax (QFX Series) on page 466</a>
<b>Syntax</b>	request system storage cleanup <dry-run>
<b>Syntax (EX Series Switches)</b>	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> >
<b>Syntax (MX Series Router)</b>	request system storage cleanup <all-members> <dry-run> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	request system storage cleanup <component ( <i>serial number</i>   <i>UUID</i>   all)> <director-group <i>name</i> > <dry-run> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <name-tag <i>name-tag</i> > <node-group <i>name</i> > <prune> <qfabric ( <i>component name</i> )   dry-run   name-tag   repository> <repository ( <i>core</i>   log)>
<b>Release Information</b>	Command introduced in Junos OS Release 7.4. <b>dry-run</b> option introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Free storage space on the router or switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion. On a QFabric system, you can delete debug files located on individual devices or on the entire QFabric system.
<b>Options</b>	<b>all-members</b> —(EX4200 switches and MX Series routers only) (Optional) Delete files on the Virtual Chassis master Routing Engine only.



**NOTE:** To delete files on the other members of the Virtual Chassis configuration, log in to each backup Routing Engine and delete the files using the **request system storage cleanup local** command.

---

**component** (*UUID | serial number | all*)—(QFabric systems only) (Optional) Delete files located on individual QFabric system devices or on the entire QFabric system.

**director-group** *name*—(QFabric systems only) (Optional) Delete files on the Director group.

**dry-run**—(Optional) List files proposed for deletion (without deleting them).

**infrastructure** *name*—(QFabric systems only) (Optional) Delete files on the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device** *name*—(QFabric systems only) (Optional) Delete files on the Interconnect device.

**local**—(EX4200 switches and MX Series routers only) (Optional) Delete files on the local Virtual Chassis member.

**member** *member-id*—(EX4200 switches and MX Series routers only) (Optional) Delete files on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**name-tag** *name-tag*—(QFabric systems only) (Optional) Delete debug files that match a specific regular expression.

**node-group** *name*—(QFabric systems only) (Optional) Delete files on the Node group.

**prune**—(QFabric systems only) (Optional) Delete debug files located in either the core or log debug repositories of a QFabric system device.

**qfabric component** *name*—(QFabric systems only) (Optional) Delete debug files located in the debug repositories of a QFabric system device.

**repository** (*core | log*)—(QFabric systems only) (Optional) Specify the repository on the QFabric system device for which you want to delete debug files.

**Additional Information** If logging is configured and being used, the **dry-run** option rotates the log files. In that case, the output displays the message “Currently rotating log files, please wait.” If no logging is currently under way, the output displays only a list of files to delete.

**Required Privilege Level** maintenance

**List of Sample Output** [request system storage cleanup dry-run on page 468](#)  
[request system storage cleanup on page 469](#)  
[request system storage cleanup director-group \(QFabric Systems\) on page 469](#)  
[request system storage cleanup infrastructure device-name \(QFabric Systems\) on page 471](#)  
[request system storage cleanup interconnect-device device-name \(QFabric Systems\) on page 472](#)  
[request system storage cleanup node-group group-name \(QFabric Systems\) on page 473](#)

[request system storage cleanup qfabric component device-name \(QFabric Systems\) on page 474](#)

[request system storage cleanup qfabric component device-name repository core \(QFabric Systems\) on page 474](#)

[request system storage cleanup qfabric component all \(QFabric Systems\) on page 474](#)

**Output Fields** Table 24 on page 468 describes the output fields for the **request system storage cleanup** command. Output fields are listed in the approximate order in which they appear.

**Table 24: request system storage cleanup Output Fields**

Field Name	Field Description
List of files to delete:	Shows list of files available for deletion.
Size	Size of the core-dump file.
Date	Last core-dump file modification date and time.
Name	Name of the core-dump file.
Directory to delete:	Shows list of directories available for deletion.
Repository scope:	Repository where core-dump files and log files are stored. The core-dump files are located in the <b>core</b> repository, and the log files are located in the <b>log</b> repository. The default <b>Repository scope</b> is shared since both the <b>core</b> and <b>log</b> repositories are shared by all of the QFabric system devices.
Repository head:	Name of the top-level repository location.
Repository name:	Name of the repository: <b>core</b> or <b>log</b> .
Creating list of debug artifacts to be removed under:	Shows location of files available for deletion.
List of debug artifacts to be removed under:	Shows list of files available for deletion.

## Sample Output

### request system storage cleanup dry-run

```
user@host> request system storage cleanup dry-run
Currently rotating log files, please wait.
This operation can take up to a minute.
```

List of files to delete:

Size	Date	Name
11.4K	Mar 8 15:00	/var/log/messages.1.gz
7245B	Feb 5 15:00	/var/log/messages.3.gz
11.8K	Feb 22 13:00	/var/log/messages.2.gz

```

3926B Mar 16 13:57 /var/log/messages.0.gz
3962B Feb 22 12:47 /var/log/sampled.1.gz
4146B Mar 8 12:20 /var/log/sampled.0.gz
4708B Dec 21 11:39 /var/log/sampled.2.gz
7068B Jan 16 18:00 /var/log/messages.4.gz
13.7K Dec 27 22:00 /var/log/messages.5.gz
890B Feb 22 17:22 /var/tmp/sampled.pkts
65.8M Oct 26 09:10 /var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
63.1M Oct 26 09:13 /var/sw/pkg/jbundle-7.4R1.7.tgz

```

### request system storage cleanup

```

user@host> request system storage cleanup
Currently rotating log files, please wait.
This operation can take up to a minute.

```

List of files to delete:

	Size	Date	Name
11.4K	Mar 8 15:00	/var/log/messages.1.gz	
7245B	Feb 5 15:00	/var/log/messages.3.gz	
11.8K	Feb 22 13:00	/var/log/messages.2.gz	
3926B	Mar 16 13:57	/var/log/messages.0.gz	
11.6K	Mar 8 15:00	/var/log/messages.5.gz	
7254B	Feb 5 15:00	/var/log/messages.6.gz	
12.9K	Feb 22 13:00	/var/log/messages.8.gz	
3726B	Mar 16 13:57	/var/log/messages.7.gz	
3962B	Feb 22 12:47	/var/log/sampled.1.gz	
4146B	Mar 8 12:20	/var/log/sampled.0.gz	
4708B	Dec 21 11:39	/var/log/sampled.2.gz	
7068B	Jan 16 18:00	/var/log/messages.4.gz	
13.7K	Dec 27 22:00	/var/log/messages.5.gz	
890B	Feb 22 17:22	/var/tmp/sampled.pkts	
65.8M	Oct 26 09:10	/var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz	
63.1M	Oct 26 09:13	/var/sw/pkg/jbundle-7.4R1.7.tgz	

Delete these files ? [yes,no] (yes)

### request system storage cleanup director-group (QFabric Systems)

```

user@switch> request system storage cleanup director-group
List of files to delete:

```

	Size	Date	Name
4.0K	2011-11-07 05:16:29	/tmp/2064.sfcauth	
4.0K	2011-11-07 05:07:34	/tmp/30804.sfcauth	
4.0K	2011-11-07 04:13:41	/tmp/26792.sfcauth	
4.0K	2011-11-07 04:13:39	/tmp/26432.sfcauth	
0	2011-11-07 07:45:40	/tmp/cluster_cleanup.log	
1.3M	2011-11-07 07:39:11	/tmp/cn_monitor.20111107-052401.log	
4.0K	2011-11-07 07:36:29	/tmp/clustat.28019.log	
4.0K	2011-11-07 07:36:29	/tmp/clustat_x.28019.log	
9.6M	2011-11-07 05:30:24	/tmp/sfc.2.log	
4.0K	2011-11-07 05:28:11	/tmp/mgd-init.1320672491.log	
248K	2011-11-07 05:19:24	/tmp/cn_monitor.20111107-045111.log	
4.0K	2011-11-07 05:17:18	/tmp/clustat.3401.log	
4.0K	2011-11-07 05:17:18	/tmp/clustat_x.3401.log	
8.0K	2011-11-07 04:58:25	/tmp/mgd-init.1320670633.log	
0	2011-11-07 04:54:01	/tmp/mysql_db_install_5.1.37.log	
4.0K	2011-11-07 04:52:08	/tmp/cn_send.log	
0	2011-11-07 04:52:00	/tmp/init_eth0.log	

```

4.0K 2011-11-07 04:49:35 /tmp/install_interfaces.sh.log
4.0K 2011-11-07 04:48:15 /tmp/bootstrap.sh.log
160K 2011-11-07 04:47:43 /tmp/bootstrap_cleanup.log
38M 2011-11-07 04:42:42 /tmp/cn_monitor.20111104-110308.log
4.0K 2011-11-07 04:38:47 /tmp/clustat.30913.log
4.0K 2011-11-07 04:38:47 /tmp/clustat_x.30913.log
4.0K 2011-11-07 04:38:03 /tmp/dcf_upgrade.sh.remove.log
4.0K 2011-11-07 04:38:03 /tmp/peer_update.log
4.0K 2011-11-07 04:38:02 /tmp/dcf_upgrade.log
4.0K 2011-11-07 04:38:02 /tmp/perl_mark_upgrade.log
8.0K 2011-11-07 04:13:42 /tmp/install_dcf_rpm.log
4.0K 2011-11-07 04:13:06 /tmp/00_cleanup.sh.1320667986.log
0 2011-11-07 04:13:06 /tmp/ccif_patch_4410_4450.sh.1320667986.log
4.0K 2011-11-07 04:13:06 /tmp/dcf-tools.sh.1320667986.log
0 2011-11-07 04:13:06 /tmp/initial.sh.1320667986.log
0 2011-11-07 04:13:06 /tmp/inventory.sh.1320667986.log
4.0K 2011-11-07 04:13:06 /tmp/qf-db.sh.1320667986.log
4.0K 2011-11-07 04:13:06 /tmp/sfc.sh.1320667986.log
8.0K 2011-11-07 04:13:05 /tmp/jinstall-qfabric.log
8.0K 2011-11-04 11:10:24 /tmp/mgd-init.1320430192.log
4.0K 2011-11-04 11:07:03 /tmp/mysql_dcf_db_install.log
8.0K 2011-11-04 10:55:07 /tmp/ccif_patch_4410_4450.sh.1320429307.log
8.0K 2011-11-04 10:55:07 /tmp/initial.sh.1320429307.log
4.0K 2011-11-04 10:55:07 /tmp/inventory.sh.1320429307.log
8.0K 2011-11-04 10:55:07 /tmp/sfc.sh.1320429307.log
4.0K 2011-11-04 10:54:09 /tmp/ks-script-Ax0tz5.log
4.0K 2011-11-07 04:13:06 /tmp//sfc.sh.1320667986.log
8.0K 2011-11-04 10:55:07 /tmp//sfc.sh.1320429307.log

```

## Directory to delete:

```

45M 2011-11-08 10:57:43 /tmp/sfc-captures

```

## List of files to delete:

	Size	Date	Name
4.0K	2011-11-08	05:47:47	/tmp/5713.sfcauth
4.0K	2011-11-08	05:14:32	/tmp/14494.sfcauth
4.0K	2011-11-08	05:11:47	/tmp/9978.sfcauth
4.0K	2011-11-08	05:09:37	/tmp/6128.sfcauth
4.0K	2011-11-08	05:04:28	/tmp/29703.sfcauth
4.0K	2011-11-07	11:59:10	/tmp/7811.sfcauth
4.0K	2011-11-07	11:36:08	/tmp/32415.sfcauth
4.0K	2011-11-07	11:30:30	/tmp/22406.sfcauth
4.0K	2011-11-07	11:24:37	/tmp/12131.sfcauth
4.0K	2011-11-07	10:48:42	/tmp/12687.sfcauth
4.0K	2011-11-07	09:27:20	/tmp/31082.sfcauth
4.0K	2011-11-07	07:33:58	/tmp/14633.sfcauth
4.0K	2011-11-07	05:08:25	/tmp/15447.sfcauth
4.0K	2011-11-07	04:12:29	/tmp/26874.sfcauth
4.0K	2011-11-07	04:12:27	/tmp/26713.sfcauth
4.0K	2011-11-07	03:49:17	/tmp/17691.sfcauth
4.0K	2011-11-05	01:32:23	/tmp/5716.sfcauth
4.0K	2011-11-07	08:00:17	/tmp/sfcsnmpd.log
4.0K	2011-11-07	07:57:50	/tmp/cluster_cleanup.log
824K	2011-11-07	07:38:37	/tmp/cn_monitor.20111107-053643.log
4.0K	2011-11-07	07:36:30	/tmp/clustat.18399.log
4.0K	2011-11-07	07:36:30	/tmp/clustat_x.18399.log
4.0K	2011-11-07	07:35:47	/tmp/command_lock.log
4.0K	2011-11-07	05:39:54	/tmp/mgd-init.1320673194.log
92K	2011-11-07	05:19:25	/tmp/cn_monitor.20111107-050412.log
4.0K	2011-11-07	05:17:20	/tmp/clustat.30115.log



```

4.0K  2011-11-07 05:17:20 /tmp/clustat_x.30115.log
8.0K  2011-11-07 05:08:07 /tmp/mgd-init.1320671241.log
4.0K  2011-11-07 05:04:57 /tmp/cn_send.log
0     2011-11-07 05:04:52 /tmp/init_eth0.log
4.0K  2011-11-07 05:02:38 /tmp/install_interfaces.sh.log
4.0K  2011-11-07 05:01:19 /tmp/bootstrap.sh.log
160K  2011-11-07 05:00:47 /tmp/bootstrap_cleanup.log
28M   2011-11-07 04:42:27 /tmp/cn_monitor.20111104-112954.log
4.0K  2011-11-07 04:38:49 /tmp/clustat.6780.log
4.0K  2011-11-07 04:38:49 /tmp/clustat_x.6780.log
4.0K  2011-11-07 04:38:05 /tmp/issue_event.log
4.0K  2011-11-07 04:38:05 /tmp/peer_upgrade_reboot.log
12K   2011-11-07 04:38:05 /tmp/primary_update.log
4.0K  2011-11-07 04:38:04 /tmp/dcf_upgrade.sh.remove.log
4.0K  2011-11-07 04:38:04 /tmp/peer_rexec_upgrade.log
4.0K  2011-11-07 04:13:42 /tmp/peer_install_dcf_rpm.log
4.0K  2011-11-07 04:11:57 /tmp/dcf-tools.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/initial.sh.1320667917.log
0     2011-11-07 04:11:57 /tmp/inventory.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/qf-db.sh.1320667917.log
4.0K  2011-11-07 04:11:57 /tmp/sfc.sh.1320667917.log
4.0K  2011-11-07 04:11:56 /tmp/00_cleanup.sh.1320667916.log
0     2011-11-07 04:11:56 /tmp/ccif_patch_4410_4450.sh.1320667916.log
8.0K  2011-11-07 04:11:56 /tmp/jinstall-qfabric.log
4.0K  2011-11-07 04:11:33 /tmp/dcf_upgrade.log
8.0K  2011-11-04 11:53:12 /tmp/mgd-init.1320432782.log
8.0K  2011-11-04 11:06:17 /tmp/ccif_patch_4410_4450.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/initial.sh.1320429977.log
4.0K  2011-11-04 11:06:17 /tmp/inventory.sh.1320429977.log
8.0K  2011-11-04 11:06:17 /tmp/sfc.sh.1320429977.log
4.0K  2011-11-04 11:05:19 /tmp/ks-script-tnWeb.log
4.0K  2011-11-07 04:11:57 /tmp//sfc.sh.1320667917.log
8.0K  2011-11-04 11:06:17 /tmp//sfc.sh.1320429977.log

```

Directory to delete:

```
49M   2011-11-08 10:45:20 /tmp/sfc-captures
```

#### request system storage cleanup infrastructure device-name (QFabric Systems)

```
user@switch> request system storage cleanup infrastructure FC-0
re0:
```

-----

List of files to delete:

	Size	Date	Name
	139B	Nov 8 19:03	/var/log/default-log-messages.0.gz
	5602B	Nov 8 19:03	/var/log/messages.0.gz
	28.4K	Nov 8 10:15	/var/log/messages.1.gz
	35.2K	Nov 7 13:45	/var/log/messages.2.gz
	207B	Nov 7 16:02	/var/log/wtmp.0.gz
	27B	Nov 7 12:14	/var/log/wtmp.1.gz
	184.4M	Nov 7 12:16	/var/sw/pkg/jinstall-dc-re-11.3I20111104_1216_dc-builder-domestic-signed.tgz
	124.0K	Nov 7 15:59	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:57	/var/tmp/gres-tp/lock
	155B	Nov 7 16:02	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 12:35	/var/tmp/last_ccif_update
	1217B	Nov 7 12:15	/var/tmp/loader.conf.preinstall
	184.4M	Nov 6 07:11	/var/tmp/mchassis-install.tgz
	10.8M	Nov 7 12:16	

```

/var/tmp/preinstall/bootstrap-install-11.3I20111104_1216_dc-builder.tar
57.4K Nov  7 12:16 /var/tmp/preinstall/configs-11.3I20111104_1216_dc-builder.tgz

259B Nov  7 12:16 /var/tmp/preinstall/install.conf
734.3K Nov  4 13:46
/var/tmp/preinstall/jboot-dc-re-11.3I20111104_1216_dc-builder.tgz
177.8M Nov  7 12:16
/var/tmp/preinstall/jbundle-dc-re-11.3I20111104_1216_dc-builder-domestic.tgz
124B Nov  7 12:15 /var/tmp/preinstall/metatags
1217B Nov  7 12:16 /var/tmp/preinstall_boot_loader.conf
0B Nov  7 16:02 /var/tmp/rtssdb/if-rtssdb

```

### request system storage cleanup interconnect-device device-name (QFabric Systems)

```

user@switch> request system storage cleanup interconnect IC-WS001
re1:
-----

```

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	128B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	9965B	Nov 8 19:06	/var/log/messages.0.gz
	15.8K	Nov 8 12:30	/var/log/messages.1.gz
	15.8K	Nov 8 11:00	/var/log/messages.2.gz
	15.7K	Nov 8 07:30	/var/log/messages.3.gz
	15.8K	Nov 8 04:00	/var/log/messages.4.gz
	15.7K	Nov 8 00:30	/var/log/messages.5.gz
	18.7K	Nov 7 21:00	/var/log/messages.6.gz
	17.6K	Nov 7 19:00	/var/log/messages.7.gz
	58.3K	Nov 7 16:00	/var/log/messages.8.gz
	20.3K	Nov 7 15:15	/var/log/messages.9.gz
	90B	Nov 7 15:41	/var/log/wtmp.0.gz
	57B	Nov 7 12:41	/var/log/wtmp.1.gz
	124.0K	Nov 7 15:42	/var/tmp/gres-tp/env.dat
	0B	Nov 7 12:40	/var/tmp/gres-tp/lock
	0B	Nov 7 12:41	/var/tmp/if-rtssdb/env.lock
	12.0K	Nov 7 15:41	/var/tmp/if-rtssdb/env.mem
	132.0K	Nov 7 15:55	/var/tmp/if-rtssdb/shm_usr1.mem
	2688.0K	Nov 7 15:41	/var/tmp/if-rtssdb/shm_usr2.mem
	2048.0K	Nov 7 15:41	/var/tmp/if-rtssdb/trace.mem
	730B	Nov 7 19:57	/var/tmp/juniper.conf+.gz
	155B	Nov 7 15:53	/var/tmp/krt_gencfg_filter.txt
	0B	Nov 7 15:41	/var/tmp/rtssdb/if-rtssdb

```

re0:
-----

```

List of files to delete:

	Size	Date	Name
	11B	Nov 7 15:55	/var/jail/tmp/alarmd.ts
	121B	Nov 8 19:06	/var/log/default-log-messages.0.gz
	16.7K	Nov 8 19:06	/var/log/messages.0.gz
	22.2K	Nov 8 17:45	/var/log/messages.1.gz
	18.4K	Nov 8 17:00	/var/log/messages.2.gz
	21.6K	Nov 8 16:00	/var/log/messages.3.gz
	17.9K	Nov 8 14:30	/var/log/messages.4.gz
	19.4K	Nov 8 13:30	/var/log/messages.5.gz
	18.2K	Nov 8 12:30	/var/log/messages.6.gz

```

20.4K Nov  8 11:30 /var/log/messages.7.gz
21.4K Nov  8 10:15 /var/log/messages.8.gz
21.0K Nov  8 09:00 /var/log/messages.9.gz
19.9K Nov  8 08:13 /var/log/snmp-traps.0.gz
203B Nov  8 15:36 /var/log/wtmp.0.gz
57B Nov  7 12:41 /var/log/wtmp.1.gz
124.0K Nov  7 15:42 /var/tmp/gres-tp/env.dat
0B Nov  7 12:40 /var/tmp/gres-tp/lock
0B Nov  7 12:41 /var/tmp/if-rtssdb/env.lck
12.0K Nov  7 15:41 /var/tmp/if-rtssdb/env.mem
132.0K Nov  7 15:55 /var/tmp/if-rtssdb/shm_usr1.mem
2688.0K Nov  7 15:41 /var/tmp/if-rtssdb/shm_usr2.mem
2048.0K Nov  7 15:41 /var/tmp/if-rtssdb/trace.mem
727B Nov  7 15:54 /var/tmp/juniper.conf+.gz
155B Nov  7 15:55 /var/tmp/krt_gencfg_filter.txt
0B Nov  7 15:41 /var/tmp/rtssdb/if-rtssdb

```

#### request system storage cleanup node-group group-name (QFabric Systems)

```

user@switch> request system storage cleanup node-group NW-NG-0
BBAK0372:

```

-----

List of files to delete:

	Size	Date	Name
	126B	Nov  8 19:07	/var/log/default-log-messages.0.gz
	179B	Nov  7 13:32	/var/log/install.0.gz
	22.9K	Nov  8 19:07	/var/log/messages.0.gz
	26.5K	Nov  8 17:30	/var/log/messages.1.gz
	20.5K	Nov  8 13:15	/var/log/messages.2.gz
	33.2K	Nov  7 17:45	/var/log/messages.3.gz
	35.5K	Nov  7 15:45	/var/log/messages.4.gz
	339B	Nov  8 17:10	/var/log/wtmp.0.gz
	58B	Nov  7 12:40	/var/log/wtmp.1.gz
	124.0K	Nov  8 17:08	/var/tmp/gres-tp/env.dat
	0B	Nov  7 12:39	/var/tmp/gres-tp/lock
	0B	Nov  7 12:59	/var/tmp/if-rtssdb/env.lck
	12.0K	Nov  8 17:09	/var/tmp/if-rtssdb/env.mem
	2688.0K	Nov  8 17:09	/var/tmp/if-rtssdb/shm_usr1.mem
	132.0K	Nov  8 17:09	/var/tmp/if-rtssdb/shm_usr2.mem
	2048.0K	Nov  8 17:09	/var/tmp/if-rtssdb/trace.mem
	1082B	Nov  8 17:09	/var/tmp/juniper.conf+.gz
	155B	Nov  7 17:39	/var/tmp/krt_gencfg_filter.txt
	0B	Nov  8 17:09	/var/tmp/rtssdb/if-rtssdb

EE3093:

-----

List of files to delete:

	Size	Date	Name
	11B	Nov  8 17:33	/var/jail/tmp/alarmd.ts
	119B	Nov  8 19:08	/var/log/default-log-messages.0.gz
	180B	Nov  7 17:41	/var/log/install.0.gz
	178B	Nov  7 13:32	/var/log/install.1.gz
	2739B	Nov  8 19:08	/var/log/messages.0.gz
	29.8K	Nov  8 18:45	/var/log/messages.1.gz
	31.8K	Nov  8 17:15	/var/log/messages.2.gz
	20.6K	Nov  8 16:00	/var/log/messages.3.gz
	15.4K	Nov  8 10:15	/var/log/messages.4.gz

```

15.4K Nov  8 02:15 /var/log/messages.5.gz
25.5K Nov  7 20:45 /var/log/messages.6.gz
48.0K Nov  7 17:45 /var/log/messages.7.gz
32.8K Nov  7 13:45 /var/log/messages.8.gz
684B Nov  8 17:02 /var/log/wtmp.0.gz
58B Nov  7 12:40 /var/log/wtmp.1.gz
124.0K Nov  7 17:34 /var/tmp/gres-tp/env.dat
  0B Nov  7 12:40 /var/tmp/gres-tp/lock
  0B Nov  7 12:59 /var/tmp/if-rtbdb/env.lck
12.0K Nov  7 17:39 /var/tmp/if-rtbdb/env.mem
2688.0K Nov  7 17:39 /var/tmp/if-rtbdb/shm_usr1.mem
132.0K Nov  7 17:40 /var/tmp/if-rtbdb/shm_usr2.mem
2048.0K Nov  7 17:39 /var/tmp/if-rtbdb/trace.mem
155B Nov  7 17:40 /var/tmp/krt_gencfg_filter.txt
  0B Nov  7 17:39 /var/tmp/rtbdb/if-rtbdb

```

### request system storage cleanup qfabric component device-name (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component A0001/YA0197
Repository type: regular
Repository head: /pbstorage
Creating list of debug artifacts to be removed under:
/pbstorage/rdumps/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.0.0.05162011123308.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.1.0.05162011123614.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/cosd.core.2.0.05162011123920.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/livekcore.05132011163930.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/tnetd.core.0.1057.05162011124500.gz ...
done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.05132011120528.gz ... done
Removing /pbstorage/rdumps/A0001/YA0197/vmcore.kz ... done
Creating list of debug artifacts to be removed under: /pbstorage/rlogs/A0001/YA0197
Removing debug artifacts ... (press control C to abort)
Removing /pbstorage/rlogs/A0001/YA0197/kdumpinfo.05132011120528 ... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.0.1039.05122011234415.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/kernel.tarball.1.1039.05132011175544.tgz
... done
Removing /pbstorage/rlogs/A0001/YA0197/tnetd.tarball.0.1057.05162011175453.tgz
... done

```

### request system storage cleanup qfabric component device-name repository core (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component EE3093 repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps/EE3093
NOTE: core repository under /pbdata/export/rdumps/EE3093 empty

```

### request system storage cleanup qfabric component all (QFabric Systems)

```

user@switch> request system storage cleanup qfabric component all
Repository scope: shared
Repository head: /pbdata/export
Creating list of debug artifacts to be removed under: /pbdata/export/rdumps
NOTE: core repository under /pbdata/export/rdumps/all empty
Creating list of debug artifacts to be removed under: /pbdata/export/rlogs
List of debug artifacts to clean up ... (press control C to abort)
/pbdata/export/rlogs/73747cd8-0710-11e1-b6a4-00e081c5297e/install-11072011125819.log
/pbdata/export/rlogs/77116f18-0710-11e1-a2a0-00e081c5297e/install-11072011125819.log

```

```
/pbdata/export/rlogs/BBAK0372/install-11072011121538.log  
/pbdata/export/rlogs/BBAK0394/install-11072011121532.log  
/pbdata/export/rlogs/EE3093/install-11072011121536.log  
/pbdata/export/rlogs/WS001/YN5999/install-11072011121644.log  
/pbdata/export/rlogs/WS001/YW3803/install-11072011122429.log  
/pbdata/export/rlogs/cd78871a-0710-11e1-878e-00e081c5297e/install-11072011125932.log  
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011125930.log  
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011133211.log  
/pbdata/export/rlogs/d0afda1e-0710-11e1-a1d0-00e081c5297e/install-11072011155302.log  
/pbdata/export/rlogs/d31ab7a6-0710-11e1-ad1b-00e081c5297e/install-11072011125931.log  
/pbdata/export/rlogs/d4d0f254-0710-11e1-90c3-00e081c5297e/install-11072011125932.log
```

## request system zeroize

---

**Syntax**    request system zeroize  
              <media>  
              <local>

**Release Information**    Command introduced before Junos OS Release 9.0.  
                              Command introduced in Junos OS Release 11.2 for EX Series switches.  
                              Option **media** added in Junos OS Release 11.4 for EX Series switches.  
                              Command introduced in Junos OS Release 12.2 for MX Series devices.  
                              Command introduced in Junos OS Release 12.3 for the QFX Series.  
                              Option *local* added in Junos OS Release 14.1.

**Description**



**NOTE:** The **media** option is not available on the QFX Series.

Remove all configuration information on the Routing Engines and reset all key values. If the device has dual Routing Engines, the command is broadcast to all Routing Engines on the device. The command removes all data files, including customized configuration and log files, by unlinking the files from their directories. The command removes all user-created files from the system including all plain-text passwords, secrets, and private keys for SSH, local encryption, local authentication, IPsec, RADIUS, TACACS+, and SNMP.

This command reboots the device and sets it to the factory default configuration. After the reboot, you cannot access the device through the management Ethernet interface. Log in through the console as **root** and start the Junos OS command-line interface (CLI) by typing **cli** at the prompt.

To completely erase user-created data so that it is unrecoverable, use the **media** option.

**Options**    **media**—(Optional) In addition to removing all configuration and log files, the **media** option causes memory and the media to be scrubbed, removing all traces of any user-created files. Every storage device attached to the system is scrubbed, including disks, flash drives, removable USBs, and the like. The duration of the scrubbing process is dependent on the size of the media being erased. As a result, the **request system zeroize media** operation can take considerably more time than the **request system zeroize** operation. However, the critical security parameters are all removed at the beginning of the process.

**local**—(Optional) Removes all the configuration information and restores all the key values on the active Routing Engine.

**Required Privilege Level**    maintenance

**Related Documentation**    • *request system snapshot*  
                                  • [request system snapshot on page 419](#)

- *Reverting to the Default Factory Configuration for the EX Series Switch*
- *Reverting to the Rescue Configuration for the EX Series Switch*
- [Reverting to the Default Factory Configuration on page 188](#)
- [Reverting to the Rescue Configuration on page 189](#)
- [Reverting to the Default Factory Configuration by Using the request system zeroize Command on page 188](#)

List of Sample Output [request system zeroize on page 477](#)  
[request system zeroize media on page 478](#)

## Sample Output

### request system zeroize

```
user@host> request system zeroize
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

0 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 5d19h20m26s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Mar 11 2011 - 04:39:06)

Board: EX4200-24T 2.11
EPLD: Version 6.0 (0x85)
DRAM: Initializing (1024 MB)
FLASH: 8 MB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.4
(user@juniper.net, Fri Mar 11 03:03:36 UTC 2011)
Memory: 1024MB
bootsequencing is enabled
bootsuccess is set
new boot device = disk0s1:
Loading /boot/defaults/loader.conf
/kernel data=0x915c84+0xa1260 syms=[0x4+0x7cbd0+0x4+0xb1c19]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
```

```
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 11.1R1.8 #0: 2011-03-09 20:14:25 UTC

user@juniper.net:/volume/build/junos/11.1/release/11.1R1.8/obj-powerpc/bsd/kernels/
JUNIPER-EX/kernel
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080
...
```

### request system zeroize media

```
user@host> request system zeroize media
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (no) yes

warning: ipsec-key-management subsystem not running - not needed by configuration.
warning: zeroizing fpc0

{master:0}
root> Waiting (max 60 seconds) for system process `vnlr' to stop...done
. . .
Syncing disks, vnodes remaining...2 4 2 4 3 2 1 1 0 0 0 done

syncing disks... All buffers synced.
Uptime: 14m50s
recorded reboot as normal shutdown
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@sv1-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
```



```

All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 50000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: HID0 80004080<EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s2a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20055 free (31 frags, 2503 blocks, 0.0% fragmentation)
zeroizing /dev/da0s1a ...
. . .
zeroizing /dev/da0s3d ...
. . .
zeroizing /dev/da0s3e ...
. . .
zeroizing /dev/da0s4d ...
. . .
zeroizing /dev/da0s4e ...
. . .

syncing disks... All buffers synced.
Uptime: 3m40s
Rebooting...

U-Boot 1.1.6 (Apr 21 2011 - 13:58:42)

Board: EX4200-48PX 1.1
EPLD: Version 8.0 (0x82)
DRAM: Initializing (512 MB)
FLASH: 8 MB
NAND: No NAND device found!!!
0 MiB

Firmware Version: --- 01.00.00 ---
USB: scanning bus for devices... 2 USB Device(s) found
      scanning bus for storage devices... 1 Storage Device(s) found

ELF file is 32 bit
Consoles: U-Boot console

FreeBSD/PowerPC U-Boot bootstrap loader, Revision 2.2
(vtseng@svl-junos-pool27.juniper.net, Fri Feb 26 17:48:51 PST 2010)
Memory: 512MB
Loading /boot/defaults/loader.conf
/kernel data=0x9abfdc+0xb06e4 syms=[0x4+0x83b30+0x4+0xbd7c6]

Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [/kernel] in 1 second... Booting [/kernel]...
Kernel entry at 0x800000e0 ...

```

```
GDB: no debug ports present
KDB: debugger backends: ddb
KDB: current backend: ddb
Copyright (c) 1996-2011, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.
JUNOS 11.4R1.2 #0: 2011-10-27 18:05:39 UTC
user@juniper.net:/volume/build/junos/11.4/release/11.4R1.2/obj-powerpc/
bsd/kernels/JUNIPER-EX/kernel
can't re-use a leaf (all_slot_serialid)!
Timecounter "decrementer" frequency 500000000 Hz quality 0
cpu0: Freescale e500v2 core revision 2.2
cpu0: H1D0 80004080 <EMCP,TBEN,EN_MAS7_UPDATE>
real memory = 511705088 (488 MB)
avail memory = 500260864 (477 MB)
ETHERNET SOCKET BRIDGE initialising
Initializing EXSERIES platform properties ...
. . .
Automatic reboot in progress...
Media check on da0 on ex platforms
** /dev/da0s1a
FILE SYSTEM CLEAN; SKIPPING CHECKS
clean, 20064 free (48 frags, 2502 blocks, 0.1% fragmentation)
zeroizing /dev/da0s2a ...
. . .
Creating initial configuration...mgd: error: Cannot open configuration file:
/config/juniper.conf
mgd: warning: activating factory configuration
mgd: commit complete
mgd: -----
mgd: Please login as 'root'. No password is required.
mgd: To start Initial Setup, type 'ezsetup' at the JUNOS prompt.
mgd: To start JUNOS CLI, type 'cli' at the JUNOS prompt.
mgd: -----
Setting initial options: debugger_on_panic=NO debugger_on_break=NO.
Starting optional daemons: .
Doing initial network setup:
. . .

Amnesiac (ttyu0)
```

## restart

### List of Syntax [Syntax on page 481](#)

[Syntax \(ACX Series Routers\) on page 481](#)  
[Syntax \(EX Series Switches\) on page 481](#)  
[Syntax \(Routing Matrix\) on page 482](#)  
[Syntax \(J Series Routing Platform\) on page 482](#)  
[Syntax \(TX Matrix Routers\) on page 482](#)  
[Syntax \(TX Matrix Plus Routers\) on page 482](#)  
[Syntax \(MX Series Routers\) on page 482](#)  
[Syntax \(J Series Routers\) on page 483](#)  
[Syntax \(QFX Series\) on page 483](#)

### Syntax `restart`

```
<adaptive-services | ancpd-service | application-identification | audit-process |
auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
class-of-service | clksyncd-service | database-replication | datapath-trace-service
| dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
ecc-error-logging | ethernet-connectivity-fault-management
| ethernet-link-fault-management | event-processing | firewall
| general-authentication-service | gracefully | iccp-service | idp-policy | immediately
| interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
| l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
| local-policy-decision-function | mac-validation | mib-process | mobile-ip | mountd-service
| mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
packet-triggered-subscribers | peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations |
root-system-domain-service | routing <logical-system logical-system-name> | sampling
| sbc-configuration-process | sdk-service | service-deployment | services | services pgcp
gateway gateway-name | snmp | soft | static-subscribers | statistics-service |
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
vrrp | web-management>
<gracefully | immediately | soft>
```

### Syntax (ACX Series Routers)

```
restart
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
| disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
| ethernet-link-fault-management | event-processing | firewall
| general-authentication-service | gracefully | immediately | interface-control |
ipsec-key-management | l2-learning | lacp | link-management | mib-process | mobile-ip |
mountd-service | mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service
| ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
sdk-service | secure-neighbor-discovery | service-deployment | services | snmp | soft
| statistics-service | subscriber-management | subscriber-management-helper | tunnel-oamd
| vrrp>
```

### Syntax (EX Series Switches)

```
restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
ethernet-switching | event-processing | firewall | general-authentication-service |
interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
```

	lldpd-service   mib-process   mountd-service   multicast-snooping   pgm   redundancy-interface-process   remote-operations   routing   secure-neighbor-discovery   service-deployment   sflow-service   snmp   vrrp   web-management>
<b>Syntax (Routing Matrix)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp> <all   all-lcc   lcc <i>number</i> > <gracefully   immediately   soft>
<b>Syntax (J Series Routing Platform)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   dhcp   dialer-services   dls   event-processing   firewall   interface-control   ipsec-key-management   isdn-signaling   l2-learning   l2tp-service   mib-process   network-access-service   pgm   ppp   pppoe   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp   usb-control   web-management> <gracefully   immediately   soft>
<b>Syntax (TX Matrix Routers)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp   statistics-service> <all-chassis   all-lcc   lcc <i>number</i>   scc> <gracefully   immediately   soft>
<b>Syntax (TX Matrix Plus Routers)</b>	restart <adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp   statistics-service> <all-chassis   all-lcc   all-sfc   lcc <i>number</i>   sfc <i>number</i> > <gracefully   immediately   soft>
<b>Syntax (MX Series Routers)</b>	restart <adaptive-services   ancpd-service   application-identification   audit-process   auto-configuration   captive-portal-content-delivery   ce-l2tp-service   chassis-control   class-of-service   clksyncd-service   database-replication   datapath-trace-service   dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging   ethernet-connectivity-fault-management   ethernet-link-fault-management   event-processing   firewall   general-authentication-service   gracefully   iccp-service   idp-policy   immediately   interface-control   ipsec-key-management   kernel-replication   l2-learning   l2cpd-service   l2tp-service   l2tp-universal-edge   lacp   license-service   link-management   local-policy-decision-function   mac-validation   mib-process   mobile-ip   mountd-service   mpls-traceroute   msp   multicast-snooping   named-service   nfsd-service

```

packet-triggered-subscribers |peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations
|root-system-domain-service | routing |routing <logical-system logical-system-name> |
sampling | sbc-configuration-process | sdk-service |service-deployment |services | services
pgcp gateway gateway-name |snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>

```

**Syntax (J Series  
Routers)**

```

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dhcp-service
| dialer-services | diameter-service | dlsr | event-processing | firewall | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

```

**Syntax (QFX Series)**

```

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsr | ethernet-connectivity | event-processing | fibre-channel | firewall
| general-authentication-service | igmp-host-services | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
redundancy-interface-process | remote-operations |logical-system-name> | routing |
sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

```

**Release Information**

Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 12.2 for ACX Series routers.  
 Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsr** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

**Description** Restart a Junos OS process.



**CAUTION:** Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

**Options** **none**—Same as **gracefully**.

**adaptive-services**—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

**all-lcc**—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

**all-members**—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

**all-sfc**—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

**ancpd-service**—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

**application-identification**—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

**audit-process**—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

**auto-configuration**—(Optional) Restart the Interface Auto-Configuration process.

**autoinstallation**—(EX Series switches only) (Optional) Restart the autoinstallation process.

**captive-portal-content-delivery**—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

**ce-l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

**chassis-control**—(Optional) Restart the chassis management process.

**class-of-service**—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

**clksyncd-service**—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

**database-replication**—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

**d datapath-trace-service**—(Optional) Restart the packet path tracing process.

**dhcp**—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

**dhcp-service**—(Optional) Restart the Dynamic Host Configuration Protocol process.

**dialer-services**—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

**diameter-service**—(Optional) Restart the diameter process.

**disk-monitoring**—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

**dlsw**—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

**dot1x-protocol**—(EX Series switches only) (Optional) Restart the port-based network access control process.

**dynamic-flow-capture**—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

**ecc-error-logging**—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

**ethernet-connectivity-fault-management**—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

**ethernet-link-fault-management**—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

**ethernet-switching**—(EX Series switches only) (Optional) Restart the Ethernet switching process.

**event-processing**—(Optional) Restart the event process (eventd).

**fibre-channel**—(QFX Series only) (Optional) Restart the Fibre Channel process.

**firewall**—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

**general-authentication-service**—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

**gracefully**—(Optional) Restart the software process.

**iccp-service**—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

**idp-policy**—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

**immediately**—(Optional) Immediately restart the software process.

**interface-control**—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

**ipsec-key-management**—(Optional) Restart the IPsec key management process.

**isdn-signaling**—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

**kernel-replication**—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

**l2-learning**—(Optional) Restart the Layer 2 address flooding and learning process.

**l2cpd-service**—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

**l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

**l2tp-universal-edge**—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

**lACP**—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG,



and to enable the transmission and reception processes for the link to function in an orderly manner.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**license-service**—(EX Series switches only) (Optional) Restart the feature license management process.

**link-management**—(TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

**lldpd-service**—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

**local**—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

**local-policy-decision-function**—(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

**mac-validation**—(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

**member *member-id***—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**mib-process**—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

**mobile-ip**—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

**mountd-service**—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

**mpls-traceroute**—(Optional) Restart the MPLS Periodic Traceroute process.

**mspd**—(Optional) Restart the Multiservice process.

**multicast-snooping**—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

**named-service**—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

**network-access-service**—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

**nfsd-service**—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

**packet-triggered-subscribers**—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

**peer-selection-service**—(Optional) Restart the Peer Selection Service process.

**pgcp-service**—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

**pgm**—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

**pic-services-logging**—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

**pki-service**—(Optional) Restart the PKI Service process.

**ppp**—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

**ppp-service**—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

**pppoe**—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

**protected-system-domain-service**—(Optional) Restart the Protected System Domain (PSD) process.

**redundancy-interface-process**—(Optional) Restart the ASP redundancy process.

**remote-operations**—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

**root-system-domain-service**—(Optional) Restart the Root System Domain (RSD) service.

**routing**—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

**routing <logical-system *logical-system-name*>**—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

**sampling**—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

**sbc-configuration-process**—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

**scc**—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

**sdk-service**—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

**secure-neighbor-discovery**—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**service-deployment**—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

**services**—(Optional) Restart a service.

**services pgcp gateway *gateway-name***—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

**sflow-service**—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

**snmp**—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

**soft**—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

**static-subscribers**—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

**statistics-service**—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

**subscriber-management**—(Optional) Restart the Subscriber Management process.

**subscriber-management-helper**—(Optional) Restart the Subscriber Management Helper process.

**tunnel-oamd**—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

**usb-control**—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

**vrrp**—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**web-management**—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

**Required Privilege Level**

reset

**Related Documentation**

- [Overview of Junos OS CLI Operational Mode Commands on page 58](#)

**List of Sample Output** [restart interfaces on page 490](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```



## rollback

---

<b>Syntax</b>	<code>rollback &lt;number   rescue&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Return to a previously committed configuration. The software saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the <b>commit</b> configuration command.</p> <p>The currently operational Junos OS configuration is stored in the file <b>juniper.conf</b>, and the last three committed configurations are stored in the files <b>juniper.conf.1</b>, <b>juniper.conf.2</b>, and <b>juniper.conf.3</b>. These four files are located in the directory <b>/config</b>, which is on the router's flash drive. The remaining 46 previous committed configurations, the files <b>juniper.conf.4</b> through <b>juniper.conf.49</b>, are stored in the directory <b>/var/db/config</b>, which is on the router's hard disk.</p> <p>During rollback, the configuration you specify is loaded from the associated file. Only objects in the rollback configuration that differ from the previously loaded configuration are marked as changed (equivalent to <b>load update</b>).</p>
<b>Options</b>	<p>none (Optional)—Return to the most recently saved configuration.</p> <p><b>number</b>—(Optional) Configuration to return to. The range of values is from <b>0</b> through <b>49</b>. The most recently saved configuration is number <b>0</b>, and the oldest saved configuration is number <b>49</b>. The default is <b>0</b>.</p> <p><b>rescue</b>—(Optional) Return to the rescue configuration.</p>
<b>Required Privilege Level</b>	rollback—To roll back to configurations other than the one most recently committed.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Returning to a Previously Committed Junos OS Configuration on page 1253</a></li><li>• <a href="#">Creating and Returning to a Rescue Configuration on page 1248</a></li></ul>

## save

<b>Syntax</b>	<code>save <i>filename</i></code>
<b>QFX Series</b>	<code>save (dhcp-snooping <i>filename</i>)</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Save the configuration to an ASCII file. The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.</p> <p>When saving a file to a remote system, the software uses the <b>scp/ssh</b> protocol.</p>
<b>Options</b>	<p><b><i>filename</i></b>—Name of the saved file. You can specify a filename in one of the following ways:</p> <ul style="list-style-type: none"> <li>• <b><i>filename</i></b>—File in the user's home directory (the current directory) on the local flash drive.</li> <li>• <b><i>path/filename</i></b>—File on the local flash drive.</li> <li>• <b><i>/var/filename</i></b> or <b><i>/var/path/filename</i></b>—File on the local hard disk.</li> <li>• <b><i>a:filename</i></b> or <b><i>a:path/filename</i></b>—File on the local drive. The default path is <b>/</b> (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.</li> <li>• <b><i>hostname:/path/filename</i></b>, <b><i>hostname:filename</i></b>, <b><i>hostname:path/filename</i></b>, or <b><i>scp://hostname/path/filename</i></b>—File on an <b>scp/ssh</b> client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify <b><i>hostname</i></b> as <b><i>username@hostname</i></b>.</li> <li>• <b><i>ftp://hostname/path/filename</i></b>—File on an FTP server. You can also specify <b><i>hostname</i></b> as <b><i>username @hostname</i></b> or <b><i>username:password @hostname</i></b>. The default path is the user's home directory. To specify an absolute path, the path must start with the string <b>%2F</b>; for example, <b><i>ftp://hostname/%2Fpath/filename</i></b>. To have the system prompt you for the password, specify <b><i>prompt</i></b> in place of the password. If a password is required, and you do not specify the password or <b><i>prompt</i></b>, an error message is displayed: <ul style="list-style-type: none"> <li><code>user@host&gt; file copy ftp://username@ftp.hostname.net//filename</code></li> <li><code>file copy ftp.hostname.net: Not logged in.</code></li> <li><code>user@host&gt; file copy ftp://username:prompt@ftphostname.net//filename</code></li> </ul> <p>Password for <b><i>username@ftp.hostname.net</i></b>:</p> </li> <li>• <b><i>http://hostname/path/filename</i></b>—File on a Hypertext Transfer Protocol (HTTP) server. You can also specify <b><i>hostname</i></b> as <b><i>username@hostname</i></b> or <b><i>username:password@hostname</i></b>. If a password is required and you omit it, you are prompted for it.</li> <li>• <b><i>re0:/path/filename</i></b> or <b><i>re1:/path/filename</i></b>—File on a local Routing Engine.</li> </ul>

**Required Privilege Level**    configure—To enter configuration mode.

**Related Documentation**    • *Deactivating and Reactivating Statements and Identifiers in a Junos OS Configuration*



## show chassis alarms

<b>List of Syntax</b>	<a href="#">Syntax on page 495</a> <a href="#">Syntax (TX Matrix Routers) on page 495</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 495</a> <a href="#">Syntax (MX Series Routers) on page 495</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 495</a> <a href="#">Syntax (QFX Series) on page 495</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 495</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 495</a>
<b>Syntax</b>	show chassis alarms
<b>Syntax (TX Matrix Routers)</b>	show chassis alarms <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis alarms <lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis alarms <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis alarms
<b>Syntax (QFX Series)</b>	show chassis alarms <interconnect-device <i>name</i> > <node-device <i>name</i> >
<b>Syntax (PTX Series Packet Transport Routers)</b>	show chassis alarms
<b>Syntax (ACX Series Universal Access Routers)</b>	show chassis alarms
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option for the TX Matrix Plus router introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** Display information about the conditions that have been configured to trigger alarms.

**Options** **none**—Display information about the conditions that have been configured to trigger alarms.

**all-members**—(MX Series routers only) (Optional) Display information about alarm conditions for all the member routers of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display information about alarm conditions for the Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display information about alarm conditions for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display information about alarm conditions for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display information about alarm conditions for the Node device.

**scc**—(TX Matrix router only) (Optional) Show information about the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus router only) (Optional) Show information about the respective TX Matrix Plus router, which is the switch-fabric chassis. Replace *number* variable with 0.

**Additional Information** You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm LED is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the

router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

In Junos OS Release 11.2 and later, the command output on EX8200 switches shows the detailed location (**Plane/FPC/PFE**) for link errors in the chassis.

In Junos OS Release 10.2 and later, an alarm is shown on T Series routers for a standby sonic clock generator (SCG) that is offline or absent.

You may often see the following error messages, in which only the error code is shown and no other information is provided:

```
Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code:
257
Apr 12 08:04:19 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code:
559
```

To understand what CM\_ALARM error codes mean, you need to first identify the structure of the CM Alarm codes. A CM\_ALARM code has the following structure:

Bits:	Error type:
1-31	Major (1)
0	Minor (0)

According to the table above, the LSB (bit 0) identifies the **Error Type** (major alarm, if the bit is set and minor alarm if the bit is unset). The rest of the bits (1 - 31) identify the actual error code.

Take an example of the following error code, which was logged on a T1600:

```
Apr 12 08:04:10 send: red alarm set, device FPC 1, reason FPC 1 Major Errors - Error code:
559
```

First, you have to convert 559 to binary; that is **100010111**. The LSB in this case is 1, which means that this is a major alarm. After removing the LSB, you are left with **10001011**, which is equal to 279 in decimal. This is the actual error code, its meaning can be found from the following list:

Chip Type: L Chip	Code
CMALARM_LCHIP_LOUT_DESRD_PARITY_ERR	1
CMALARM_LCHIP_LOUT_DESRD_UNINIT_ERR	2
CMALARM_LCHIP_LOUT_DESRD_ILLEGALLINK_ERR	3
CMALARM_LCHIP_LOUT_DESRD_ILLEGALSIZERR	4

CMALARM_LCHIP_LOUT_HDRF_TOERR_ERR	5
CMALARM_LCHIP_LOUT_HDRF_PARITY_ERR	6
CMALARM_LCHIP_LOUT_HDRF_UCERR_ERR	7
CMALARM_LCHIP_LOUT_NLIF_CRCDROP_ERR	8
CMALARM_LCHIP_LOUT_NLIF_CRCERR_ERR	9
CMALARM_LCHIP_UCODE_TIMEOUT_ERR	10
CMALARM_LCHIP_LIN_SRCTL_ACCT_DROP_ERR	11
CMALARM_LCHIP_LIN_SRCTL_ACCT_ADDR_SIZE_ERR	12
CMALARM_LCHIP_SRAM_PARITY_ERR	13
CMALARM_LCHIP_UCODE_OVFLW_ERR	14
CMALARM_LCHIP_LOUT_HDRF_MTU_ERR	15

Chip Type: M Chip	Code
CMALARM_MCHIP_ECC_UNCORRECT_ERR	128

Chip Type: N Chip	Code
CMALARM_NCHIP_RDDMA_JBUS_TIMEOUT_ERR	256
CMALARM_NCHIP_RDDMA_FIFO_OVFLW_ERR	257
CMALARM_NCHIP_RDDMA_FIFO_UNFLW_ERR	258
CMALARM_NCHIP_RDDMA_SIZE_ERR	259
CMALARM_NCHIP_RDDMA_JBUS_CRC_ERR	260
CMALARM_NCHIP_WRDMA_PKTR_ERR	261
CMALARM_NCHIP_WRDMA_PKT_CRC_ERR	262
CMALARM_NCHIP_WRDMA_JBUS_TIMEOUT_ERR	263
CMALARM_NCHIP_WRDMA_FIFO_OVFLW_ERR	264
CMALARM_NCHIP_WRDMA_FIFO_UNFLW_ERR	265
CMALARM_NCHIP_WRDMA_PKT_LEN_ERR	266

CMALARM_NCHIP_WRDMA_JBUS_CRC_ERR	267
CMALARM_NCHIP_PKTR_DMA_AGE_ERR	268
CMALARM_NCHIP_PKTR_ICELLSIG_ERR	269
CMALARM_NCHIP_PKTR_FTTL_ERR	270
CMALARM_NCHIP_RODR_OFFSET_OVFLW_ERR	271
CMALARM_NCHIP_PKTR_TMO_CELL_ERR	272
CMALARM_NCHIP_PKTR_TMO_OUTRANGE_ERR	273
CMALARM_NCHIP_PKTR_MD_REQUEST_Q_OVFLW_ERR	274
CMALARM_NCHIP_PKTR_DMA_BUFFER_OVFLW_ERR	275
CMALARM_NCHIP_PKTR_GRT_OVFLW_ERR	276
CMALARM_NCHIP_FRQ_ERR	277
CMALARM_NCHIP_RODR_IN_Q_OVFLW_ERR	278
CMALARM_NCHIP_DBUF_CRC_ERR	279

Chip Type: R Chip	Code
CMALARM_RCHIP_SRAM_PARITY_ERR	512

Chip Type: R Chip	Code
CMALARM_ICHIP_WO_DESRD_ID_ERR	601
CMALARM_ICHIP_WO_DESRD_DATA_ERR	602
CMALARM_ICHIP_WO_DESRD_OFLOW_ERR	603
CMALARM_ICHIP_WO_HDRF_UCERR_ERR	604
CMALARM_ICHIP_WO_HDRF_MTUERR_ERR	605
CMALARM_ICHIP_WO_HDRF_PARITY_ERR	606
CMALARM_ICHIP_WO_HDRF_TOERR_ERR	607
CMALARM_ICHIP_WO_IP_CRC_ERR	608
CMALARM_ICHIP_WO_IP_INTER_ERR	609

CMALARM_ICHIP_WI_WAN_TIMEOUT_ERR	625
CMALARM_ICHIP_WI_FAB_TIMEOUT_ERR	626
CMALARM_ICHIP_RLDRAM_BIST_ERR	630
CMALARM_ICHIP_SDRAM_BIST_ERR	631
CMALARM_ICHIP_RLDRAM_PARITY_ERR	632
CMALARM_ICHIP_SDRAM_UNCORRECT_ERR	633
CMALARM_ICHIP_SDRAM_CORRECT_ERR	634
CMALARM_ICHIP_FUSE_DONE_ERR	635

According to the table above, the **279** error code corresponds to **CMALARM\_NCHIP\_DBUF\_CRC\_ERR**; this means that new CRC errors were seen on the NCHIP of this particular FPC, which is FPC as per the logs.

If you do not want to convert decimal to binary and vice versa, you may use the following shortcut:

For major alarms, the **Actual Error Code = (Error Code - 1)/2**, where **Error Code** is the code that you get in the log message. For example, if you get the following log:

Apr 12 08:04:10 send: red alarm set, device FPC 6, reason FPC 6 Major Errors - Error code: 257

Actual Error Code =  $(257-1)/2 = 128$ . Similarly, for minor alarms, Actual Error Code =  $(\text{Error Code})/2$

**Required Privilege Level**

view

**Related Documentation**

- *Configuring an RMON Alarm Entry and Its Attributes*
- *Chassis Conditions That Trigger Alarms*

**List of Sample Output**

[show chassis alarms \(Alarms Active\) on page 501](#)  
[show chassis alarms \(No Alarms Active\) on page 501](#)  
[show chassis alarms \(Fan Tray\) on page 502](#)  
[show chassis alarms \(MX104 Router\) on page 502](#)  
[show chassis alarms \(MX2010 Router\) on page 502](#)  
[show chassis alarms \(MX2020 Router\) on page 502](#)  
[show chassis alarms \(T4000 Router\) on page 502](#)  
[show chassis alarms \(Unreachable Destinations Present on a T Series Router\) on page 502](#)  
[show chassis alarms \(FPC Offline Due to Unreachable Destinations on a T Series Router\) on page 503](#)

[show chassis alarms \(SCG Absent on a T Series Router\) on page 503](#)  
[show chassis alarms \(Alarms Active on a TX Matrix Router\) on page 503](#)  
[show chassis alarms \(TX Matrix Plus router with 3D SIBs\) on page 504](#)  
[show chassis alarms \(Alarms on a T4000 Router After the enhanced-mode Statement is Enabled\) on page 505](#)  
[show chassis alarms \(Backup Routing Engine\) on page 506](#)  
[show chassis alarms \(EX Series Switch\) on page 506](#)  
[show chassis alarms \(Alarms Active on the QFX Series\) on page 506](#)  
[show chassis alarms node-device \(Alarms Active on the QFabric System\) on page 506](#)  
[show chassis alarms \(Alarms Active on the QFabric System\) on page 506](#)  
[show chassis alarms \(Alarms Active on an EX8200 Switch\) on page 507](#)  
[show chassis alarms \(Alarms Active on a PTX5000 Packet Transport Router\) on page 507](#)  
[show chassis alarms \(Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 507](#)  
[show chassis alarms \(PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 507](#)  
[show chassis alarms \(No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 508](#)  
[show chassis alarms \(Alarms Active on an ACX2000 Universal Access Router\) on page 508](#)  
[show chassis alarms \(Active Alarm to Indicate Status of the Bad SCB Clock on MX Series\) on page 508](#)

**Output Fields** [Table 25 on page 501](#) lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

**Table 25: show chassis alarms Output Fields**

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: <b>Minor</b> or <b>Major</b> .
Description	Information about the alarm.

## Sample Output

### show chassis alarms (Alarms Active)

```

user@host> show chassis alarms
3 alarms are currently active
Alarm time           Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed

```

### show chassis alarms (No Alarms Active)

```

user@host> show chassis alarms
No alarms are currently active

```

### show chassis alarms (Fan Tray)

```
user@host> show chassis alarms
4 alarms currently active
Alarm time      Class  Description
2010-11-11 20:27:38 UTC Major Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC Minor Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC Major Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC Major Side Fan Tray 0 Failure
```

### show chassis alarms (MX104 Router)

```
user@host >show chassis alarms
1 alarms currently active
Alarm time      Class  Description
2013-06-05 14:43:31 IST Minor Backup RE Active
```

### show chassis alarms (MX2010 Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time      Class  Description
2012-08-07 00:46:06 PDT Major Fan Tray 2 Failure
2012-08-06 18:24:36 PDT Minor Redundant feed missing for PSM 6
2012-08-06 07:41:04 PDT Minor Redundant feed missing for PSM 8
2012-08-04 02:42:06 PDT Minor Redundant feed missing for PSM 5
2012-08-03 21:14:24 PDT Minor Loss of communication with Backup RE
2012-08-03 12:26:03 PDT Minor Redundant feed missing for PSM 4
2012-08-03 10:40:18 PDT Minor Redundant feed missing for PSM 7
```

### show chassis alarms (MX2020 Router)

```
user@host> show chassis alarms
1 alarms currently active
Alarm time Class Description
2012-10-03 12:14:59 PDT Minor Plane 0 not online
```

### show chassis alarms (T4000 Router)

```
user@host> show chassis alarms
9 alarms currently active
Alarm time      Class  Description
2007-06-02 01:41:10 UTC Minor RE 0 Not Supported
2007-06-02 01:41:10 UTC Minor CB 0 Not Supported
2007-06-02 01:41:10 UTC Minor Mixed Master and Backup RE types
2007-05-30 19:37:33 UTC Major SPMB 1 not online
2007-05-30 19:37:29 UTC Minor Front Bottom Fan Tray Absent
2007-05-30 19:37:13 UTC Major PEM 1 Input Failure
2007-05-30 19:37:13 UTC Major PEM 0 Not OK
2007-05-30 19:37:03 UTC Major PEM 0 Improper for Platform
2007-05-30 19:37:03 UTC Minor Backup RE Active
```

### show chassis alarms (Unreachable Destinations Present on a T Series Router)

```
user@host> show chassis alarms
10 alarms currently active
Alarm time      Class  Description
2011-08-30 18:43:53 PDT Major FPC 7 has unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 has unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 has unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
```



```

2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

#### show chassis alarms (FPC Offline Due to Unreachable Destinations on a T Series Router)

```

user@host> show chassis alarms
10 alarms currently active
Alarm time      Class Description
2011-08-30 18:43:53 PDT Major FPC 7 offline due to unreachable destinations
2011-08-30 18:43:53 PDT Major FPC 5 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 3 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Major FPC 2 offline due to unreachable destinations
2011-08-30 18:43:52 PDT Minor SIB 0 Not Online
2011-08-30 18:43:33 PDT Minor SIB 4 Not Online
2011-08-30 18:43:28 PDT Minor SIB 3 Not Online
2011-08-30 18:43:05 PDT Minor SIB 2 Not Online
2011-08-30 18:43:28 PDT Minor SIB 1 Not Online
2011-08-30 18:43:05 PDT Major PEM 1 Not Ok

```

#### show chassis alarms (SCG Absent on a T Series Router)

```

user@host> show chassis alarms
4 alarms currently active
Alarm time      Class Description
2011-01-23 21:42:46 PST Major SCG 0 NO EXT CLK MEAS-BKUP SCG ABS

```

#### show chassis alarms (Alarms Active on a TX Matrix Router)

```

user@host> show chassis alarms
scc-re0:
-----
8 alarms currently active
Alarm time      Class Description
2004-08-05 18:43:53 PDT Minor LCC 0 Minor Errors
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:52 PDT Major SIB 2 Absent
2004-08-05 18:43:52 PDT Major SIB 1 Absent
2004-08-05 18:43:52 PDT Major SIB 0 Absent
2004-08-05 18:43:33 PDT Major LCC 2 Major Errors
2004-08-05 18:43:28 PDT Major LCC 0 Major Errors
2004-08-05 18:43:05 PDT Minor LCC 2 Minor Errors
lcc0-re0:
-----
5 alarms currently active
Alarm time      Class Description
2004-08-05 18:43:53 PDT Minor SIB 3 Not Online
2004-08-05 18:43:49 PDT Major SIB 2 Absent
2004-08-05 18:43:49 PDT Major SIB 1 Absent
2004-08-05 18:43:49 PDT Major SIB 0 Absent
2004-08-05 18:43:28 PDT Major PEM 0 Not OK
lcc2-re0:
-----
5 alarms currently active
Alarm time      Class Description
2004-08-05 18:43:35 PDT Minor SIB 3 Not Online
2004-08-05 18:43:33 PDT Major SIB 2 Absent
2004-08-05 18:43:33 PDT Major SIB 1 Absent
2004-08-05 18:43:33 PDT Major SIB 0 Absent
2004-08-05 18:43:05 PDT Minor PEM 1 Absent

```

## show chassis alarms (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis alarms
sfc0-re0:
```

Alarm time	Class	Description
2014-04-08 14:35:13 IST	Minor	FPM 0 SFC Config Size Changed
2014-04-08 14:32:58 IST	Major	Fan Tray Failure
2014-04-08 14:31:53 IST	Major	SIB F13 6 Fault
2014-04-08 14:31:43 IST	Major	SIB F13 11 Fault
2014-04-08 14:31:08 IST	Minor	Check SIB F13 12 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 12 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 12 CXP 3 Fbr Cbl
2014-04-08 14:31:08 IST	Major	SIB F13 12 CXP 15 fault
2014-04-08 14:31:08 IST	Minor	SIB F13 12 CXP 14 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 12 CXP 14
2014-04-08 14:31:08 IST	Major	SIB F13 12 CXP 10 fault
2014-04-08 14:31:08 IST	Minor	SIB F13 12 CXP 8 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 12 CXP 8
2014-04-08 14:31:08 IST	Major	SIB F13 12 CXP 7 fault
2014-04-08 14:31:08 IST	Major	SIB F13 12 CXP 4 fault
2014-04-08 14:31:08 IST	Minor	SIB F13 12 CXP 3 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 12 CXP 3
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 14 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 12 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 8 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 6 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 4 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 2 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 0 Fbr Cbl
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 14 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 14
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 12 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 12
2014-04-08 14:31:08 IST	Major	SIB F13 6 CXP 10 fault
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 8 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 8
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 6 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 6
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 4 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 4
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 2 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 2
2014-04-08 14:31:08 IST	Minor	SIB F13 6 CXP 0 LOL
2014-04-08 14:31:08 IST	Minor	Check SIB F13 6 CXP 0
2014-04-08 14:31:08 IST	Minor	SIB F13 12 CXP 14 XC HSL Link Error
2014-04-08 14:29:27 IST	Minor	LCC 0 Minor Errors
2014-04-08 14:28:37 IST	Major	LCC 0 Major Errors
2014-04-08 14:28:37 IST	Major	LCC 2 Major Errors
2014-04-08 14:28:37 IST	Minor	LCC 2 Minor Errors
2014-04-08 14:28:24 IST	Major	SIB F2S 4/6 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 4/4 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 4/2 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 4/0 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 3/6 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 3/4 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 3/2 Absent
2014-04-08 14:28:24 IST	Major	SIB F2S 3/0 Absent
2014-04-08 14:28:24 IST	Major	SIB F13 9 Absent
2014-04-08 14:28:24 IST	Major	SIB F13 8 Absent

```

2014-04-08 14:28:24 IST Major SIB F13 7 Absent
2014-04-08 14:28:24 IST Major SIB F13 4 Absent
2014-04-08 14:28:24 IST Major SIB F13 1 Absent
2014-04-08 14:28:22 IST Major PEM 0 Input Failure
2014-04-08 14:28:22 IST Major PEM 0 Not OK

```

```
lcc0-re0:
```

```
-----
12 alarms currently active
```

Alarm time	Class	Description
2014-04-08 14:36:08 IST	Minor	CB 1 M/S Switch Changed
2014-04-08 14:36:08 IST	Minor	CB 1 CHASSIS ID Changed
2014-04-08 14:35:43 IST	Minor	CB 0 M/S Switch Changed
2014-04-08 14:35:43 IST	Minor	CB 0 CHASSIS ID Changed
2014-04-08 14:29:30 IST	Minor	SIB 4 Not Online
2014-04-08 14:29:30 IST	Minor	SIB 3 Not Online
2014-04-08 14:29:30 IST	Minor	SIB 2 Not Online
2014-04-08 14:29:24 IST	Major	Rear Fan Tray Failure
2014-04-08 14:29:24 IST	Major	Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:24 IST	Major	Front Top Fan Tray Improper for Platform
2014-04-08 14:28:37 IST	Major	SIB 4 Absent
2014-04-08 14:28:37 IST	Major	SIB 3 Absent

```
lcc2-re0:
```

```
-----
12 alarms currently active
```

Alarm time	Class	Description
2014-04-08 14:36:02 IST	Minor	CB 1 M/S Switch Changed
2014-04-08 14:36:02 IST	Minor	CB 1 CHASSIS ID Changed
2014-04-08 14:35:42 IST	Minor	CB 0 M/S Switch Changed
2014-04-08 14:34:42 IST	Minor	CB 0 CHASSIS ID Changed
2014-04-08 14:29:29 IST	Minor	SIB 0 CXP 7 Unsupported Optics
2014-04-08 14:29:27 IST	Major	Front Bottom Fan Tray Improper for Platform
2014-04-08 14:29:27 IST	Major	Front Top Fan Tray Improper for Platform
2014-04-08 14:29:25 IST	Minor	SIB 4 Not Online
2014-04-08 14:29:25 IST	Minor	SIB 3 Not Online
2014-04-08 14:28:47 IST	Major	PEM 0 Not OK
2014-04-08 14:28:36 IST	Major	SIB 2 Absent
2014-04-08 14:28:36 IST	Minor	Host 0 Boot from alternate media

```
lcc6-re0:
```

```
-----
2 alarms currently active
```

Alarm time	Class	Description
2013-11-06 04:03:56 PST	Minor	SIB 1 CXP 0 XC HSL Link Error
2013-11-06 03:49:32 PST	Major	PEM 1 Not OK

### show chassis alarms (Alarms on a T4000 Router After the enhanced-mode Statement is Enabled)

To enable improved virtual private LAN service (VPLS) MAC address learning on T4000 routers, you must include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and reboot the router. When router reboots, only the T4000 Type 5 FPCs are required to be present on the router. If there are any other FPCs (apart from T4000 Type 5 FPCs) on the T4000 router, such FPCs become offline, and FPC misconfiguration alarms are generated. The **show chassis alarm** command output displays FPC misconfiguration (**FPC *fpc-slot* misconfig**) as the reason for the generation of the alarms.

```
user@host> show chassis alarms
```

```
2 alarms currently active
Alarm time      Class  Description
2011-10-22 10:10:47 PDT Major FPC 1 misconfig
2011-10-22 10:10:46 PDT Major FPC 0 misconfig
```

#### show chassis alarms (Backup Routing Engine)

```
user@host> show chassis alarms
2 alarms are currently active
Alarm time      Class  Description
2005-04-07 10:12:22 PDT Minor Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT Major Host 1 compact-flash missing in Boot List
```

#### show chassis alarms (EX Series Switch)

```
user@switch> show chassis alarms
4 alarms currently active
Alarm time      Class  Description
2014-03-12 15:36:09 UTC Minor Require a Fan Tray upgrade
2014-03-12 15:00:02 UTC Major PEM 0 Input Failure
2014-03-12 15:00:02 UTC Major PEM 0 Not OK
2014-03-12 14:59:51 UTC Minor Host 1 Boot from alternate media
```

#### show chassis alarms (Alarms Active on the QFX Series)

```
user@switch> show chassis alarms
1 alarms currently active
Alarm time      Class  Description
2012-03-05 2:10:24 UTC Major FPC 0 PEM 0 Airflow not matching Chassis Airflow
```

#### show chassis alarms node-device (Alarms Active on the QFabric System)

```
user@switch> show chassis alarms node-device ED3691
node-device ED3694
3 alarms currently active
Alarm time      Class  Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered
```

#### show chassis alarms (Alarms Active on the QFabric System)

```
user@switch> show chassis alarms
IC-A0001:
-----
1 alarms currently active
Alarm time      Class  Description
2011-08-24 16:04:15 UTC Minor Backup RE Active

ED3694:
-----
3 alarms currently active
Alarm time      Class  Description
2011-08-24 16:04:15 UTC Major ED3694:fte-0/1/2: Link down
2011-08-24 16:04:14 UTC Major ED3694:fte-0/1/0: Link down
2011-08-24 14:21:14 UTC Major ED3694 PEM 0 is not supported/powered

SNG-0:
-----

NW-NG-0:
-----
```

```

1 alarms currently active
Alarm time          Class  Description
2011-08-24 15:49:27 UTC Major  ED3691 PEM 0 is not supported/powered

```

#### show chassis alarms (Alarms Active on an EX8200 Switch)

```

user@switch> show chassis alarms

6 alarms currently active
Alarm time          Class  Description
2010-12-02 19:15:22 UTC Major  Fan Tray Failure
2010-12-02 19:15:22 UTC Major  Fan Tray Failure
2010-12-02 19:15:14 UTC Minor  Check CB 0 Fabric Chip 1 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:15:14 UTC Minor  Check CB 0 Fabric Chip 0 on Plane/FPC/PFE: 1/5/0,
1/5/1, 1/5/2, 1/5/3, 1/7/0, 1/7/1, 1/7/2, 1/7/3, 2/5/0, 2/5/1, ...
2010-12-02 19:14:18 UTC Major  PSU 1 Output Failure
2010-12-02 19:14:18 UTC Minor  Loss of communication with Backup RE

```

#### show chassis alarms (Alarms Active on a PTX5000 Packet Transport Router)

```

user@host> show chassis alarms

23 alarms currently active
Alarm time          Class  Description
2011-07-12 16:22:05 PDT Minor  No Redundant Power for Rear Chassis
2011-07-12 16:22:05 PDT Major  PDU 0 PSM 1 Not OK
2011-07-12 16:21:57 PDT Minor  No Redundant Power for Fan 0-2
2011-07-12 16:21:57 PDT Major  PDU 0 PSM 0 Not OK
2011-07-12 15:56:06 PDT Major  PDU 1 PSM 2 Not OK
2011-07-12 15:56:06 PDT Minor  No Redundant Power for FPC 0-7
2011-07-12 15:56:06 PDT Major  PDU 0 PSM 3 Not OK
2011-07-12 15:28:20 PDT Major  PDU 0 PSM 2 Not OK
2011-07-12 15:19:14 PDT Minor  Backup RE Active

```

#### show chassis alarms (Mix of PDUs Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA)

All PDUs installed on a PTX5000 router must be of the same type. The **Mix of PDUs** or **Power Manager Non Operational** alarm is raised when different types of PDUs are installed on a PTX5000 router.

```

user@host> show chassis alarms

15 alarms currently active
Alarm time          Class  Description
2013-03-19 23:03:53 PDT Minor  No Redundant Power
2013-03-19 23:03:48 PDT Minor Mix of PDUs
2013-03-19 23:03:47 PDT Minor  PDU 1 PSM 3 Absent
2013-03-19 23:03:47 PDT Minor  PDU 1 PSM 2 Absent
2013-03-19 23:03:47 PDT Minor  PDU 1 PSM 1 Absent
2013-03-19 23:03:47 PDT Minor  PDU 1 PSM 0 Absent
2013-03-19 23:03:46 PDT Major  No CG Online

```

#### show chassis alarms (PDU Converter Failed Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-PIA)

The **PDU Converter Failed** alarm is raised when one or more 36 V booster converter of a DC PDU fails. If two or more 36 V booster converter fails, fan trays fail and the router might get over heated. Therefore, when this alarm is raised, check the PDU and replace it, if required.

```
user@host> show chassis alarms
11 alarms currently active
Alarm time          Class Description
2013-12-11 22:14:13 PST Minor No Redundant Power for System
2013-12-11 22:14:10 PST Major PDU 0 PSM 7 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 6 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 5 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 4 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 3 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 2 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 1 Not OK
2013-12-11 22:14:10 PST Major PDU 0 PSM 0 Not OK
2013-12-11 22:14:10 PST Major PDU 0 Not OK
2013-12-11 22:14:01 PST Major PDU 0 Converter Failed
```

#### show chassis alarms (No Power for System Alarm on a PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis alarms
8 alarms currently active
Alarm time          Class Description
2013-11-19 01:58:41 PST Major No Power for System
2013-11-19 01:58:37 PST Major PDU 0 PSM 1 Not OK
2013-11-19 01:56:46 PST Major PDU 0 PSM 2 Not OK
2013-11-19 01:54:26 PST Major PDU 0 PSM 3 Not OK
2013-11-19 01:53:30 PST Major PDU 1 PSM 3 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 2 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 1 Not OK
2013-11-19 01:53:29 PST Major PDU 1 PSM 0 Not OK
```

#### show chassis alarms (Alarms Active on an ACX2000 Universal Access Router)

```
user@host> show chassis alarms
7 alarms currently active
Alarm time          Class Description
2012-05-22 11:19:09 UTC Major xe-0/3/1: Link down
2012-05-22 11:19:09 UTC Major xe-0/3/0: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/7: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/6: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/3: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/2: Link down
2012-05-22 11:19:09 UTC Major ge-0/1/1: Link down
```

#### show chassis alarms (Active Alarm to Indicate Status of the Bad SCB Clock on MX Series)

```
user@host> show chassis alarms
1 alarm currently active
Alarm time          Class Description
2013-08-06 07:48:35 PDT Major CB 0 19.44 MHz clock failure
```

## show chassis beacon

**show chassis beacon**  
(QFX Series)

```
show chassis beacon
<cb slot-number>
<fpc slot-number>
<interconnect-device name (cb slot-number | fpc slot-number)>
<node-device name>
```

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display the beacon LED status on a QFX3500 standalone switch, Node device, and an Interconnect device. You can also display the beacon LED status of the Control Boards and Flexible PIC Concentrators on the Interconnect device.

**Options**

**cb slot-number**— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Control Board on the Interconnect device.

**fpc slot-number**— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Flexible PIC Concentrator (FPC) on the Interconnect device. (QFX3500 switches only) (Optional) Display the status of the beacon LEDs for the Flexible PIC Concentrator on the standalone switch.

**interconnect-device name**— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Interconnect device.

**node-device name**— (QFabric systems only) (Optional) Display the status of the beacon LEDs for the Node device.

**Required Privilege Level** view

**Related Documentation**

- [request chassis beacon on page 382](#)

**List of Sample Output**

[show chassis beacon \(QFX Series\) on page 510](#)  
[show chassis beacon interconnect-device \(QFabric System\) on page 510](#)  
[show chassis beacon interconnect-device fpc \(QFabric System\) on page 510](#)  
[show chassis beacon node-device \(QFabric System\) on page 510](#)  
[show chassis beacon node-device fpc \(QFabric System\) on page 510](#)

**Output Fields** [Table 26 on page 509](#) lists the output fields for the **show chassis beacon** command. Output fields are listed in the approximate order in which they appear.

**Table 26: show chassis led Output Fields**

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always 0.

Table 26: show chassis led Output Fields (*continued*)

Field Name	Field Description
<b>Beacon State</b>	Status of the beacon state: <ul style="list-style-type: none"> <li>• Off—The beacon is <b>OFF</b>.</li> <li>• On—The beacon is <b>ON</b>.</li> </ul>

## Sample Output

### show chassis beacon (QFX Series)

```
user@switch> show chassis beacon
Slot          Beacon State
FPC          0          OFF
```

### show chassis beacon interconnect-device (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1
Chassis              OFF
CB 0                  OFF
CB 1                  OFF
FC 0 FPC 0           OFF
FC 1 FPC 1           OFF
RC 0 FPC 8           OFF
RC 1 FPC 9           OFF
```

### show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1 fpc 0
FPC 0                ON
```

### show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device node1
node1                ON
```

### show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device node1 fpc 0
FPC 0                ON
```



## show chassis environment

<b>List of Syntax</b>	<a href="#">Syntax on page 511</a> <a href="#">Syntax (T320, T640, T1600, and T4000 Routers) on page 511</a> <a href="#">Syntax (TX Matrix Routers) on page 511</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 511</a> <a href="#">Syntax (MX Series Routers) on page 511</a> <a href="#">Syntax (MX104 3D Universal Edge Routers) on page 511</a> <a href="#">Syntax (MX2010 and MX2020 3D Universal Edge Routers) on page 512</a> <a href="#">Syntax (EX8200 Switches) on page 512</a> <a href="#">Syntax (EX Series Switches except EX8200) on page 512</a> <a href="#">Syntax (QFX Series) on page 512</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 512</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 512</a>
<b>Syntax</b>	<b>show chassis environment</b>
<b>Syntax (T320, T640, T1600, and T4000 Routers)</b>	show chassis environment <cb <i>cb-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> > <scg <i>scg-slot-number</i> > <sib <i>sib-slot-number</i> >
<b>Syntax (TX Matrix Routers)</b>	show chassis environment <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis environment <cb <i>cb-slot-number</i> > <cip <i>cip-slot-number</i> > <fpc <i>fpc-slot-number</i> > <fpm> <lcc <i>number</i> > <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> > <scg <i>scg-slot-number</i> > <sfc <i>number</i> > <sib <i>sib-slot-number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis environment <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104 3D Universal Edge Routers)</b>	show chassis environment <cb> <pem <i>pem-slot-number</i> > <routing-engine <i>re-slot-number</i> >

Syntax (MX2010 and MX2020 3D Universal Edge Routers)	<pre>show chassis environment &lt;adc <i>adc-slot-number</i>&gt; &lt;cb <i>cb-slot-number</i>&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;fpm&gt; &lt;monitored&gt; &lt;psm <i>psm-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt; &lt;sfb <i>sfb-slot-number</i>&gt;</pre>
Syntax (EX8200 Switches)	<pre>show chassis environment &lt;all-members&gt; &lt;cb <i>cb-slot-number</i>&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;psu <i>psu-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt;</pre>
Syntax (EX Series Switches except EX8200)	<pre>show chassis environment &lt;all-members&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;power-supply-unit&gt; &lt;routing-engine&gt;</pre>
Syntax (QFX Series)	<pre>show chassis environment &lt;cb <i>slot-number</i> &lt;interconnect-device name&gt;&gt; &lt;fpc <i>slot-number</i> &lt;interconnect-device name&gt;&gt; &lt;interconnect-device name &lt;slot-number&gt; &lt;node-device name&gt; &lt;pem <i>slot-number</i> (interconnect-device name <i>slot-number</i>)   (node-device name)&gt; &lt;routing-engine name &lt;interconnect-device name slot-number&gt;&gt;</pre>
Syntax (PTX Series Packet Transport Routers)	<pre>show chassis environment &lt;cb <i>cb-slot-number</i>&gt; &lt;ccg <i>ccg-slot-number</i>&gt; &lt;fpc <i>fpc-slot-number</i>&gt; &lt;fpm&gt; &lt;monitored&gt; &lt;pdu <i>pdu-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt; &lt;sib <i>sib-slot-number</i>&gt;</pre>
Syntax (ACX Series Universal Access Routers)	<pre>show chassis environment &lt;cb <i>cb-slot-number</i>&gt; &lt;pem <i>pem-slot-number</i>&gt; &lt;routing-engine <i>re-slot-number</i>&gt;</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p>

Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

**monitored** option added in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.1 for T4000 Core Routers.

Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.

Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

**pem** option introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers.

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

In addition, on ACX4000 routers, display temperature information about the different channels of a Modular Interface Card (MIC). The number of channels displayed depends on the type of MIC installed.

Starting with Junos OS Release 14.1, the **show chassis environment cb cb-slot-number | ccg ccg-slot-number | fpc fpc-slot-number | fpm | monitored | pdu pdu-slot-number | routing-engine re-slot-number | sib sib-slot-number** operational mode command output displays environmental information for the the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-PIA) and other components in a PTX5000 Packet Transport Router.

**Options** **none**—Display environmental information about the router or switch chassis. On a TX Matrix router, display environmental information about the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for all the members of the Virtual Chassis configuration.

**adc adc-slot-number**—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the adapter cards. For MX2020 routers, replace **adc-slot-number** with a value from 0 through 19. For MX2010 routers, replace **adc-slot-number** with a value from 0 through 9.

**cb cb-slot-number**—(ACX Series Universal Access Routers, EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2020 routers, MX2010 routers, PTX Series Packet Transport Routers, QFX Series, and T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Control Board. On devices other than EX Series switches, replace **cb-slot** with 0 or 1. For the EX Series switches, see *EX Series Switches Hardware and CLI Terminology Mapping* for information on CB slot numbering.

**cip cip-slot-number**—(TX Matrix Plus routers only) (Optional) Display chassis environmental information for the Connection Interface Panel (CIP). Replace the **cip-slot-number** variable with a value of 0 or 1.

**cb interconnect-device name**—(QFabric systems only) (Optional) Display chassis environmental information for the Control Board on an Interconnect device.

**ccg ccg-slot-number**—(PTX Series only) (Optional) Display chassis environmental information for the Centralized Clock Generator. Replace **cb-slot** with a value of 0 or 1.

**fpc fpc-slot**—(EX Series switches, M120, M320, and M40e routers, MX Series routers, MX2010 routers, MX2020 routers, PTX Series Packet Transport Routers, QFX Series, QFX3500 switches, QFabric systems, T Series routers, and TX Matrix Plus routers) (Optional) Display chassis environmental information for a specified Flexible PIC Concentrator. For MX2010 routers, replace **fpc-slot** with a value from 0 through 9. For MX2020 routers, replace **fpc-slot** with a value from 0 through 19. For information about FPC numbering, see [show chassis environment fpc](#). On a QFabric system, display chassis environmental information for a specified Flexible PIC Concentrator on an Interconnect device. On an EX Series switch, display chassis environmental information for a specified Flexible PIC Concentrator; see *EX Series Switches Hardware and CLI Terminology Mapping* for information on FPC numbering. On a TX Matrix Plus router with 3D SIBs replace **fpc-slot** with a value from 0 through 63.

**fpm**—(M120, M320, and M40e routers, MX2010 routers, MX2020 routers, PTX Series, Packet Transport Routers, T Series routers, and TX Matrix Plus routers only) (Optional) Display chassis environmental information for the craft interface (FPM).

**interconnect-device name**—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

**monitored**—(MX2020 routers and PTX Series Packet Transport Routers only) (Optional) Display chassis environmental information for monitored temperatures only. Temperatures that are not included in temperature alarm computations are not displayed.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers and EX Series switches) (Optional) Display chassis environmental information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display chassis environmental information for the specified member of the Virtual Chassis configuration. On MX Series routers, replace *member-id* variable with a value of **0** or **1**. For EX Series switches, see [member](#) for member ID values.

**node-device *name***—(QFabric systems only) (Optional) Display chassis environmental information for the Node device.

**pdu *pdu-slot-number***—(PTX Series only) (Optional) Display chassis environmental information for the specified power distribution unit.

**pem**—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Interconnect device or Node device.

**pem *pem-slot-number***—(ACX Series Universal Access Routers, M120, M320, and M40e routers, MX Series routers, MX104 routers, QFX Series, and T Series routers only) (Optional) Display chassis environmental information for the Power Entry Module on the specified Power Entry Module. For information about the options, see [show chassis environment pem](#).

**psm *psm-slot-number***—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. For MX2020 routers, replace *psm-slot-number* with a value from **0** through **17**. For MX2010 routers, replace *psm-slot-number* with a value from **0** through **8**.

**psu *psu-slot-number***—(EX Series switches only) (Optional) Display chassis environmental information for a specified power supply. See *EX Series Switches Hardware and CLI Terminology Mapping* for detailed information.

**routing-engine**—(QFX3500 switches and QFabric systems only) (Optional) Display chassis environmental information for the Routing Engine on the specified Interconnect device.

**routing-engine *re-slot-number***—(Optional) Display chassis environmental information for the specified Routing Engine. For information about the options, see [show chassis environment routing-engine](#).

**scg**—(T Series routers only) (Optional) Display chassis environmental information about the SONET Clock Generator.

**scc**—(TX Matrix routers only) (Optional) Display chassis environmental information about the TX Matrix router (switch-card chassis).

**sfb *sfb-slot-number***—(MX2020 and MX2010 routers only) (Optional) Display chassis environmental information for the power supply module. Replace *sfb-slot-number* with a value from **0** through **7**.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display chassis environmental information about the respective TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with **0**.

**sib *sib-slot-number***—(M320 routers, PTX Series Packet Transport Routers, and T Series routers only) (Optional) Display chassis environmental information about the specified switch interface board. For information about the options, see *show chassis environment sib*.

**Required Privilege Level** view

**Related Documentation**

- *show chassis environment adc*
- *show chassis environment cb*
- *show chassis environment ccg*
- *show chassis environment cip*
- [show chassis environment fpc on page 575](#)
- *show chassis environment fpm*
- *show chassis environment lcc*
- *show chassis environment mcs*
- *show chassis environment monitored*
- *show chassis environment pcg*
- *show chassis environment pdu*
- [show chassis environment pem on page 601](#)
- *show chassis environment psm*
- *show chassis environment psu*
- [show chassis environment routing-engine on page 610](#)
- *show chassis environment scg*
- *show chassis environment sfb*
- *show chassis environment sib*
- *show chassis environment sfc*

**List of Sample Output**

[show chassis environment \(J2300 Router\) on page 519](#)  
[show chassis environment \(J4300 or J6300 Router\) on page 519](#)  
[show chassis environment \(M5 Router\) on page 519](#)  
[show chassis environment \(M7i Router\) on page 520](#)  
[show chassis environment \(M10 Router\) on page 520](#)  
[show chassis environment \(M10i Router\) on page 520](#)  
[show chassis environment \(M20 Router\) on page 521](#)  
[show chassis environment \(M40 Router\) on page 521](#)  
[show chassis environment \(M40e Router\) on page 521](#)  
[show chassis environment \(M120 Router\) on page 522](#)  
[show chassis environment \(M160 Router\) on page 523](#)  
[show chassis environment \(M320 Router\) on page 523](#)

[show chassis environment \(MX104 Router\) on page 524](#)  
[show chassis environment \(MX240 Router\) on page 525](#)  
[show chassis environment \(MX240 Router with SCBE\) on page 526](#)  
[show chassis environment \(MX480 Router\) on page 526](#)  
[show chassis environment \(MX480 Router with SCBE\) on page 527](#)  
[show chassis environment \(MX960 Router\) on page 528](#)  
[show chassis environment \(MX960 Router with SCBE\) on page 529](#)  
[show chassis environment \(MX960 Router with MPC5EQ\) on page 532](#)  
[show chassis environment \(MX2020 Router\) on page 536](#)  
[show chassis environment \(MX2020 Router with MPC5EQ and MPC6E\) on page 545](#)  
[show chassis environment \(MX2010 Router\) on page 549](#)  
[show chassis environment \(T320 Router\) on page 554](#)  
[show chassis environment \(T640 Router\) on page 555](#)  
[show chassis environment \(T4000 Router\) on page 556](#)  
[show chassis environment \(TX Matrix Router\) on page 558](#)  
[show chassis environment \(T1600 Router\) on page 559](#)  
[show chassis environment \(TX Matrix Plus Router\) on page 560](#)  
[show chassis environment \(TX Matrix Plus router with 3D SIBs\) on page 562](#)  
[show chassis environment \(EX4200 Standalone Switch\) on page 565](#)  
[show chassis environment \(EX8216 Switch\) on page 566](#)  
[show chassis environment \(EX9200 Switch\) on page 566](#)  
[show chassis environment \(QFX Series\) on page 567](#)  
[show chassis environment interconnect-device \(QFabric System\) on page 567](#)  
[show chassis environment node-device \(QFabric System\) on page 569](#)  
[show chassis environment pem node-device \(QFabric System\) on page 569](#)  
[show chassis environment \(PTX5000 Packet Transport Router\) on page 570](#)  
[show chassis environment \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 572](#)  
[show chassis environment \(ACX2000 Universal Access Router\) on page 573](#)  
[show chassis environment \(ACX4000 Universal Access Router\) on page 573](#)

**Output Fields** [Table 27 on page 518](#) lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.

Table 27: show chassis environment Output Fields

Field Name	Field Description
<b>Class</b>	<p>Information about the category or class of chassis component:</p> <ul style="list-style-type: none"> <li>• <b>Power:</b> Power information: <ul style="list-style-type: none"> <li>• (M5, M10, M20, and M40 routers and EX Series switches only) Power supply status: <b>OK</b>, <b>Testing</b>, (during initial power-on), <b>Failed</b>, or <b>Absent</b>.</li> <li>• (M7i, M10i, M40e, M120, M160, M320, and T Series routers and EX Series switches only) Power Entry Modules status: <b>OK</b>, <b>Testing</b>, (during initial power-on), <b>Check</b>, <b>Failed</b>, or <b>Absent</b>.</li> <li>• (PTX Series only) Power information is reported in PDU or PSM combinations. The status is: <b>OK</b>, <b>Testing</b>, (during initial power-on), <b>Check</b>, <b>Failed</b>, or <b>Absent</b>.</li> </ul> </li> <li>• <b>Temp:</b> Temperature of air flowing through the chassis in degrees Celsius (C) and Fahrenheit (F). <ul style="list-style-type: none"> <li>• On PTX Series Packet Transport Routers and MX2010 and MX2020 Routers, multiple cooling zones are supported. FRU temperatures in each zone are coordinated with the fan speed of fan trays in those zones.</li> <li>• EX2200 switches have a side-to-rear cooling system. The <b>Local Intake</b> temperature is measured by the sensor on the right side of the chassis, and the <b>Remote Intake</b> temperature is measured by the sensor on the left side of the chassis.</li> </ul> </li> <li>• <b>Pic:</b> On ACX4000 Routers, multiple temperature channels on a MIC. The status is: <b>OK</b> and the <b>Measurement</b> is in degrees Celsius (C) and Fahrenheit (F).</li> <li>• <b>Fan:</b> Fan status: <b>OK</b>, <b>Testing</b> (during initial power-on), <b>Failed</b>, or <b>Absent</b>. On PTX Series Packet Transport Routers and MX2010 and MX2020 Routers, multiple fan trays are supported. Fan status is reported in Fan Tray or Fan combinations. <b>Measurement</b> indicates actual fan RPM (PTX and MX2010 and MX2020 Routers only).</li> <li>• <b>Misc:</b> Information about other components of the chassis. <ul style="list-style-type: none"> <li>• On some routers, this field indicates the status of one or more additional components.</li> <li>• On the M40e, M160, and M320 router, <b>Misc</b> includes <b>CIP</b> (Connector Interface Panel). <b>OK</b> indicates that the CIP is present. <b>Absent</b> indicates that the CIP is not present.</li> <li>• On T Series routers, <b>Misc</b> includes <b>CIP</b> and <b>SPMB</b> (Switch Processor Mezzanine Board). <b>OK</b> indicates that the <b>CIP</b> or <b>SPMB</b> is present. <b>Absent</b> indicates that the <b>CIP</b> or <b>SPMB</b> is not present.</li> <li>• On PTX Series Packet Transport Routers, <b>Misc</b> includes the <b>SPMB</b> (Switch Processor Mezzanine Board). The SPMB is located on the control boards. <b>OK</b> indicates that the control board is present. <b>Absent</b> indicates that the control board is not present.</li> </ul> </li> </ul>
<b>Item</b>	<p>(MX2010 and MX2020 Routers) Information about the chassis component: Routing Engines, Controls Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).</p> <p>(MX104 Routers) Information about the chassis components: Routing Engines, Control Board (CB), Power Entry Module (PEM), and Compact Forwarding Engine Board (AFEB).</p> <p>(QFabric Systems) Information about the chassis component: Control Boards, Routing Engines, Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs), Node Devices, and Interconnect Devices.</p> <p>(QFX Series) Information about the chassis component: Flexible PIC Concentrators (FPCs), and Power Entry Modules (PEMs).</p>



Table 27: show chassis environment Output Fields (*continued*)

Field Name	Field Description
<b>Status</b>	<p>(MX104, MX2010, and MX2020 Routers) Status of the specified chassis component. For example, if the Class is Fan, the fan status can be:</p> <ul style="list-style-type: none"> <li>• <b>OK:</b> The fans are operational.</li> <li>• <b>Testing:</b> The fans are being tested during initial power-on.</li> <li>• <b>Failed:</b> The fans have failed or the fans are not spinning.</li> <li>• <b>Absent:</b> The fan tray is not installed.</li> </ul> <p>If the Class is Power, the power supply status can be:</p> <ul style="list-style-type: none"> <li>• <b>OK:</b> The power component is operational.</li> <li>• <b>Testing:</b> The power component is being tested during initial power-on.</li> <li>• <b>Check:</b> There is insufficient power---that is, fewer than the minimum required feeds are connected.</li> <li>• <b>Failed:</b> The inputs leads have failed.</li> <li>• <b>Absent:</b> The power component is not installed.</li> </ul>
<b>Measurement</b>	<p>(MX104, MX2010, and MX2020 Routers) Dependant on the Class. For example, if the Class is Temp, indicates the temperature in degree Celsius and degrees Fahrenheit. If the Class is Fan, indicates actual fan RPM.</p>

## Sample Output

### show chassis environment (J2300 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Temp Routing Engine   OK      40 degrees C / 104 degrees F
Fan  Fan              OK

```

### show chassis environment (J4300 or J6300 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Temp Routing Engine   OK      41 degrees C / 105 degrees F
Fan  Fan 0            OK
     Fan 1            OK

```

### show chassis environment (M5 Router)

```

user@host> show chassis environment
Class Item           Status Measurement
Power Power Supply A   OK
Power Power Supply B   Absent
Temp  FPC 0            OK      30 degrees C / 86 degrees F
     FEB              OK      33 degrees C / 91 degrees F
     PS Intake         OK      27 degrees C / 80 degrees F
     PS Exhaust        OK      27 degrees C / 80 degrees F
     Routing Engine     OK      34 degrees C / 93 degrees F
Fans  Left Fan 1        OK      Spinning at normal speed
     Left Fan 2         OK      Spinning at normal speed
     Left Fan 3         OK      Spinning at normal speed
     Left Fan 4         OK      Spinning at normal speed
Misc  Craft Interface   OK

```

**show chassis environment (M7i Router)**

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply 0       OK
       Power Supply 1      Absent
Temp  Intake               OK          22 degrees C / 71 degrees F
       FPC 0                OK          23 degrees C / 73 degrees F
       Power Supplies       OK          23 degrees C / 73 degrees F
       CFEB Intake          OK          24 degrees C / 75 degrees F
       CFEB Exhaust         OK          29 degrees C / 84 degrees F
       Routing Engine       OK          26 degrees C / 78 degrees F
Fans  Fan 1                 OK          Spinning at normal speed
       Fan 2                 OK          Spinning at normal speed
       Fan 3                 OK          Spinning at normal speed
       Fan 4                 OK          Spinning at normal speed

```

**show chassis environment (M10 Router)**

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply A       OK
       Power Supply B      Failed
Temp  FPC 0                 OK          36 degrees C / 96 degrees F
       FPC 1                 OK          35 degrees C / 95 degrees F
       FEB                   OK          34 degrees C / 93 degrees F
       PS Intake             OK          31 degrees C / 87 degrees F
       PS Exhaust            OK          34 degrees C / 93 degrees F
       Routing Engine        OK          35 degrees C / 95 degrees F
Fans  Left Fan 1            OK          Spinning at normal speed
       Left Fan 2            OK          Spinning at normal speed
       Left Fan 3            OK          Spinning at normal speed
       Left Fan 4            OK          Spinning at normal speed
Misc  Craft Interface       OK

```

**show chassis environment (M10i Router)**

```

user@host> show chassis environment
Class Item                Status      Measurement
Power Power Supply 0       OK
       Power Supply 1       OK
       Power Supply 2       Absent
       Power Supply 3       Absent
Temp  Intake               OK          26 degrees C / 78 degrees F
       FPC 0                OK          27 degrees C / 80 degrees F
       FPC 1                OK          28 degrees C / 82 degrees F
       Lower Power Supplies  OK          29 degrees C / 84 degrees F
       Upper Power Supplies  OK          28 degrees C / 82 degrees F
       CFEB Intake           OK          27 degrees C / 80 degrees F
       CFEB Exhaust          OK          36 degrees C / 96 degrees F
       Routing Engine 0      OK          31 degrees C / 87 degrees F
       Routing Engine 1      OK          27 degrees C / 80 degrees F
Fans  Fan Tray 0 Fan 1      OK          Spinning at normal speed
       Fan Tray 0 Fan 2      OK          Spinning at normal speed
       Fan Tray 0 Fan 3      OK          Spinning at normal speed
       Fan Tray 0 Fan 4      OK          Spinning at normal speed
       Fan Tray 0 Fan 5      OK          Spinning at normal speed
       Fan Tray 0 Fan 6      OK          Spinning at normal speed
       Fan Tray 0 Fan 7      OK          Spinning at normal speed

```

Fan Tray 0 Fan 8	OK	Spinning at normal speed
Fan Tray 1 Fan 1	Absent	
Fan Tray 1 Fan 2	Absent	
Fan Tray 1 Fan 3	Absent	
Fan Tray 1 Fan 4	Absent	
Fan Tray 1 Fan 5	Absent	
Fan Tray 1 Fan 6	Absent	
Fan Tray 1 Fan 7	Absent	
Fan Tray 1 Fan 8	Absent	

### show chassis environment (M20 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 0	OK	28 degrees C / 82 degrees F
	FPC 1	OK	27 degrees C / 80 degrees F
	Power Supply A	OK	22 degrees C / 71 degrees F
	Power Supply B	Absent	
	SSB 0	OK	30 degrees C / 86 degrees F
	Backplane	OK	22 degrees C / 71 degrees F
Fans	Routing Engine 0	OK	26 degrees C / 78 degrees F
	Routing Engine 1	Testing	
	Rear Fan	OK	Spinning at normal speed
	Front Upper Fan	OK	Spinning at normal speed
	Front Middle Fan	OK	Spinning at normal speed
	Front Bottom Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

### show chassis environment (M40 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	Power Supply A	OK	
	Power Supply B	Absent	
Temp	FPC 3	OK	24 degrees C / 75 degrees F
	FPC 6	OK	26 degrees C / 78 degrees F
	SCB	OK	26 degrees C / 78 degrees F
	Backplane @ A1	OK	28 degrees C / 82 degrees F
	Backplane @ A2	OK	23 degrees C / 73 degrees F
	Routing Engine	OK	26 degrees C / 78 degrees F
Fans	Top Impeller	OK	Spinning at normal speed
	Bottom impeller	OK	Spinning at normal speed
	Rear Left Fan	OK	Spinning at normal speed
	Rear Center Fan	OK	Spinning at normal speed
	Rear Right Fan	OK	Spinning at normal speed
Misc	Craft Interface	OK	

### show chassis environment (M40e Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	PCG 0	OK	44 degrees C / 111 degrees F
	PCG 1	OK	47 degrees C / 116 degrees F
	Routing Engine 0	OK	40 degrees C / 104 degrees F
	Routing Engine 1	OK	37 degrees C / 98 degrees F

	MCS 0	OK	45 degrees C / 113 degrees F
	MCS 1	OK	42 degrees C / 107 degrees F
	SFM 0 SPP	OK	40 degrees C / 104 degrees F
	SFM 0 SPR	OK	44 degrees C / 111 degrees F
	SFM 1 SPP	OK	43 degrees C / 109 degrees F
	SFM 1 SPR	OK	45 degrees C / 113 degrees F
	FPC 0	OK	38 degrees C / 100 degrees F
	FPC 1	OK	40 degrees C / 104 degrees F
	FPC 2	OK	38 degrees C / 100 degrees F
	FPC 4	OK	34 degrees C / 93 degrees F
	FPC 5	OK	43 degrees C / 109 degrees F
	FPC 6	OK	41 degrees C / 105 degrees F
	FPC 7	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	28 degrees C / 82 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

**show chassis environment (M120 Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	
	PEM 1	OK	
	Routing Engine 0	OK	43 degrees C / 109 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	33 degrees C / 91 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	35 degrees C / 95 degrees F
	CB 1 Intake	OK	34 degrees C / 93 degrees F
	CB 1 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 1 Exhaust B	OK	35 degrees C / 95 degrees F
	FEB 3 Intake	OK	35 degrees C / 95 degrees F
	FEB 3 Exhaust A	OK	37 degrees C / 98 degrees F
	FEB 3 Exhaust B	OK	39 degrees C / 102 degrees F
	FEB 4 Intake	OK	33 degrees C / 91 degrees F
	FEB 4 Exhaust A	OK	39 degrees C / 102 degrees F
	FEB 4 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 2 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 2 Exhaust B	OK	31 degrees C / 87 degrees F
	FPC 3 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 3 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 4 Exhaust A	OK	32 degrees C / 89 degrees F
	FPC 4 Exhaust B	OK	30 degrees C / 86 degrees F
Fans	Front Top Tray Fan 1	OK	Spinning at normal speed
	Front Top Tray Fan 2	OK	Spinning at normal speed
	Front Top Tray Fan 3	OK	Spinning at normal speed
	Front Top Tray Fan 4	OK	Spinning at normal speed
	Front Top Tray Fan 5	OK	Spinning at normal speed
	Front Top Tray Fan 6	OK	Spinning at normal speed
	Front Top Tray Fan 7	OK	Spinning at normal speed
	Front Top Tray Fan 8	OK	Spinning at normal speed
	Front Bottom Tray Fan 1	OK	Spinning at normal speed
	Front Bottom Tray Fan 2	OK	Spinning at normal speed

Front Bottom Tray Fan 3	OK	Spinning at normal speed
Front Bottom Tray Fan 4	OK	Spinning at normal speed
Front Bottom Tray Fan 5	OK	Spinning at normal speed
Front Bottom Tray Fan 6	OK	Spinning at normal speed
Front Bottom Tray Fan 7	OK	Spinning at normal speed
Front Bottom Tray Fan 8	OK	Spinning at normal speed
Rear Top Tray Fan 1	OK	Spinning at normal speed
Rear Top Tray Fan 2	OK	Spinning at normal speed
Rear Top Tray Fan 3	OK	Spinning at normal speed
Rear Top Tray Fan 4	OK	Spinning at normal speed
Rear Top Tray Fan 5	OK	Spinning at normal speed
Rear Top Tray Fan 6	OK	Spinning at normal speed
Rear Top Tray Fan 7	OK	Spinning at normal speed
Rear Top Tray Fan 8	OK	Spinning at normal speed
Rear Bottom Tray Fan 1	OK	Spinning at normal speed
Rear Bottom Tray Fan 2	OK	Spinning at normal speed
Rear Bottom Tray Fan 3	OK	Spinning at normal speed
Rear Bottom Tray Fan 4	OK	Spinning at normal speed
Rear Bottom Tray Fan 5	OK	Spinning at normal speed
Rear Bottom Tray Fan 6	OK	Spinning at normal speed
Rear Bottom Tray Fan 7	OK	Spinning at normal speed
Rear Bottom Tray Fan 8	OK	Spinning at normal speed

#### show chassis environment (M160 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	PEM 1
Temp	PCG 0	OK	45 degrees C / 113 degrees F
	PCG 1	Absent	
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 1	Absent	
	MCS 0	OK	50 degrees C / 122 degrees F
	SFM 0 SPP	OK	47 degrees C / 116 degrees F
	SFM 0 SPR	OK	49 degrees C / 120 degrees F
	SFM 1 SPP	OK	50 degrees C / 122 degrees F
	SFM 1 SPR	OK	50 degrees C / 122 degrees F
	SFM 2 SPP	OK	51 degrees C / 123 degrees F
	SFM 2 SPR	OK	52 degrees C / 125 degrees F
	SFM 3 SPP	OK	52 degrees C / 125 degrees F
	SFM 3 SPR	OK	48 degrees C / 118 degrees F
	FPC 0	OK	45 degrees C / 113 degrees F
	FPC 6	OK	43 degrees C / 109 degrees F
	FPM CMB	OK	31 degrees C / 87 degrees F
	FPM Display	OK	33 degrees C / 91 degrees F
Fans	Rear Bottom Blower	OK	Spinning at normal speed
	Rear Top Blower	OK	Spinning at normal speed
	Front Top Blower	OK	Spinning at normal speed
	Fan Tray Rear Left	OK	Spinning at normal speed
	Fan Tray Rear Right	OK	Spinning at normal speed
	Fan Tray Front Left	OK	Spinning at normal speed
	Fan Tray Front Right	OK	Spinning at normal speed
Misc	CIP	OK	

#### show chassis environment (M320 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	

	PEM 2	OK	
	PEM 3	OK	
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	CB 0	OK	36 degrees C / 96 degrees F
	CB 1	OK	36 degrees C / 96 degrees F
	SIB 0	OK	38 degrees C / 100 degrees F
	SIB 1	OK	29 degrees C / 84 degrees F
	SIB 2	OK	38 degrees C / 100 degrees F
	SIB 3	OK	41 degrees C / 105 degrees F
	FPC 0 Intake	OK	28 degrees C / 82 degrees F
	FPC 0 Exhaust	OK	40 degrees C / 104 degrees F
	FPC 1 Intake	OK	29 degrees C / 84 degrees F
	FPC 1 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 2 Intake	OK	28 degrees C / 82 degrees F
	FPC 2 Exhaust	OK	38 degrees C / 100 degrees F
	FPC 3 Intake	OK	28 degrees C / 82 degrees F
	FPC 3 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 6 Intake	OK	27 degrees C / 80 degrees F
	FPC 6 Exhaust	OK	39 degrees C / 102 degrees F
	FPC 7 Intake	OK	27 degrees C / 80 degrees F
	FPC 7 Exhaust	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
Fan	Top Left Front fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Rear Fan 1 (TOP)	OK	Spinning at normal speed
	Rear Fan 2	OK	Spinning at normal speed
	Rear Fan 3	OK	Spinning at normal speed
	Rear Fan 4	OK	Spinning at normal speed
	Rear Fan 5	OK	Spinning at normal speed
	Rear Fan 6	OK	Spinning at normal speed
	Rear Fan 7 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	

**show chassis environment (MX104 Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	34 degrees C / 93 degrees F
	PEM 1	Absent	
	ABB 0 Intake	OK	33 degrees C / 91 degrees F
	ABB 0 Exhaust A	OK	42 degrees C / 107 degrees F
	ABB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	ABB 1 Intake	Absent	
	ABB 1 Exhaust A	Absent	
	ABB 1 Exhaust B	Absent	
	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 0 CPU	OK	46 degrees C / 114 degrees F
Fans	Routing Engine 1	Absent	
	Routing Engine 1 CPU	Absent	
	AFEB 0 AFEB Processor	OK	33 degrees C / 91 degrees F
	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed
	Fan 3	OK	Spinning at normal speed

Fan 4	OK	Spinning at normal speed
Fan 5	OK	Spinning at normal speed

### show chassis environment (MX240 Router)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp PEM 0                             OK          40 degrees C / 104 degrees F
      PEM 1                             OK          45 degrees C / 113 degrees F
      PEM 2                             Absent
      PEM 3                             Absent
      Routing Engine 0                   OK          39 degrees C / 102 degrees F
      Routing Engine 1                   OK          37 degrees C / 98 degrees F
      CB 0 Intake                         OK          36 degrees C / 96 degrees F
      CB 0 Exhaust A                     OK          34 degrees C / 93 degrees F
      CB 0 Exhaust B                     OK          38 degrees C / 100 degrees F
      CB 0 ACBC                           OK          37 degrees C / 98 degrees F
      CB 0 SF A                           OK          49 degrees C / 120 degrees F
      CB 0 SF B                           OK          41 degrees C / 105 degrees F
      CB 1 Intake                         OK          37 degrees C / 98 degrees F
      CB 1 Exhaust A                     OK          34 degrees C / 93 degrees F
      CB 1 Exhaust B                     OK          39 degrees C / 102 degrees F
      CB 1 ACBC                           OK          38 degrees C / 100 degrees F
      CB 1 SF A                           OK          47 degrees C / 116 degrees F
      CB 1 SF B                           OK          41 degrees C / 105 degrees F
      FPC 1 Intake                       OK          33 degrees C / 91 degrees F
      FPC 1 Exhaust A                     OK          38 degrees C / 100 degrees F
      FPC 1 Exhaust B                     OK          53 degrees C / 127 degrees F
      FPC 1 I3 0 TSensor                  OK          50 degrees C / 122 degrees F
      FPC 1 I3 0 Chip                     OK          53 degrees C / 127 degrees F
      FPC 1 I3 1 TSensor                  OK          49 degrees C / 120 degrees F
      FPC 1 I3 1 Chip                     OK          52 degrees C / 125 degrees F
      FPC 1 I3 2 TSensor                  OK          47 degrees C / 116 degrees F
      FPC 1 I3 2 Chip                     OK          49 degrees C / 120 degrees F
      FPC 1 I3 3 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 1 I3 3 Chip                     OK          46 degrees C / 114 degrees F
      FPC 1 IA 0 TSensor                  OK          45 degrees C / 113 degrees F
      FPC 1 IA 0 Chip                     OK          44 degrees C / 111 degrees F
      FPC 1 IA 1 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 1 IA 1 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 Intake                       OK          32 degrees C / 89 degrees F
      FPC 2 Exhaust A                     OK          40 degrees C / 104 degrees F
      FPC 2 Exhaust B                     OK          52 degrees C / 125 degrees F
      FPC 2 I3 0 TSensor                  OK          52 degrees C / 125 degrees F
      FPC 2 I3 0 Chip                     OK          56 degrees C / 132 degrees F
      FPC 2 I3 1 TSensor                  OK          52 degrees C / 125 degrees F
      FPC 2 I3 1 Chip                     OK          55 degrees C / 131 degrees F
      FPC 2 I3 2 TSensor                  OK          49 degrees C / 120 degrees F
      FPC 2 I3 2 Chip                     OK          52 degrees C / 125 degrees F
      FPC 2 I3 3 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 2 I3 3 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 IA 0 TSensor                  OK          50 degrees C / 122 degrees F
      FPC 2 IA 0 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 IA 1 TSensor                  OK          47 degrees C / 116 degrees F
      FPC 2 IA 1 Chip                     OK          53 degrees C / 127 degrees F
Fans  Front Fan                         OK          Spinning at normal speed
      Middle Fan                         OK          Spinning at normal speed
      Rear Fan                           OK          Spinning at normal speed

```

## show chassis environment (MX240 Router with SCBE)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp  PEM 0                             OK          40 degrees C / 104 degrees F
      PEM 1                             OK          45 degrees C / 113 degrees F
      PEM 2                             Absent
      PEM 3                             Absent
      Routing Engine 0                   OK          39 degrees C / 102 degrees F
      Routing Engine 1                   OK          37 degrees C / 98 degrees F
      CB 0 Intake                         OK          36 degrees C / 96 degrees F
      CB 0 Exhaust A                     OK          34 degrees C / 93 degrees F
      CB 0 Exhaust B                     OK          38 degrees C / 100 degrees F
      CB 0 ACBC                           OK          37 degrees C / 98 degrees F
      CB 0 XF A                           OK          49 degrees C / 120 degrees F
      CB 0 XF B                           OK          41 degrees C / 105 degrees F
      CB 1 Intake                         OK          37 degrees C / 98 degrees F
      CB 1 Exhaust A                     OK          34 degrees C / 93 degrees F
      CB 1 Exhaust B                     OK          39 degrees C / 102 degrees F
      CB 1 ACBC                           OK          38 degrees C / 100 degrees F
      CB 1 XF A                           OK          47 degrees C / 116 degrees F
      CB 1 XF B                           OK          41 degrees C / 105 degrees F
      FPC 1 Intake                       OK          33 degrees C / 91 degrees F
      FPC 1 Exhaust A                     OK          38 degrees C / 100 degrees F
      FPC 1 Exhaust B                     OK          53 degrees C / 127 degrees F
      FPC 1 I3 0 TSensor                  OK          50 degrees C / 122 degrees F
      FPC 1 I3 0 Chip                     OK          53 degrees C / 127 degrees F
      FPC 1 I3 1 TSensor                  OK          49 degrees C / 120 degrees F
      FPC 1 I3 1 Chip                     OK          52 degrees C / 125 degrees F
      FPC 1 I3 2 TSensor                  OK          47 degrees C / 116 degrees F
      FPC 1 I3 2 Chip                     OK          49 degrees C / 120 degrees F
      FPC 1 I3 3 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 1 I3 3 Chip                     OK          46 degrees C / 114 degrees F
      FPC 1 IA 0 TSensor                  OK          45 degrees C / 113 degrees F
      FPC 1 IA 0 Chip                     OK          44 degrees C / 111 degrees F
      FPC 1 IA 1 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 1 IA 1 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 Intake                       OK          32 degrees C / 89 degrees F
      FPC 2 Exhaust A                     OK          40 degrees C / 104 degrees F
      FPC 2 Exhaust B                     OK          52 degrees C / 125 degrees F
      FPC 2 I3 0 TSensor                  OK          52 degrees C / 125 degrees F
      FPC 2 I3 0 Chip                     OK          56 degrees C / 132 degrees F
      FPC 2 I3 1 TSensor                  OK          52 degrees C / 125 degrees F
      FPC 2 I3 1 Chip                     OK          55 degrees C / 131 degrees F
      FPC 2 I3 2 TSensor                  OK          49 degrees C / 120 degrees F
      FPC 2 I3 2 Chip                     OK          52 degrees C / 125 degrees F
      FPC 2 I3 3 TSensor                  OK          44 degrees C / 111 degrees F
      FPC 2 I3 3 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 IA 0 TSensor                  OK          50 degrees C / 122 degrees F
      FPC 2 IA 0 Chip                     OK          48 degrees C / 118 degrees F
      FPC 2 IA 1 TSensor                  OK          47 degrees C / 116 degrees F
      FPC 2 IA 1 Chip                     OK          53 degrees C / 127 degrees F
Fans  Front Fan                         OK          Spinning at normal speed
      Middle Fan                         OK          Spinning at normal speed
      Rear Fan                           OK          Spinning at normal speed

```

## show chassis environment (MX480 Router)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp  PEM 0                             OK          35 degrees C / 95 degrees F

```



	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	
	PEM 3	Absent	
	Routing Engine 0	OK	44 degrees C / 111 degrees F
	Routing Engine 1	OK	45 degrees C / 113 degrees F
	CB 0 Intake	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
	CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
	CB 0 ACBC	OK	37 degrees C / 98 degrees F
	CB 0 SF A	OK	51 degrees C / 123 degrees F
	CB 0 SF B	OK	44 degrees C / 111 degrees F
	CB 1 Intake	OK	36 degrees C / 96 degrees F
	CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
	CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
	CB 1 ACBC	OK	37 degrees C / 98 degrees F
	CB 1 SF A	OK	50 degrees C / 122 degrees F
	CB 1 SF B	OK	43 degrees C / 109 degrees F
	FPC 0 Intake	OK	36 degrees C / 96 degrees F
	FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
	FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
	FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
	FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
	FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
	FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
	FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 Intake	OK	37 degrees C / 98 degrees F
	FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
	FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
	FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
	FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
	FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
	FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
	FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
	FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
	FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
	FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
	FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans	Top Rear Fan	OK	Spinning at normal speed
	Bottom Rear Fan	OK	Spinning at normal speed
	Top Middle Fan	OK	Spinning at normal speed
	Bottom Middle Fan	OK	Spinning at normal speed
	Top Front Fan	OK	Spinning at normal speed
	Bottom Front Fan	OK	Spinning at normal speed

#### show chassis environment (MX480 Router with SCBE)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	OK	35 degrees C / 95 degrees F
	PEM 1	OK	40 degrees C / 104 degrees F
	PEM 2	Absent	

PEM 3	Absent	
Routing Engine 0	OK	44 degrees C / 111 degrees F
Routing Engine 1	OK	45 degrees C / 113 degrees F
CB 0 Intake	OK	36 degrees C / 96 degrees F
CB 0 Exhaust A	OK	38 degrees C / 100 degrees F
CB 0 Exhaust B	OK	39 degrees C / 102 degrees F
CB 0 ACBC	OK	37 degrees C / 98 degrees F
CB 0 XF A	OK	51 degrees C / 123 degrees F
CB 0 XF B	OK	44 degrees C / 111 degrees F
CB 1 Intake	OK	36 degrees C / 96 degrees F
CB 1 Exhaust A	OK	39 degrees C / 102 degrees F
CB 1 Exhaust B	OK	40 degrees C / 104 degrees F
CB 1 ACBC	OK	37 degrees C / 98 degrees F
CB 1 XF A	OK	50 degrees C / 122 degrees F
CB 1 XF B	OK	43 degrees C / 109 degrees F
FPC 0 Intake	OK	36 degrees C / 96 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 0 I3 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 0 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 0 I3 3 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	37 degrees C / 98 degrees F
FPC 1 Exhaust A	OK	41 degrees C / 105 degrees F
FPC 1 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 1 I3 0 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 I3 0 Chip	OK	57 degrees C / 134 degrees F
FPC 1 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 1 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 1 I3 2 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 I3 2 Chip	OK	50 degrees C / 122 degrees F
FPC 1 I3 3 TSensor	OK	42 degrees C / 107 degrees F
FPC 1 I3 3 Chip	OK	46 degrees C / 114 degrees F
FPC 1 IA 0 TSensor	OK	49 degrees C / 120 degrees F
FPC 1 IA 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 IA 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 1 IA 1 Chip	OK	50 degrees C / 122 degrees F
Fans		
Top Rear Fan	OK	Spinning at normal speed
Bottom Rear Fan	OK	Spinning at normal speed
Top Middle Fan	OK	Spinning at normal speed
Bottom Middle Fan	OK	Spinning at normal speed
Top Front Fan	OK	Spinning at normal speed
Bottom Front Fan	OK	Spinning at normal speed

### show chassis environment (MX960 Router)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	Absent	
	PEM 2	Check	
	PEM 3	OK	35 degrees C / 95 degrees F
	Routing Engine 0	OK	37 degrees C / 98 degrees F

	Routing Engine 1	Absent	
	CB 0 Intake	OK	24 degrees C / 75 degrees F
	CB 0 Exhaust A	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 Intake	Absent	
	CB 1 Exhaust A	Absent	
	CB 1 Exhaust B	Absent	
	CB 1 ACBC	Absent	
	CB 1 SF A	Absent	
	CB 1 SF B	Absent	
	CB 2 Intake	Absent	
	CB 2 Exhaust A	Absent	
	CB 2 Exhaust B	Absent	
	CB 2 ACBC	Absent	
	CB 2 SF A	Absent	
	CB 2 SF B	Absent	
	FPC 4 Intake	OK	24 degrees C / 75 degrees F
	FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 4 Exhaust B	OK	38 degrees C / 100 degrees F
	FPC 7 Intake	OK	24 degrees C / 75 degrees F
	FPC 7 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust B	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	Failed	
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	Failed	
	Bottom Tray Fan 1	OK	Spinning at normal speed
	Bottom Tray Fan 2	OK	Spinning at normal speed
	Bottom Tray Fan 3	OK	Spinning at normal speed
	Bottom Tray Fan 4	OK	Spinning at normal speed
	Bottom Tray Fan 5	OK	Spinning at normal speed
	Bottom Tray Fan 6	OK	Spinning at normal speed

#### show chassis environment (MX960 Router with SCBE)

user@host> show chassis environment			
Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	50 degrees C / 122 degrees F
	PEM 2	OK	50 degrees C / 122 degrees F
	PEM 3	OK	50 degrees C / 122 degrees F
	Routing Engine 0	OK	42 degrees C / 107 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	39 degrees C / 102 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	35 degrees C / 95 degrees F
	CB 0 Exhaust A	OK	36 degrees C / 96 degrees F
	CB 0 Exhaust B	OK	43 degrees C / 109 degrees F
	CB 0 ACBC	OK	38 degrees C / 100 degrees F
	CB 0 XF A	OK	53 degrees C / 127 degrees F
	CB 0 XF B	OK	47 degrees C / 116 degrees F
	CB 1 Intake	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust A	OK	35 degrees C / 95 degrees F
	CB 1 Exhaust B	OK	41 degrees C / 105 degrees F
	CB 1 ACBC	OK	38 degrees C / 100 degrees F
	CB 1 XF A	OK	52 degrees C / 125 degrees F
	CB 1 XF B	OK	47 degrees C / 116 degrees F

CB 2 Intake	OK	32 degrees C / 89 degrees F
CB 2 Exhaust A	OK	30 degrees C / 86 degrees F
CB 2 Exhaust B	OK	35 degrees C / 95 degrees F
CB 2 ACBC	OK	33 degrees C / 91 degrees F
CB 2 XF A	OK	51 degrees C / 123 degrees F
CB 2 XF B	OK	50 degrees C / 122 degrees F
FPC 0 Intake	OK	35 degrees C / 95 degrees F
FPC 0 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 0 Exhaust B	OK	50 degrees C / 122 degrees F
FPC 0 I3 0 TSensor	OK	50 degrees C / 122 degrees F
FPC 0 I3 0 Chip	OK	56 degrees C / 132 degrees F
FPC 0 I3 1 TSensor	OK	47 degrees C / 116 degrees F
FPC 0 I3 1 Chip	OK	50 degrees C / 122 degrees F
FPC 0 I3 2 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 I3 2 Chip	OK	48 degrees C / 118 degrees F
FPC 0 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 0 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 0 IA 0 TSensor	OK	45 degrees C / 113 degrees F
FPC 0 IA 0 Chip	OK	45 degrees C / 113 degrees F
FPC 0 IA 1 TSensor	OK	44 degrees C / 111 degrees F
FPC 0 IA 1 Chip	OK	48 degrees C / 118 degrees F
FPC 1 Intake	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 1 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 1 LU 0 TCAM TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 TCAM Chip	OK	57 degrees C / 134 degrees F
FPC 1 LU 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 1 MQ 0 TSensor	OK	53 degrees C / 127 degrees F
FPC 1 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TCAM TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 TCAM Chip	OK	52 degrees C / 125 degrees F
FPC 1 LU 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 LU 1 Chip	OK	53 degrees C / 127 degrees F
FPC 1 MQ 1 TSensor	OK	51 degrees C / 123 degrees F
FPC 1 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 2 Intake	OK	35 degrees C / 95 degrees F
FPC 2 Exhaust A	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust B	OK	54 degrees C / 129 degrees F
FPC 2 I3 0 TSensor	OK	52 degrees C / 125 degrees F
FPC 2 I3 0 Chip	OK	59 degrees C / 138 degrees F
FPC 2 I3 1 TSensor	OK	48 degrees C / 118 degrees F
FPC 2 I3 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 I3 2 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 I3 2 Chip	OK	49 degrees C / 120 degrees F
FPC 2 I3 3 TSensor	OK	41 degrees C / 105 degrees F
FPC 2 I3 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 IA 0 TSensor	OK	47 degrees C / 116 degrees F
FPC 2 IA 0 Chip	OK	46 degrees C / 114 degrees F
FPC 2 IA 1 TSensor	OK	45 degrees C / 113 degrees F
FPC 2 IA 1 Chip	OK	49 degrees C / 120 degrees F
FPC 3 Intake	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 3 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 3 I3 0 TSensor	OK	48 degrees C / 118 degrees F
FPC 3 I3 0 Chip	OK	52 degrees C / 125 degrees F
FPC 3 I3 1 TSensor	OK	46 degrees C / 114 degrees F
FPC 3 I3 1 Chip	OK	48 degrees C / 118 degrees F
FPC 3 IA 0 TSensor	OK	41 degrees C / 105 degrees F
FPC 3 IA 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 Intake	OK	42 degrees C / 107 degrees F

	FPC 5 Exhaust A	OK	42 degrees C / 107 degrees F
	FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
	FPC 5 LU 0 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 0 Chip	OK	54 degrees C / 129 degrees F
	FPC 5 LU 1 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 1 Chip	OK	61 degrees C / 141 degrees F
	FPC 5 LU 2 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 2 Chip	OK	51 degrees C / 123 degrees F
	FPC 5 LU 3 TSensor	OK	53 degrees C / 127 degrees F
	FPC 5 LU 3 Chip	OK	53 degrees C / 127 degrees F
	FPC 5 MQ 0 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
	FPC 5 MQ 1 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 5 MQ 2 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 2 Chip	OK	46 degrees C / 114 degrees F
	FPC 5 MQ 3 TSensor	OK	47 degrees C / 116 degrees F
	FPC 5 MQ 3 Chip	OK	45 degrees C / 113 degrees F
	FPC 7 Intake	OK	36 degrees C / 96 degrees F
	FPC 7 Exhaust A	OK	35 degrees C / 95 degrees F
	FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
	FPC 7 QX 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 QX 0 Chip	OK	47 degrees C / 116 degrees F
	FPC 7 LU 0 TCAM TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 TCAM Chip	OK	44 degrees C / 111 degrees F
	FPC 7 LU 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 LU 0 Chip	OK	46 degrees C / 114 degrees F
	FPC 7 MQ 0 TSensor	OK	42 degrees C / 107 degrees F
	FPC 7 MQ 0 Chip	OK	45 degrees C / 113 degrees F
	FPC 8 Intake	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust A	OK	33 degrees C / 91 degrees F
	FPC 8 Exhaust B	OK	36 degrees C / 96 degrees F
	FPC 8 I3 0 TSensor	OK	38 degrees C / 100 degrees F
	FPC 8 I3 0 Chip	OK	43 degrees C / 109 degrees F
	FPC 8 BDS 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 BDS 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 8 IA 0 TSensor	OK	37 degrees C / 98 degrees F
	FPC 8 IA 0 Chip	OK	37 degrees C / 98 degrees F
	FPC 10 Intake	OK	38 degrees C / 100 degrees F
	FPC 10 Exhaust A	OK	36 degrees C / 96 degrees F
	FPC 10 Exhaust B	OK	41 degrees C / 105 degrees F
	FPC 10 I3 0 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 0 Chip	OK	42 degrees C / 107 degrees F
	FPC 10 I3 1 TSensor	OK	40 degrees C / 104 degrees F
	FPC 10 I3 1 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 I3 2 TSensor	OK	42 degrees C / 107 degrees F
	FPC 10 I3 2 Chip	OK	43 degrees C / 109 degrees F
	FPC 10 I3 3 TSensor	OK	39 degrees C / 102 degrees F
	FPC 10 I3 3 Chip	OK	44 degrees C / 111 degrees F
	FPC 10 IA 0 TSensor	OK	36 degrees C / 96 degrees F
	FPC 10 IA 0 Chip	OK	36 degrees C / 96 degrees F
	FPC 10 IA 1 TSensor	OK	43 degrees C / 109 degrees F
	FPC 10 IA 1 Chip	OK	42 degrees C / 107 degrees F
Fans	Top Fan Tray Temp	OK	37 degrees C / 98 degrees F
	Top Tray Fan 1	OK	Spinning at normal speed
	Top Tray Fan 2	OK	Spinning at normal speed
	Top Tray Fan 3	OK	Spinning at normal speed
	Top Tray Fan 4	OK	Spinning at normal speed
	Top Tray Fan 5	OK	Spinning at normal speed
	Top Tray Fan 6	OK	Spinning at normal speed
	Bottom Fan Tray Temp	OK	28 degrees C / 82 degrees F

Bottom Tray Fan 1	OK	Spinning at normal speed
Bottom Tray Fan 2	OK	Spinning at normal speed
Bottom Tray Fan 3	OK	Spinning at normal speed
Bottom Tray Fan 4	OK	Spinning at normal speed
Bottom Tray Fan 5	OK	Spinning at normal speed
Bottom Tray Fan 6	OK	Spinning at normal speed

#### show chassis environment (MX960 Router with MPC5EQ)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	50 degrees C / 122 degrees F
	PEM 1	OK	45 degrees C / 113 degrees F
	PEM 2	OK	45 degrees C / 113 degrees F
	PEM 3	Absent	
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 0 CPU	OK	30 degrees C / 86 degrees F
	Routing Engine 1	Present	
	Routing Engine 1 CPU	Present	
	CB 0 Intake	OK	29 degrees C / 84 degrees F
	CB 0 Exhaust A	OK	29 degrees C / 84 degrees F
	CB 0 Exhaust B	OK	34 degrees C / 93 degrees F
	CB 0 ACBC	OK	32 degrees C / 89 degrees F
	CB 0 XF A	OK	49 degrees C / 120 degrees F
	CB 0 XF B	OK	45 degrees C / 113 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	27 degrees C / 80 degrees F
	CB 1 ACBC	OK	26 degrees C / 78 degrees F
	CB 1 XF A	OK	32 degrees C / 89 degrees F
	CB 1 XF B	OK	32 degrees C / 89 degrees F
	CB 2 Intake	OK	28 degrees C / 82 degrees F
	CB 2 Exhaust A	OK	27 degrees C / 80 degrees F
	CB 2 Exhaust B	OK	33 degrees C / 91 degrees F
	CB 2 ACBC	OK	30 degrees C / 86 degrees F
	CB 2 XF A	OK	48 degrees C / 118 degrees F
	CB 2 XF B	OK	46 degrees C / 114 degrees F
	FPC 0 Intake	OK	38 degrees C / 100 degrees F
	FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
	FPC 0 Exhaust B	OK	49 degrees C / 120 degrees F
	FPC 0 XL TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XL Chip	OK	50 degrees C / 122 degrees F
	FPC 0 XL_XR0 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XL_XR0 Chip	OK	53 degrees C / 127 degrees F
	FPC 0 XL_XR1 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XL_XR1 Chip	OK	54 degrees C / 129 degrees F
	FPC 0 XQ TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XQ Chip	OK	52 degrees C / 125 degrees F
	FPC 0 XQ_XR0 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XQ_XR0 Chip	OK	62 degrees C / 143 degrees F
	FPC 0 XQ_XR1 TSen	OK	48 degrees C / 118 degrees F
	FPC 0 XQ_XR1 Chip	OK	62 degrees C / 143 degrees F
	FPC 0 XM 0 TSen	OK	53 degrees C / 127 degrees F
	FPC 0 XM 0 Chip	OK	63 degrees C / 145 degrees F
	FPC 0 XM 1 TSen	OK	53 degrees C / 127 degrees F
	FPC 0 XM 1 Chip	OK	46 degrees C / 114 degrees F
	FPC 0 PLX PCIe Switch TSe	OK	53 degrees C / 127 degrees F
	FPC 0 PLX PCIe Switch Chi	OK	66 degrees C / 150 degrees F
	FPC 1 Intake	OK	31 degrees C / 87 degrees F
	FPC 1 Exhaust A	OK	38 degrees C / 100 degrees F
	FPC 1 Exhaust B	OK	49 degrees C / 120 degrees F

FPC 1 LU 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 0 Chip	OK	47 degrees C / 116 degrees F
FPC 1 LU 1 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 1 Chip	OK	42 degrees C / 107 degrees F
FPC 1 LU 2 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 2 Chip	OK	46 degrees C / 114 degrees F
FPC 1 LU 3 TSen	OK	41 degrees C / 105 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 XM 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 XM 0 Chip	OK	49 degrees C / 120 degrees F
FPC 1 XF 0 TSen	OK	41 degrees C / 105 degrees F
FPC 1 XF 0 Chip	OK	63 degrees C / 145 degrees F
FPC 1 PLX Switch TSen	OK	41 degrees C / 105 degrees F
FPC 1 PLX Switch Chip	OK	43 degrees C / 109 degrees F
FPC 3 Intake	OK	31 degrees C / 87 degrees F
FPC 3 Exhaust A	OK	37 degrees C / 98 degrees F
FPC 3 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 3 LU 0 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 0 Chip	OK	43 degrees C / 109 degrees F
FPC 3 LU 1 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 1 Chip	OK	46 degrees C / 114 degrees F
FPC 3 LU 2 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 2 Chip	OK	40 degrees C / 104 degrees F
FPC 3 LU 3 TSen	OK	42 degrees C / 107 degrees F
FPC 3 LU 3 Chip	OK	41 degrees C / 105 degrees F
FPC 3 MQ 0 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 0 Chip	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 Chip	OK	40 degrees C / 104 degrees F
FPC 3 MQ 2 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 2 Chip	OK	36 degrees C / 96 degrees F
FPC 3 MQ 3 TSen	OK	37 degrees C / 98 degrees F
FPC 3 MQ 3 Chip	OK	38 degrees C / 100 degrees F
FPC 4 Intake	OK	34 degrees C / 93 degrees F
FPC 4 Exhaust A	OK	45 degrees C / 113 degrees F
FPC 4 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 4 XL TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL Chip	OK	47 degrees C / 116 degrees F
FPC 4 XL_XR0 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL_XR0 Chip	OK	48 degrees C / 118 degrees F
FPC 4 XL_XR1 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XL_XR1 Chip	OK	47 degrees C / 116 degrees F
FPC 4 XQ TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ Chip	OK	47 degrees C / 116 degrees F
FPC 4 XQ_XR0 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 4 XQ_XR1 TSen	OK	44 degrees C / 111 degrees F
FPC 4 XQ_XR1 Chip	OK	58 degrees C / 136 degrees F
FPC 4 XM 0 TSen	OK	51 degrees C / 123 degrees F
FPC 4 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 4 XM 1 TSen	OK	51 degrees C / 123 degrees F
FPC 4 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 4 PLX PCIe Switch TSe	OK	51 degrees C / 123 degrees F
FPC 4 PLX PCIe Switch Chi	OK	60 degrees C / 140 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	45 degrees C / 113 degrees F
FPC 5 Exhaust B	OK	47 degrees C / 116 degrees F
FPC 5 XL TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL Chip	OK	47 degrees C / 116 degrees F
FPC 5 XL_XR0 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL_XR0 Chip	OK	49 degrees C / 120 degrees F

FPC 5 XL_XR1 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XL_XR1 Chip	OK	49 degrees C / 120 degrees F
FPC 5 XQ TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ Chip	OK	48 degrees C / 118 degrees F
FPC 5 XQ_XR0 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ_XR0 Chip	OK	60 degrees C / 140 degrees F
FPC 5 XQ_XR1 TSen	OK	45 degrees C / 113 degrees F
FPC 5 XQ_XR1 Chip	OK	58 degrees C / 136 degrees F
FPC 5 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 5 XM 0 Chip	OK	48 degrees C / 118 degrees F
FPC 5 XM 1 TSen	OK	50 degrees C / 122 degrees F
FPC 5 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 5 PLX PCIe Switch TSe	OK	50 degrees C / 122 degrees F
FPC 5 PLX PCIe Switch Chi	OK	59 degrees C / 138 degrees F
FPC 7 Intake	OK	32 degrees C / 89 degrees F
FPC 7 Exhaust A	OK	32 degrees C / 89 degrees F
FPC 7 Exhaust B	OK	33 degrees C / 91 degrees F
FPC 7 LU 0 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 0 Chip	OK	44 degrees C / 111 degrees F
FPC 7 LU 1 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 1 Chip	OK	47 degrees C / 116 degrees F
FPC 7 LU 2 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 2 Chip	OK	39 degrees C / 102 degrees F
FPC 7 LU 3 TSen	OK	49 degrees C / 120 degrees F
FPC 7 LU 3 Chip	OK	43 degrees C / 109 degrees F
FPC 7 XM 0 TSen	OK	49 degrees C / 120 degrees F
FPC 7 XM 0 Chip	OK	57 degrees C / 134 degrees F
FPC 7 XM 1 TSen	OK	49 degrees C / 120 degrees F
FPC 7 XM 1 Chip	OK	48 degrees C / 118 degrees F
FPC 7 PLX Switch TSen	OK	49 degrees C / 120 degrees F
FPC 7 PLX Switch Chip	OK	45 degrees C / 113 degrees F
FPC 8 Intake	OK	36 degrees C / 96 degrees F
FPC 8 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 8 Exhaust B	OK	46 degrees C / 114 degrees F
FPC 8 XL TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL Chip	OK	47 degrees C / 116 degrees F
FPC 8 XL_XR0 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL_XR0 Chip	OK	53 degrees C / 127 degrees F
FPC 8 XL_XR1 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XL_XR1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 XQ TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ Chip	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR0 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR0 Chip	OK	59 degrees C / 138 degrees F
FPC 8 XQ_XR1 TSen	OK	46 degrees C / 114 degrees F
FPC 8 XQ_XR1 Chip	OK	57 degrees C / 134 degrees F
FPC 8 XM 0 TSen	OK	52 degrees C / 125 degrees F
FPC 8 XM 0 Chip	OK	61 degrees C / 141 degrees F
FPC 8 XM 1 TSen	OK	52 degrees C / 125 degrees F
FPC 8 XM 1 Chip	OK	47 degrees C / 116 degrees F
FPC 8 PLX PCIe Switch TSe	OK	52 degrees C / 125 degrees F
FPC 8 PLX PCIe Switch Chi	OK	63 degrees C / 145 degrees F
FPC 9 Intake	OK	31 degrees C / 87 degrees F
FPC 9 Exhaust A	OK	34 degrees C / 93 degrees F
FPC 9 Exhaust B	OK	35 degrees C / 95 degrees F
FPC 9 QX 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 QX 0 Chip	OK	45 degrees C / 113 degrees F
FPC 9 LU 0 TCAM TSen	OK	42 degrees C / 107 degrees F
FPC 9 LU 0 TCAM Chip	OK	41 degrees C / 105 degrees F
FPC 9 LU 0 TSen	OK	42 degrees C / 107 degrees F
FPC 9 LU 0 Chip	OK	43 degrees C / 109 degrees F



	FPC 9 MQ 0 TSen	OK	42 degrees C / 107 degrees F
	FPC 9 MQ 0 Chip	OK	43 degrees C / 109 degrees F
	FPC 9 QX 1 TSen	OK	38 degrees C / 100 degrees F
	FPC 9 QX 1 Chip	OK	40 degrees C / 104 degrees F
	FPC 9 LU 1 TCAM TSen	OK	38 degrees C / 100 degrees F
	FPC 9 LU 1 TCAM Chip	OK	38 degrees C / 100 degrees F
	FPC 9 LU 1 TSen	OK	38 degrees C / 100 degrees F
	FPC 9 LU 1 Chip	OK	41 degrees C / 105 degrees F
	FPC 9 MQ 1 TSen	OK	38 degrees C / 100 degrees F
	FPC 9 MQ 1 Chip	OK	41 degrees C / 105 degrees F
	FPC 10 Intake	OK	35 degrees C / 95 degrees F
	FPC 10 Exhaust A	OK	51 degrees C / 123 degrees F
	FPC 10 Exhaust B	OK	46 degrees C / 114 degrees F
	FPC 10 XL TSen	OK	42 degrees C / 107 degrees F
	FPC 10 XL Chip	OK	44 degrees C / 111 degrees F
	FPC 10 XL_XR0 TSen	OK	42 degrees C / 107 degrees F
	FPC 10 XL_XR0 Chip	OK	47 degrees C / 116 degrees F
	FPC 10 XL_XR1 TSen	OK	42 degrees C / 107 degrees F
	FPC 10 XL_XR1 Chip	OK	48 degrees C / 118 degrees F
	FPC 10 XQ TSen	OK	42 degrees C / 107 degrees F
	FPC 10 XQ Chip	OK	46 degrees C / 114 degrees F
	FPC 10 XQ_XR0 TSen	OK	42 degrees C / 107 degrees F
	FPC 10 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
	FPC 10 XQ_XR1 TSen	OK	42 degrees C / 107 degrees F
	FPC 10 XQ_XR1 Chip	OK	53 degrees C / 127 degrees F
	FPC 10 XM 0 TSen	OK	51 degrees C / 123 degrees F
	FPC 10 XM 0 Chip	OK	61 degrees C / 141 degrees F
	FPC 10 XM 1 TSen	OK	51 degrees C / 123 degrees F
	FPC 10 XM 1 Chip	OK	49 degrees C / 120 degrees F
	FPC 10 PLX PCIe Switch TSe	OK	51 degrees C / 123 degrees F
	FPC 10 PLX PCIe Switch Chi	OK	61 degrees C / 141 degrees F
	FPC 11 Intake	OK	33 degrees C / 91 degrees F
	FPC 11 Exhaust A	OK	33 degrees C / 91 degrees F
	FPC 11 Exhaust B	OK	34 degrees C / 93 degrees F
	FPC 11 LU 0 TSen	OK	50 degrees C / 122 degrees F
	FPC 11 LU 0 Chip	OK	48 degrees C / 118 degrees F
	FPC 11 LU 1 TSen	OK	50 degrees C / 122 degrees F
	FPC 11 LU 1 Chip	OK	50 degrees C / 122 degrees F
	FPC 11 LU 2 TSen	OK	50 degrees C / 122 degrees F
	FPC 11 LU 2 Chip	OK	41 degrees C / 105 degrees F
	FPC 11 LU 3 TSen	OK	50 degrees C / 122 degrees F
	FPC 11 LU 3 Chip	OK	48 degrees C / 118 degrees F
	FPC 11 XM 0 TSen	OK	50 degrees C / 122 degrees F
	FPC 11 XM 0 Chip	OK	57 degrees C / 134 degrees F
	FPC 11 XM 1 TSen	OK	50 degrees C / 122 degrees F
	FPC 11 XM 1 Chip	OK	52 degrees C / 125 degrees F
	FPC 11 PLX Switch TSen	OK	50 degrees C / 122 degrees F
	FPC 11 PLX Switch Chip	OK	45 degrees C / 113 degrees F
Fans	Top Fan Tray Temp	OK	42 degrees C / 107 degrees F
	Top Tray Fan 1	OK	Spinning at high speed
Top Tray Fan 2	OK	Spinning at high speed	
	Top Tray Fan 3	OK	Spinning at high speed
	Top Tray Fan 4	OK	Spinning at high speed
	Top Tray Fan 5	OK	Spinning at high speed
	Top Tray Fan 6	OK	Spinning at high speed
	Top Tray Fan 7	OK	Spinning at high speed
	Top Tray Fan 8	OK	Spinning at high speed
	Top Tray Fan 9	OK	Spinning at high speed
	Top Tray Fan 10	OK	Spinning at high speed
	Top Tray Fan 11	OK	Spinning at high speed
	Top Tray Fan 12	OK	Spinning at high speed

Bottom Fan Tray Temp	OK	33 degrees C / 91 degrees F
Bottom Tray Fan 1	OK	Spinning at high speed
Bottom Tray Fan 2	OK	Spinning at high speed
Bottom Tray Fan 3	OK	Spinning at high speed
Bottom Tray Fan 4	OK	Spinning at high speed
Bottom Tray Fan 5	OK	Spinning at high speed
Bottom Tray Fan 6	OK	Spinning at high speed
Bottom Tray Fan 7	OK	Spinning at high speed
Bottom Tray Fan 8	OK	Spinning at high speed
Bottom Tray Fan 9	OK	Spinning at high speed
Bottom Tray Fan 10	OK	Spinning at high speed
Bottom Tray Fan 11	OK	Spinning at high speed
Bottom Tray Fan 12	OK	Spinning at high speed

### show chassis environment (MX2020 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PSM 0	Absent	
	PSM 1	Absent	
	PSM 2	OK	41 degrees C / 105 degrees F
	PSM 3	OK	39 degrees C / 102 degrees F
	PSM 4	OK	39 degrees C / 102 degrees F
	PSM 5	OK	38 degrees C / 100 degrees F
	PSM 6	OK	38 degrees C / 100 degrees F
	PSM 7	OK	38 degrees C / 100 degrees F
	PSM 8	OK	37 degrees C / 98 degrees F
	PSM 9	Absent	
	PSM 10	Absent	
	PSM 11	OK	47 degrees C / 116 degrees F
	PSM 12	OK	45 degrees C / 113 degrees F
	PSM 13	OK	44 degrees C / 111 degrees F
	PSM 14	OK	44 degrees C / 111 degrees F
	PSM 15	OK	43 degrees C / 109 degrees F
	PSM 16	OK	42 degrees C / 107 degrees F
	PSM 17	OK	41 degrees C / 105 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	PDM 2	Absent	
	PDM 3	OK	
	CB 0 IntakeA-Zone0	OK	45 degrees C / 113 degrees F
	CB 0 IntakeB-Zone1	OK	34 degrees C / 93 degrees F
	CB 0 IntakeC-Zone0	OK	48 degrees C / 118 degrees F
	CB 0 ExhaustA-Zone0	OK	45 degrees C / 113 degrees F
	CB 0 ExhaustB-Zone1	OK	37 degrees C / 98 degrees F
	CB 0 TCBC-Zone0	OK	41 degrees C / 105 degrees F
	CB 1 IntakeA-Zone0	OK	46 degrees C / 114 degrees F
	CB 1 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
	CB 1 IntakeC-Zone0	OK	49 degrees C / 120 degrees F
	CB 1 ExhaustA-Zone0	OK	46 degrees C / 114 degrees F
	CB 1 ExhaustB-Zone1	OK	41 degrees C / 105 degrees F
	CB 1 TCBC-Zone0	OK	46 degrees C / 114 degrees F
	SPMB 0 Intake	OK	33 degrees C / 91 degrees F
	SPMB 1 Intake	OK	42 degrees C / 107 degrees F
	Routing Engine 0	OK	35 degrees C / 95 degrees F
	Routing Engine 0 CPU	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	44 degrees C / 111 degrees F
	Routing Engine 1 CPU	OK	42 degrees C / 107 degrees F
	SFB 0 Intake-Zone0	OK	55 degrees C / 131 degrees F
	SFB 0 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
	SFB 0 IntakeA-Zone0	OK	50 degrees C / 122 degrees F

SFB 0 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 0 Exhaust-Zone0	OK	52 degrees C / 125 degrees F
SFB 0 SFB-XF2-Zone1	OK	61 degrees C / 141 degrees F
SFB 0 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 0 SFB-XF0-Zone0	OK	68 degrees C / 154 degrees F
SFB 1 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 1 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 1 Exhaust-Zone0	OK	51 degrees C / 123 degrees F
SFB 1 SFB-XF2-Zone1	OK	62 degrees C / 143 degrees F
SFB 1 SFB-XF1-Zone0	OK	67 degrees C / 152 degrees F
SFB 1 SFB-XF0-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 Intake-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Exhaust-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 IntakeA-Zone0	OK	51 degrees C / 123 degrees F
SFB 2 IntakeB-Zone1	OK	40 degrees C / 104 degrees F
SFB 2 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 2 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 2 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 2 SFB-XF0-Zone0	OK	70 degrees C / 158 degrees F
SFB 3 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 3 Exhaust-Zone1	OK	48 degrees C / 118 degrees F
SFB 3 IntakeA-Zone0	OK	52 degrees C / 125 degrees F
SFB 3 IntakeB-Zone1	OK	41 degrees C / 105 degrees F
SFB 3 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 3 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 3 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 3 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 4 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 4 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 4 IntakeA-Zone0	OK	54 degrees C / 129 degrees F
SFB 4 IntakeB-Zone1	OK	42 degrees C / 107 degrees F
SFB 4 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 4 SFB-XF2-Zone1	OK	64 degrees C / 147 degrees F
SFB 4 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 4 SFB-XF0-Zone0	OK	71 degrees C / 159 degrees F
SFB 5 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 5 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 5 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 5 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 5 SFB-XF2-Zone1	OK	66 degrees C / 150 degrees F
SFB 5 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 5 SFB-XF0-Zone0	OK	74 degrees C / 165 degrees F
SFB 6 Intake-Zone0	OK	58 degrees C / 136 degrees F
SFB 6 Exhaust-Zone1	OK	49 degrees C / 120 degrees F
SFB 6 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 Exhaust-Zone0	OK	53 degrees C / 127 degrees F
SFB 6 SFB-XF2-Zone1	OK	65 degrees C / 149 degrees F
SFB 6 SFB-XF1-Zone0	OK	68 degrees C / 154 degrees F
SFB 6 SFB-XF0-Zone0	OK	72 degrees C / 161 degrees F
SFB 7 Intake-Zone0	OK	57 degrees C / 134 degrees F
SFB 7 Exhaust-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 IntakeA-Zone0	OK	53 degrees C / 127 degrees F
SFB 7 IntakeB-Zone1	OK	43 degrees C / 109 degrees F
SFB 7 Exhaust-Zone0	OK	54 degrees C / 129 degrees F
SFB 7 SFB-XF2-Zone1	OK	68 degrees C / 154 degrees F
SFB 7 SFB-XF1-Zone0	OK	69 degrees C / 156 degrees F
SFB 7 SFB-XF0-Zone0	OK	73 degrees C / 163 degrees F

FPC 0 Intake	OK	41 degrees C / 105 degrees F
FPC 0 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 0 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 0 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 0 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 0 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 0 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 0 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 0 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 0 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 0 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 0 MQ 3 Chip	OK	45 degrees C / 113 degrees F
FPC 1 Intake	OK	40 degrees C / 104 degrees F
FPC 1 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 1 Exhaust B	OK	58 degrees C / 136 degrees F
FPC 1 LU 0 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 1 LU 1 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 1 Chip	OK	58 degrees C / 136 degrees F
FPC 1 LU 2 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 2 Chip	OK	49 degrees C / 120 degrees F
FPC 1 LU 3 TSen	OK	55 degrees C / 131 degrees F
FPC 1 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 1 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 1 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 1 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 1 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 1 MQ 3 Chip	OK	44 degrees C / 111 degrees F
FPC 2 Intake	OK	39 degrees C / 102 degrees F
FPC 2 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 2 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 2 LU 0 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 2 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 2 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 2 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 2 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 2 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 2 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 2 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 2 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 3 Intake	OK	40 degrees C / 104 degrees F
FPC 3 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 3 Exhaust B	OK	61 degrees C / 141 degrees F
FPC 3 LU 0 TSen	OK	58 degrees C / 136 degrees F

FPC 3 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 3 LU 1 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 3 LU 2 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 3 LU 3 TSen	OK	58 degrees C / 136 degrees F
FPC 3 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 3 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 3 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 3 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 3 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 3 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 4 Intake	OK	40 degrees C / 104 degrees F
FPC 4 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 4 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 4 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 4 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 4 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 2 Chip	OK	51 degrees C / 123 degrees F
FPC 4 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 4 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 4 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 4 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 4 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 5 Intake	OK	41 degrees C / 105 degrees F
FPC 5 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 5 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 5 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 0 Chip	OK	63 degrees C / 145 degrees F
FPC 5 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 5 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 5 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 5 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 5 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 5 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 5 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 5 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 5 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 6 Intake	OK	42 degrees C / 107 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 6 LU 0 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 6 LU 1 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 1 Chip	OK	66 degrees C / 150 degrees F
FPC 6 LU 2 TSen	OK	61 degrees C / 141 degrees F

FPC 6 LU 2 Chip	OK	56 degrees C / 132 degrees F
FPC 6 LU 3 TSen	OK	61 degrees C / 141 degrees F
FPC 6 LU 3 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 0 Chip	OK	56 degrees C / 132 degrees F
FPC 6 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 6 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 6 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 6 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 7 Intake	OK	41 degrees C / 105 degrees F
FPC 7 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 7 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 7 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 0 Chip	OK	61 degrees C / 141 degrees F
FPC 7 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 1 Chip	OK	65 degrees C / 149 degrees F
FPC 7 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 7 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 7 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 0 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 0 Chip	OK	53 degrees C / 127 degrees F
FPC 7 MQ 1 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 7 MQ 2 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 7 MQ 3 TSen	OK	50 degrees C / 122 degrees F
FPC 7 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 8 Intake	OK	41 degrees C / 105 degrees F
FPC 8 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 8 Exhaust B	OK	62 degrees C / 143 degrees F
FPC 8 LU 0 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 0 Chip	OK	62 degrees C / 143 degrees F
FPC 8 LU 1 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 1 Chip	OK	64 degrees C / 147 degrees F
FPC 8 LU 2 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 2 Chip	OK	55 degrees C / 131 degrees F
FPC 8 LU 3 TSen	OK	59 degrees C / 138 degrees F
FPC 8 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 8 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 8 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 8 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 8 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 8 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 9 Intake	OK	42 degrees C / 107 degrees F
FPC 9 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 9 Exhaust B	OK	63 degrees C / 145 degrees F
FPC 9 LU 0 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 0 Chip	OK	65 degrees C / 149 degrees F
FPC 9 LU 1 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 1 Chip	OK	67 degrees C / 152 degrees F
FPC 9 LU 2 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 2 Chip	OK	54 degrees C / 129 degrees F
FPC 9 LU 3 TSen	OK	60 degrees C / 140 degrees F
FPC 9 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 9 MQ 0 TSen	OK	51 degrees C / 123 degrees F

FPC 9 MQ 0 Chip	OK	55 degrees C / 131 degrees F
FPC 9 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 1 Chip	OK	59 degrees C / 138 degrees F
FPC 9 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 2 Chip	OK	49 degrees C / 120 degrees F
FPC 9 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 9 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 10 Intake	OK	44 degrees C / 111 degrees F
FPC 10 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 10 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 10 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 0 Chip	OK	55 degrees C / 131 degrees F
FPC 10 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 1 Chip	OK	59 degrees C / 138 degrees F
FPC 10 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 2 Chip	OK	52 degrees C / 125 degrees F
FPC 10 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 10 LU 3 Chip	OK	51 degrees C / 123 degrees F
FPC 10 MQ 0 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 0 Chip	OK	49 degrees C / 120 degrees F
FPC 10 MQ 1 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 1 Chip	OK	52 degrees C / 125 degrees F
FPC 10 MQ 2 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 10 MQ 3 TSen	OK	48 degrees C / 118 degrees F
FPC 10 MQ 3 Chip	OK	47 degrees C / 116 degrees F
FPC 11 Intake	OK	30 degrees C / 86 degrees F
FPC 11 Exhaust A	OK	35 degrees C / 95 degrees F
FPC 11 Exhaust B	OK	30 degrees C / 86 degrees F
FPC 11 LU 0 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 0 Chip	OK	58 degrees C / 136 degrees F
FPC 11 LU 1 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 11 LU 2 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 11 LU 3 TSen	OK	57 degrees C / 134 degrees F
FPC 11 LU 3 Chip	OK	54 degrees C / 129 degrees F
FPC 11 MQ 0 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 0 Chip	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 1 Chip	OK	57 degrees C / 134 degrees F
FPC 11 MQ 2 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 2 Chip	OK	48 degrees C / 118 degrees F
FPC 11 MQ 3 TSen	OK	52 degrees C / 125 degrees F
FPC 11 MQ 3 Chip	OK	52 degrees C / 125 degrees F
FPC 12 Intake	OK	40 degrees C / 104 degrees F
FPC 12 Exhaust A	OK	47 degrees C / 116 degrees F
FPC 12 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 12 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 12 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 12 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 12 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 12 LU 3 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 12 MQ 2 TSen	OK	46 degrees C / 114 degrees F

FPC 12 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 12 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 12 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 13 Intake	OK	40 degrees C / 104 degrees F
FPC 13 Exhaust A	OK	48 degrees C / 118 degrees F
FPC 13 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 13 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 13 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 13 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 13 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 13 LU 3 Chip	OK	48 degrees C / 118 degrees F
FPC 13 MQ 0 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 1 Chip	OK	50 degrees C / 122 degrees F
FPC 13 MQ 2 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 2 Chip	OK	44 degrees C / 111 degrees F
FPC 13 MQ 3 TSen	OK	46 degrees C / 114 degrees F
FPC 13 MQ 3 Chip	OK	46 degrees C / 114 degrees F
FPC 14 Intake	OK	40 degrees C / 104 degrees F
FPC 14 Exhaust A	OK	50 degrees C / 122 degrees F
FPC 14 Exhaust B	OK	51 degrees C / 123 degrees F
FPC 14 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 0 Chip	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 1 Chip	OK	54 degrees C / 129 degrees F
FPC 14 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 2 Chip	OK	47 degrees C / 116 degrees F
FPC 14 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 14 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 14 MQ 0 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 0 Chip	OK	46 degrees C / 114 degrees F
FPC 14 MQ 1 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 1 Chip	OK	51 degrees C / 123 degrees F
FPC 14 MQ 2 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 2 Chip	OK	45 degrees C / 113 degrees F
FPC 14 MQ 3 TSen	OK	47 degrees C / 116 degrees F
FPC 14 MQ 3 Chip	OK	48 degrees C / 118 degrees F
FPC 15 Intake	OK	44 degrees C / 111 degrees F
FPC 15 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 15 Exhaust B	OK	60 degrees C / 140 degrees F
FPC 15 LU 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 LU 1 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 1 Chip	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 2 Chip	OK	58 degrees C / 136 degrees F
FPC 15 LU 3 TSen	OK	50 degrees C / 122 degrees F
FPC 15 LU 3 Chip	OK	63 degrees C / 145 degrees F
FPC 15 XM 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XM 0 Chip	OK	56 degrees C / 132 degrees F
FPC 15 XF 0 TSen	OK	50 degrees C / 122 degrees F
FPC 15 XF 0 Chip	OK	68 degrees C / 154 degrees F
FPC 15 PLX Switch TSen	OK	50 degrees C / 122 degrees F
FPC 15 PLX Switch Chip	OK	56 degrees C / 132 degrees F
FPC 16 Intake	OK	42 degrees C / 107 degrees F
FPC 16 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 16 Exhaust B	OK	53 degrees C / 127 degrees F



FPC 16 LU 0 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 0 Chip	OK	52 degrees C / 125 degrees F
FPC 16 LU 1 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 1 Chip	OK	55 degrees C / 131 degrees F
FPC 16 LU 2 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 2 Chip	OK	48 degrees C / 118 degrees F
FPC 16 LU 3 TSen	OK	51 degrees C / 123 degrees F
FPC 16 LU 3 Chip	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 0 Chip	OK	48 degrees C / 118 degrees F
FPC 16 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 1 Chip	OK	53 degrees C / 127 degrees F
FPC 16 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 2 Chip	OK	46 degrees C / 114 degrees F
FPC 16 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 16 MQ 3 Chip	OK	49 degrees C / 120 degrees F
FPC 17 Intake	OK	43 degrees C / 109 degrees F
FPC 17 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 17 Exhaust B	OK	55 degrees C / 131 degrees F
FPC 17 LU 0 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 17 LU 1 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 1 Chip	OK	60 degrees C / 140 degrees F
FPC 17 LU 2 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 17 LU 3 TSen	OK	54 degrees C / 129 degrees F
FPC 17 LU 3 Chip	OK	53 degrees C / 127 degrees F
FPC 17 MQ 0 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 0 Chip	OK	50 degrees C / 122 degrees F
FPC 17 MQ 1 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 1 Chip	OK	54 degrees C / 129 degrees F
FPC 17 MQ 2 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 2 Chip	OK	47 degrees C / 116 degrees F
FPC 17 MQ 3 TSen	OK	49 degrees C / 120 degrees F
FPC 17 MQ 3 Chip	OK	51 degrees C / 123 degrees F
FPC 18 Intake	OK	44 degrees C / 111 degrees F
FPC 18 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 18 Exhaust B	OK	57 degrees C / 134 degrees F
FPC 18 LU 0 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 0 Chip	OK	57 degrees C / 134 degrees F
FPC 18 LU 1 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 1 Chip	OK	62 degrees C / 143 degrees F
FPC 18 LU 2 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 2 Chip	OK	53 degrees C / 127 degrees F
FPC 18 LU 3 TSen	OK	56 degrees C / 132 degrees F
FPC 18 LU 3 Chip	OK	55 degrees C / 131 degrees F
FPC 18 MQ 0 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 0 Chip	OK	54 degrees C / 129 degrees F
FPC 18 MQ 1 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 1 Chip	OK	58 degrees C / 136 degrees F
FPC 18 MQ 2 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 2 Chip	OK	50 degrees C / 122 degrees F
FPC 18 MQ 3 TSen	OK	51 degrees C / 123 degrees F
FPC 18 MQ 3 Chip	OK	53 degrees C / 127 degrees F
FPC 19 Intake	OK	48 degrees C / 118 degrees F
FPC 19 Exhaust A	OK	56 degrees C / 132 degrees F
FPC 19 Exhaust B	OK	64 degrees C / 147 degrees F
FPC 19 LU 0 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 0 Chip	OK	64 degrees C / 147 degrees F
FPC 19 LU 1 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 1 Chip	OK	70 degrees C / 158 degrees F

FPC 19 LU 2 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 2 Chip	OK	61 degrees C / 141 degrees F
FPC 19 LU 3 TSen	OK	63 degrees C / 145 degrees F
FPC 19 LU 3 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 0 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 0 Chip	OK	60 degrees C / 140 degrees F
FPC 19 MQ 1 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 1 Chip	OK	62 degrees C / 143 degrees F
FPC 19 MQ 2 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 2 Chip	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 TSen	OK	56 degrees C / 132 degrees F
FPC 19 MQ 3 Chip	OK	57 degrees C / 134 degrees F
ADC 0 Intake	OK	40 degrees C / 104 degrees F
ADC 0 Exhaust	OK	52 degrees C / 125 degrees F
ADC 0 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 0 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 1 Intake	OK	38 degrees C / 100 degrees F
ADC 1 Exhaust	OK	50 degrees C / 122 degrees F
ADC 1 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 1 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 2 Intake	OK	37 degrees C / 98 degrees F
ADC 2 Exhaust	OK	52 degrees C / 125 degrees F
ADC 2 ADC-XF1	OK	53 degrees C / 127 degrees F
ADC 2 ADC-XF0	OK	61 degrees C / 141 degrees F
ADC 3 Intake	OK	40 degrees C / 104 degrees F
ADC 3 Exhaust	OK	51 degrees C / 123 degrees F
ADC 3 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 3 ADC-XF0	OK	64 degrees C / 147 degrees F
ADC 4 Intake	OK	39 degrees C / 102 degrees F
ADC 4 Exhaust	OK	51 degrees C / 123 degrees F
ADC 4 ADC-XF1	OK	60 degrees C / 140 degrees F
ADC 4 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 5 Intake	OK	38 degrees C / 100 degrees F
ADC 5 Exhaust	OK	54 degrees C / 129 degrees F
ADC 5 ADC-XF1	OK	56 degrees C / 132 degrees F
ADC 5 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 6 Intake	OK	39 degrees C / 102 degrees F
ADC 6 Exhaust	OK	52 degrees C / 125 degrees F
ADC 6 ADC-XF1	OK	59 degrees C / 138 degrees F
ADC 6 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 7 Intake	OK	39 degrees C / 102 degrees F
ADC 7 Exhaust	OK	54 degrees C / 129 degrees F
ADC 7 ADC-XF1	OK	62 degrees C / 143 degrees F
ADC 7 ADC-XF0	OK	70 degrees C / 158 degrees F
ADC 8 Intake	OK	39 degrees C / 102 degrees F
ADC 8 Exhaust	OK	52 degrees C / 125 degrees F
ADC 8 ADC-XF1	OK	61 degrees C / 141 degrees F
ADC 8 ADC-XF0	OK	65 degrees C / 149 degrees F
ADC 9 Intake	OK	41 degrees C / 105 degrees F
ADC 9 Exhaust	OK	51 degrees C / 123 degrees F
ADC 9 ADC-XF1	OK	63 degrees C / 145 degrees F
ADC 9 ADC-XF0	OK	63 degrees C / 145 degrees F
ADC 10 Intake	OK	48 degrees C / 118 degrees F
ADC 10 Exhaust	OK	53 degrees C / 127 degrees F
ADC 10 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 10 ADC-XF0	OK	66 degrees C / 150 degrees F
ADC 12 Intake	OK	49 degrees C / 120 degrees F
ADC 12 Exhaust	OK	54 degrees C / 129 degrees F
ADC 12 ADC-XF1	OK	67 degrees C / 152 degrees F
ADC 12 ADC-XF0	OK	67 degrees C / 152 degrees F
ADC 13 Intake	OK	49 degrees C / 120 degrees F

	ADC 13 Exhaust	OK	57 degrees C / 134 degrees F
	ADC 13 ADC-XF1	OK	66 degrees C / 150 degrees F
	ADC 13 ADC-XF0	OK	69 degrees C / 156 degrees F
	ADC 14 Intake	OK	51 degrees C / 123 degrees F
	ADC 14 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 14 ADC-XF1	OK	69 degrees C / 156 degrees F
	ADC 14 ADC-XF0	OK	74 degrees C / 165 degrees F
	ADC 15 Intake	OK	50 degrees C / 122 degrees F
	ADC 15 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 15 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 15 ADC-XF0	OK	69 degrees C / 156 degrees F
	ADC 16 Intake	OK	52 degrees C / 125 degrees F
	ADC 16 Exhaust	OK	58 degrees C / 136 degrees F
	ADC 16 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 16 ADC-XF0	OK	70 degrees C / 158 degrees F
	ADC 17 Intake	OK	52 degrees C / 125 degrees F
	ADC 17 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 17 ADC-XF1	OK	69 degrees C / 156 degrees F
	ADC 17 ADC-XF0	OK	71 degrees C / 159 degrees F
	ADC 18 Intake	OK	53 degrees C / 127 degrees F
	ADC 18 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 18 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 18 ADC-XF0	OK	73 degrees C / 163 degrees F
	ADC 19 Intake	OK	50 degrees C / 122 degrees F
	ADC 19 Exhaust	OK	59 degrees C / 138 degrees F
	ADC 19 ADC-XF1	OK	68 degrees C / 154 degrees F
	ADC 19 ADC-XF0	OK	72 degrees C / 161 degrees F
Fans	Fan Tray 0 Fan 1	OK	7440 RPM
	Fan Tray 0 Fan 2	OK	7200 RPM
	Fan Tray 0 Fan 3	OK	6960 RPM
	Fan Tray 0 Fan 4	OK	7200 RPM
	Fan Tray 0 Fan 5	OK	7080 RPM
	Fan Tray 0 Fan 6	OK	6840 RPM
	Fan Tray 1 Fan 1	OK	6840 RPM
	Fan Tray 1 Fan 2	OK	6960 RPM
	Fan Tray 1 Fan 3	OK	6960 RPM
	Fan Tray 1 Fan 4	OK	7080 RPM
	Fan Tray 1 Fan 5	OK	6960 RPM
	Fan Tray 1 Fan 6	OK	6960 RPM
	Fan Tray 2 Fan 1	OK	8640 RPM
	Fan Tray 2 Fan 2	OK	8640 RPM
	Fan Tray 2 Fan 3	OK	8760 RPM
	Fan Tray 2 Fan 4	OK	8760 RPM
	Fan Tray 2 Fan 5	OK	8640 RPM
	Fan Tray 2 Fan 6	OK	8640 RPM
	Fan Tray 3 Fan 1	OK	8520 RPM
	Fan Tray 3 Fan 2	OK	8520 RPM
	Fan Tray 3 Fan 3	OK	8640 RPM
	Fan Tray 3 Fan 4	OK	8640 RPM
	Fan Tray 3 Fan 5	OK	8520 RPM
	Fan Tray 3 Fan 6	OK	8520 RPM

#### show chassis environment (MX2020 Router with MPC5EQ and MPC6E)

Class	Item	Status	Measurement
Temp	PSM 0	OK	32 degrees C / 89 degrees F
	PSM 1	OK	32 degrees C / 89 degrees F
	PSM 2	OK	32 degrees C / 89 degrees F
	PSM 3	OK	32 degrees C / 89 degrees F
	PSM 4	OK	32 degrees C / 89 degrees F
	PSM 5	OK	33 degrees C / 91 degrees F

PSM 6	OK	32 degrees C / 89 degrees F
PSM 7	OK	32 degrees C / 89 degrees F
PSM 8	OK	32 degrees C / 89 degrees F
PSM 9	Absent	
PSM 10	Absent	
PSM 11	Absent	
PSM 12	OK	33 degrees C / 91 degrees F
PSM 13	OK	33 degrees C / 91 degrees F
PSM 14	OK	34 degrees C / 93 degrees F
PSM 15	OK	34 degrees C / 93 degrees F
PSM 16	OK	33 degrees C / 91 degrees F
PSM 17	OK	33 degrees C / 91 degrees F
PDM 0	OK	
PDM 1	OK	
PDM 2	OK	
PDM 3	OK	
CB 0 IntakeA-Zone0	OK	34 degrees C / 93 degrees F
CB 0 IntakeB-Zone1	OK	26 degrees C / 78 degrees F
CB 0 IntakeC-Zone0	OK	38 degrees C / 100 degrees F
CB 0 ExhaustA-Zone0	OK	34 degrees C / 93 degrees F
CB 0 ExhaustB-Zone1	OK	27 degrees C / 80 degrees F
CB 0 TCBC-Zone0	OK	32 degrees C / 89 degrees F
CB 1 IntakeA-Zone0	OK	24 degrees C / 75 degrees F
CB 1 IntakeB-Zone1	OK	22 degrees C / 71 degrees F
CB 1 IntakeC-Zone0	OK	34 degrees C / 93 degrees F
CB 1 ExhaustA-Zone0	OK	31 degrees C / 87 degrees F
CB 1 ExhaustB-Zone1	OK	24 degrees C / 75 degrees F
CB 1 TCBC-Zone0	OK	27 degrees C / 80 degrees F
SPMB 0 Intake	OK	25 degrees C / 77 degrees F
SPMB 1 Intake	OK	23 degrees C / 73 degrees F
Routing Engine 0	OK	28 degrees C / 82 degrees F
Routing Engine 0 CPU	OK	25 degrees C / 77 degrees F
Routing Engine 1	OK	25 degrees C / 77 degrees F
Routing Engine 1 CPU	OK	24 degrees C / 75 degrees F
SFB 0 Intake-Zone0	OK	45 degrees C / 113 degrees F
SFB 0 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 0 IntakeA-Zone0	OK	32 degrees C / 89 degrees F
SFB 0 IntakeB-Zone1	OK	28 degrees C / 82 degrees F
SFB 0 Exhaust-Zone0	OK	36 degrees C / 96 degrees F
SFB 0 SFB-XF2-Zone1	OK	46 degrees C / 114 degrees F
SFB 0 SFB-XF1-Zone0	OK	48 degrees C / 118 degrees F
SFB 0 SFB-XF0-Zone0	OK	60 degrees C / 140 degrees F
SFB 1 Intake-Zone0	OK	44 degrees C / 111 degrees F
SFB 1 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 1 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 1 IntakeB-Zone1	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone0	OK	37 degrees C / 98 degrees F
SFB 1 SFB-XF2-Zone1	OK	47 degrees C / 116 degrees F
SFB 1 SFB-XF1-Zone0	OK	49 degrees C / 120 degrees F
SFB 1 SFB-XF0-Zone0	OK	56 degrees C / 132 degrees F
SFB 2 Intake-Zone0	OK	41 degrees C / 105 degrees F
SFB 2 Exhaust-Zone1	OK	34 degrees C / 93 degrees F
SFB 2 IntakeA-Zone0	OK	35 degrees C / 95 degrees F
SFB 2 IntakeB-Zone1	OK	28 degrees C / 82 degrees F
SFB 2 Exhaust-Zone0	OK	37 degrees C / 98 degrees F
SFB 2 SFB-XF2-Zone1	OK	47 degrees C / 116 degrees F
SFB 2 SFB-XF1-Zone0	OK	55 degrees C / 131 degrees F
SFB 2 SFB-XF0-Zone0	OK	55 degrees C / 131 degrees F
SFB 3 Intake-Zone0	OK	43 degrees C / 109 degrees F
SFB 3 Exhaust-Zone1	OK	33 degrees C / 91 degrees F
SFB 3 IntakeA-Zone0	OK	35 degrees C / 95 degrees F

SFB 3 IntakeB-Zone1	OK	27 degrees C / 80 degrees F
SFB 3 Exhaust-Zone0	OK	36 degrees C / 96 degrees F
SFB 3 SFB-XF2-Zone1	OK	46 degrees C / 114 degrees F
SFB 3 SFB-XF1-Zone0	OK	46 degrees C / 114 degrees F
SFB 3 SFB-XF0-Zone0	OK	57 degrees C / 134 degrees F
SFB 4 Intake-Zone0	OK	36 degrees C / 96 degrees F
SFB 4 Exhaust-Zone1	OK	32 degrees C / 89 degrees F
SFB 4 IntakeA-Zone0	OK	31 degrees C / 87 degrees F
SFB 4 IntakeB-Zone1	OK	26 degrees C / 78 degrees F
SFB 4 Exhaust-Zone0	OK	32 degrees C / 89 degrees F
SFB 4 SFB-XF2-Zone1	OK	44 degrees C / 111 degrees F
SFB 4 SFB-XF1-Zone0	OK	45 degrees C / 113 degrees F
SFB 4 SFB-XF0-Zone0	OK	52 degrees C / 125 degrees F
SFB 5 Intake-Zone0	OK	31 degrees C / 87 degrees F
SFB 5 Exhaust-Zone1	OK	30 degrees C / 86 degrees F
SFB 5 IntakeA-Zone0	OK	26 degrees C / 78 degrees F
SFB 5 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 5 Exhaust-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF2-Zone1	OK	43 degrees C / 109 degrees F
SFB 5 SFB-XF1-Zone0	OK	47 degrees C / 116 degrees F
SFB 5 SFB-XF0-Zone0	OK	49 degrees C / 120 degrees F
SFB 6 Intake-Zone0	OK	30 degrees C / 86 degrees F
SFB 6 Exhaust-Zone1	OK	29 degrees C / 84 degrees F
SFB 6 IntakeA-Zone0	OK	25 degrees C / 77 degrees F
SFB 6 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 Exhaust-Zone0	OK	29 degrees C / 84 degrees F
SFB 6 SFB-XF2-Zone1	OK	43 degrees C / 109 degrees F
SFB 6 SFB-XF1-Zone0	OK	44 degrees C / 111 degrees F
SFB 6 SFB-XF0-Zone0	OK	45 degrees C / 113 degrees F
SFB 7 Intake-Zone0	OK	31 degrees C / 87 degrees F
SFB 7 Exhaust-Zone1	OK	30 degrees C / 86 degrees F
SFB 7 IntakeA-Zone0	OK	26 degrees C / 78 degrees F
SFB 7 IntakeB-Zone1	OK	24 degrees C / 75 degrees F
SFB 7 Exhaust-Zone0	OK	28 degrees C / 82 degrees F
SFB 7 SFB-XF2-Zone1	OK	50 degrees C / 122 degrees F
SFB 7 SFB-XF1-Zone0	OK	43 degrees C / 109 degrees F
SFB 7 SFB-XF0-Zone0	OK	47 degrees C / 116 degrees F
FPC 0 Intake	OK	31 degrees C / 87 degrees F
FPC 0 Exhaust A	OK	49 degrees C / 120 degrees F
FPC 0 Exhaust B	OK	43 degrees C / 109 degrees F
FPC 0 XL TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL Chip	OK	46 degrees C / 114 degrees F
FPC 0 XL_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL_XR0 Chip	OK	48 degrees C / 118 degrees F
FPC 0 XL_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XL_XR1 Chip	OK	48 degrees C / 118 degrees F
FPC 0 XQ TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ Chip	OK	44 degrees C / 111 degrees F
FPC 0 XQ_XR0 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ_XR0 Chip	OK	57 degrees C / 134 degrees F
FPC 0 XQ_XR1 TSen	OK	42 degrees C / 107 degrees F
FPC 0 XQ_XR1 Chip	OK	55 degrees C / 131 degrees F
FPC 0 XM 0 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XM 0 Chip	OK	62 degrees C / 143 degrees F
FPC 0 XM 1 TSen	OK	48 degrees C / 118 degrees F
FPC 0 XM 1 Chip	OK	44 degrees C / 111 degrees F
FPC 0 PLX PCIe Switch TSe	OK	48 degrees C / 118 degrees F
FPC 0 PLX PCIe Switch Chi	OK	57 degrees C / 134 degrees F
FPC 1 Intake	OK	29 degrees C / 84 degrees F
FPC 1 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 1 Exhaust B	OK	44 degrees C / 111 degrees F

FPC 1 LU 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 0 Chip	OK	45 degrees C / 113 degrees F
FPC 1 LU 1 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 1 Chip	OK	38 degrees C / 100 degrees F
FPC 1 LU 2 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 2 Chip	OK	42 degrees C / 107 degrees F
FPC 1 LU 3 TSen	OK	38 degrees C / 100 degrees F
FPC 1 LU 3 Chip	OK	47 degrees C / 116 degrees F
FPC 1 XM 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 XM 0 Chip	OK	44 degrees C / 111 degrees F
FPC 1 XF 0 TSen	OK	38 degrees C / 100 degrees F
FPC 1 XF 0 Chip	OK	54 degrees C / 129 degrees F
FPC 1 PLX Switch TSen	OK	38 degrees C / 100 degrees F
FPC 1 PLX Switch Chip	OK	41 degrees C / 105 degrees F
FPC 2 Intake	OK	28 degrees C / 82 degrees F
FPC 2 Exhaust A	OK	28 degrees C / 82 degrees F
FPC 2 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 2 LU 0 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 0 Chip	OK	40 degrees C / 104 degrees F
FPC 2 LU 1 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 1 Chip	OK	41 degrees C / 105 degrees F
FPC 2 LU 2 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 2 Chip	OK	34 degrees C / 93 degrees F
FPC 2 LU 3 TSen	OK	40 degrees C / 104 degrees F
FPC 2 LU 3 Chip	OK	38 degrees C / 100 degrees F
FPC 2 XM 0 TSen	OK	40 degrees C / 104 degrees F
FPC 2 XM 0 Chip	OK	47 degrees C / 116 degrees F
FPC 2 XM 1 TSen	OK	40 degrees C / 104 degrees F
FPC 2 XM 1 Chip	OK	42 degrees C / 107 degrees F
FPC 2 PLX Switch TSen	OK	40 degrees C / 104 degrees F
FPC 2 PLX Switch Chip	OK	39 degrees C / 102 degrees F
FPC 3 Intake	OK	27 degrees C / 80 degrees F
FPC 3 Exhaust A	OK	38 degrees C / 100 degrees F
FPC 3 Exhaust B	OK	31 degrees C / 87 degrees F
FPC 3 QX 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 QX 0 Chip	OK	42 degrees C / 107 degrees F
FPC 3 LU 0 TCAM TSen	OK	38 degrees C / 100 degrees F
FPC 3 LU 0 TCAM Chip	OK	43 degrees C / 109 degrees F
FPC 3 LU 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 LU 0 Chip	OK	42 degrees C / 107 degrees F
FPC 3 MQ 0 TSen	OK	38 degrees C / 100 degrees F
FPC 3 MQ 0 Chip	OK	39 degrees C / 102 degrees F
FPC 3 QX 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 QX 1 Chip	OK	36 degrees C / 96 degrees F
FPC 3 LU 1 TCAM TSen	OK	32 degrees C / 89 degrees F
FPC 3 LU 1 TCAM Chip	OK	35 degrees C / 95 degrees F
FPC 3 LU 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 LU 1 Chip	OK	37 degrees C / 98 degrees F
FPC 3 MQ 1 TSen	OK	32 degrees C / 89 degrees F
FPC 3 MQ 1 Chip	OK	36 degrees C / 96 degrees F
FPC 4 Intake	OK	29 degrees C / 84 degrees F
FPC 4 Exhaust A	OK	36 degrees C / 96 degrees F
FPC 4 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 4 XL TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL Chip	OK	42 degrees C / 107 degrees F
FPC 4 XL_XR0 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL_XR0 Chip	OK	45 degrees C / 113 degrees F
FPC 4 XL_XR1 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XL_XR1 Chip	OK	46 degrees C / 114 degrees F
FPC 4 XQ TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ Chip	OK	42 degrees C / 107 degrees F

FPC 4 XQ_XR0 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ_XR0 Chip	OK	54 degrees C / 129 degrees F
FPC 4 XQ_XR1 TSen	OK	39 degrees C / 102 degrees F
FPC 4 XQ_XR1 Chip	OK	53 degrees C / 127 degrees F
FPC 4 XM 0 TSen	OK	45 degrees C / 113 degrees F
FPC 4 XM 0 Chip	OK	59 degrees C / 138 degrees F
FPC 4 XM 1 TSen	OK	45 degrees C / 113 degrees F
FPC 4 XM 1 Chip	OK	41 degrees C / 105 degrees F
FPC 4 PLX PCIe Switch TSe	OK	45 degrees C / 113 degrees F
FPC 4 PLX PCIe Switch Chi	OK	58 degrees C / 136 degrees F
FPC 5 Intake	OK	29 degrees C / 84 degrees F
FPC 5 Exhaust A	OK	33 degrees C / 91 degrees F
FPC 5 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 5 LU 0 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 0 Chip	OK	40 degrees C / 104 degrees F
FPC 5 LU 1 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 1 Chip	OK	45 degrees C / 113 degrees F
FPC 5 LU 2 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 2 Chip	OK	40 degrees C / 104 degrees F
FPC 5 LU 3 TSen	OK	40 degrees C / 104 degrees F
FPC 5 LU 3 Chip	OK	46 degrees C / 114 degrees F
FPC 5 MQ 0 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 0 Chip	OK	33 degrees C / 91 degrees F
FPC 5 MQ 1 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 1 Chip	OK	35 degrees C / 95 degrees F
FPC 5 MQ 2 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 2 Chip	OK	32 degrees C / 89 degrees F
FPC 5 MQ 3 TSen	OK	32 degrees C / 89 degrees F
FPC 5 MQ 3 Chip	OK	32 degrees C / 89 degrees F
FPC 9 Intake	OK	25 degrees C / 77 degrees F
FPC 9 Exhaust A	OK	37 degrees C / 98 degrees F
FPC 9 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 9 XL 0 TSen	OK	40 degrees C / 104 degrees F

...

**show chassis environment (MX2010 Router)**

user@host&gt; show chassis environment

Class	Item	Status	Measurement
Temp	PSM 0	OK	7 degrees C / 44 degrees F
	PSM 1	OK	7 degrees C / 44 degrees F
	PSM 2	OK	7 degrees C / 44 degrees F
	PSM 3	OK	6 degrees C / 42 degrees F
	PSM 4	OK	6 degrees C / 42 degrees F
	PSM 5	OK	6 degrees C / 42 degrees F
	PSM 6	OK	6 degrees C / 42 degrees F
	PSM 7	OK	7 degrees C / 44 degrees F
	PSM 8	OK	7 degrees C / 44 degrees F
	PDM 0	OK	
	PDM 1	Absent	
	CB 0 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 IntakeB-Zone1	OK	7 degrees C / 44 degrees F
	CB 0 IntakeC-Zone0	OK	22 degrees C / 71 degrees F
	CB 0 ExhaustA-Zone0	OK	14 degrees C / 57 degrees F
	CB 0 ExhaustB-Zone1	OK	9 degrees C / 48 degrees F
	CB 0 TCBC-Zone0	OK	11 degrees C / 51 degrees F
	CB 1 IntakeA-Zone0	OK	9 degrees C / 48 degrees F
	CB 1 IntakeB-Zone1	OK	5 degrees C / 41 degrees F
	CB 1 IntakeC-Zone0	OK	20 degrees C / 68 degrees F
	CB 1 ExhaustA-Zone0	OK	12 degrees C / 53 degrees F
	CB 1 ExhaustB-Zone1	OK	7 degrees C / 44 degrees F

CB 1 TCBC-Zone0	OK	10 degrees C / 50 degrees F
SPMB 0 Intake	OK	5 degrees C / 41 degrees F
SPMB 1 Intake	OK	4 degrees C / 39 degrees F
Routing Engine 0	OK	9 degrees C / 48 degrees F
Routing Engine 0 CPU	OK	9 degrees C / 48 degrees F
Routing Engine 1	OK	6 degrees C / 42 degrees F
Routing Engine 1 CPU	OK	6 degrees C / 42 degrees F
SFB 0 Intake-Zone0	OK	26 degrees C / 78 degrees F
SFB 0 Exhaust-Zone1	OK	17 degrees C / 62 degrees F
SFB 0 IntakeA-Zone0	OK	16 degrees C / 60 degrees F
SFB 0 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 0 Exhaust-Zone0	OK	18 degrees C / 64 degrees F
SFB 0 SFB-XF2-Zone1	OK	25 degrees C / 77 degrees F
SFB 0 SFB-XF1-Zone0	OK	23 degrees C / 73 degrees F
SFB 0 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 1 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 Exhaust-Zone1	OK	15 degrees C / 59 degrees F
SFB 1 IntakeA-Zone0	OK	20 degrees C / 68 degrees F
SFB 1 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 1 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 1 SFB-XF2-Zone1	OK	26 degrees C / 78 degrees F
SFB 1 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 1 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 2 Intake-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 2 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 2 IntakeB-Zone1	OK	9 degrees C / 48 degrees F
SFB 2 Exhaust-Zone0	OK	16 degrees C / 60 degrees F
SFB 2 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 2 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 2 SFB-XF0-Zone0	OK	26 degrees C / 78 degrees F
SFB 4 Intake-Zone0	OK	28 degrees C / 82 degrees F
SFB 4 Exhaust-Zone1	OK	16 degrees C / 60 degrees F
SFB 4 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 4 IntakeB-Zone1	OK	11 degrees C / 51 degrees F
SFB 4 Exhaust-Zone0	OK	19 degrees C / 66 degrees F
SFB 4 SFB-XF2-Zone1	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF1-Zone0	OK	27 degrees C / 80 degrees F
SFB 4 SFB-XF0-Zone0	OK	32 degrees C / 89 degrees F
SFB 5 Intake-Zone0	OK	22 degrees C / 71 degrees F
SFB 5 Exhaust-Zone1	OK	14 degrees C / 57 degrees F
SFB 5 IntakeA-Zone0	OK	18 degrees C / 64 degrees F
SFB 5 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 5 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 5 SFB-XF2-Zone1	OK	22 degrees C / 71 degrees F
SFB 5 SFB-XF1-Zone0	OK	29 degrees C / 84 degrees F
SFB 5 SFB-XF0-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Intake-Zone0	OK	27 degrees C / 80 degrees F
SFB 6 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 6 IntakeA-Zone0	OK	19 degrees C / 66 degrees F
SFB 6 IntakeB-Zone1	OK	10 degrees C / 50 degrees F
SFB 6 Exhaust-Zone0	OK	20 degrees C / 68 degrees F
SFB 6 SFB-XF2-Zone1	OK	24 degrees C / 75 degrees F
SFB 6 SFB-XF1-Zone0	OK	32 degrees C / 89 degrees F
SFB 6 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
SFB 7 Intake-Zone0	OK	25 degrees C / 77 degrees F
SFB 7 Exhaust-Zone1	OK	13 degrees C / 55 degrees F
SFB 7 IntakeA-Zone0	OK	14 degrees C / 57 degrees F
SFB 7 IntakeB-Zone1	OK	8 degrees C / 46 degrees F
SFB 7 Exhaust-Zone0	OK	17 degrees C / 62 degrees F
SFB 7 SFB-XF2-Zone1	OK	21 degrees C / 69 degrees F



SFB 7 SFB-XF1-Zone0	OK	21 degrees C / 69 degrees F
SFB 7 SFB-XF0-Zone0	OK	33 degrees C / 91 degrees F
FPC 0 Intake	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust A	OK	13 degrees C / 55 degrees F
FPC 0 Exhaust B	OK	14 degrees C / 57 degrees F
FPC 0 LU 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 0 LU 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 0 LU 2 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 0 LU 3 TSen	OK	28 degrees C / 82 degrees F
FPC 0 LU 3 Chip	OK	23 degrees C / 73 degrees F
FPC 0 XM 0 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 0 Chip	OK	33 degrees C / 91 degrees F
FPC 0 XM 1 TSen	OK	28 degrees C / 82 degrees F
FPC 0 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 0 PLX Switch TSen	OK	28 degrees C / 82 degrees F
FPC 0 PLX Switch Chip	OK	26 degrees C / 78 degrees F
FPC 1 Intake	OK	10 degrees C / 50 degrees F
FPC 1 Exhaust A	OK	24 degrees C / 75 degrees F
FPC 1 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 1 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 0 Chip	OK	31 degrees C / 87 degrees F
FPC 1 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 1 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 2 Chip	OK	25 degrees C / 77 degrees F
FPC 1 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 1 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 1 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 1 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 1 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 1 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 1 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 2 Intake	OK	9 degrees C / 48 degrees F
FPC 2 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 2 Exhaust B	OK	10 degrees C / 50 degrees F
FPC 2 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 0 Chip	OK	25 degrees C / 77 degrees F
FPC 2 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 2 Chip	OK	17 degrees C / 62 degrees F
FPC 2 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 2 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 2 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 2 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 2 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 2 PLX Switch Chip	OK	20 degrees C / 68 degrees F
FPC 3 Intake	OK	12 degrees C / 53 degrees F
FPC 3 Exhaust A	OK	16 degrees C / 60 degrees F
FPC 3 Exhaust B	OK	26 degrees C / 78 degrees F
FPC 3 LU 0 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 3 LU 1 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 3 LU 2 TSen	OK	23 degrees C / 73 degrees F

FPC 3 LU 2 Chip	OK	22 degrees C / 71 degrees F
FPC 3 LU 3 TSen	OK	23 degrees C / 73 degrees F
FPC 3 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 3 MQ 0 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 0 Chip	OK	18 degrees C / 64 degrees F
FPC 3 MQ 1 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 3 MQ 2 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 3 MQ 3 TSen	OK	15 degrees C / 59 degrees F
FPC 3 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 4 Intake	OK	11 degrees C / 51 degrees F
FPC 4 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 4 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 4 LU 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 4 LU 1 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 1 Chip	OK	21 degrees C / 69 degrees F
FPC 4 LU 2 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 4 LU 3 TSen	OK	22 degrees C / 71 degrees F
FPC 4 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 4 XM 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XM 0 Chip	OK	30 degrees C / 86 degrees F
FPC 4 XF 0 TSen	OK	22 degrees C / 71 degrees F
FPC 4 XF 0 Chip	OK	37 degrees C / 98 degrees F
FPC 4 PLX Switch TSen	OK	22 degrees C / 71 degrees F
FPC 4 PLX Switch Chip	OK	23 degrees C / 73 degrees F
FPC 5 Intake	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust A	OK	12 degrees C / 53 degrees F
FPC 5 Exhaust B	OK	12 degrees C / 53 degrees F
FPC 5 LU 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 0 Chip	OK	28 degrees C / 82 degrees F
FPC 5 LU 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 1 Chip	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 5 LU 3 TSen	OK	27 degrees C / 80 degrees F
FPC 5 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 5 XM 0 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 0 Chip	OK	36 degrees C / 96 degrees F
FPC 5 XM 1 TSen	OK	27 degrees C / 80 degrees F
FPC 5 XM 1 Chip	OK	26 degrees C / 78 degrees F
FPC 5 PLX Switch TSen	OK	27 degrees C / 80 degrees F
FPC 5 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 6 Intake	OK	12 degrees C / 53 degrees F
FPC 6 Exhaust A	OK	17 degrees C / 62 degrees F
FPC 6 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 6 LU 0 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 0 Chip	OK	29 degrees C / 84 degrees F
FPC 6 LU 1 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 1 Chip	OK	30 degrees C / 86 degrees F
FPC 6 LU 2 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 2 Chip	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 TSen	OK	24 degrees C / 75 degrees F
FPC 6 LU 3 Chip	OK	22 degrees C / 71 degrees F
FPC 6 MQ 0 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 0 Chip	OK	19 degrees C / 66 degrees F
FPC 6 MQ 1 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 1 Chip	OK	20 degrees C / 68 degrees F
FPC 6 MQ 2 TSen	OK	16 degrees C / 60 degrees F

FPC 6 MQ 2 Chip	OK	17 degrees C / 62 degrees F
FPC 6 MQ 3 TSen	OK	16 degrees C / 60 degrees F
FPC 6 MQ 3 Chip	OK	16 degrees C / 60 degrees F
FPC 7 Intake	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust A	OK	10 degrees C / 50 degrees F
FPC 7 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 7 LU 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 0 Chip	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 1 Chip	OK	29 degrees C / 84 degrees F
FPC 7 LU 2 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 2 Chip	OK	19 degrees C / 66 degrees F
FPC 7 LU 3 TSen	OK	26 degrees C / 78 degrees F
FPC 7 LU 3 Chip	OK	24 degrees C / 75 degrees F
FPC 7 XM 0 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 0 Chip	OK	34 degrees C / 93 degrees F
FPC 7 XM 1 TSen	OK	26 degrees C / 78 degrees F
FPC 7 XM 1 Chip	OK	32 degrees C / 89 degrees F
FPC 7 PLX Switch TSen	OK	26 degrees C / 78 degrees F
FPC 7 PLX Switch Chip	OK	22 degrees C / 71 degrees F
FPC 8 Intake	OK	10 degrees C / 50 degrees F
FPC 8 Exhaust A	OK	22 degrees C / 71 degrees F
FPC 8 Exhaust B	OK	28 degrees C / 82 degrees F
FPC 8 LU 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 0 Chip	OK	33 degrees C / 91 degrees F
FPC 8 LU 1 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 1 Chip	OK	23 degrees C / 73 degrees F
FPC 8 LU 2 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 2 Chip	OK	26 degrees C / 78 degrees F
FPC 8 LU 3 TSen	OK	20 degrees C / 68 degrees F
FPC 8 LU 3 Chip	OK	33 degrees C / 91 degrees F
FPC 8 XM 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XM 0 Chip	OK	29 degrees C / 84 degrees F
FPC 8 XF 0 TSen	OK	20 degrees C / 68 degrees F
FPC 8 XF 0 Chip	OK	38 degrees C / 100 degrees F
FPC 8 PLX Switch TSen	OK	20 degrees C / 68 degrees F
FPC 8 PLX Switch Chip	OK	24 degrees C / 75 degrees F
FPC 9 Intake	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust A	OK	11 degrees C / 51 degrees F
FPC 9 Exhaust B	OK	11 degrees C / 51 degrees F
FPC 9 LU 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 0 Chip	OK	24 degrees C / 75 degrees F
FPC 9 LU 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 1 Chip	OK	26 degrees C / 78 degrees F
FPC 9 LU 2 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 2 Chip	OK	16 degrees C / 60 degrees F
FPC 9 LU 3 TSen	OK	25 degrees C / 77 degrees F
FPC 9 LU 3 Chip	OK	21 degrees C / 69 degrees F
FPC 9 XM 0 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 0 Chip	OK	32 degrees C / 89 degrees F
FPC 9 XM 1 TSen	OK	25 degrees C / 77 degrees F
FPC 9 XM 1 Chip	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch TSen	OK	25 degrees C / 77 degrees F
FPC 9 PLX Switch Chip	OK	21 degrees C / 69 degrees F
ADC 0 Intake	OK	12 degrees C / 53 degrees F
ADC 0 Exhaust	OK	20 degrees C / 68 degrees F
ADC 0 ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 0 ADC-XF0	OK	32 degrees C / 89 degrees F
ADC 1 Intake	OK	11 degrees C / 51 degrees F
ADC 1 Exhaust	OK	21 degrees C / 69 degrees F
ADC 1 ADC-XF1	OK	24 degrees C / 75 degrees F

ADC 1	ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 2	Intake	OK	14 degrees C / 57 degrees F
ADC 2	Exhaust	OK	21 degrees C / 69 degrees F
ADC 2	ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 2	ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 3	Intake	OK	13 degrees C / 55 degrees F
ADC 3	Exhaust	OK	19 degrees C / 66 degrees F
ADC 3	ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 3	ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 4	Intake	OK	9 degrees C / 48 degrees F
ADC 4	Exhaust	OK	22 degrees C / 71 degrees F
ADC 4	ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 4	ADC-XF0	OK	35 degrees C / 95 degrees F
ADC 5	Intake	OK	12 degrees C / 53 degrees F
ADC 5	Exhaust	OK	22 degrees C / 71 degrees F
ADC 5	ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 5	ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 6	Intake	OK	11 degrees C / 51 degrees F
ADC 6	Exhaust	OK	21 degrees C / 69 degrees F
ADC 6	ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 6	ADC-XF0	OK	35 degrees C / 95 degrees F
ADC 7	Intake	OK	14 degrees C / 57 degrees F
ADC 7	Exhaust	OK	22 degrees C / 71 degrees F
ADC 7	ADC-XF1	OK	26 degrees C / 78 degrees F
ADC 7	ADC-XF0	OK	34 degrees C / 93 degrees F
ADC 8	Intake	OK	14 degrees C / 57 degrees F
ADC 8	Exhaust	OK	21 degrees C / 69 degrees F
ADC 8	ADC-XF1	OK	24 degrees C / 75 degrees F
ADC 8	ADC-XF0	OK	31 degrees C / 87 degrees F
ADC 9	Intake	OK	10 degrees C / 50 degrees F
ADC 9	Exhaust	OK	22 degrees C / 71 degrees F
ADC 9	ADC-XF1	OK	28 degrees C / 82 degrees F
ADC 9	ADC-XF0	OK	36 degrees C / 96 degrees F
Fans	Fan Tray 0 Fan 1	OK	3480 RPM
	Fan Tray 0 Fan 2	OK	3480 RPM
	Fan Tray 0 Fan 3	OK	3480 RPM
	Fan Tray 0 Fan 4	OK	3360 RPM
	Fan Tray 0 Fan 5	OK	3360 RPM
	Fan Tray 0 Fan 6	OK	3480 RPM
	Fan Tray 1 Fan 1	OK	3360 RPM
	Fan Tray 1 Fan 2	OK	3360 RPM
	Fan Tray 1 Fan 3	OK	3360 RPM
	Fan Tray 1 Fan 4	OK	3480 RPM
	Fan Tray 1 Fan 5	OK	3480 RPM
	Fan Tray 1 Fan 6	OK	3480 RPM
	Fan Tray 2 Fan 1	OK	3360 RPM
	Fan Tray 2 Fan 2	OK	3360 RPM
	Fan Tray 2 Fan 3	OK	3480 RPM
	Fan Tray 2 Fan 4	OK	3480 RPM
	Fan Tray 2 Fan 5	OK	3360 RPM
	Fan Tray 2 Fan 6	OK	3480 RPM
	Fan Tray 3 Fan 1	OK	3360 RPM
	Fan Tray 3 Fan 2	OK	3360 RPM
	Fan Tray 3 Fan 3	OK	3480 RPM
	Fan Tray 3 Fan 4	OK	3480 RPM
	Fan Tray 3 Fan 5	OK	3480 RPM
	Fan Tray 3 Fan 6	OK	3360 RPM

**show chassis environment (T320 Router)**

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Power	PEM 0	OK	
	PEM 1	Absent	
Temp	SCG 0	OK	28 degrees C / 82 degrees F
	SCG 1	OK	28 degrees C / 82 degrees F
	Routing Engine 0	OK	31 degrees C / 87 degrees F
	Routing Engine 1	OK	30 degrees C / 86 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	33 degrees C / 91 degrees F
	SIB 1	OK	33 degrees C / 91 degrees F
	SIB 2	OK	34 degrees C / 93 degrees F
	FPC 0 Top	OK	38 degrees C / 100 degrees F
	FPC 0 Bottom	OK	32 degrees C / 89 degrees F
	FPC 1 Top	OK	38 degrees C / 100 degrees F
	FPC 1 Bottom	OK	33 degrees C / 91 degrees F
	FPC 2 Top	OK	36 degrees C / 96 degrees F
	FPC 2 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	26 degrees C / 78 degrees F
	FPM Display	OK	29 degrees C / 84 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Middle fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

### show chassis environment (T640 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	22 degrees C / 71 degrees F
	SCG 0	OK	30 degrees C / 86 degrees F
	SCG 1	OK	30 degrees C / 86 degrees F
	Routing Engine 0	Present	
	Routing Engine 1	OK	27 degrees C / 80 degrees F
	CB 0	Present	
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	Absent	
	SIB 1	Absent	
	SIB 2	Absent	
	SIB 3	Absent	
	SIB 4	Absent	
	FPC 4 Top	Testing	
	FPC 4 Bottom	Testing	

	FPC 5 Top	Testing	
	FPC 5 Bottom	Testing	
	FPC 6 Top	Testing	
	FPC 6 Bottom	Testing	
	FPM GBUS	OK	23 degrees C / 73 degrees F
	FPM Display	Absent	
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Fourth Blower from top	OK	Spinning at normal speed
	Bottom Blower	OK	Spinning at normal speed
	Middle Blower	OK	Spinning at normal speed
	Top Blower	OK	Spinning at normal speed
Misc	Second Blower from top	OK	Spinning at normal speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

### show chassis environment (T4000 Router)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	33 degrees C / 91 degrees F
	PEM 1	Absent	
	SCG 0	OK	33 degrees C / 91 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU	OK	50 degrees C / 122 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	46 degrees C / 114 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	33 degrees C / 91 degrees F
	SIB 0	OK	42 degrees C / 107 degrees F
	SIB 1	OK	42 degrees C / 107 degrees F
	SIB 2	OK	42 degrees C / 107 degrees F
	SIB 3	OK	43 degrees C / 109 degrees F
	SIB 4	OK	45 degrees C / 113 degrees F
	FPC 0 Fan Intake	OK	34 degrees C / 93 degrees F
	FPC 0 Fan Exhaust	OK	48 degrees C / 118 degrees F
	FPC 0 PMB	OK	47 degrees C / 116 degrees F
	FPC 0 LMB0	OK	50 degrees C / 122 degrees F
	FPC 0 LMB1	OK	41 degrees C / 105 degrees F
	FPC 0 LMB2	OK	35 degrees C / 95 degrees F
	FPC 0 PFE1 LU2	OK	46 degrees C / 114 degrees F
	FPC 0 PFE1 LU0	OK	41 degrees C / 105 degrees F
	FPC 0 PFE0 LU0	OK	57 degrees C / 134 degrees F
	FPC 0 XF1	OK	46 degrees C / 114 degrees F
	FPC 0 XF0	OK	52 degrees C / 125 degrees F
	FPC 0 XM1	OK	41 degrees C / 105 degrees F
	FPC 0 XM0	OK	50 degrees C / 122 degrees F
	FPC 0 PFE0 LU1	OK	56 degrees C / 132 degrees F

	FPC 0 PFE0 LU2	OK	45 degrees C / 113 degrees F
	FPC 0 PFE1 LU1	OK	37 degrees C / 98 degrees F
	FPC 3 Fan Intake	OK	36 degrees C / 96 degrees F
	FPC 3 Fan Exhaust	OK	51 degrees C / 123 degrees F
	FPC 3 PMB	OK	43 degrees C / 109 degrees F
	FPC 3 LMB0	OK	57 degrees C / 134 degrees F
	FPC 3 LMB1	OK	54 degrees C / 129 degrees F
	FPC 3 LMB2	OK	38 degrees C / 100 degrees F
	FPC 3 PFE1 LU2	OK	63 degrees C / 145 degrees F
	FPC 3 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 3 PFE0 LU0	OK	69 degrees C / 156 degrees F
	FPC 3 XF1	OK	62 degrees C / 143 degrees F
	FPC 3 XF0	OK	63 degrees C / 145 degrees F
	FPC 3 XM1	OK	43 degrees C / 109 degrees F
	FPC 3 XM0	OK	67 degrees C / 152 degrees F
	FPC 3 PFE0 LU1	OK	63 degrees C / 145 degrees F
	FPC 3 PFE0 LU2	OK	66 degrees C / 150 degrees F
	FPC 3 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPC 5 Top	OK	39 degrees C / 102 degrees F
	FPC 5 Bottom	OK	38 degrees C / 100 degrees F
	FPC 6 Fan Intake	OK	33 degrees C / 91 degrees F
	FPC 6 Fan Exhaust	OK	49 degrees C / 120 degrees F
	FPC 6 PMB	OK	40 degrees C / 104 degrees F
	FPC 6 LMB0	OK	60 degrees C / 140 degrees F
	FPC 6 LMB1	OK	58 degrees C / 136 degrees F
	FPC 6 LMB2	OK	40 degrees C / 104 degrees F
	FPC 6 PFE1 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 6 PFE0 LU0	OK	71 degrees C / 159 degrees F
	FPC 6 XF1	OK	58 degrees C / 136 degrees F
	FPC 6 XF0	OK	65 degrees C / 149 degrees F
	FPC 6 XM1	OK	39 degrees C / 102 degrees F
	FPC 6 XM0	OK	66 degrees C / 150 degrees F
	FPC 6 PFE0 LU1	OK	69 degrees C / 156 degrees F
	FPC 6 PFE0 LU2	OK	69 degrees C / 156 degrees F
	FPC 6 PFE1 LU1	OK	42 degrees C / 107 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	27 degrees C / 80 degrees F
Fans	Top Left Front fan	OK	Spinning at high speed
	Top Left Middle fan	OK	Spinning at high speed
	Top Left Rear fan	OK	Spinning at high speed
	Top Right Front fan	OK	Spinning at high speed
	Top Right Middle fan	OK	Spinning at high speed
	Top Right Rear fan	OK	Spinning at high speed
	Bottom Left Front fan	OK	Spinning at high speed
	Bottom Left Middle fan	OK	Spinning at high speed
	Bottom Left Rear fan	OK	Spinning at high speed
	Bottom Right Front fan	OK	Spinning at high speed
	Bottom Right Middle fan	OK	Spinning at high speed
	Bottom Right Rear fan	OK	Spinning at high speed
	Rear Tray Top fan	OK	Spinning at high speed
	Rear Tray Second fan	OK	Spinning at high speed
	Rear Tray Third fan	OK	Spinning at high speed
	Rear Tray Fourth fan	OK	Spinning at high speed
Misc	Rear Tray Fifth fan	OK	Spinning at high speed
	Rear Tray Sixth fan	OK	Spinning at high speed
	Rear Tray Seventh fan	OK	Spinning at high speed
	Rear Tray Bottom fan	OK	Spinning at high speed
	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

## show chassis environment (TX Matrix Router)

```
user@host> show chassis environment
scc-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	Absent	
	PEM 1	OK	29 degrees C / 84 degrees F
	Routing Engine 0	OK	34 degrees C / 93 degrees F
	Routing Engine 1	OK	34 degrees C / 93 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	44 degrees C / 111 degrees F
	SIB 0 (B)	OK	44 degrees C / 111 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	32 degrees C / 89 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP 0	OK	
	CIP 1	OK	
	SPMB 0	OK	
	SPMB 1	OK	

```
1cc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Absent	
	SCG 0	OK	35 degrees C / 95 degrees F
	SCG 1	Absent	
	Routing Engine 0	OK	39 degrees C / 102 degrees F
	Routing Engine 1	OK	36 degrees C / 96 degrees F
	CB 0	OK	32 degrees C / 89 degrees F
	CB 1	OK	32 degrees C / 89 degrees F
	SIB 0	OK	40 degrees C / 104 degrees F
	SIB 0 (B)	OK	51 degrees C / 123 degrees F
	FPC 0 Top	OK	45 degrees C / 113 degrees F
	FPC 0 Bottom	OK	31 degrees C / 87 degrees F
	FPC 1 Top	OK	34 degrees C / 93 degrees F
	FPC 1 Bottom	OK	31 degrees C / 87 degrees F
	FPM GBUS	OK	30 degrees C / 86 degrees F
	FPM Display	OK	34 degrees C / 93 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed



```

Top Left Middle fan    OK      Spinning at normal speed
Top Left Rear fan      OK      Spinning at normal speed
Top Right Front fan    OK      Spinning at normal speed
Top Right Middle fan   OK      Spinning at normal speed
Top Right Rear fan     OK      Spinning at normal speed
Bottom Left Front fan  OK      Spinning at normal speed
Bottom Left Middle fan OK      Spinning at normal speed
Bottom Left Rear fan   OK      Spinning at normal speed
Bottom Right Front fan OK      Spinning at normal speed
Bottom Right Middle fan OK     Spinning at normal speed
Bottom Right Rear fan  OK      Spinning at normal speed
Rear Tray Top fan      OK      Spinning at normal speed
Rear Tray Second fan   OK      Spinning at normal speed
Rear Tray Third fan    OK      Spinning at normal speed
Rear Tray Fourth fan   OK      Spinning at normal speed
Rear Tray Fifth fan    OK      Spinning at normal speed
Rear Tray Sixth fan    OK      Spinning at normal speed
Rear Tray Seventh fan  OK      Spinning at normal speed
Rear Tray Bottom fan   OK      Spinning at normal speed
Misc CIP               OK
SPMB 0                 OK
SPMB 1                 OK

```

lcc2-re0:

```

-----
Class Item              Status      Measurement
Temp PEM 0              OK          29 degrees C / 84 degrees F
      PEM 1              Absent
      SCG 0              OK          32 degrees C / 89 degrees F
      SCG 1              Absent
      Routing Engine 0    OK          31 degrees C / 87 degrees F
      Routing Engine 1    OK          32 degrees C / 89 degrees F
      CB 0                OK          30 degrees C / 86 degrees F
      SIB 0               OK          38 degrees C / 100 degrees F
      SIB 0 (B)           OK          49 degrees C / 120 degrees F
      FPC 0 Top           OK          45 degrees C / 113 degrees F
      FPC 0 Bottom        OK          33 degrees C / 91 degrees F
      FPC 1 Top           OK          37 degrees C / 98 degrees F
      FPC 1 Bottom        OK          33 degrees C / 91 degrees F
      FPM GBUS            OK          30 degrees C / 86 degrees F
      FPM Display         OK          34 degrees C / 93 degrees F
Fans  Top Left Front fan  OK          Spinning at normal speed
      Top Left Middle fan OK          Spinning at normal speed
...

```

#### show chassis environment (T1600 Router)

```

user@host> show chassis environment
Class Item              Status      Measurement
Temp PEM 0              OK          27 degrees C / 80 degrees F
      PEM 1              Absent
      SCG 0              OK          31 degrees C / 87 degrees F
      SCG 1              OK          35 degrees C / 95 degrees F
      Routing Engine 0    OK          30 degrees C / 86 degrees F
      Routing Engine 1    OK          30 degrees C / 86 degrees F
      CB 0                OK          31 degrees C / 87 degrees F
      CB 1                OK          31 degrees C / 87 degrees F
      SIB 0               OK          41 degrees C / 105 degrees F
      SIB 0 (B)           OK          34 degrees C / 93 degrees F
      SIB 1               OK          0 degrees C / 32 degrees F
      SIB 1 (B)           OK          0 degrees C / 32 degrees F

```

	SIB 2	OK	0 degrees C / 32 degrees F
	SIB 2 (B)	OK	0 degrees C / 32 degrees F
	SIB 3	OK	0 degrees C / 32 degrees F
	SIB 3 (B)	OK	0 degrees C / 32 degrees F
	SIB 4	OK	0 degrees C / 32 degrees F
	SIB 4 (B)	OK	0 degrees C / 32 degrees F
	FPC 0 Top	OK	49 degrees C / 120 degrees F
	FPC 0 Bottom	OK	50 degrees C / 122 degrees F
	FPC 1 Top	OK	48 degrees C / 118 degrees F
	FPC 1 Bottom	OK	49 degrees C / 120 degrees F
	FPM GBUS	OK	27 degrees C / 80 degrees F
	FPM Display	OK	30 degrees C / 86 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray Top fan	OK	Spinning at normal speed
	Rear Tray Second fan	OK	Spinning at normal speed
	Rear Tray Third fan	OK	Spinning at normal speed
	Rear Tray Fourth fan	OK	Spinning at normal speed
	Rear Tray Fifth fan	OK	Spinning at normal speed
	Rear Tray Sixth fan	OK	Spinning at normal speed
	Rear Tray Seventh fan	OK	Spinning at normal speed
	Rear Tray Bottom fan	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

**show chassis environment (TX Matrix Plus Router)**

```
user@host> show chassis environment
sfc0-re0:
```

Class	Item	Status	Measurement
Temp	PEM 0	OK	28 degrees C / 82 degrees F
	PEM 1	Absent	
	Routing Engine 0	OK	27 degrees C / 80 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	CB 0 Intake	OK	26 degrees C / 78 degrees F
	CB 0 Exhaust A	OK	25 degrees C / 77 degrees F
	CB 0 Exhaust B	OK	25 degrees C / 77 degrees F
	CB 1 Intake	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust A	OK	26 degrees C / 78 degrees F
	CB 1 Exhaust B	OK	26 degrees C / 78 degrees F
	SIB F13 0	OK	47 degrees C / 116 degrees F
	SIB F13 0 (B)	OK	48 degrees C / 118 degrees F
	SIB F13 1	OK	38 degrees C / 100 degrees F
	SIB F13 1 (B)	OK	37 degrees C / 98 degrees F
	SIB F2S 0/0	OK	27 degrees C / 80 degrees F
	SIB F2S 0/2	OK	28 degrees C / 82 degrees F
	SIB F2S 0/4	OK	27 degrees C / 80 degrees F
	SIB F2S 0/6	OK	28 degrees C / 82 degrees F
	SIB F2S 1/0	OK	26 degrees C / 78 degrees F

	SIB F2S 1/2	OK	26 degrees C / 78 degrees F
	SIB F2S 1/4	OK	26 degrees C / 78 degrees F
	SIB F2S 1/6	OK	26 degrees C / 78 degrees F
	SIB F2S 2/0	OK	25 degrees C / 77 degrees F
	SIB F2S 2/2	OK	25 degrees C / 77 degrees F
	SIB F2S 2/4	OK	23 degrees C / 73 degrees F
	CIP 0 Intake	OK	23 degrees C / 73 degrees F
	CIP 0 Exhaust A	OK	24 degrees C / 75 degrees F
	CIP 0 Exhaust B	OK	24 degrees C / 75 degrees F
	CIP 1 Intake	OK	24 degrees C / 75 degrees F
	CIP 1 Exhaust A	OK	25 degrees C / 77 degrees F
	CIP 1 Exhaust B	OK	25 degrees C / 77 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed
	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed
	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed
	Fan Tray 3 Fan 8	OK	Spinning at normal speed
	Fan Tray 3 Fan 9	OK	Spinning at normal speed
	Fan Tray 4 Fan 1	OK	Spinning at normal speed
	Fan Tray 4 Fan 2	OK	Spinning at normal speed
	Fan Tray 4 Fan 3	OK	Spinning at normal speed
	Fan Tray 4 Fan 4	OK	Spinning at normal speed
	Fan Tray 4 Fan 5	OK	Spinning at normal speed
	Fan Tray 4 Fan 6	OK	Spinning at normal speed
	Fan Tray 4 Fan 7	OK	Spinning at normal speed
	Fan Tray 4 Fan 8	OK	Spinning at normal speed
	Fan Tray 4 Fan 9	OK	Spinning at normal speed
	Fan Tray 5 Fan 1	OK	Spinning at normal speed
	Fan Tray 5 Fan 2	OK	Spinning at normal speed
	Fan Tray 5 Fan 3	OK	Spinning at normal speed
	Fan Tray 5 Fan 4	OK	Spinning at normal speed
	Fan Tray 5 Fan 5	OK	Spinning at normal speed
	Fan Tray 5 Fan 6	OK	Spinning at normal speed
	Fan Tray 5 Fan 7	OK	Spinning at normal speed
	Fan Tray 5 Fan 8	OK	Spinning at normal speed
	Fan Tray 5 Fan 9	OK	Spinning at normal speed
Misc	SPMB 0	OK	

```

SPMB 1                                OK

lcc0-re0:
-----
Class Item                               Status Measurement
Temp PEM 0                             OK          27 degrees C / 80 degrees F
    PEM 1                             Absent
    SCG 0                             OK          31 degrees C / 87 degrees F
    SCG 1                             OK          35 degrees C / 95 degrees F
    Routing Engine 0                  OK          30 degrees C / 86 degrees F
    Routing Engine 1                  OK          30 degrees C / 86 degrees F
    CB 0                             OK          31 degrees C / 87 degrees F
    CB 1                             OK          31 degrees C / 87 degrees F
    SIB 0                             OK          41 degrees C / 105 degrees F
    SIB 0 (B)                         OK          34 degrees C / 93 degrees F
    SIB 1                             OK          0 degrees C / 32 degrees F
    SIB 1 (B)                         OK          0 degrees C / 32 degrees F
    SIB 2                             OK          0 degrees C / 32 degrees F
    SIB 2 (B)                         OK          0 degrees C / 32 degrees F
    SIB 3                             OK          0 degrees C / 32 degrees F
    SIB 3 (B)                         OK          0 degrees C / 32 degrees F
    SIB 4                             OK          0 degrees C / 32 degrees F
    SIB 4 (B)                         OK          0 degrees C / 32 degrees F
    FPC 0 Top                         OK          49 degrees C / 120 degrees F
    FPC 0 Bottom                     OK          50 degrees C / 122 degrees F
    FPC 1 Top                         OK          48 degrees C / 118 degrees F
    FPC 1 Bottom                     OK          49 degrees C / 120 degrees F
    FPM GBUS                         OK          27 degrees C / 80 degrees F
    FPM Display                       OK          30 degrees C / 86 degrees F
Fans Top Left Front fan              OK          Spinning at normal speed
    Top Left Middle fan              OK          Spinning at normal speed
    Top Left Rear fan                OK          Spinning at normal speed
    Top Right Front fan              OK          Spinning at normal speed
    Top Right Middle fan             OK          Spinning at normal speed
    Top Right Rear fan               OK          Spinning at normal speed
    Bottom Left Front fan            OK          Spinning at normal speed
    Bottom Left Middle fan           OK          Spinning at normal speed
    Bottom Left Rear fan             OK          Spinning at normal speed
    Bottom Right Front fan           OK          Spinning at normal speed
    Bottom Right Middle fan          OK          Spinning at normal speed
    Bottom Right Rear fan            OK          Spinning at normal speed
    Rear Tray Top fan                OK          Spinning at normal speed
    Rear Tray Second fan             OK          Spinning at normal speed
    Rear Tray Third fan              OK          Spinning at normal speed
    Rear Tray Fourth fan             OK          Spinning at normal speed
    Rear Tray Fifth fan              OK          Spinning at normal speed
    Rear Tray Sixth fan              OK          Spinning at normal speed
    Rear Tray Seventh fan            OK          Spinning at normal speed
    Rear Tray Bottom fan             OK          Spinning at normal speed
Misc CIP                             OK
    SPMB 0                           OK
    SPMB 1                           OK

```

#### show chassis environment (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis environment
sfc0-re0:
-----
Class Item                               Status Measurement
Temp PEM 0                             Check       30 degrees C / 86 degrees F
    PEM 1                             OK          33 degrees C / 91 degrees F

```

	Routing Engine 0	OK	28 degrees C / 82 degrees F
	Routing Engine 0 CPU	OK	42 degrees C / 107 degrees F
	Routing Engine 1	OK	29 degrees C / 84 degrees F
	Routing Engine 1 CPU	OK	44 degrees C / 111 degrees F
	CB 0 Intake	OK	30 degrees C / 86 degrees F
	CB 0 Exhaust A	OK	28 degrees C / 82 degrees F
	CB 0 Exhaust B	OK	30 degrees C / 86 degrees F
	CB 1 Intake	OK	31 degrees C / 87 degrees F
	CB 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CB 1 Exhaust B	OK	31 degrees C / 87 degrees F
	SIB F13 0 Board	OK	44 degrees C / 111 degrees F
	SIB F13 0 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F13 3 Board	OK	45 degrees C / 113 degrees F
	SIB F13 3 XF Junction	OK	60 degrees C / 140 degrees F
	SIB F13 6 Board	OK	47 degrees C / 116 degrees F
	SIB F13 6 XF Junction	OK	62 degrees C / 143 degrees F
	SIB F2S 0/0 Board	OK	32 degrees C / 89 degrees F
	SIB F2S 0/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/2 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/2 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 0/4 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 0/6 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 0/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/0 Board	OK	31 degrees C / 87 degrees F
	SIB F2S 1/0 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 1/2 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 1/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 1/4 XF Junction	OK	35 degrees C / 95 degrees F
	SIB F2S 1/6 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 1/6 XF Junction	OK	41 degrees C / 105 degrees F
	SIB F2S 2/0 Board	OK	30 degrees C / 86 degrees F
	SIB F2S 2/0 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/2 Board	OK	28 degrees C / 82 degrees F
	SIB F2S 2/2 XF Junction	OK	39 degrees C / 102 degrees F
	SIB F2S 2/4 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/4 XF Junction	OK	42 degrees C / 107 degrees F
	SIB F2S 2/6 Board	OK	29 degrees C / 84 degrees F
	SIB F2S 2/6 XF Junction	OK	41 degrees C / 105 degrees F
	CIP 0 Intake	OK	25 degrees C / 77 degrees F
	CIP 0 Exhaust A	OK	26 degrees C / 78 degrees F
	CIP 0 Exhaust B	OK	26 degrees C / 78 degrees F
	CIP 1 Intake	OK	26 degrees C / 78 degrees F
	CIP 1 Exhaust A	OK	27 degrees C / 80 degrees F
	CIP 1 Exhaust B	OK	27 degrees C / 80 degrees F
Fans	Fan Tray 0 Fan 1	OK	Spinning at normal speed
	Fan Tray 0 Fan 2	OK	Spinning at normal speed
	Fan Tray 0 Fan 3	OK	Spinning at normal speed
	Fan Tray 0 Fan 4	OK	Spinning at normal speed
	Fan Tray 0 Fan 5	OK	Spinning at normal speed
	Fan Tray 0 Fan 6	OK	Spinning at normal speed
	Fan Tray 1 Fan 1	OK	Spinning at normal speed
	Fan Tray 1 Fan 2	OK	Spinning at normal speed
	Fan Tray 1 Fan 3	OK	Spinning at normal speed
	Fan Tray 1 Fan 4	OK	Spinning at normal speed
	Fan Tray 1 Fan 5	OK	Spinning at normal speed
	Fan Tray 1 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 1	OK	Spinning at normal speed
	Fan Tray 2 Fan 2	OK	Spinning at normal speed
	Fan Tray 2 Fan 3	OK	Spinning at normal speed

	Fan Tray 2 Fan 4	OK	Spinning at normal speed
	Fan Tray 2 Fan 5	OK	Spinning at normal speed
	Fan Tray 2 Fan 6	OK	Spinning at normal speed
	Fan Tray 2 Fan 7	OK	Spinning at normal speed
	Fan Tray 2 Fan 8	OK	Spinning at normal speed
	Fan Tray 2 Fan 9	OK	Spinning at normal speed
	Fan Tray 3 Fan 1	OK	Spinning at normal speed
	Fan Tray 3 Fan 2	OK	Spinning at normal speed
	Fan Tray 3 Fan 3	OK	Spinning at normal speed
	Fan Tray 3 Fan 4	OK	Spinning at normal speed
	Fan Tray 3 Fan 5	OK	Spinning at normal speed
	Fan Tray 3 Fan 6	OK	Spinning at normal speed
	Fan Tray 3 Fan 7	OK	Spinning at normal speed
	Fan Tray 3 Fan 8	OK	Spinning at normal speed
	Fan Tray 3 Fan 9	OK	Spinning at normal speed
	Fan Tray 4 Fan 1	OK	Spinning at normal speed
	Fan Tray 4 Fan 2	OK	Spinning at normal speed
	Fan Tray 4 Fan 3	OK	Spinning at normal speed
	Fan Tray 4 Fan 4	OK	Spinning at normal speed
	Fan Tray 4 Fan 5	OK	Spinning at normal speed
	Fan Tray 4 Fan 6	OK	Spinning at normal speed
	Fan Tray 4 Fan 7	OK	Spinning at normal speed
	Fan Tray 4 Fan 8	OK	Spinning at normal speed
	Fan Tray 4 Fan 9	OK	Spinning at normal speed
	Fan Tray 5 Fan 1	OK	Spinning at normal speed
	Fan Tray 5 Fan 2	OK	Spinning at normal speed
	Fan Tray 5 Fan 3	OK	Spinning at normal speed
	Fan Tray 5 Fan 4	OK	Spinning at normal speed
	Fan Tray 5 Fan 5	OK	Spinning at normal speed
	Fan Tray 5 Fan 6	OK	Spinning at normal speed
	Fan Tray 5 Fan 7	OK	Spinning at normal speed
	Fan Tray 5 Fan 8	OK	Spinning at normal speed
	Fan Tray 5 Fan 9	Check	
Misc	SPMB 0	OK	
	SPMB 1	OK	

## 1cc0-re0:

Class	Item	Status	Measurement
Temp	PEM 0	OK	29 degrees C / 84 degrees F
	PEM 1	Check	29 degrees C / 84 degrees F
	SCG 0	OK	32 degrees C / 89 degrees F
	SCG 1	OK	33 degrees C / 91 degrees F
	Routing Engine 0	OK	32 degrees C / 89 degrees F
	Routing Engine 0 CPU	OK	51 degrees C / 123 degrees F
	Routing Engine 1	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU	OK	49 degrees C / 120 degrees F
	CB 0	OK	34 degrees C / 93 degrees F
	CB 1	OK	34 degrees C / 93 degrees F
	SIB 0	OK	39 degrees C / 102 degrees F
	SIB 0 (B)	Absent	
	SIB 1	OK	39 degrees C / 102 degrees F
	SIB 1 (B)	Absent	
	SIB 2	OK	39 degrees C / 102 degrees F
	SIB 2 (B)	Absent	
	FPC 4 Top	OK	43 degrees C / 109 degrees F
	FPC 4 Bottom	OK	43 degrees C / 109 degrees F
	FPC 7 Fan Intake	OK	35 degrees C / 95 degrees F
	FPC 7 Fan Exhaust	OK	50 degrees C / 122 degrees F
	FPC 7 PMB	OK	50 degrees C / 122 degrees F
	FPC 7 LMB0	OK	55 degrees C / 131 degrees F

	FPC 7 LMB1	OK	49 degrees C / 120 degrees F
	FPC 7 LMB2	OK	39 degrees C / 102 degrees F
	FPC 7 PFE1 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU0	OK	45 degrees C / 113 degrees F
	FPC 7 PFE0 LU0	OK	62 degrees C / 143 degrees F
	FPC 7 XF1	OK	52 degrees C / 125 degrees F
	FPC 7 XF0	OK	61 degrees C / 141 degrees F
	FPC 7 XM1	OK	39 degrees C / 102 degrees F
	FPC 7 XM0	OK	56 degrees C / 132 degrees F
	FPC 7 PFE0 LU1	OK	60 degrees C / 140 degrees F
	FPC 7 PFE0 LU2	OK	55 degrees C / 131 degrees F
	FPC 7 PFE1 LU1	OK	41 degrees C / 105 degrees F
	FPM GBUS	OK	24 degrees C / 75 degrees F
	FPM Display	OK	28 degrees C / 82 degrees F
Fans	Top Left Front fan	OK	Spinning at normal speed
	Top Left Middle fan	OK	Spinning at normal speed
	Top Left Rear fan	OK	Spinning at normal speed
	Top Right Front fan	OK	Spinning at normal speed
	Top Right Middle fan	OK	Spinning at normal speed
	Top Right Rear fan	OK	Spinning at normal speed
	Bottom Left Front fan	OK	Spinning at normal speed
	Bottom Left Middle fan	OK	Spinning at normal speed
	Bottom Left Rear fan	OK	Spinning at normal speed
	Bottom Right Front fan	OK	Spinning at normal speed
	Bottom Right Middle fan	OK	Spinning at normal speed
	Bottom Right Rear fan	OK	Spinning at normal speed
	Rear Tray fan 1 (Top)	OK	Spinning at normal speed
	Rear Tray fan 2	OK	Spinning at normal speed
	Rear Tray fan 3	OK	Spinning at normal speed
	Rear Tray fan 4	OK	Spinning at normal speed
	Rear Tray fan 5	OK	Spinning at normal speed
	Rear Tray fan 6	OK	Spinning at normal speed
	Rear Tray fan 7	OK	Spinning at normal speed
	Rear Tray fan 8	OK	Spinning at normal speed
	Rear Tray fan 9	OK	Spinning at normal speed
	Rear Tray fan 10	OK	Spinning at normal speed
	Rear Tray fan 11	OK	Spinning at normal speed
	Rear Tray fan 12	OK	Spinning at normal speed
	Rear Tray fan 13	OK	Spinning at normal speed
	Rear Tray fan 14	OK	Spinning at normal speed
	Rear Tray fan 15	OK	Spinning at normal speed
	Rear Tray fan 16 (Bottom)	OK	Spinning at normal speed
Misc	CIP	OK	
	SPMB 0	OK	
	SPMB 1	OK	

### show chassis environment (EX4200 Standalone Switch)

user@switch> show chassis environment			
Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	Absent	
Temp	FPC 0 CPU	OK	41 degrees C / 105 degrees F
	FPC 0 EX-PFE1	OK	42 degrees C / 107 degrees F
	FPC 0 EX-PFE2	OK	46 degrees C / 114 degrees F
	FPC 0 GEPHY Front Left	OK	25 degrees C / 77 degrees F
	FPC 0 GEPHY Front Right	OK	27 degrees C / 80 degrees F
	FPC 0 Uplink Conn	OK	29 degrees C / 84 degrees F
Fans	FPC 0 Fan 1	OK	Spinning at normal speed
	FPC 0 Fan 2	OK	Spinning at normal speed
	FPC 0 Fan 3	OK	Spinning at normal speed

## show chassis environment (EX8216 Switch)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Power PSU 0                             OK
      PSU 1                             OK
      PSU 2                             OK
      PSU 3                             Check
      PSU 4                             Absent
      PSU 5                             Absent
Temp  CB 0 Intake                         OK          23 degrees C / 73 degrees F
      CB 0 Exhaust                       OK          26 degrees C / 78 degrees F
      CB 1 Intake                         OK          22 degrees C / 71 degrees F
      CB 1 Exhaust                       OK          25 degrees C / 77 degrees F
      FPC 4 Intake                       OK          49 degrees C / 120 degrees F
      FPC 4 Exhaust                     OK          59 degrees C / 138 degrees F
      SIB 5 Intake                       OK          25 degrees C / 77 degrees F
      SIB 5 Exhaust                     OK          35 degrees C / 95 degrees F
      SIB 6 Intake                       OK          25 degrees C / 77 degrees F
      SIB 6 Exhaust                     OK          38 degrees C / 100 degrees F
Fans  Top Fan 1                         OK          Spinning at normal speed
      Top Fan 2                         OK          Spinning at normal speed
      Top Fan 3                         OK          Spinning at normal speed
      Top Fan 4                         OK          Spinning at normal speed
      Top Fan 5                         OK          Spinning at normal speed
      Top Fan 6                         OK          Spinning at normal speed
      Top Fan 7                         OK          Spinning at normal speed
      Top Fan 8                         OK          Spinning at normal speed
      Top Fan 9                         OK          Spinning at normal speed
      Bottom Fan 1                     OK          Spinning at normal speed
      Bottom Fan 2                     OK          Spinning at normal speed
      Bottom Fan 3                     OK          Spinning at normal speed
      Bottom Fan 4                     OK          Spinning at normal speed
      Bottom Fan 5                     OK          Spinning at normal speed
      Bottom Fan 6                     OK          Spinning at normal speed
      Bottom Fan 7                     OK          Spinning at normal speed
      Bottom Fan 8                     OK          Spinning at normal speed
      Bottom Fan 9                     OK          Spinning at normal speed

```

## show chassis environment (EX9200 Switch)

```

user@switch> show chassis environment
Class Item                               Status      Measurement
Temp PEM 0                             Check
      PEM 1                             OK          40 degrees C / 104 degrees F
      PEM 2                             OK          40 degrees C / 104 degrees F
      PEM 3                             Absent
      Routing Engine 0                 OK          35 degrees C / 95 degrees F
      Routing Engine 0 CPU              OK          33 degrees C / 91 degrees F
      Routing Engine 1                 OK          38 degrees C / 100 degrees F
      Routing Engine 1 CPU              OK          33 degrees C / 91 degrees F
      CB 0 Intake                       OK          35 degrees C / 95 degrees F
      CB 0 Exhaust A                    OK          33 degrees C / 91 degrees F
      CB 0 Exhaust B                    OK          40 degrees C / 104 degrees F
      CB 0 ACBC                         OK          39 degrees C / 102 degrees F
      CB 0 XF A                         OK          49 degrees C / 120 degrees F
      CB 0 XF B                         OK          46 degrees C / 114 degrees F
      CB 1 Intake                       OK          37 degrees C / 98 degrees F
      CB 1 Exhaust A                    OK          32 degrees C / 89 degrees F
      CB 1 Exhaust B                    OK          39 degrees C / 102 degrees F
      CB 1 ACBC                         OK          41 degrees C / 105 degrees F

```



CB 1 XF A	OK	49 degrees C / 120 degrees F
CB 1 XF B	OK	49 degrees C / 120 degrees F
FPC 2 Intake	OK	37 degrees C / 98 degrees F
FPC 2 Exhaust A	OK	40 degrees C / 104 degrees F
FPC 2 Exhaust B	OK	34 degrees C / 93 degrees F
FPC 2 LU 0 TCAM TSen	OK	44 degrees C / 111 degrees F
FPC 2 LU 0 TCAM Chip	OK	48 degrees C / 118 degrees F
FPC 2 LU 0 TSen	OK	44 degrees C / 111 degrees F
FPC 2 LU 0 Chip	OK	60 degrees C / 140 degrees F
FPC 2 MQ 0 TSen	OK	44 degrees C / 111 degrees F
FPC 2 MQ 0 Chip	OK	51 degrees C / 123 degrees F
FPC 3 Intake	OK	39 degrees C / 102 degrees F
FPC 3 Exhaust A	OK	51 degrees C / 123 degrees F

[...Output truncated...]

Fans	Top Rear Fan	OK	Spinning at intermediate-speed
	Bottom Rear Fan	OK	Spinning at intermediate-speed
	Top Middle Fan	OK	Spinning at intermediate-speed
	Bottom Middle Fan	OK	Spinning at intermediate-speed
	Top Front Fan	OK	Spinning at intermediate-speed
	Bottom Front Fan	OK	Spinning at intermediate-speed

#### show chassis environment (QFX Series)

```
user@switch> show chassis environment
```

Class	Item	Status	Measurement
Power	FPC 0 Power Supply 0	OK	
	FPC 0 Power Supply 1	OK	
Temp	FPC 0 Sensor TopLeft I	OK	26 degrees C / 78 degrees F
	FPC 0 Sensor TopRight I	OK	24 degrees C / 75 degrees F
	FPC 0 Sensor TopLeft E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopRight E	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle I	OK	30 degrees C / 86 degrees F
	FPC 0 Sensor TopMiddle E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Bottom I	OK	34 degrees C / 93 degrees F
	FPC 0 Sensor Bottom E	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Die Temp	OK	38 degrees C / 100 degrees F
	FPC 0 Sensor Mgmt Brd I	OK	24 degrees C / 75 degrees F
	FPC 0 Sensor Switch I	OK	28 degrees C / 82 degrees F
Fans	FPC 0 Fan 1 (left)	Failed	
	FPC 0 Fan 2 (right)	OK	Spinning at normal speed
	FPC 0 Fan 3 (middle)	OK	Spinning at normal speed

#### show chassis environment interconnect-device (QFabric System)

```
user@switch> show chassis environment interconnect-device IC-A0004
```

Class	Item	Status	Measurement
CB 0			
	CB 0 L Intake	OK	30 degrees C / 86 degrees F
	CB 0 R Intake	OK	31 degrees C / 87 degrees F
	CB 0 L Exhaust	OK	32 degrees C / 89 degrees F
	CB 0 R Exhaust	OK	33 degrees C / 91 degrees F
	Routing Engine 0 CPU temp	OK	51 degrees C / 123 degrees F
CB 1			
	CB 1 L Intake	OK	27 degrees C / 80 degrees F
	CB 1 R Intake	OK	29 degrees C / 84 degrees F
	CB 1 L Exhaust	OK	31 degrees C / 87 degrees F
	CB 1 R Exhaust	OK	32 degrees C / 89 degrees F
	Routing Engine 1 CPU temp	OK	40 degrees C / 104 degrees F
FC 0	FPC 0		

FPC 0 L Intake	OK	25 degrees C / 77 degrees F
FPC 0 R Intake	OK	28 degrees C / 82 degrees F
FPC 0 L Exhaust	OK	28 degrees C / 82 degrees F
FPC 0 R Exhaust	OK	29 degrees C / 84 degrees F
FC 7 FPC 7		
FPC 7 L Intake	OK	25 degrees C / 77 degrees F
FPC 7 R Intake	OK	26 degrees C / 78 degrees F
FPC 7 L Exhaust	OK	28 degrees C / 82 degrees F
FPC 7 R Exhaust	OK	29 degrees C / 84 degrees F
RC 0 FPC 8		
FPC 8 L Intake	OK	25 degrees C / 77 degrees F
FPC 8 R Intake	OK	26 degrees C / 78 degrees F
FPC 8 L Exhaust	OK	32 degrees C / 89 degrees F
FPC 8 R Exhaust	OK	30 degrees C / 86 degrees F
RC 7 FPC 15		
FPC 15 L Intake	OK	24 degrees C / 75 degrees F
FPC 15 R Intake	OK	25 degrees C / 77 degrees F
FPC 15 L Exhaust	OK	33 degrees C / 91 degrees F
FPC 15 R Exhaust	OK	31 degrees C / 87 degrees F
Fans TFT 0 Fan 0	OK	Spinning at normal speed
Fans TFT 0 Fan 1	OK	Spinning at normal speed
Fans TFT 0 Fan 2	OK	Spinning at normal speed
Fans TFT 0 Fan 3	OK	Spinning at normal speed
Fans TFT 0 Fan 4	OK	Spinning at normal speed
Fans TFT 0 Fan 5	OK	Spinning at normal speed
Fans BFT 1 Fan 0	OK	Spinning at normal speed
Fans BFT 1 Fan 1	OK	Spinning at normal speed
Fans BFT 1 Fan 2	OK	Spinning at normal speed
Fans BFT 1 Fan 3	Check	
Fans BFT 1 Fan 4	OK	Spinning at normal speed
Fans BFT 1 Fan 5	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 0 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 1 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 2 Rotor 1	OK	Spinning at normal speed
Fans SFT 2 Fan 3 Rotor 0	OK	Spinning at normal speed
Fans SFT 2 Fan 3 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 0 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 0 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 1 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 1 Rotor 1	OK	Spinning at normal speed
Fans SFT 3 Fan 2 Rotor 0	OK	Spinning at normal speed
Fans SFT 3 Fan 2 Rotor 1	OK	Spinning at normal speed

Fans	SFT 3	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 3	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 4	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 4	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 5	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 6	Fan 3	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 0	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 1	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 2	Rotor 1	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 0	OK	Spinning at normal speed
Fans	SFT 7	Fan 3	Rotor 1	OK	Spinning at normal speed
Power	PEM 0			OK	30 degrees C / 86 degrees F
Power	PEM 1			OK	30 degrees C / 86 degrees F
Power	PEM 2			OK	30 degrees C / 86 degrees F
Power	PEM 3			Absent	
Power	PEM 4			Absent	
Power	PEM 5			Absent	

#### show chassis environment node-device (QFabric System)

```

user@switch> show chassis environment node-device node1
Class Item                               Status Measurement
Power node1 Power Supply 0              Absent
      node1 Power Supply 1              Absent
Fans  node1 Fan Tray 0                  Testing
      node1 Fan Tray 1                  Testing
      node1 Fan Tray 2                  Testing

```

#### show chassis environment pem node-device (QFabric System)

```

user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
  State          Check
  Airflow        Front to Back
  Temperature     OK
  AC Input:      OK
  DC Output      Voltage(V) Current(A) Power(W) Load(%)
                  12          10        120      18

```

```

FPC 0 PEM 1 status:
  State           Online
  Airflow         Back to Front
  Temperature      OK
  AC Input:        OK
  DC Output        Voltage(V) Current(A) Power(W) Load(%)
                   11          10       110     17

```

### show chassis environment (PTX5000 Packet Transport Router)

```

user@host> show chassis environment
Class Item                               Status      Measurement
Temp  PDU 0                               OK
      PDU 0 PSM 0                         OK          36 degrees C / 96 degrees F
      PDU 0 PSM 1                         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 2                         OK          38 degrees C / 100 degrees F
      PDU 0 PSM 3                         OK          37 degrees C / 98 degrees F
      PDU 1                               Absent
      CCG 0                               OK          44 degrees C / 111 degrees F
      CCG 1                               OK          44 degrees C / 111 degrees F
      Routing Engine 0                     OK          62 degrees C / 143 degrees F
      Routing Engine 0 CPU                  OK          75 degrees C / 167 degrees F
      Routing Engine 1                     OK          51 degrees C / 123 degrees F
      Routing Engine 1 CPU                  OK          64 degrees C / 147 degrees F
      CB 0 Intake                          OK          38 degrees C / 100 degrees F
      CB 0 Exhaust A                      OK          46 degrees C / 114 degrees F
      CB 0 Exhaust B                      OK          42 degrees C / 107 degrees F
      CB 1 Intake                          OK          35 degrees C / 95 degrees F
      CB 1 Exhaust A                      OK          39 degrees C / 102 degrees F
      CB 1 Exhaust B                      OK          36 degrees C / 96 degrees F
      SIB 0 Exhaust                       OK          47 degrees C / 116 degrees F
      SIB 0 Junction                      OK          45 degrees C / 113 degrees F
      SIB 1 Exhaust                       OK          44 degrees C / 111 degrees F
      SIB 1 Junction                      OK          43 degrees C / 109 degrees F
      SIB 2 Exhaust                       OK          47 degrees C / 116 degrees F
      SIB 2 Junction                      OK          42 degrees C / 107 degrees F
      SIB 3 Exhaust                       OK          43 degrees C / 109 degrees F
      SIB 3 Junction                      OK          43 degrees C / 109 degrees F
      SIB 4 Exhaust                       OK          47 degrees C / 116 degrees F
      SIB 4 Junction                      OK          42 degrees C / 107 degrees F
      SIB 5 Exhaust                       OK          42 degrees C / 107 degrees F
      SIB 5 Junction                      OK          40 degrees C / 104 degrees F
      SIB 6 Exhaust                       OK          46 degrees C / 114 degrees F
      SIB 6 Junction                      OK          42 degrees C / 107 degrees F
      SIB 7 Exhaust                       OK          43 degrees C / 109 degrees F
      SIB 7 Junction                      OK          39 degrees C / 102 degrees F
      SIB 8 Exhaust                       OK          44 degrees C / 111 degrees F
      SIB 8 Junction                      OK          41 degrees C / 105 degrees F
      FPC 0 PMB                           OK          35 degrees C / 95 degrees F
      FPC 0 Intake                        OK          33 degrees C / 91 degrees F
      FPC 0 Exhaust A                    OK          51 degrees C / 123 degrees F
      FPC 0 Exhaust B                    OK          43 degrees C / 109 degrees F
      FPC 0 TL0                          OK          48 degrees C / 118 degrees F
      FPC 0 TQ0                          OK          53 degrees C / 127 degrees F
      FPC 0 TL1                          OK          56 degrees C / 132 degrees F
      FPC 0 TQ1                          OK          58 degrees C / 136 degrees F
      FPC 0 TL2                          OK          55 degrees C / 131 degrees F
      FPC 0 TQ2                          OK          56 degrees C / 132 degrees F
      FPC 0 TL3                          OK          59 degrees C / 138 degrees F
      FPC 0 TQ3                          OK          59 degrees C / 138 degrees F
      FPC 2 PMB                           OK          35 degrees C / 95 degrees F

```

FPC 2 Intake	OK	34 degrees C / 93 degrees F
FPC 2 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 2 Exhaust B	OK	52 degrees C / 125 degrees F
FPC 2 TL0	OK	53 degrees C / 127 degrees F
FPC 2 TQ0	OK	53 degrees C / 127 degrees F
FPC 2 TL1	OK	57 degrees C / 134 degrees F
FPC 2 TQ1	OK	58 degrees C / 136 degrees F
FPC 2 TL2	OK	54 degrees C / 129 degrees F
FPC 2 TQ2	OK	59 degrees C / 138 degrees F
FPC 2 TL3	OK	60 degrees C / 140 degrees F
FPC 2 TQ3	OK	64 degrees C / 147 degrees F
PIC 2/0 Ambient	OK	49 degrees C / 120 degrees F
FPC 3 PMB	OK	34 degrees C / 93 degrees F
FPC 3 Intake	OK	35 degrees C / 95 degrees F
FPC 3 Exhaust A	OK	54 degrees C / 129 degrees F
FPC 3 Exhaust B	OK	49 degrees C / 120 degrees F
FPC 3 TL0	OK	49 degrees C / 120 degrees F
FPC 3 TQ0	OK	55 degrees C / 131 degrees F
FPC 3 TL1	OK	56 degrees C / 132 degrees F
FPC 3 TQ1	OK	58 degrees C / 136 degrees F
FPC 3 TL2	OK	56 degrees C / 132 degrees F
FPC 3 TQ2	OK	59 degrees C / 138 degrees F
FPC 3 TL3	OK	62 degrees C / 143 degrees F
FPC 3 TQ3	OK	63 degrees C / 145 degrees F
PIC 3/1	Absent	
FPC 5 PMB	OK	35 degrees C / 95 degrees F
FPC 5 Intake	OK	34 degrees C / 93 degrees F
FPC 5 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 5 Exhaust B	OK	53 degrees C / 127 degrees F
FPC 5 TL0	OK	54 degrees C / 129 degrees F
FPC 5 TQ0	OK	52 degrees C / 125 degrees F
FPC 5 TL1	OK	61 degrees C / 141 degrees F
FPC 5 TQ1	OK	60 degrees C / 140 degrees F
FPC 5 TL2	OK	55 degrees C / 131 degrees F
FPC 5 TQ2	OK	55 degrees C / 131 degrees F
FPC 5 TL3	OK	59 degrees C / 138 degrees F
FPC 5 TQ3	OK	58 degrees C / 136 degrees F
PIC 5/0 Ambient	OK	51 degrees C / 123 degrees F
PIC 5/1 Ambient	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/0	OK	34 degrees C / 93 degrees F
PIC 5/1 cfp-5/1/1	OK	36 degrees C / 96 degrees F
FPC 6 PMB	OK	36 degrees C / 96 degrees F
FPC 6 Intake	OK	33 degrees C / 91 degrees F
FPC 6 Exhaust A	OK	51 degrees C / 123 degrees F
FPC 6 Exhaust B	OK	39 degrees C / 102 degrees F
FPC 6 TL0	OK	44 degrees C / 111 degrees F
FPC 6 TQ0	OK	54 degrees C / 129 degrees F
FPC 6 TL1	OK	59 degrees C / 138 degrees F
FPC 6 TQ1	OK	58 degrees C / 136 degrees F
FPC 6 TL2	OK	60 degrees C / 140 degrees F
FPC 6 TQ2	OK	57 degrees C / 134 degrees F
FPC 6 TL3	OK	65 degrees C / 149 degrees F
FPC 6 TQ3	OK	60 degrees C / 140 degrees F
FPC 7 PMB	OK	35 degrees C / 95 degrees F
FPC 7 Intake	OK	33 degrees C / 91 degrees F
FPC 7 Exhaust A	OK	53 degrees C / 127 degrees F
FPC 7 Exhaust B	OK	40 degrees C / 104 degrees F
FPC 7 TL0	OK	46 degrees C / 114 degrees F
FPC 7 TQ0	OK	58 degrees C / 136 degrees F
FPC 7 TL1	OK	53 degrees C / 127 degrees F
FPC 7 TQ1	OK	59 degrees C / 138 degrees F

	FPC 7 TL2	OK	56 degrees C / 132 degrees F
	FPC 7 TQ2	OK	61 degrees C / 141 degrees F
	FPC 7 TL3	OK	63 degrees C / 145 degrees F
	FPC 7 TQ3	OK	63 degrees C / 145 degrees F
	FPM I2CS	OK	37 degrees C / 98 degrees F
Fans	Fan Tray 0 Fan 1	OK	3042 RPM
	Fan Tray 0 Fan 2	OK	3042 RPM
	Fan Tray 0 Fan 3	OK	3000 RPM
	Fan Tray 0 Fan 4	OK	3042 RPM
	Fan Tray 0 Fan 5	OK	3000 RPM
	Fan Tray 0 Fan 6	OK	3042 RPM
	Fan Tray 0 Fan 7	OK	3085 RPM
	Fan Tray 0 Fan 8	OK	3042 RPM
	Fan Tray 0 Fan 9	OK	3042 RPM
	Fan Tray 0 Fan 10	OK	3085 RPM
	Fan Tray 0 Fan 11	OK	3085 RPM
	Fan Tray 0 Fan 12	OK	3128 RPM
	Fan Tray 0 Fan 13	OK	3128 RPM
	Fan Tray 0 Fan 14	OK	3042 RPM
	Fan Tray 1 Fan 1	OK	2299 RPM
	Fan Tray 1 Fan 2	OK	2399 RPM
	Fan Tray 1 Fan 3	OK	2299 RPM
	Fan Tray 1 Fan 4	OK	2266 RPM
	Fan Tray 1 Fan 5	OK	2266 RPM
	Fan Tray 1 Fan 6	OK	2366 RPM
Misc	Fan Tray 2 Fan 1	OK	2199 RPM
	Fan Tray 2 Fan 2	OK	2133 RPM
	Fan Tray 2 Fan 3	OK	2366 RPM
	Fan Tray 2 Fan 4	OK	2233 RPM
	Fan Tray 2 Fan 5	OK	2399 RPM
	Fan Tray 2 Fan 6	OK	2233 RPM
	SPMB 0 Intake	OK	50 degrees C / 122 degrees F
	SPMB 1 Intake	OK	40 degrees C / 104 degrees F

#### show chassis environment (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis environment
```

Class	Item	Status	Measurement
Temp	PDU 0	OK	
	PDU 0 PSM 0	OK	41 degrees C / 105 degrees F
	PDU 0 PSM 1	Absent	
	PDU 0 PSM 2	OK	43 degrees C / 109 degrees F
	PDU 0 PSM 3	Absent	
	PDU 0 PSM 4	OK	44 degrees C / 111 degrees F
	PDU 0 PSM 5	Absent	
	PDU 0 PSM 6	OK	45 degrees C / 113 degrees F
	PDU 0 PSM 7	Absent	
	PDU 1	OK	
	PDU 1 PSM 0	Absent	
	PDU 1 PSM 1	OK	45 degrees C / 113 degrees F
	PDU 1 PSM 2	Absent	
	PDU 1 PSM 3	OK	43 degrees C / 109 degrees F
	PDU 1 PSM 4	Absent	
	PDU 1 PSM 5	OK	46 degrees C / 114 degrees F
	PDU 1 PSM 6	Absent	
	PDU 1 PSM 7	OK	46 degrees C / 114 degrees F
	CCG 0	OK	27 degrees C / 80 degrees F
	CCG 1	OK	29 degrees C / 84 degrees F
	...		

### show chassis environment (ACX2000 Universal Access Router)

```

user@host> show chassis environment
Class Item                               Status      Measurement
PCB Left                                OK          44 degrees C / 111 degrees F
SFP+ Xcvr                               OK          50 degrees C / 122 degrees F
FEB                                      OK          70 degrees C / 158 degrees F
PCB Up                                  OK          63 degrees C / 145 degrees F
PCB Mid                                 OK          66 degrees C / 150 degrees F
Telecom Mod                             OK          65 degrees C / 149 degrees F
Routing Engine                           OK          54 degrees C / 129 degrees F
Heater off

```

### show chassis environment (ACX4000 Universal Access Router)

On the ACX4000 router, the MIC output of the **show chassis environment** command varies depending on the number of temperature channels present in the installed MIC.

```

user@host> show chassis environment

Class Item                               Status      Measurement
Temp PEM 0                              OK          33 degrees C / 91 degrees F
      PEM 1                              Absent
      PCB Bottom                          OK          30 degrees C / 86 degrees F
      PCB Middle                          OK          34 degrees C / 93 degrees F
      BCM56445                             OK          33 degrees C / 91 degrees F
      SFP+ Xcvr                             OK          32 degrees C / 89 degrees F
      Fan tray inlet                       OK          39 degrees C / 102 degrees F
      Exhaust                             OK          30 degrees C / 86 degrees F
      Routing Engine                       OK          32 degrees C / 89 degrees F
      Heater off
Pic PIC 0/0 Channel 0                     OK          28 degrees C / 82 degrees F
    PIC 0/0 Channel 1                     OK          29 degrees C / 84 degrees F
    PIC 0/0 Channel 2                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 3                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 4                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 5                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 6                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 7                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 8                     OK          0 degrees C / 32 degrees F
    PIC 0/0 Channel 9                     OK          0 degrees C / 32 degrees F
    PIC 1/0 Channel 0                     OK          33 degrees C / 91 degrees F
    PIC 1/0 Channel 1                     OK          31 degrees C / 87 degrees F
    PIC 1/0 Channel 2                     OK          30 degrees C / 86 degrees F
    PIC 1/0 Channel 3                     OK          0 degrees C / 32 degrees F
    PIC 1/0 Channel 4                     OK          0 degrees C / 32 degrees F
    PIC 1/0 Channel 5                     OK          0 degrees C / 32 degrees F
    PIC 1/0 Channel 6                     OK          0 degrees C / 32 degrees F
    PIC 1/0 Channel 7                     OK          0 degrees C / 32 degrees F
    PIC 1/0 Channel 8                     OK          0 degrees C / 32 degrees F
    PIC 1/1 Channel 0                     OK          31 degrees C / 87 degrees F
    PIC 1/1 Channel 1                     OK          29 degrees C / 84 degrees F
    PIC 1/1 Channel 2                     OK          28 degrees C / 82 degrees F
    PIC 1/1 Channel 3                     OK          0 degrees C / 32 degrees F
    PIC 1/1 Channel 4                     OK          0 degrees C / 32 degrees F
    PIC 1/1 Channel 5                     OK          0 degrees C / 32 degrees F
    PIC 1/1 Channel 6                     OK          0 degrees C / 32 degrees F
    PIC 1/1 Channel 7                     OK          0 degrees C / 32 degrees F
    PIC 1/1 Channel 8                     OK          0 degrees C / 32 degrees F

```

Fans	Fan 1	OK	Spinning at normal speed
	Fan 2	OK	Spinning at normal speed



## show chassis environment fpc

<b>List of Syntax</b>	<a href="#">Syntax on page 575</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 575</a> <a href="#">Syntax (MX Series Routers) on page 575</a> <a href="#">Syntax (MX2010 3D Universal Edge Routers) on page 575</a> <a href="#">Syntax (MX2020 3D Universal Edge Routers) on page 575</a> <a href="#">Syntax (QFX Series) on page 575</a>
<b>Syntax</b>	show chassis environment fpc <slot>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	show chassis environment fpc <lcc number> <slot>
<b>Syntax (MX Series Routers)</b>	show chassis environment fpc <slot> <all-members> <local> <member member-id>
<b>Syntax (MX2010 3D Universal Edge Routers)</b>	show chassis environment fpc <slot>
<b>Syntax (MX2020 3D Universal Edge Routers)</b>	show chassis environment fpc <slot>
<b>Syntax (QFX Series)</b>	show chassis environment fpc <fpc-slot> interconnect-device <i>name</i>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.1 for T4000 Core Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
<b>Description</b>	(M40e, M120, M160, M320, MX Series, T Series routers, EX Series, QFX Series, and PTX Series routers only) Display environmental information about Flexible PIC Concentrators (FPCs).
<b>Options</b>	<b>none</b> —Display environmental information about all FPCs. On a TX Matrix router, display environmental information about all FPCs on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display environmental information about all FPCs on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display environmental information for the FPCs in all the members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display chassis environmental information for the Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display environmental information for the FPCs in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display environmental information for the FPCs in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**slot or fpc-slot**—(Optional) Display environmental information about an individual FPC:

- (TX Matrix and TX Matrix Plus routers only) On a TX Matrix router, if you specify the number of the T640 router by using only the **lcc *number*** option (the recommended method), replace **slot** with a value from 0 through 7. Similarly, on a TX Matrix Plus router, if you specify the number of the router by using only the **lcc *number*** option (the recommended method), replace **slot** with a value from 0 through 7. Otherwise, replace **slot** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis environment fpc 1 lcc 1
user@host> show chassis environment fpc 9
```

- M120 router—Replace **slot** with a value from 0 through 5.
- MX240 router—Replace **slot** with a value from 0 through 2.
- MX480 router—Replace **slot** with a value from 0 through 5.
- MX960 router—Replace **slot** with a value from 0 through 11.
- MX2010 router—Replace **slot** with a value from 0 through 9.
- MX2020 router—Replace **slot** with a value from 0 through 19.
- Other routers—Replace **slot** with a value from 0 through 7.
- EX Series switches:

- EX3200 switches and EX4200 standalone switches—Replace **slot** with 0.
- EX4200 switches in a Virtual Chassis configuration—Replace **slot** with a value from 0 through 9 (switch's member ID).
- EX6210 switches—Replace **slot** with a value from 0 through 3 (line card only), 4 or 5 (line card or Switch Fabric and Rotuing Engine (SRE) module), or 6 through 9 (line card only).
- EX8208 switches—Replace **slot** with a value from 0 through 7 (line card).
- EX8216 switches—Replace **slot** with a value from 0 through 15 (line card).
- QFX3500 switches —Replace **fpc-slot** with 0 through 15.
- PTX5000 Packet Transport Router—Replace **fpc-slot** with 0 through 7.

**Required Privilege Level** view

- Related Documentation**
- [request chassis fpc on page 384](#)
  - [show chassis fpc on page 639](#)
  - *show chassis fpc-feb-connectivity*
  - *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
  - *MX960 Flexible PIC Concentrator Description*

- List of Sample Output**
- [show chassis environment fpc \(M120 Router\) on page 579](#)
  - [show chassis environment fpc \(M160 Router\) on page 580](#)
  - [show chassis environment fpc \(M320 Router\) on page 580](#)
  - [show chassis environment fpc \(MX2020 Router\) on page 581](#)
  - [show chassis environment fpc \(MX2010 Router\) on page 584](#)
  - [show chassis environment fpc \(MX240 Router\) on page 586](#)
  - [show chassis environment fpc \(MX480 Router\) on page 587](#)
  - [show chassis environment fpc \(MX960 Router\) on page 588](#)
  - [show chassis environment fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 589](#)
  - [show chassis environment fpc \(MX240, MX480, MX960 with Application Services Modular Line Card on page 590](#)
  - [show chassis environment fpc \(T320, T640, and T1600 Routers\) on page 591](#)
  - [show chassis environment fpc \(T4000 Router\) on page 591](#)
  - [show chassis environment fpc lcc \(TX Matrix Router\) on page 596](#)
  - [show chassis environment fpc lcc \(TX Matrix Plus Router\) on page 597](#)
  - [show chassis environment fpc \(QFX Series\) on page 598](#)
  - [show chassis environment fpc interconnect-device \(QFabric Systems\) on page 598](#)
  - [show chassis environment fpc 0 \(PTX5000 Packet Transport Router\) on page 598](#)
  - [show chassis environment fpc 07 \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 599](#)

[show chassis environment FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 600](#)

**Output Fields** [Table 28 on page 578](#) lists the output fields for the **show chassis environment fpc** command. Output fields are listed in the approximate order in which they appear.

**Table 28: show chassis environment fpc Output Fields**

Field Name	Field Description
<b>State</b>	<p>Status of the FPC:</p> <ul style="list-style-type: none"> <li>• <b>Unknown</b>—FPC is not detected by the router.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online.</li> <li>• <b>Ready</b>—FPC is in intermediate or transition state.</li> <li>• <b>Announce online</b>—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative.</li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Offline</b>—FPC is powered down.</li> <li>• <b>Diagnostics</b>—FPC is set to operate in diagnostics mode.</li> </ul>
<b>Temperature</b>	(M40e and M160 routers and QFX Series only) Temperature of the air flowing past the FPC.
<b>PMB Temperature</b>	<p>(PTX Series only) Temperature of the air flowing past the PMB (bottom of the FPC).</p> <p>The PTX5000 Packet Transport Router with FPC2-PTX-PIA include multiple temperatures for PMB (<b>TEMPO</b> and <b>TEMPI</b>).</p>
<b>PMB CPU Temperature</b>	(PTX5000 Packet Transport Router with FPC2-PTX-PIA only) Temperature of the air flowing past the PMB CPU.
<b>Temperature Intake</b>	(M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing into the chassis.
<b>Temperature Top</b>	(T Series routers only) Temperature of the air flowing past the top of the FPC.
<b>Temperature Exhaust</b>	<p>(M120 and M320 routers, MX2010 routers, MX2020 routers, and PTX Series only) Temperature of the air flowing out of the chassis.</p> <p>The PTX Series Packet Transport Routers, and the MX2010 and MX2020 routers include exhaust temperatures for multiple zones (<b>Exhaust A</b> and <b>Exhaust B</b>).</p>
<b>Temperature Bottom</b>	(T Series routers only) Temperature of the air flowing past the bottom of the FPC.
<b>TL <i>n</i> Temperature</b>	(PTX Series only) Temperature of the air flowing past the specified TL area of the packet forwarding engine (PFE) on the FPC.
<b>TQ <i>n</i> Temperature</b>	(PTX Series only) Temperature of the air flowing past the specified TQ area of the packet forwarding engine (PFE) on the FPC.
<b>Temperature MMBO</b>	(T640 router only) Temperature of the air flowing past the type 3 FPC.

Table 28: show chassis environment fpc Output Fields (*continued*)

Field Name	Field Description
<b>Temperature MMB1</b>	(M320 and T Series routers only) Temperature of the air flowing past the type 1, type 2, and type 3 FPC.
<b>Power</b>	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.
<b>CMB Revision or BUS revision</b>	Revision level of the chassis management bus device (M Series router) or bus (T Series routers).

## Sample Output

### show chassis environment fpc (M120 Router)

```

user@host> show chassis environment fpc
FPC 2 status:
  State                               Online
  Temperature Exhaust A               32 degrees C / 89 degrees F
  Temperature Exhaust B               31 degrees C / 87 degrees F
  Power A-Board
    1.2 V                             1202 mV
    1.5 V                             1508 mV
    1.8 V                             1798 mV
    2.5 V                             2507 mV
    3.3 V                             3351 mV
    5.0 V                             4995 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 3 status:
  State                               Online
  Temperature Exhaust A               31 degrees C / 87 degrees F
  Temperature Exhaust B               33 degrees C / 91 degrees F
  Power A-Board
    1.2 V                             1211 mV
    1.5 V                             1501 mV
    1.8 V                             1798 mV
    2.5 V                             2471 mV
    3.3 V                             3293 mV
    5.0 V                             4930 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  Power B-Board
    1.2 V                             1214 mV
    1.5 V                             1501 mV
    2.5 V                             2471 mV
    3.3 V                             3300 mV
    5.0 V                             4943 mV
    3.3 V bias                         3296 mV
    1.2 V Rocket IO                   1205 mV
    1.5 V Rocket IO                   1501 mV
  I2C Slave Revision                 12
FPC 4 status:
  State                               Online

```

Temperature Exhaust A	32 degrees C / 89 degrees F
Temperature Exhaust B	30 degrees C / 86 degrees F
Power A-Board	
1.2 V	1195 mV
1.5 V	1504 mV
1.8 V	1801 mV
2.5 V	2504 mV
3.3 V	3293 mV
5.0 V	4917 mV
3.3 V bias	3296 mV
1.2 V Rocket IO	1202 mV
1.5 V Rocket IO	1492 mV
I2C Slave Revision	12

#### show chassis environment fpc (M160 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
State                Online
Temperature          42 degrees C / 107 degrees F
Power:
  1.5 V              1500 mV
  2.5 V              2509 mV
  3.3 V              3308 mV
  5.0 V              4991 mV
  5.0 V bias         4952 mV
  8.0 V bias         8307 mV
CMB Revision         12
FPC 1 status:
State                Online
Temperature          45 degrees C / 113 degrees F
Power:
  1.5 V              1498 mV
  2.5 V              2501 mV
  3.3 V              3319 mV
  5.0 V              5020 mV
  5.0 V bias         5025 mV
  8.0 V bias         8307 mV
CMB Revision         12
```

#### show chassis environment fpc (M320 Router)

```
user@host> show chassis environment fpc
FPC 0 status:
State                Online
Temperature Intake    27 degrees C / 80 degrees F
Temperature Exhaust   38 degrees C / 100 degrees F
Temperature MMB1      31 degrees C / 87 degrees F
Power:
  1.5 V              1487 mV
  1.5 V *            1494 mV
  1.8 V              1821 mV
  2.5 V              2533 mV
  3.3 V              3323 mV
  5.0 V              5028 mV
  3.3 V bias         3296 mV
  5.0 V bias         4984 mV
CMB Revision         16
FPC 1 status:
State                Online
Temperature Intake    27 degrees C / 80 degrees F
```

```

Temperature Exhaust      37 degrees C / 98 degrees F
Temperature MMB1         32 degrees C / 89 degrees F
Power:
  1.5 V                  1504 mV
  1.5 V *                1499 mV
  1.8 V                  1820 mV
  2.5 V                  2529 mV
  3.3 V                  3328 mV
  5.0 V                  5013 mV
  3.3 V bias             3294 mV
  5.0 V bias             4984 mV
CMB Revision             16
FPC 2 status:
State                    Online
Temperature Intake        28 degrees C / 82 degrees F
Temperature Exhaust       38 degrees C / 100 degrees F
Temperature MMB1         32 degrees C / 89 degrees F
Power:
  1.5 V                  1498 mV
  1.5 V *                1487 mV
  1.8 V                  1816 mV
  2.5 V                  2531 mV
  3.3 V                  3324 mV
  5.0 V                  5025 mV
  3.3 V bias             3277 mV
  5.0 V bias             5013 mV
CMB Revision             17
FPC 3 status:
...
```

### show chassis environment fpc (MX2020 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                    Online
Temperature Intake        41 degrees C / 105 degrees F
Temperature Exhaust A     48 degrees C / 118 degrees F
Temperature Exhaust B     60 degrees C / 140 degrees F
Temperature LU 0 TSen     56 degrees C / 132 degrees F
Temperature LU 0 Chip     59 degrees C / 138 degrees F
Temperature LU 1 TSen     56 degrees C / 132 degrees F
Temperature LU 1 Chip     61 degrees C / 141 degrees F
Temperature LU 2 TSen     56 degrees C / 132 degrees F
Temperature LU 2 Chip     52 degrees C / 125 degrees F
Temperature LU 3 TSen     56 degrees C / 132 degrees F
Temperature LU 3 Chip     52 degrees C / 125 degrees F
Temperature MQ 0 TSen     49 degrees C / 120 degrees F
Temperature MQ 0 Chip     49 degrees C / 120 degrees F
Temperature MQ 1 TSen     49 degrees C / 120 degrees F
Temperature MQ 1 Chip     52 degrees C / 125 degrees F
Temperature MQ 2 TSen     49 degrees C / 120 degrees F
Temperature MQ 2 Chip     45 degrees C / 113 degrees F
Temperature MQ 3 TSen     49 degrees C / 120 degrees F
Temperature MQ 3 Chip     46 degrees C / 114 degrees F
Power
AS-BIAS3V3-z12105        3299 mV
AS-VDD1V8-z12006         1807 mV
AS-VDD2V5-z12006         2512 mV
AS-AVDD1V0-z12004         997 mV
AS-PCIE_1V0-z12004        996 mV
AS-VDD3V3-z12004         3294 mV
```

```

AS-VDD_1V5A-z12004      1501 mV
AS-VDD_1V5B-z12004      1498 mV
AS-LU0_1V0-z12004        998 mV
AS-LU1_1V0-z12004       1002 mV
AS-MQ0_1V0-z12004        999 mV
AS-MQ1_1V0-z12004        994 mV
AS-LU2_1V0-z12004       1000 mV
AS-LU3_1V0-z12004        998 mV
AS-MQ2_1V0-z12004       1002 mV
AS-MQ3_1V0-z12004        999 mV
AS-PMB_1V1-z12006       1096 mV
I2C Slave Revision      68
FPC 1 status:
State                    Online
Temperature Intake       39 degrees C / 102 degrees F
Temperature Exhaust A    48 degrees C / 118 degrees F
Temperature Exhaust B    55 degrees C / 131 degrees F
Temperature LU 0 TSen     52 degrees C / 125 degrees F
Temperature LU 0 Chip     54 degrees C / 129 degrees F
Temperature LU 1 TSen     52 degrees C / 125 degrees F
Temperature LU 1 Chip     56 degrees C / 132 degrees F
Temperature LU 2 TSen     52 degrees C / 125 degrees F
Temperature LU 2 Chip     49 degrees C / 120 degrees F
Temperature LU 3 TSen     52 degrees C / 125 degrees F
Temperature LU 3 Chip     50 degrees C / 122 degrees F
Temperature MQ 0 TSen     48 degrees C / 118 degrees F
Temperature MQ 0 Chip     48 degrees C / 118 degrees F
Temperature MQ 1 TSen     48 degrees C / 118 degrees F
Temperature MQ 1 Chip     51 degrees C / 123 degrees F
Temperature MQ 2 TSen     48 degrees C / 118 degrees F
Temperature MQ 2 Chip     45 degrees C / 113 degrees F
Temperature MQ 3 TSen     48 degrees C / 118 degrees F
Temperature MQ 3 Chip     45 degrees C / 113 degrees F
Power
AS-BIAS3V3-z12105       3291 mV
AS-VDD1V8-z12006       1786 mV
AS-VDD2V5-z12006       2496 mV
AS-AVDD1V0-z12004       1000 mV
AS-PCIE_1V0-z12004       1000 mV
AS-VDD3V3-z12004       3294 mV
AS-VDD_1V5A-z12004      1500 mV
AS-VDD_1V5B-z12004      1498 mV
AS-LU0_1V0-z12004       1003 mV
AS-LU1_1V0-z12004       1000 mV
AS-MQ0_1V0-z12004       1000 mV
AS-MQ1_1V0-z12004        995 mV
AS-LU2_1V0-z12004       1002 mV
AS-LU3_1V0-z12004        997 mV
AS-MQ2_1V0-z12004       1000 mV
AS-MQ3_1V0-z12004        998 mV
AS-PMB_1V1-z12006       1096 mV
I2C Slave Revision      68
FPC 2 status:
State                    Online
Temperature Intake       39 degrees C / 102 degrees F
Temperature Exhaust A    48 degrees C / 118 degrees F
Temperature Exhaust B    58 degrees C / 136 degrees F
Temperature LU 0 TSen     55 degrees C / 131 degrees F
Temperature LU 0 Chip     57 degrees C / 134 degrees F
Temperature LU 1 TSen     55 degrees C / 131 degrees F
Temperature LU 1 Chip     63 degrees C / 145 degrees F

```



```

Temperature LU 2 TSen      55 degrees C / 131 degrees F
Temperature LU 2 Chip      51 degrees C / 123 degrees F
Temperature LU 3 TSen      55 degrees C / 131 degrees F
Temperature LU 3 Chip      52 degrees C / 125 degrees F
Temperature MQ 0 TSen      48 degrees C / 118 degrees F
Temperature MQ 0 Chip      50 degrees C / 122 degrees F
Temperature MQ 1 TSen      48 degrees C / 118 degrees F
Temperature MQ 1 Chip      52 degrees C / 125 degrees F
Temperature MQ 2 TSen      48 degrees C / 118 degrees F
Temperature MQ 2 Chip      47 degrees C / 116 degrees F
Temperature MQ 3 TSen      48 degrees C / 118 degrees F
Temperature MQ 3 Chip      47 degrees C / 116 degrees F
Power
AS-BIAS3V3-z12105         3299 mV
AS-VDD1V8-z12006          1805 mV
AS-VDD2V5-z12006          2510 mV
AS-AVDD1V0-z12004          999 mV
AS-PCIE_1V0-z12004          998 mV
AS-VDD3V3-z12004          3296 mV
AS-VDD_1V5A-z12004         1492 mV
AS-VDD_1V5B-z12004         1497 mV
AS-LU0_1V0-z12004          997 mV
AS-LU1_1V0-z12004         1000 mV
AS-MQ0_1V0-z12004          998 mV
AS-MQ1_1V0-z12004         1001 mV
AS-LU2_1V0-z12004          996 mV
AS-LU3_1V0-z12004          995 mV
AS-MQ2_1V0-z12004          998 mV
AS-MQ3_1V0-z12004          997 mV
AS-PMB_1V1-z12006         1100 mV
I2C Slave Revision        68
FPC 3 status:
State                      Online
Temperature Intake          41 degrees C / 105 degrees F
Temperature Exhaust A       48 degrees C / 118 degrees F
Temperature Exhaust B       58 degrees C / 136 degrees F
Temperature LU 0 TSen       56 degrees C / 132 degrees F
Temperature LU 0 Chip       59 degrees C / 138 degrees F
Temperature LU 1 TSen       56 degrees C / 132 degrees F
Temperature LU 1 Chip       61 degrees C / 141 degrees F
Temperature LU 2 TSen       56 degrees C / 132 degrees F
Temperature LU 2 Chip       51 degrees C / 123 degrees F
Temperature LU 3 TSen       56 degrees C / 132 degrees F
Temperature LU 3 Chip       53 degrees C / 127 degrees F
Temperature MQ 0 TSen       50 degrees C / 122 degrees F
Temperature MQ 0 Chip       51 degrees C / 123 degrees F
Temperature MQ 1 TSen       50 degrees C / 122 degrees F
Temperature MQ 1 Chip       55 degrees C / 131 degrees F
Temperature MQ 2 TSen       50 degrees C / 122 degrees F
Temperature MQ 2 Chip       47 degrees C / 116 degrees F
Temperature MQ 3 TSen       50 degrees C / 122 degrees F
Temperature MQ 3 Chip       50 degrees C / 122 degrees F
Power
AS-BIAS3V3-z12105         3305 mV
AS-VDD1V8-z12006          1810 mV
AS-VDD2V5-z12006          2508 mV
AS-AVDD1V0-z12004          999 mV
AS-PCIE_1V0-z12004         1001 mV
AS-VDD3V3-z12004          3294 mV
AS-VDD_1V5A-z12004         1500 mV
AS-VDD_1V5B-z12004         1498 mV

```

```

AS-LU0_1V0-z12004      998 mV
AS-LU1_1V0-z12004      998 mV
AS-MQ0_1V0-z12004      999 mV
AS-MQ1_1V0-z12004      998 mV
AS-LU2_1V0-z12004      1000 mV
AS-LU3_1V0-z12004      1001 mV
AS-MQ2_1V0-z12004      996 mV
AS-MQ3_1V0-z12004      998 mV
AS-PMB_1V1-z12006      1098 mV
I2C Slave Revision      68
FPC 4 status:
...

```

### show chassis environment fpc (MX2010 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
State                               Online
Temperature Intake                  36 degrees C / 96 degrees F
Temperature Exhaust A                42 degrees C / 107 degrees F
Temperature Exhaust B                51 degrees C / 123 degrees F
Temperature LU 0 TSen                49 degrees C / 120 degrees F
Temperature LU 0 Chip                50 degrees C / 122 degrees F
Temperature LU 1 TSen                49 degrees C / 120 degrees F
Temperature LU 1 Chip                54 degrees C / 129 degrees F
Temperature LU 2 TSen                49 degrees C / 120 degrees F
Temperature LU 2 Chip                45 degrees C / 113 degrees F
Temperature LU 3 TSen                49 degrees C / 120 degrees F
Temperature LU 3 Chip                46 degrees C / 114 degrees F
Temperature MQ 0 TSen                40 degrees C / 104 degrees F
Temperature MQ 0 Chip                41 degrees C / 105 degrees F
Temperature MQ 1 TSen                40 degrees C / 104 degrees F
Temperature MQ 1 Chip                44 degrees C / 111 degrees F
Temperature MQ 2 TSen                40 degrees C / 104 degrees F
Temperature MQ 2 Chip                38 degrees C / 100 degrees F
Temperature MQ 3 TSen                40 degrees C / 104 degrees F
Temperature MQ 3 Chip                41 degrees C / 105 degrees F
Power
AS-BIAS3V3-z12105                  3300 mV
AS-VDD1V8-z12006                    1805 mV
AS-VDD2V5-z12006                    2505 mV
AS-AVDD1V0-z12004                    998 mV
AS-PCIE_1V0-z12004                    999 mV
AS-VDD3V3-z12004                      3303 mV
AS-VDD_1V5A-z12004                    1497 mV
AS-VDD_1V5B-z12004                    1497 mV
AS-LU0_1V0-z12004                      998 mV
AS-LU1_1V0-z12004                    1003 mV
AS-MQ0_1V0-z12004                      998 mV
AS-MQ1_1V0-z12004                      998 mV
AS-LU2_1V0-z12004                      997 mV
AS-LU3_1V0-z12004                    1001 mV
AS-MQ2_1V0-z12004                      996 mV
AS-MQ3_1V0-z12004                      994 mV
AS-PMB_1V1-z12006                    1097 mV
I2C Slave Revision                    68
FPC 1 status:
State                               Online
Temperature Intake                  34 degrees C / 93 degrees F
Temperature Exhaust A                46 degrees C / 114 degrees F
Temperature Exhaust B                54 degrees C / 129 degrees F

```

```

Temperature LU 0 TSen      45 degrees C / 113 degrees F
Temperature LU 0 Chip      55 degrees C / 131 degrees F
Temperature LU 1 TSen      45 degrees C / 113 degrees F
Temperature LU 1 Chip      44 degrees C / 111 degrees F
Temperature LU 2 TSen      45 degrees C / 113 degrees F
Temperature LU 2 Chip      50 degrees C / 122 degrees F
Temperature LU 3 TSen      45 degrees C / 113 degrees F
Temperature LU 3 Chip      58 degrees C / 136 degrees F
Temperature XM 0 TSen      45 degrees C / 113 degrees F
Temperature XM 0 Chip      51 degrees C / 123 degrees F
Temperature XF 0 TSen      45 degrees C / 113 degrees F
Temperature XF 0 Chip      63 degrees C / 145 degrees F
Temperature PLX Switch TSen45 degrees C / 113 degrees F
Temperature PLX Switch Chip47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105        3300 mV
MPC-VDD3V3-z16100         3294 mV
MPC-VDD2V5-z16100         2505 mV
MPC-VDD1V8-z12004         1796 mV
MPC-AVDD1V0-z12004         991 mV
MPC-VDD1V2-z16100         1196 mV
MPC-VDD1V5A-z12004         1491 mV
MPC-VDD1V5B-z12004         1492 mV
MPC-XF_OV9-z12004          996 mV
MPC-PCIE_1V0-z16100        1003 mV
MPC-LU0_1V0-z12004          996 mV
MPC-LU1_1V0-z12004          996 mV
MPC-LU2_1V0-z12004          998 mV
MPC-LU3_1V0-z12004          994 mV
MPC-12VA-BMR453            12031 mV
MPC-12VB-BMR453            12003 mV
MPC-PMB_1V1-z12006         1104 mV
MPC-PMB_1V2-z12106         1194 mV
MPC-XM_OV9-vt273m          911 mV
I2C Slave Revision         110
FPC 8 status:
State                       Online
Temperature Intake           32 degrees C / 89 degrees F
Temperature Exhaust A        44 degrees C / 111 degrees F
Temperature Exhaust B        37 degrees C / 98 degrees F
Temperature LU 0 TCAM TSen   41 degrees C / 105 degrees F
Temperature LU 0 TCAM Chip   49 degrees C / 120 degrees F
Temperature LU 0 TSen        41 degrees C / 105 degrees F
Temperature LU 0 Chip        52 degrees C / 125 degrees F
Temperature MQ 0 TSen        41 degrees C / 105 degrees F
Temperature MQ 0 Chip        47 degrees C / 116 degrees F
Temperature LU 1 TCAM TSen   39 degrees C / 102 degrees F
Temperature LU 1 TCAM Chip   42 degrees C / 107 degrees F
Temperature LU 1 TSen        39 degrees C / 102 degrees F
Temperature LU 1 Chip        46 degrees C / 114 degrees F
Temperature MQ 1 TSen        39 degrees C / 102 degrees F
Temperature MQ 1 Chip        45 degrees C / 113 degrees F
Power
MPC-BIAS3V3-z12105        3296 mV
MPC-VDD3V3-z12006         3298 mV
MPC-VDD2V5-z12006         2505 mV
MPC-TCAM_1V0-z12004         997 mV
MPC-AVDD1V0-z12006         1007 mV
MPC-VDD1V8-z12006         1803 mV
MPC-PCIE_1V0-z12006         1004 mV
MPC-LU0_1V0-z12004         1000 mV

```

```

MPC-MQ0_1V0-z12004      999 mV
MPC-VDD_1V5-z12004      1498 mV
MPC-PMB_1V1-z12006      1102 mV
MPC-9VA-BMR453          9009 mV
MPC-9VB-BMR453          8960 mV
MPC-PMB_1V2-z12105      1202 mV
MPC-LU1_1V0-z12004      1005 mV
MPC-MQ1_1V0-z12004      1000 mV
I2C Slave Revision      70
FPC 9 status:
State                   Online
Temperature Intake      34 degrees C / 93 degrees F
Temperature Exhaust A   41 degrees C / 105 degrees F
Temperature Exhaust B   54 degrees C / 129 degrees F
Temperature LU 0 TSen   51 degrees C / 123 degrees F
Temperature LU 0 Chip   52 degrees C / 125 degrees F
Temperature LU 1 TSen   51 degrees C / 123 degrees F
Temperature LU 1 Chip   55 degrees C / 131 degrees F
Temperature LU 2 TSen   51 degrees C / 123 degrees F
Temperature LU 2 Chip   47 degrees C / 116 degrees F
Temperature LU 3 TSen   51 degrees C / 123 degrees F
Temperature LU 3 Chip   47 degrees C / 116 degrees F
Temperature MQ 0 TSen   40 degrees C / 104 degrees F
Temperature MQ 0 Chip   42 degrees C / 107 degrees F
Temperature MQ 1 TSen   40 degrees C / 104 degrees F
Temperature MQ 1 Chip   44 degrees C / 111 degrees F
Temperature MQ 2 TSen   40 degrees C / 104 degrees F
Temperature MQ 2 Chip   38 degrees C / 100 degrees F
Temperature MQ 3 TSen   40 degrees C / 104 degrees F
Temperature MQ 3 Chip   40 degrees C / 104 degrees F
Power
AS-BIAS3V3-z12105      3302 mV
AS-VDD1V8-z12006      1808 mV
AS-VDD2V5-z12006      2513 mV
AS-AVDD1V0-z12004      997 mV
AS-PCIE_1V0-z12004      999 mV
AS-VDD3V3-z12004      3294 mV
AS-VDD_1V5A-z12004      1503 mV
AS-VDD_1V5B-z12004      1502 mV
AS-LU0_1V0-z12004      996 mV
AS-LU1_1V0-z12004      999 mV
AS-MQ0_1V0-z12004      997 mV
AS-MQ1_1V0-z12004      999 mV
AS-LU2_1V0-z12004      997 mV
AS-LU3_1V0-z12004      998 mV
AS-MQ2_1V0-z12004      1000 mV
AS-MQ3_1V0-z12004      1000 mV
AS-PMB_1V1-z12006      1102 mV
I2C Slave Revision      68

```

### show chassis environment fpc (MX240 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                   Online
Temperature Intake      34 degrees C / 93 degrees F
Temperature Exhaust A   39 degrees C / 102 degrees F
Temperature Exhaust B   53 degrees C / 127 degrees F
Temperature I3 0 TSensor 51 degrees C / 123 degrees F
Temperature I3 0 Chip   54 degrees C / 129 degrees F
Temperature I3 1 TSensor 50 degrees C / 122 degrees F

```

```

Temperature I3 1 Chip      53 degrees C / 127 degrees F
Temperature I3 2 TSensor   48 degrees C / 118 degrees F
Temperature I3 2 Chip      51 degrees C / 123 degrees F
Temperature I3 3 TSensor   45 degrees C / 113 degrees F
Temperature I3 3 Chip      48 degrees C / 118 degrees F
Temperature IA 0 TSensor   45 degrees C / 113 degrees F
Temperature IA 0 Chip      45 degrees C / 113 degrees F
Temperature IA 1 TSensor   45 degrees C / 113 degrees F
Temperature IA 1 Chip      49 degrees C / 120 degrees F
Power
  1.5 V                    1492 mV
  2.5 V                    2507 mV
  3.3 V                    3306 mV
  1.8 V PFE 0              1801 mV
  1.8 V PFE 1              1804 mV
  1.8 V PFE 2              1798 mV
  1.8 V PFE 3              1798 mV
  1.2 V PFE 0              1169 mV
  1.2 V PFE 1              1189 mV
  1.2 V PFE 2              1182 mV
  1.2 V PFE 3              1176 mV
I2C Slave Revision        42
FPC 2 status:
State                      Online
Temperature Intake          33 degrees C / 91 degrees F
Temperature Exhaust A       41 degrees C / 105 degrees F
Temperature Exhaust B       53 degrees C / 127 degrees F
Temperature I3 0 TSensor    53 degrees C / 127 degrees F
Temperature I3 0 Chip       58 degrees C / 136 degrees F
Temperature I3 1 TSensor    52 degrees C / 125 degrees F
Temperature I3 1 Chip       56 degrees C / 132 degrees F
Temperature I3 2 TSensor    50 degrees C / 122 degrees F
Temperature I3 2 Chip       52 degrees C / 125 degrees F
Temperature I3 3 TSensor    46 degrees C / 114 degrees F
Temperature I3 3 Chip       49 degrees C / 120 degrees F
Temperature IA 0 TSensor    51 degrees C / 123 degrees F
Temperature IA 0 Chip       49 degrees C / 120 degrees F
Temperature IA 1 TSensor    48 degrees C / 118 degrees F
Temperature IA 1 Chip       53 degrees C / 127 degrees F
Power
  1.5 V                    1492 mV
  2.5 V                    2445 mV
  3.3 V                    3293 mV
  1.8 V PFE 0              1827 mV
  1.8 V PFE 1              1775 mV
  1.8 V PFE 2              1788 mV
  1.8 V PFE 3              1798 mV
  1.2 V PFE 0              1250 mV
  1.2 V PFE 1              1234 mV
  1.2 V PFE 2              1231 mV
  1.2 V PFE 3              1192 mV
I2C Slave Revision        42

```

#### show chassis environment fpc (MX480 Router)

```

user@host> show chassis environment fpc
FPC 1 status:
State                      Online
Temperature Intake          36 degrees C / 96 degrees F
Temperature Exhaust A       41 degrees C / 105 degrees F
Temperature Exhaust B       55 degrees C / 131 degrees F

```

```

Temperature I3 0 TSensor 55 degrees C / 131 degrees F
Temperature I3 0 Chip    57 degrees C / 134 degrees F
Temperature I3 1 TSensor 53 degrees C / 127 degrees F
Temperature I3 1 Chip    53 degrees C / 127 degrees F
Temperature I3 2 TSensor 52 degrees C / 125 degrees F
Temperature I3 2 Chip    49 degrees C / 120 degrees F
Temperature I3 3 TSensor 47 degrees C / 116 degrees F
Temperature I3 3 Chip    47 degrees C / 116 degrees F
Temperature IA 0 TSensor 54 degrees C / 129 degrees F
Temperature IA 0 Chip    58 degrees C / 136 degrees F
Temperature IA 1 TSensor 48 degrees C / 118 degrees F
Temperature IA 1 Chip    53 degrees C / 127 degrees F
Power
  1.5 V      1479 mV
  2.5 V      2542 mV
  3.3 V      3319 mV
  1.8 V PFE 0 1811 mV
  1.8 V PFE 1 1804 mV
  1.8 V PFE 2 1804 mV
  1.8 V PFE 3 1814 mV
  1.2 V PFE 0 1192 mV
  1.2 V PFE 1 1202 mV
  1.2 V PFE 2 1205 mV
  1.2 V PFE 3 1189 mV
I2C Slave Revision 40

```

#### show chassis environment fpc (MX960 Router)

```

user@host> show chassis environment fpc
FPC 5 status:
State Online
Temperature Intake 27 degrees C / 80 degrees F
Temperature Exhaust A 34 degrees C / 93 degrees F
Temperature Exhaust B 40 degrees C / 104 degrees F
Temperature I3 0 TSensor 39 degrees C / 102 degrees F
Temperature I3 0 Chip 41 degrees C / 105 degrees F
Temperature I3 1 TSensor 38 degrees C / 100 degrees F
Temperature I3 1 Chip 37 degrees C / 98 degrees F
Temperature I3 2 TSensor 37 degrees C / 98 degrees F
Temperature I3 2 Chip 34 degrees C / 93 degrees F
Temperature I3 3 TSensor 32 degrees C / 89 degrees F
Temperature I3 3 Chip 33 degrees C / 91 degrees F
Temperature IA 0 TSensor 39 degrees C / 102 degrees F
Temperature IA 0 Chip 44 degrees C / 111 degrees F
Temperature IA 1 TSensor 36 degrees C / 96 degrees F
Temperature IA 1 Chip 44 degrees C / 111 degrees F
Power
  1.5 V      1479 mV
  2.5 V      2523 mV
  3.3 V      3254 mV
  1.8 V PFE 0 1798 mV
  1.8 V PFE 1 1798 mV
  1.8 V PFE 2 1807 mV
  1.8 V PFE 3 1791 mV
  1.2 V PFE 0 1173 mV
  1.2 V PFE 1 1179 mV
  1.2 V PFE 2 1179 mV
  1.2 V PFE 3 1185 mV
I2C Slave Revision 6
FPC 6 status:
State Online

```

```

Temperature Intake          25 degrees C / 77 degrees F
Temperature Exhaust A      38 degrees C / 100 degrees F
Temperature Exhaust B      38 degrees C / 100 degrees F
Temperature I3 0 TSensor   40 degrees C / 104 degrees F
Temperature I3 0 Chip      40 degrees C / 104 degrees F
Temperature I3 1 TSensor   40 degrees C / 104 degrees F
Temperature I3 1 Chip      38 degrees C / 100 degrees F
Temperature I3 2 TSensor   37 degrees C / 98 degrees F
Temperature I3 2 Chip      32 degrees C / 89 degrees F
Temperature I3 3 TSensor   34 degrees C / 93 degrees F
Temperature I3 3 Chip      33 degrees C / 91 degrees F
Temperature IA 0 TSensor   45 degrees C / 113 degrees F
Temperature IA 0 Chip      47 degrees C / 116 degrees F
Temperature IA 1 TSensor   37 degrees C / 98 degrees F
Temperature IA 1 Chip      42 degrees C / 107 degrees F
Power
  1.5 V                    1485 mV
  2.5 V                    2510 mV
  3.3 V                    3332 mV
  1.8 V PFE 0              1801 mV
  1.8 V PFE 1              1814 mV
  1.8 V PFE 2              1804 mV
  1.8 V PFE 3              1820 mV
  1.2 V PFE 0              1192 mV
  1.2 V PFE 1              1189 mV
  1.2 V PFE 2              1202 mV
  1.2 V PFE 3              1156 mV
I2C Slave Revision        40

```

#### show chassis environment fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis environment fpc
FPC 0 status:
State          Online
Temperature Intake          32 degrees C / 89 degrees F
Temperature Exhaust A      39 degrees C / 102 degrees F
Temperature Exhaust B      37 degrees C / 98 degrees F
Temperature QX 0 TSen      44 degrees C / 111 degrees F
Temperature QX 0 Chip      48 degrees C / 118 degrees F
Temperature LU 0 TCAM TSen 44 degrees C / 111 degrees F
Temperature LU 0 TCAM Chip 47 degrees C / 116 degrees F
Temperature LU 0 TSen      44 degrees C / 111 degrees F
Temperature LU 0 Chip      48 degrees C / 118 degrees F
Temperature MQ 0 TSen      44 degrees C / 111 degrees F
Temperature MQ 0 Chip      47 degrees C / 116 degrees F
Power
MPC-BIAS3V3-z12105        3297 mV
MPC-VDD3V3-z12105        3306 mV
MPC-VDD2V5-z12105        2498 mV
MPC-TCAM_1V0-z12004       999 mV
MPC-AVDD1V0-z12006       999 mV
MPC-VDD1V8-z12006        1796 mV
MPC-PCIE_1V0-z12006       1002 mV
MPC-LU0_1V0-z12004        997 mV
MPC-MQ0_1V0-z12004        995 mV
MPC-VDD_1V5-z12004        1496 mV
MPC-PMB_1V1-z12006        1094 mV
MPC-9VA-BMR453            9054 mV
MPC-9VB-BMR453            9037 mV
MPC-PMB_1V2-z12106        1191 mV
MPC-QXM0_1V0-z12006       1000 mV

```

```

I2C Slave Revision          66
FPC 1 status:
State                       Online
Temperature Intake          35 degrees C / 95 degrees F
Temperature Exhaust A      50 degrees C / 122 degrees F
Temperature Exhaust B      56 degrees C / 132 degrees F
Temperature LU 0 TSen      46 degrees C / 114 degrees F
Temperature LU 0 Chip      59 degrees C / 138 degrees F
Temperature LU 1 TSen      46 degrees C / 114 degrees F
Temperature LU 1 Chip      45 degrees C / 113 degrees F
Temperature LU 2 TSen      46 degrees C / 114 degrees F
Temperature LU 2 Chip      60 degrees C / 140 degrees F
Temperature LU 3 TSen      46 degrees C / 114 degrees F
Temperature LU 3 Chip      71 degrees C / 159 degrees F
Temperature XM 0 TSen      46 degrees C / 114 degrees F
Temperature XM 0 Chip      -18 degrees C / 0 degrees F
Temperature XF 0 TSen      46 degrees C / 114 degrees F
Temperature XF 0 Chip      76 degrees C / 168 degrees F
Power
MPC-BIAS3V3-z12105        3292 mV
MPC-VDD3V3-z16100         3303 mV
MPC-VDD2V5-z16100         2501 mV
MPC-VDD1V8-z12004         1801 mV
MPC-AVDD1V0-z12006         996 mV
MPC-VDD1V2-z16100         1199 mV
MPC-VDD1V5A-z12004        1493 mV
MPC-VDD1V5B-z12004        1498 mV
MPC-XF_0V9-z12006         996 mV
MPC-PCIE_1V0-z16100       1000 mV
MPC-LU0_1V0-z12004        994 mV
MPC-LU1_1V0-z12004        994 mV
MPC-LU2_1V0-z12004        992 mV
MPC-LU3_1V0-z12004        993 mV
MPC-12VA-BMR453           12003 mV
MPC-12VB-BMR453           12043 mV
MPC-PMB_1V1-z12006        1091 mV
MPC-PMB_1V2-z12106        1196 mV
MPC-XM_0V9-vt273m         899 mV
I2C Slave Revision        106

```

#### show chassis environment fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host>show chassis environment fpc 1
FPC 1 status:
State                       Online
Temperature Intake          36 degrees C / 96 degrees F
Temperature Exhaust A      39 degrees C / 102 degrees F
Temperature LU TSen        52 degrees C / 125 degrees F
Temperature LU Chip        54 degrees C / 129 degrees F
Temperature XM TSen        52 degrees C / 125 degrees F
Temperature XM Chip        60 degrees C / 140 degrees F
Temperature PCIE TSen      52 degrees C / 125 degrees F
Temperature PCIE Chip      69 degrees C / 156 degrees F
Power
MPC-BIAS3V3-z12106        3302 mV
MPC-VDD3V3-z16100         3325 mV
MPC-AVDD1V0-z16100        1007 mV
MPC-PCIE_1V0-z16100        904 mV
MPC-LU0_1V0-z12004        996 mV
MPC-VDD_1V5-z12004        1498 mV
MPC-12VA-BMR453           11733 mV

```



```

MPC-12VB-BMR453      11728 mV
MPC-XM_0V9-vt273m    900 mV
I2C Slave Revision    81

```

### show chassis environment fpc (T320, T640, and T1600 Routers)

```

user@host> show chassis environment fpc
FPC 0 status:
  State                Online
  Temperature Top      42 degrees C / 107 degrees F
  Temperature Bottom   36 degrees C / 96 degrees F
  Temperature MMB1     39 degrees C / 102 degrees F
  Power:
    1.8 V              1959 mV
    2.5 V              2495 mV
    3.3 V              3344 mV
    5.0 V              5047 mV
    1.8 V bias         1787 mV
    3.3 V bias         3291 mV
    5.0 V bias         4998 mV
    8.0 V bias         7343 mV
  BUS Revision        40
FPC 1 status:
  State                Online
  Temperature Top      42 degrees C / 107 degrees F
  Temperature Bottom   39 degrees C / 102 degrees F
  Temperature MMB1     40 degrees C / 104 degrees F
  Power:
    1.8 V              1956 mV
    2.5 V              2498 mV
    3.3 V              3340 mV
    5.0 V              5023 mV
    1.8 V bias         1782 mV
    3.3 V bias         3277 mV
    5.0 V bias         4989 mV
    8.0 V bias         7289 mV
  BUS Revision        40
FPC 2 status:
  State                Online
  Temperature Top      43 degrees C / 109 degrees F
  Temperature Bottom   39 degrees C / 102 degrees F
  Temperature MMB1     41 degrees C / 105 degrees F
  Power:
    1.8 V              1963 mV
    2.5 V              2503 mV
    3.3 V              3340 mV
    5.0 V              5042 mV
    1.8 V bias         1797 mV
    3.3 V bias         3311 mV
    5.0 V bias         5013 mV
    8.0 V bias         7221 mV
  BUS Revision        40

```

### show chassis environment fpc (T4000 Router)

```

user@host> show chassis environment fpc
FPC 0 status:
  State                Online
  Fan Intake           34 degrees C / 93 degrees F
  Fan Exhaust          48 degrees C / 118 degrees F
  PMB                  47 degrees C / 116 degrees F

```

LMB0	50 degrees C / 122 degrees F
LMB1	41 degrees C / 105 degrees F
LMB2	35 degrees C / 95 degrees F
PFE1 LU2	46 degrees C / 114 degrees F
PFE1 LU0	41 degrees C / 105 degrees F
PFE0 LU0	57 degrees C / 134 degrees F
XF1	47 degrees C / 116 degrees F
XF0	52 degrees C / 125 degrees F
XM1	41 degrees C / 105 degrees F
XM0	50 degrees C / 122 degrees F
PFE0 LU1	56 degrees C / 132 degrees F
PFE0 LU2	45 degrees C / 113 degrees F
PFE1 LU1	37 degrees C / 98 degrees F
Power 1	
1.0 V	991 mV
1.2 V bias	1195 mV
1.8 V	1788 mV
2.5 V	2483 mV
3.3 V	3289 mV
3.3 V bias	3299 mV
12.0 V A	10608 mV
12.0 V B	10637 mV
Power 2	
0.9 V	881 mV
0.9 V PFE0	916 mV
0.9 V PFE1	903 mV
1.0 V PFE0	1012 mV
1.0 V PFE1	1002 mV
1.1 V	1095 mV
1.5 V_0	1494 mV
1.5 V_1	1479 mV
Power 3	
1.0 V PFE0	1000 mV
1.0 V PFE1	1002 mV
1.0 V PFE0 *	995 mV
1.0 V PFE1 *	995 mV
1.8 V PFE 0	1788 mV
1.8 V PFE 1	1789 mV
2.5 V	2482 mV
12.0 V	11614 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1003 mV
1.0 V PFE1 LU2	1004 mV
1.0 V PFE0 LU0 *	995 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	996 mV
12.0 V	11643 mV
12.0 V C	11711 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2488 mV
LMB0 VDD1V8	1788 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	1002 mV
LMB0 PFE0 LU0 VDD1V0	1000 mV
LMB0 VDD12V0	10752 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1792 mV
LMB1 VDD1V5	1480 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1002 mV

LMB1 VDD12V0	10800 mV
LMB2 VDD2V5	2472 mV
LMB2 VDD1V8	1792 mV
LMB2 VDD1V5	1486 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	998 mV
LMB2 VDD12V0	10704 mV
PMB 1.05v	1049 mV
PMB 1.5v	1500 mV
PMB 2.5v	2500 mV
PMB 3.3v	3299 mV
Bus Revision	113
FPC 3 status:	
State	Online
Fan Intake	37 degrees C / 98 degrees F
Fan Exhaust	51 degrees C / 123 degrees F
PMB	43 degrees C / 109 degrees F
LMB0	57 degrees C / 134 degrees F
LMB1	54 degrees C / 129 degrees F
LMB2	38 degrees C / 100 degrees F
PFE1 LU2	63 degrees C / 145 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	69 degrees C / 156 degrees F
XF1	62 degrees C / 143 degrees F
XF0	63 degrees C / 145 degrees F
XM1	43 degrees C / 109 degrees F
XM0	67 degrees C / 152 degrees F
PFE0 LU1	63 degrees C / 145 degrees F
PFE0 LU2	66 degrees C / 150 degrees F
PFE1 LU1	41 degrees C / 105 degrees F
Power 1	
1.0 V	1002 mV
1.2 V bias	1201 mV
1.8 V	1785 mV
2.5 V	2485 mV
3.3 V	3288 mV
3.3 V bias	3285 mV
12.0 V A	10412 mV
12.0 V B	10515 mV
Power 2	
0.9 V	882 mV
0.9 V PFE0	920 mV
0.9 V PFE1	905 mV
1.0 V PFE0	1015 mV
1.0 V PFE1	1001 mV
1.1 V	1094 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
0.92 V PFE1	998 mV
1.0 V PFE0	997 mV
1.0 V PFE0 *	992 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1780 mV
1.8 V PFE 1	1797 mV
2.5 V	2492 mV
12.0 V	11604 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1004 mV
1.0 V PFE1 LU2	1003 mV

```

1.0 V PFE0 LU0 *      1000 mV
1.0 V PFE1 LU0 *      1001 mV
1.0 V PFE1 LU2 *      1003 mV
12.0 V                 11653 mV
12.0 V C               11672 mV
Power (Base/PMB/MMB)
LMB0 VDD2V5           2512 mV
LMB0 VDD1V8           1790 mV
LMB0 VDD1V5           1500 mV
LMB0 PFE0 LU0 AVDD1V0 1004 mV
LMB0 PFE0 LU0 VDD1V0   1002 mV
LMB0 VDD12V0          10608 mV
LMB1 VDD2V5           2472 mV
LMB1 VDD1V8           1788 mV
LMB1 VDD1V5           1480 mV
LMB1 PFE0 LU2 AVDD1V0 1000 mV
LMB1 PFE0 LU2 VDD1V0   1004 mV
LMB1 VDD12V0          10672 mV
LMB2 VDD2V5           2488 mV
LMB2 VDD1V8           1798 mV
LMB2 VDD1V5           1494 mV
LMB2 PFE1 LU1 AVDD1V0 1000 mV
LMB2 PFE1 LU1 VDD1V0   1004 mV
LMB2 VDD12V0          10528 mV
PMB 1.05v             1050 mV
PMB 1.5v              1500 mV
PMB 2.5v              2499 mV
PMB 3.3v              3299 mV
Bus Revision           113
FPC 5 status:
State                  Online
Temperature Top        39 degrees C / 102 degrees F
Temperature Bottom     38 degrees C / 100 degrees F
Power
1.8 V                 1804 mV
1.8 V bias            1802 mV
3.3 V                 3294 mV
3.3 V bias            3277 mV
5.0 V bias            5008 mV
5.0 V TOP             5067 mV
8.0 V bias            6642 mV
Power (Base/PMB/MMB)
1.2 V                 1202 mV
1.5 V                 1504 mV
5.0 V BOT             5079 mV
12.0 V TOP Base       11848 mV
12.0 V BOT Base       11780 mV
1.1 V PMB             1111 mV
1.2 V PMB             1189 mV
1.5 V PMB             1494 mV
1.8 V PMB             1819 mV
2.5 V PMB             2503 mV
3.3 V PMB             3294 mV
5.0 V PMB             5035 mV
12.0 V PMB            11788 mV
0.75 MMB TOP          766 mV
1.5 V MMB TOP         1484 mV
1.8 V MMB TOP         1772 mV
2.5 V MMB TOP         2485 mV
1.2 V MMB TOP         1137 mV
5.0 V MMB TOP         4946 mV

```

12.0 V MMB TOP	11772 mV
3.3 V MMB TOP	3289 mV
0.75 MMB BOT	759 mV
1.5 V MMB BOT	1482 mV
1.8 V MMB BOT	1792 mV
2.5 V MMB BOT	2490 mV
1.2 V MMB BOT	1145 mV
5.0 V MMB BOT	4922 mV
12.0 V MMB BOT	11625 mV
3.3 V MMB BOT	3282 mV
APS 00	2495 mV
APS 01	3308 mV
APS 02	3301 mV
5.0 V PIC 0	4967 mV
APS 10	2512 mV
APS 11	3316 mV
APS 12	3304 mV
5.0 V PIC 1	5081 mV
Bus Revision	49
FPC 6 status:	
State	Online
Fan Intake	34 degrees C / 93 degrees F
Fan Exhaust	49 degrees C / 120 degrees F
PMB	40 degrees C / 104 degrees F
LMB0	60 degrees C / 140 degrees F
LMB1	58 degrees C / 136 degrees F
LMB2	40 degrees C / 104 degrees F
PFE1 LU2	69 degrees C / 156 degrees F
PFE1 LU0	45 degrees C / 113 degrees F
PFE0 LU0	71 degrees C / 159 degrees F
XF1	58 degrees C / 136 degrees F
XF0	65 degrees C / 149 degrees F
XM1	40 degrees C / 104 degrees F
XM0	66 degrees C / 150 degrees F
PFE0 LU1	69 degrees C / 156 degrees F
PFE0 LU2	68 degrees C / 154 degrees F
PFE1 LU1	42 degrees C / 107 degrees F
Power 1	
1.0 V	998 mV
1.2 V bias	1191 mV
1.8 V	1781 mV
2.5 V	2487 mV
3.3 V	3302 mV
3.3 V bias	3300 mV
12.0 V A	10388 mV
12.0 V B	10388 mV
Power 2	
0.9 V	902 mV
0.9 V PFE0	921 mV
0.9 V PFE1	907 mV
1.0 V PFE0	996 mV
1.0 V PFE1	974 mV
1.1 V	1095 mV
1.5 V_0	1495 mV
1.5 V_1	1478 mV
Power 3	
1.0 V PFE0	997 mV
1.0 V PFE1	998 mV
1.0 V PFE0 *	993 mV
1.0 V PFE1 *	991 mV
1.8 V PFE 0	1796 mV

1.8 V PFE 1	1789 mV
2.5 V	2465 mV
12.0 V	11609 mV
Power 4	
1.0 V PFE0 LU0	1003 mV
1.0 V PFE1 LU0	1006 mV
1.0 V PFE1 LU2	1002 mV
1.0 V PFE0 LU0 *	1000 mV
1.0 V PFE1 LU0 *	998 mV
1.0 V PFE1 LU2 *	998 mV
12.0 V	11638 mV
12.0 V C	11702 mV
Power (Base/PMB/MMB)	
LMB0 VDD2V5	2484 mV
LMB0 VDD1V8	1780 mV
LMB0 VDD1V5	1496 mV
LMB0 PFE0 LU0 AVDD1V0	998 mV
LMB0 PFE0 LU0 VDD1V0	1004 mV
LMB0 VDD12V0	10528 mV
LMB1 VDD2V5	2472 mV
LMB1 VDD1V8	1776 mV
LMB1 VDD1V5	1474 mV
LMB1 PFE0 LU2 AVDD1V0	994 mV
LMB1 PFE0 LU2 VDD1V0	1004 mV
LMB1 VDD12V0	10544 mV
LMB2 VDD2V5	2476 mV
LMB2 VDD1V8	1790 mV
LMB2 VDD1V5	1492 mV
LMB2 PFE1 LU1 AVDD1V0	996 mV
LMB2 PFE1 LU1 VDD1V0	1010 mV
LMB2 VDD12V0	10528 mV
PMB 1.05v	1050 mV
PMB 1.5v	1499 mV
PMB 2.5v	2500 mV
PMB 3.3v	3300 mV
Bus Revision	80

### show chassis environment fpc lcc (TX Matrix Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

#### FPC 1 status:

State	Online
Temperature Top	30 degrees C / 86 degrees F
Temperature Bottom	25 degrees C / 77 degrees F
Temperature MMB0	Absent
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1813 mV
2.5 V	2504 mV
3.3 V	3338 mV
5.0 V	5037 mV
1.8 V bias	1797 mV
3.3 V bias	3301 mV
5.0 V bias	5013 mV
8.0 V bias	7345 mV
BUS Revision	40

#### FPC 2 status:

State	Online
Temperature Top	37 degrees C / 98 degrees F

Temperature Bottom	26 degrees C / 78 degrees F
Temperature MMB0	32 degrees C / 89 degrees F
Temperature MMB1	27 degrees C / 80 degrees F
Power:	
1.8 V	1791 mV
2.5 V	2517 mV
3.3 V	3308 mV
5.0 V	5052 mV
1.8 V bias	1797 mV
3.3 V bias	3289 mV
5.0 V bias	4991 mV
8.0 V bias	7477 mV
BUS Revision	40

### show chassis environment fpc lcc (TX Matrix Plus Router)

```
user@host> show chassis environment fpc lcc 0
lcc0-re0:
```

```
-----
FPC 1 status:
State                               Online
Temperature Top                     46 degrees C / 114 degrees F
Temperature Bottom                   47 degrees C / 116 degrees F
Power
  1.8 V                             1788 mV
  1.8 V bias                         1787 mV
  3.3 V                             3321 mV
  3.3 V bias                         3306 mV
  5.0 V bias                         5018 mV
  5.0 V TOP                          5037 mV
  8.0 V bias                         7223 mV
Power (Base/PMB/MMB)
  1.2 V                             1205 mV
  1.5 V                             1503 mV
  5.0 V BOT                          5084 mV
  12.0 V TOP Base                    11775 mV
  12.0 V BOT Base                    11794 mV
  1.1 V PMB                          1108 mV
  1.2 V PMB                          1196 mV
  1.5 V PMB                          1499 mV
  1.8 V PMB                          1811 mV
  2.5 V PMB                          2515 mV
  3.3 V PMB                          3318 mV
  5.0 V PMB                          5030 mV
  12.0 V PMB                         11832 mV
  0.75 MMB TOP                       752 mV
  1.5 V MMB TOP                      1489 mV
  1.8 V MMB TOP                      1782 mV
  2.5 V MMB TOP                      2498 mV
  1.2 V MMB TOP                      1155 mV
  5.0 V MMB TOP                      4902 mV
  12.0 V MMB TOP                     11721 mV
  3.3 V MMB TOP                      3316 mV
  0.75 MMB BOT                       754 mV
  1.5 V MMB BOT                      1482 mV
  1.8 V MMB BOT                      1758 mV
  2.5 V MMB BOT                      2488 mV
  1.2 V MMB BOT                      1157 mV
  5.0 V MMB BOT                      4962 mV
  12.0 V MMB BOT                     11691 mV
  3.3 V MMB BOT                      3308 mV
```

APS 00	1484 mV
APS 01	2503 mV
APS 02	3313 mV
5.0 V PIC 0	5025 mV
APS 10	1501 mV
APS 11	2466 mV
APS 12	3311 mV
5.0 V PIC 1	5081 mV
Bus Revision	49

**show chassis environment fpc (QFX Series)**

```
user@switch> show chassis environment fpc 0
FPC 0 status:
  State                Online
  Temperature          42 degrees C / 107 degrees F
```

**show chassis environment fpc interconnect-device (QFabric Systems)**

```
user@switch> show chassis environment fpc interconnect-device interconnect1 0
FC 0 FPC 0 status:
  State                Online
  Left Intake Temperature 24 degrees C / 75 degrees F
  Right Intake Temperature 24 degrees C / 75 degrees F
  Left Exhaust Temperature 27 degrees C / 80 degrees F
  Right Exhaust Temperature 27 degrees C / 80 degrees F
  Power
    BIAS 3V3            3330 mV
    VDD 3V3             3300 mV
    VDD 2V5             2502 mV
    VDD 1V5             1496 mV
    VDD 1V2             1194 mV
    VDD 1V0             1000 mV
    SW0 VDD 1V0         1020 mV
    SW0 CVDD 1V025      1032 mV
    SW1 VDD 1V0         1022 mV
    SW1 CVDD 1V025      1030 mV
    VDD 12V0 DIV3_33    3414 mV
```

**show chassis environment fpc 0 (PTX5000 Packet Transport Router)**

```
user@host> show chassis environment fpc 0
FPC 0 status:
  State                Online
  PMB Temperature      35 degrees C / 95 degrees F
  Intake Temperature   33 degrees C / 91 degrees F
  Exhaust A Temperature 51 degrees C / 123 degrees F
  Exhaust B Temperature 43 degrees C / 109 degrees F
  TL0 Temperature      48 degrees C / 118 degrees F
  TQ0 Temperature      53 degrees C / 127 degrees F
  TL1 Temperature      56 degrees C / 132 degrees F
  TQ1 Temperature      58 degrees C / 136 degrees F
  TL2 Temperature      55 degrees C / 131 degrees F
  TQ2 Temperature      57 degrees C / 134 degrees F
  TL3 Temperature      59 degrees C / 138 degrees F
  TQ3 Temperature      59 degrees C / 138 degrees F
  Power
    PMB 1.05v          1049 mV
    PMB 1.5v           1500 mV
    PMB 2.5v           2500 mV
    PMB 3.3v           3299 mV
    PFE0 1.5v          1500 mV
```



PFE0	1.0v	999 mV
TQ0	0.9v	900 mV
TL0	0.9v	900 mV
PFE1	1.5v	1499 mV
PFE1	1.0v	999 mV
TQ1	0.9v	899 mV
TL1	0.9v	900 mV
PFE2	1.5v	1500 mV
PFE2	1.0v	1000 mV
TQ2	0.9v	900 mV
TL2	0.9v	900 mV
PFE3	1.5v	1499 mV
PFE3	1.0v	1000 mV
TQ3	0.9v	900 mV
TL3	0.9v	900 mV
Bias	3.3v	3327 mV
FPC	3.3v	3300 mV
FPC	2.5v	2500 mV
SAM	0.9v	900 mV
A	12.0v	2014 mV
B	12.0v	2030 mV

#### show chassis environment fpc 07 (PTX5000 Packet Transport Router with FPC2-PTX-PIA)

```
user@host> show chassis environment fpc 07
```

```
FPC 7 status:
```

State	Online
PMB TEMP0 Temperature	32 degrees C / 89 degrees F
PMB TEMP1 Temperature	28 degrees C / 82 degrees F
PMB CPU Temperature	46 degrees C / 114 degrees F
Intake Temperature	35 degrees C / 95 degrees F
Exhaust A Temperature	55 degrees C / 131 degrees F
Exhaust B Temperature	54 degrees C / 129 degrees F
TL5 Temperature	59 degrees C / 138 degrees F
TQ5 Temperature	57 degrees C / 134 degrees F
TL6 Temperature	57 degrees C / 134 degrees F
TQ6 Temperature	51 degrees C / 123 degrees F
TL1 Temperature	76 degrees C / 168 degrees F
TQ1 Temperature	58 degrees C / 136 degrees F
TL2 Temperature	75 degrees C / 167 degrees F
TQ2 Temperature	57 degrees C / 134 degrees F
TL4 Temperature	52 degrees C / 125 degrees F
TQ4 Temperature	66 degrees C / 150 degrees F
TL7 Temperature	52 degrees C / 125 degrees F
TQ7 Temperature	60 degrees C / 140 degrees F
TL0 Temperature	72 degrees C / 161 degrees F
TQ0 Temperature	73 degrees C / 163 degrees F
TL3 Temperature	64 degrees C / 147 degrees F
TQ3 Temperature	70 degrees C / 158 degrees F
Power	
PMB 1.05v	1049 mV
PMB 3.3v	3299 mV
PMB 1.1v-a	1100 mV
PMB 1.5v	1499 mV
PMB 1.1v-b	1100 mV
Base 3.3v	3300 mV
FPC Base 2.5v	2499 mV
TL1 0.9v	897 mV
TQ1 0.9v	897 mV
PFE1 1.0v	999 mV
PFE1 1.5v	1499 mV

TL2	0.9v	897 mV
TQ2	0.9v	897 mV
PFE2	1.0v	999 mV
PFE2	1.5v	1499 mV
FPC Base	1.0v	1000 mV
FPC Base	1.2v	1199 mV
TL5	0.9v	898 mV
TQ5	0.9v	898 mV
PFE5	1.0v	1000 mV
PFE5	1.5v	1500 mV
TL6	0.9v	897 mV
TQ6	0.9v	897 mV
PFE6	1.0v	1000 mV
PFE6	1.5v	1499 mV
Mezz Base	2.5v	2500 mV
TL0	0.9v	896 mV
TQ0	0.9v	896 mV
PFE0	1.0v	999 mV
PFE0	1.5v	1499 mV

#### show chassis environment FPC 1 (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis environment fpc 1
```

```
FPC 1 status:
```

State	Online
Temperature Intake	36 degrees C / 96 degrees F
Temperature Exhaust A	39 degrees C / 102 degrees F
Temperature LU TSen	52 degrees C / 125 degrees F
Temperature LU Chip	54 degrees C / 129 degrees F
Temperature XM TSen	52 degrees C / 125 degrees F
Temperature XM Chip	60 degrees C / 140 degrees F
Temperature PCIE TSen	52 degrees C / 125 degrees F
Temperature PCIE Chip	69 degrees C / 156 degrees F
Power	
MPC-BIAS3V3-z12106	3302 mV
MPC-VDD3V3-z16100	3325 mV
MPC-AVDD1V0-z16100	1007 mV
MPC-PCIE_1V0-z16100	904 mV
MPC-LU0_1V0-z12004	996 mV
MPC-VDD_1V5-z12004	1498 mV
MPC-12VA-BMR453	11733 mV
MPC-12VB-BMR453	11728 mV
MPC-XM_0V9-vt273m	900 mV
I2C Slave Revision	81

## show chassis environment pem

<b>List of Syntax</b>	<a href="#">Syntax on page 601</a> <a href="#">Syntax (ACX4000 Router) on page 601</a> <a href="#">Syntax (TX Matrix Routers) on page 601</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 601</a> <a href="#">Syntax (MX Series Router) on page 601</a> <a href="#">Syntax (MX104 3D Universal Edge Routers) on page 601</a> <a href="#">Syntax (QFX Series) on page 601</a>
<b>Syntax</b>	show chassis environment pem <slot>
<b>Syntax (ACX4000 Router)</b>	show chassis environment pem
<b>Syntax (TX Matrix Routers)</b>	show chassis environment pem <lcc number   scc> <slot>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis environment pem <lcc number   sfc number> <slot>
<b>Syntax (MX Series Router)</b>	show chassis environment pem <slot> <all-members> <local> <member member-id>
<b>Syntax (MX104 3D Universal Edge Routers)</b>	show chassis environment pem <slot>
<b>Syntax (QFX Series)</b>	show chassis environment pem <slot (interconnect-device name slot )   (node-device name)>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS 11.3 for the QFX Series. Command introduced in Junos OS 12.3R2 for EX Series. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	Display Power Entry Module (PEM) environmental status information.



**NOTE:** The new high-capacity (4100W) enhanced DC PEM on MX960 routers includes a new design that can condition the input voltage. This results in the output voltage differing from the input voltage. The earlier generation of DC PEMs coupled the input power directly to the output, thereby making it safe to assume that the output voltage was equal to the input voltage.

- Options**    **none**—Display environmental information about both PEMs. For the TX Matrix router, display environmental information about the PEMs, the TX Matrix router, and its attached T640 routers. For the TX Matrix Plus router, display environmental information about the PEMs, the TX Matrix Plus router, and its attached routers.
- all-members**—(MX Series routers only) (Optional) Display environmental information about the PEMs in all the member routers of the Virtual Chassis configuration.
- interconnect-device *name***—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Interconnect device.
- lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.  
              Replace *number* with the following values depending on the LCC configuration:
- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
  - 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
  - 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
  - 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- local**—(MX Series routers only) (Optional) Display environmental information about the PEM in the local Virtual Chassis member.
- member *member-id***—(MX Series routers only) (Optional) Display environmental information about the PEM in the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.
- node-device *name***—(QFabric systems only) (Optional) Display chassis environmental information about the PEMs in the Node device.
- scc**—(TX Matrix routers only) (Optional) Display environmental information about the PEM in the TX Matrix router (switch-card chassis).
- sfc**—(TX Matrix Plus routers only) (Optional) Display environmental information about the PEM in the TX Matrix Plus router (or switch-fabric chassis).
- slot** —(Optional) Display environmental information about an individual PEM. Replace *slot* with 0 or 1.

**Required Privilege Level**    view

**Related Documentation**    • [show chassis hardware on page 676](#)

**List of Sample Output**    [show chassis environment pem \(M40e Router\) on page 604](#)

[show chassis environment pem \(M120 Router\) on page 604](#)  
[show chassis environment pem \(M160 Router\) on page 604](#)  
[show chassis environment pem \(M320 Router\) on page 605](#)  
[show chassis environment pem \(MX104 Router\) on page 605](#)  
[show chassis environment pem \(MX240 Router\) on page 605](#)  
[show chassis environment pem \(MX480 Router\) on page 605](#)  
[show chassis environment pem \(MX960 Router\) on page 606](#)  
[show chassis environment pem \(T320 Router\) on page 606](#)  
[show chassis environment pem \(T640 Router\) on page 606](#)  
[show chassis environment pem \(T4000 Router\) on page 606](#)  
[show chassis environment pem \(T640/T1600/T4000 Routers With Six-Input DC Power Supply\) on page 607](#)  
[show chassis environment pem lcc \(TX Matrix Routing Matrix\) on page 607](#)  
[show chassis environment pem scc \(TX Matrix Routing Matrix\) on page 607](#)  
[show chassis environment pem sfc \(TX Matrix Plus Routing Matrix\) on page 608](#)  
[show chassis environment pem lcc \(TX Matrix Plus Routing Matrix\) on page 608](#)  
[show chassis environment pem node-device \(QFabric System\) on page 608](#)  
[show chassis environment pem \(QFX Series\) on page 609](#)  
[show chassis environment pem interconnect-device \(QFabric System\) on page 609](#)

**Output Fields** Table 29 on page 603 lists the output fields for the **show chassis environment pem** command. Output fields are listed in the approximate order in which they appear.

**Table 29: show chassis environment pem Output Fields**

Field Name	Field Description
<b>PEM slot status</b>	Number of the PEM slot.
<b>State</b>	Status of the PEM.
<b>Temperature</b>	Temperature of the air flowing past the PEM.
<b>AC Input</b>	Status of the AC input for the specified component
<b>AC Output</b>	Status of the AC output for the specified component.
<b>DC input</b>	Status of the DC input for the specified component.
<b>DC output</b>	Status of the DC output for the specified component.
<b>Load</b>	(Not available on M40e or M160 routers) Information about the load on supply, in percentage of rated current being used.
<b>Voltage</b>	(M120, M160, M320, T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about voltage supplied to the PEM.  (MX104 routers only) Information about voltage supplied by the PEM to the system.
<b>Current</b>	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM current.
<b>Power</b>	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) Information about the PEM power.

Table 29: show chassis environment pem Output Fields (*continued*)

Field Name	Field Description
<b>SCG/CB/SIB</b>	(T640, T1600, TX Matrix, and TX Matrix Plus routers only) SONET Clock Generator/Control Board/Switch Interface Board.
<b>FAN</b>	(T640, T1600, and T4000 routers with six-input DC power supply only) Information about the DC output to the fan.

## Sample Output

### show chassis environment pem (M40e Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  AC input              OK
  DC output             OK
```

### show chassis environment pem (M120 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       52864 mV
    48.0 V fan supply  41655 mV
    3.3 V              3399 mV
PEM 1 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       54537 mV
    48.0 V fan supply  42910 mV
    3.3 V              3506 mV
```

### show chassis environment pem (M160 Router)

```
user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC input              OK
  DC output             OK
  Load                Less than 20 percent
  Voltage:
    48.0 V input       54833 mV
    48.0 V fan supply  50549 mV
    8.0 V bias         8239 mV
    5.0 V bias         5006 mV
```

**show chassis environment pem (M320 Router)**

```

user@host> show chassis environment pem
PEM 2 status:
  State                Online
  Temperature           OK
  DC input              OK
  Load                 Less than 40 percent
    48.0 V input        51853 mV
    48.0 V fan supply   48877 mV
    8.0 V bias          8449 mV
    5.0 V bias          4998 mV
PEM 3 status:
  State                Online
  Temperature           OK
  DC input              OK
  Load                 Less than 40 percent
    48.0 V input        51717 mV
    48.0 V fan supply   49076 mV
    8.0 V bias          8442 mV
    5.0 V bias          4998 mV

```

**show chassis environment pem (MX104 Router)**

```

user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC Output:           OK
  Voltage:
    12.0 V output       12281 mV
    3.3 V output        3353 mV
PEM 1 status:
  State                Empty

```

**show chassis environment pem (MX240 Router)**

```

user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC Output:           OK
PEM 1 status:
  State                Online
  Temperature           OK
  DC Output:           OK

```

**show chassis environment pem (MX480 Router)**

```

user@host> show chassis environment pem
PEM 0 status:
  State                Online
  Temperature           OK
  DC Input:            OK
  DC Output:           OK
  Voltage:
PEM 1 status:
  State                Online
  Temperature           OK
  DC Input:            OK

```

```

DC Output:          OK
Voltage:

```

#### show chassis environment pem (MX960 Router)

```

user@host> show chassis environment pem
PEM 2 status:
  State          Present
PEM 3 status:
  State          Online
  Temperature     OK
  DC Output:      OK

```

#### show chassis environment pem (T320 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State          Online
  Temperature     OK
  DC input:       OK

```

#### show chassis environment pem (T640 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State          Online
  Temperature     22 degrees C / 71 degrees F
  AC input:       OK
  DC output:
    Voltage      Current      Power      Load
    FPC 0        56875        606        34        4
    FPC 1        57016        525        29        3
    FPC 2         0         0         0         0
    FPC 3         0         0         0         0
    FPC 4         0         0         0         0
    FPC 5         0         0         0         0
    FPC 6        57158       1581        90       12
    FPC 7         0         0         0         0
  SCG/CB/SIB     56750       1125        63        5

```

#### show chassis environment pem (T4000 Router)

```

user@host> show chassis environment pem
PEM 0 status:
  State          Online
  Temperature     33 degrees C / 91 degrees F
  DC Input:       OK
    Voltage(V)    Current(A)    Power(W)    Load(%)
  INPUT 0        54.625      9.812      535        22
  INPUT 1        54.625     10.250     559        23
  INPUT 2        55.125      0.125        6         0
  INPUT 3        54.500     10.062     548        22
  INPUT 4        54.750      9.375     513        21
  INPUT 5        54.750     10.187     557        23
  DC Output      Voltage(V)    Current(A)    Power(W)    Load(%)
  FPC 0         55.750     10.125     564        37
  FPC 1         51.625      0.000        0         0
  FPC 2         52.000      0.000        0         0
  FPC 3         55.062     10.437     574        38
  FPC 4         52.125      0.000        0         0
  FPC 5         55.000      9.375     515        34
  FPC 6         55.187      9.687     534        35
  FPC 7         51.437      0.000        0         0

```



SCG/CB/SIB	55.375	15.750	872	35
FAN	54.562	14.750	804	42

### show chassis environment pem (T640/T1600/T4000 Routers With Six-Input DC Power Supply)

```
user@host> show chassis environment pem
PEM 1 status:
State                Online
Temperature          36 degrees C / 96 degrees F
DC Input:            OK

```

	Voltage(V)	Current(A)	Power(W)	Load(%)
INPUT 0	0.000	0.000	0	0
INPUT 1	54.875	3.812	209	27
INPUT 2	55.375	3.937	218	29
INPUT 3	54.625	3.750	204	27
INPUT 4	55.125	3.375	186	24
INPUT 5	55.125	3.375	186	24

```
DC Output
```

	Voltage(V)	Current(A)	Power(W)	Load(%)
FPC 0	52.312	0.000	0	0
FPC 1	52.687	0.000	0	0
FPC 2	52.812	0.000	0	0
FPC 3	55.812	7.062	394	52
FPC 4	52.625	0.000	0	0
FPC 5	52.625	0.000	0	0
FPC 6	52.750	0.000	0	0
FPC 7	52.750	0.000	0	0
SCG/CB/SIB	55.937	11.937	667	55
FAN	55.812	4.937	275	36

### show chassis environment pem lcc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem 0 lcc 0
lcc0-re0:
-----
PEM 0 status:
State                Present
Temperature          27 degrees C / 80 degrees F
DC input:            Check
DC output:           Voltage Current Power Load
FPC 0                0      0      0      0
FPC 1                0      0      0      0
FPC 2                0      0      0      0
FPC 3                0      0      0      0
FPC 4                0      0      0      0
FPC 5                0      0      0      0
FPC 6                0      0      0      0
FPC 7                0      0      0      0
SCG/CB/SIB           0      0      0      0
```

### show chassis environment pem scc (TX Matrix Routing Matrix)

```
user@host> show chassis environment pem scc
scc-re0:
-----
PEM 1 status:
State                Online
Temperature          24 degrees C / 75 degrees F
DC input:            OK
DC output:           Voltage Current Power Load
SIB 0                0      0      0      0
SIB 1                0      0      0      0
SIB 2                0      0      0      0
```

SIB 3	56550	0	0	0
SIB 4	55958	6912	386	51

### show chassis environment pem sfc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment pem sfc 0
sfc0-re0:
```

```
-----
PEM 0 status:
State                Online
Temperature          35 degrees C / 95 degrees F
DC Input:            OK
DC Output             Voltage    Current    Power    Load
Channel 0            53820    14140    761      59
Channel 1            53550    12720    681      53
Channel 2            53840    12930    696      54
Channel 3            53690    14990    804      63
Channel 4            53620    15070    808      63
Channel 5            53900    14820    798      62
Channel 6            54120    5020     271      21
```

### show chassis environment pem lcc (TX Matrix Plus Routing Matrix)

```
user@host> show chassis environment lcc 0
```

```
lcc0-re1:
```

```
-----
PEM 0 status:
State                Online
Temperature          38 degrees C / 100 degrees F
DC Input:            OK
DC Output             Voltage    Current    Power    Load
FPC 0                0         0         0         0
FPC 1                0         0         0         0
FPC 2                0         0         0         0
FPC 3                0         0         0         0
FPC 4                56408    7575     427      56
FPC 5                0         0         0         0
FPC 6                56266    7956     447      59
FPC 7                56283    6100     343      45
SCG/CB/SIB           55916    8950     500      41

PEM 1 status:
State                Present
Temperature          35 degrees C / 95 degrees F
DC Input:            Check
DC Output             Voltage    Current    Power    Load
FPC 0                0         0         0         0
FPC 1                0         0         0         0
FPC 2                0         0         0         0
FPC 3                0         0         0         0
FPC 4                0         0         0         0
FPC 5                0         0         0         0
FPC 6                0         0         0         0
FPC 7                0         0         0         0
SCG/CB/SIB           0         0         0         0
```

### show chassis environment pem node-device (QFabric System)

```
user@switch> show chassis environment pem node-device node1
FPC 0 PEM 0 status:
```

State	Check
Airflow	Front to Back

```

Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                      12         10         120      18
FPC 0 PEM 1 status:
State                Online
Airflow              Back to Front
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                      11         10         110      17

```

#### show chassis environment pem (QFX Series)

```

user@switch> show chassis environment pem
FPC 0 PEM 1 status:
State                Online
Airflow              Front to Back
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                      12         17         204      31

```

#### show chassis environment pem interconnect-device (QFabric System)

```

user@switch> show chassis environment pem interconnect-device IC11
IC1 PEM 1 status:
State                Online
Airflow              Front to Back
Temperature          OK
AC Input:            OK
DC Output            Voltage(V) Current(A) Power(W) Load(%)
                      12         18         216      33

```

## show chassis environment routing-engine

---

<b>List of Syntax</b>	<a href="#">Syntax on page 610</a> <a href="#">Syntax (TX Matrix Routers) on page 610</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 610</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 610</a> <a href="#">Syntax (MX Series Routers) on page 610</a> <a href="#">Syntax (QFX Series) on page 610</a>
<b>Syntax</b>	show chassis environment routing-engine <slot>
<b>Syntax (TX Matrix Routers)</b>	show chassis environment routing-engine <lcc number   scc> <slot>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis environment routing-engine <lcc number   sfc number> <slot>
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis environment routing-engine <slot>
<b>Syntax (MX Series Routers)</b>	show chassis environment routing-engine <slot> <all-members> <local> <member member-id>
<b>Syntax (QFX Series)</b>	show chassis environment routing-engine interconnect-device <i>name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers. Command introduced in Junos OS Release 12.1 for the T4000 Core Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	Display Routing Engine environmental status information.
<b>Options</b>	<b>none</b> —Display environmental information about all Routing Engines. For a TX Matrix router, display environmental information about all Routing Engines on the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display

environmental information about all Routing Engines on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in all member routers in the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display environmental information about the Routing Engines for the Interconnect device.

**lcc *number***—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display environmental information about the Routing Engines in the specified member in the Virtual Chassis configuration. Replace *member-id* with the value of 0 or 1.

**scc**—(TX Matrix router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix router (switch-card chassis).

**sfc**—(TX Matrix Plus router only) (Optional) Display environmental information about the Routing Engine in the TX Matrix Plus router (or switch-fabric chassis).

**slot**—(Optional) Display environmental information about an individual Routing Engine. On M10i, M20, M40e, M120, M160, M320, MX Series, MX104 routers, MX2010 routers, MX2020 routers, and T Series routers, replace *slot* with **0** or **1**. On M5, M7i, M10, and M40 routers and on the J Series router, replace *slot* with **0**. On EX3200 and EX4200 standalone switches, replace *slot* with **0**. On EX4200 switches in a Virtual Chassis configuration and on EX8208 and EX8216 switches, replace *slot* with **0** or **1**. On the QFX3500 switch, there is only one Routing Engine, so you do not need to specify the slot number. On PTX Series Packet Transport Routers, replace *slot* with **0** or **1**.

**Required Privilege Level** view

**Related Documentation**

- [request chassis routing-engine master on page 392](#)
- [show chassis routing-engine on page 905](#)

**List of Sample Output**

- [show chassis environment routing-engine \(Nonredundant\) on page 612](#)
- [show chassis environment routing-engine \(Redundant\) on page 612](#)
- [show chassis environment routing-engine \(MX104 Router\) on page 612](#)
- [show chassis environment routing-engine \(MX2010 Router\) on page 613](#)
- [show chassis environment routing-engine \(MX2020 Router\) on page 613](#)
- [show chassis environment routing-engine \(TX Matrix Plus Router\) on page 613](#)
- [show chassis environment routing-engine \(T4000 Core Router\) on page 613](#)
- [show chassis environment routing-engine \(QFX Series\) on page 614](#)
- [show chassis environment routing-engine interconnect-device \(QFabric System\) on page 614](#)
- [show chassis environment routing-engine \(PTX5000 Packet Transport Router\) on page 614](#)

**Output Fields** [Table 30 on page 612](#) lists the output fields for the **show chassis environment routing-engine** command. Output fields are listed in the approximate order in which they appear.

**Table 30: show chassis environment routing-engine Output Fields**

Field Name	Field Description
<b>Routing engine slot status</b>	Number of the Routing Engine slot: 0 or 1.
<b>State</b>	Status of the Routing Engine: <ul style="list-style-type: none"> <li>• <b>Online Master</b>—Routing Engine is online, operating as Master.</li> <li>• <b>Online Standby</b>—Routing Engine is online, operating as Standby.</li> <li>• <b>Offline</b>—Routing Engine is offline.</li> </ul>
<b>Temperature</b>	Temperature of the air flowing past the Routing Engine.
<b>CPU Temperature</b>	(PTX Series and T4000 Core Routers only) Temperature of the air flowing past the Routing Engine CPU.

## Sample Output

### show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                Online Master
  Temperature          27 degrees C / 80 degrees
```

### show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State                Online Master
  Temperature          26 degrees C / 78 degrees F
Route Engine 1 status:
  State                Online Standby
  Temperature          26 degrees C / 78 degrees F
```

### show chassis environment routing-engine (MX104 Router)

```
user@ host >show chassis environment routing-engine
```

```

Routing Engine 0 status:
  State           Online Master
  Temperature      34 degrees C / 93 degrees F
  CPU Temperature  43 degrees C / 109 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  39 degrees C / 102 degrees F

```

#### show chassis environment routing-engine (MX2010 Router)

```

user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      37 degrees C / 98 degrees F
  CPU Temperature  37 degrees C / 98 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      35 degrees C / 95 degrees F
  CPU Temperature  34 degrees C / 93 degrees F

```

#### show chassis environment routing-engine (MX2020 Router)

```

user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      35 degrees C / 95 degrees F
  CPU Temperature  34 degrees C / 93 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      44 degrees C / 111 degrees F
  CPU Temperature  43 degrees C / 109 degrees F

```

#### show chassis environment routing-engine (TX Matrix Plus Router)

```

user@host> show chassis environment routing-engine
sfc0-re0:
-----
Routing Engine 0 status:
  State           Online Master
  Temperature      26 degrees C / 78 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      28 degrees C / 82 degrees F

lcc0-re0:
-----
Routing Engine 0 status:
  State           Online Master
  Temperature      30 degrees C / 86 degrees F
Routing Engine 1 status:
  State           Online Standby
  Temperature      29 degrees C / 84 degrees F

```

#### show chassis environment routing-engine (T4000 Core Router)

```

user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State           Online Master
  Temperature      33 degrees C / 91 degrees F
  CPU Temperature  50 degrees C / 122 degrees F
Routing Engine 1 status:

```

State	Online Standby
Temperature	33 degrees C / 91 degrees F
CPU Temperature	46 degrees C / 114 degrees F

#### show chassis environment routing-engine (QFX Series)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State      Online Master
  Temperature 42 degrees C / 107 degrees F
```

#### show chassis environment routing-engine interconnect-device (QFabric System)

```
user@switch> show chassis environment routing-engine interconnect-device interconnect1
routing-engine interconnect-device interconnect1
Routing Engine 0 status:
  State      Online Standby
  Temperature 52 degrees C / 125 degrees F
Routing Engine 1 status:
  State      Online Master
  Temperature 57 degrees C / 134 degrees F
```

#### show chassis environment routing-engine (PTX5000 Packet Transport Router)

```
user@switch> show chassis environment routing-engine
Routing Engine 0 status:
  State      Online Master
  Temperature 55 degrees C / 131 degrees F
  CPU Temperature 66 degrees C / 150 degrees F
Routing Engine 1 status:
  State      Online Standby
  Temperature 52 degrees C / 125 degrees F
  CPU Temperature 64 degrees C / 147 degrees F
```



## show chassis fan

<b>List of Syntax</b>	<a href="#">Syntax on page 615</a> <a href="#">Syntax (ACX4000 Series Router) on page 615</a> <a href="#">Syntax (MX Series Router) on page 615</a> <a href="#">Syntax (T Series Routers) on page 615</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Router) on page 615</a> <a href="#">Syntax (QFabric Systems) on page 615</a> <a href="#">Syntax (TX Matrix Router) on page 615</a> <a href="#">Syntax (TX Matrix Plus Router) on page 615</a>
<b>Syntax</b>	show chassis fan
<b>Syntax (ACX4000 Series Router)</b>	show chassis fan
<b>Syntax (MX Series Router)</b>	show chassis fan <all-members> <local> <member <i>member-id</i> >
<b>Syntax (T Series Routers)</b>	show chassis fan
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Router)</b>	show chassis fan
<b>Syntax (QFabric Systems)</b>	show chassis fan <interconnect-device <i>name</i> >
<b>Syntax (TX Matrix Router)</b>	show chassis fan <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis fan <lcc <i>number</i>   sfc <i>number</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 on MX Series 3D Universal Edge Routers, M120 routers, and M320 routers, T320 routers, T640 routers, T1600 routers, TX Matrix Routers, and TX Matrix Plus routers. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.4 for EX Series switches. Command introduced in Junos OS Release 12.3 for PTX5000 Packet Transport Routers. Command introduced in Junos OS Release 12.1 for T4000 routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for ACX Series Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	(T Series routers, TX Matrix routers, TX Matrix Plus routers, M120 routers, M320 routers, MX104 routers, MX2010 routers, MX2020 routers, MX Series 3D Universal Edge Routers,

QFX3008-I Interconnect devices, EX Series switches, and PTX Series Packet Transport Routers only) Show information about the fan tray and fans.

**Options**    **all-members**—(MX Series routers only) (Optional) Display information about the fan tray and fans for all members of the Virtual Chassis configuration.

**local**—(MX Series routers only) (Optional) Display information about the fan tray and fans for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display information about the fan tray and fans for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* variable with a value 0 or 1.

**interconnect-device *name***—(QFX3000-G QFabric systems only) (Optional) Display information about the fan tray and fans for the specified QFX3008-I Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display information about the fan tray and fans for the specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display information about the fan tray and fans for the specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**scc**—(TX Matrix routers only) (Optional) Display information about the fan tray and fans for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display information about the fan tray and fans for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

**Required Privilege Level**

view

**List of Sample Output**

[show chassis fan on page 617](#)  
[show chassis fan \(QFabric Systems\) on page 618](#)  
[show chassis fan \(EX Series Switches\) on page 619](#)  
[show chassis fan \(T320 Router\) on page 619](#)  
[show chassis fan \(T640 Router\) on page 620](#)  
[show chassis fan \(T1600 Router\) on page 620](#)

[show chassis fan \(T4000 Core Router\) on page 621](#)  
[show chassis fan \(TX Matrix Router\) on page 621](#)  
[show chassis fan \(TX Matrix Plus Router\) on page 622](#)  
[show chassis fan \(TX Matrix Plus Router with 3D SIBs\) on page 623](#)  
[show chassis fan \(PTX5000 Packet Transport Router\) on page 625](#)  
[show chassis fan \(MX104 Router\) on page 626](#)  
[show chassis fan \(MX2010 Router\) on page 626](#)  
[show chassis fan \(MX2020 Router\) on page 626](#)  
[show chassis fan \(ACX4000 Router\) on page 627](#)  
[show chassis fan \(QFX5100 Switch\) on page 627](#)

**Output Fields** Table 31 on page 617 lists the output fields for the **show chassis fan** command. Output fields are listed in the approximate order in which they appear.

**Table 31: show chassis fan Output Fields**

Field Name	Field Description
<b>Item</b>	Fan item identifier.
<b>Status</b>	Status of the fan: <ul style="list-style-type: none"> <li>• <b>OK</b>—Fan is running properly and within the normal range.</li> <li>• <b>Check</b>—Fan is in <b>Check</b> state because of some fault or alarm condition.</li> </ul>
<b>RPM</b>	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed in revolutions per minute (RPM).
<b>% RPM</b>	(MX2010 routers, MX2020 routers, and PTX Series Packet Transport Routers only) Percentage of the fan speed being used.
<b>Measurement</b>	(T Series routers, TX Matrix routers, TX Matrix Plus routers, MX Series 3D Universal Edge Routers, QFX3108 Interconnect devices, and EX Series switches only) Fan speed status based on different chassis cooling requirements: <ul style="list-style-type: none"> <li>• Spinning at high speed</li> <li>• Spinning at intermediate speed</li> <li>• Spinning at normal speed</li> <li>• Spinning at low speed (except EX Series switches)</li> </ul> (MX2010 routers, MX2020 routers, and PTX Series Packet Transport Routers only) Fan speed in revolutions per minute (RPM) for each fan in the fan tray.

## Sample Output

show chassis fan

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
------	--------	-----	-------------

Top Tray Fan 1	OK	3790	Spinning at normal speed
Top Tray Fan 2	OK	3769	Spinning at normal speed
Top Tray Fan 3	OK	3769	Spinning at normal speed
Top Tray Fan 4	OK	3790	Spinning at normal speed
Top Tray Fan 5	OK	3790	Spinning at normal speed
Top Tray Fan 6	OK	3769	Spinning at normal speed
Top Tray Fan 7	OK	3790	Spinning at normal speed
Top Tray Fan 8	OK	3769	Spinning at normal speed
Top Tray Fan 9	OK	3769	Spinning at normal speed
Top Tray Fan 10	OK	3790	Spinning at normal speed
Top Tray Fan 11	OK	3790	Spinning at normal speed
Top Tray Fan 12	OK	3769	Spinning at normal speed
Bottom Tray Fan 1	OK	2880	Spinning at normal speed
Bottom Tray Fan 2	OK	2912	Spinning at normal speed
Bottom Tray Fan 3	OK	2928	Spinning at normal speed
Bottom Tray Fan 4	OK	2896	Spinning at normal speed
Bottom Tray Fan 5	OK	2896	Spinning at normal speed
Bottom Tray Fan 6	OK	2928	Spinning at normal speed

### show chassis fan (QFabric Systems)

```
user@host> show chassis fan interconnect-device interconnect1
```

Item	Status	RPM	Measurement
TFT 0 Fan 0	OK	2849	Spinning at normal speed
TFT 0 Fan 1	OK	2821	Spinning at normal speed
TFT 0 Fan 2	OK	2735	Spinning at normal speed
TFT 0 Fan 3	OK	2815	Spinning at normal speed
TFT 0 Fan 4	OK	2828	Spinning at normal speed
TFT 0 Fan 5	OK	2863	Spinning at normal speed
BFT 1 Fan 0	OK	2941	Spinning at normal speed
BFT 1 Fan 1	OK	3008	Spinning at normal speed
BFT 1 Fan 2	OK	3073	Spinning at normal speed
BFT 1 Fan 3	OK	2925	Spinning at normal speed
BFT 1 Fan 4	OK	2863	Spinning at normal speed
BFT 1 Fan 5	OK	2933	Spinning at normal speed
SFT 0 Fan 0 Rotor 0	OK	15472	Spinning at normal speed
SFT 0 Fan 0 Rotor 1	OK	14477	Spinning at normal speed
SFT 0 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 0 Fan 1 Rotor 1	OK	14210	Spinning at normal speed
SFT 0 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 0 Fan 2 Rotor 1	OK	14248	Spinning at normal speed
SFT 0 Fan 3 Rotor 0	OK	16463	Spinning at normal speed
SFT 0 Fan 3 Rotor 1	OK	14099	Spinning at normal speed
SFT 1 Fan 0 Rotor 0	OK	15083	Spinning at normal speed
SFT 1 Fan 0 Rotor 1	OK	13533	Spinning at normal speed
SFT 1 Fan 1 Rotor 0	OK	16071	Spinning at normal speed
SFT 1 Fan 1 Rotor 1	OK	14400	Spinning at normal speed
SFT 1 Fan 2 Rotor 0	OK	15517	Spinning at normal speed
SFT 1 Fan 2 Rotor 1	OK	14210	Spinning at normal speed
SFT 1 Fan 3 Rotor 0	OK	16413	Spinning at normal speed
SFT 1 Fan 3 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 2 Fan 0 Rotor 1	OK	14634	Spinning at normal speed
SFT 2 Fan 1 Rotor 0	OK	15561	Spinning at normal speed
SFT 2 Fan 1 Rotor 1	OK	14285	Spinning at normal speed
SFT 2 Fan 2 Rotor 0	OK	15835	Spinning at normal speed
SFT 2 Fan 2 Rotor 1	OK	14400	Spinning at normal speed
SFT 2 Fan 3 Rotor 0	OK	15789	Spinning at normal speed
SFT 2 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 0 Rotor 0	OK	16314	Spinning at normal speed

SFT 3 Fan 0 Rotor 1	OK	14876	Spinning at normal speed
SFT 3 Fan 1 Rotor 0	OK	15835	Spinning at normal speed
SFT 3 Fan 1 Rotor 1	OK	14323	Spinning at normal speed
SFT 3 Fan 2 Rotor 0	OK	16265	Spinning at normal speed
SFT 3 Fan 2 Rotor 1	OK	14594	Spinning at normal speed
SFT 3 Fan 3 Rotor 0	OK	16071	Spinning at normal speed
SFT 3 Fan 3 Rotor 1	OK	14323	Spinning at normal speed
SFT 4 Fan 0 Rotor 0	OK	15652	Spinning at normal speed
SFT 4 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 4 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 4 Fan 1 Rotor 1	OK	14555	Spinning at normal speed
SFT 4 Fan 2 Rotor 0	OK	16023	Spinning at normal speed
SFT 4 Fan 2 Rotor 1	OK	14361	Spinning at normal speed
SFT 4 Fan 3 Rotor 0	OK	16216	Spinning at normal speed
SFT 4 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 5 Fan 0 Rotor 0	OK	15297	Spinning at normal speed
SFT 5 Fan 0 Rotor 1	OK	14173	Spinning at normal speed
SFT 5 Fan 1 Rotor 0	OK	15472	Spinning at normal speed
SFT 5 Fan 1 Rotor 1	OK	13846	Spinning at normal speed
SFT 5 Fan 2 Rotor 0	OK	15340	Spinning at normal speed
SFT 5 Fan 2 Rotor 1	OK	13917	Spinning at normal speed
SFT 5 Fan 3 Rotor 0	OK	15835	Spinning at normal speed
SFT 5 Fan 3 Rotor 1	OK	13917	Spinning at normal speed
SFT 6 Fan 0 Rotor 0	OK	15743	Spinning at normal speed
SFT 6 Fan 0 Rotor 1	OK	14594	Spinning at normal speed
SFT 6 Fan 1 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 1 Rotor 1	OK	14634	Spinning at normal speed
SFT 6 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 6 Fan 2 Rotor 1	OK	14516	Spinning at normal speed
SFT 6 Fan 3 Rotor 0	OK	16666	Spinning at normal speed
SFT 6 Fan 3 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 0 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 0 Rotor 1	OK	14438	Spinning at normal speed
SFT 7 Fan 1 Rotor 0	OK	15517	Spinning at normal speed
SFT 7 Fan 1 Rotor 1	OK	14361	Spinning at normal speed
SFT 7 Fan 2 Rotor 0	OK	16167	Spinning at normal speed
SFT 7 Fan 2 Rotor 1	OK	14555	Spinning at normal speed
SFT 7 Fan 3 Rotor 0	OK	15697	Spinning at normal speed
SFT 7 Fan 3 Rotor 1	OK	14361	Spinning at normal speed

#### show chassis fan (EX Series Switches)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	3477	Spinning at normal speed
Fan 2	OK	3477	Spinning at normal speed
Fan 3	OK	3479	Spinning at normal speed
Fan 4	OK	3508	Spinning at normal speed
Fan 5	OK	3517	Spinning at normal speed
Fan 6	OK	3531	Spinning at normal speed
Fan 7	OK	3439	Spinning at normal speed
Fan 8	OK	3424	Spinning at normal speed
Fan 9	OK	3413	Spinning at normal speed
Fan 10	OK	3439	Spinning at normal speed
Fan 11	OK	3446	Spinning at normal speed
Fan 12	OK	3432	Spinning at normal speed

#### show chassis fan (T320 Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	2850	Spinning at normal speed
Top Left Middle fan	OK	2820	Spinning at normal speed
Top Left Rear fan	OK	2970	Spinning at normal speed
Top Right Front fan	OK	2790	Spinning at normal speed
Top Right Middle fan	OK	2640	Spinning at normal speed
Top Right Rear fan	OK	2790	Spinning at normal speed
Bottom Left Front fan	OK	2520	Spinning at normal speed
Bottom Left Middle fan	OK	2610	Spinning at normal speed
Bottom Left Rear fan	OK	2550	Spinning at normal speed
Bottom Right Front fan	OK	2610	Spinning at normal speed
Bottom Right Middle fan	OK	2880	Spinning at normal speed
Bottom Right Rear fan	OK	2790	Spinning at normal speed
Rear Tray Top fan	OK	2130	Spinning at normal speed
Rear Tray Second fan	OK	2190	Spinning at normal speed
Rear Tray Middle fan	OK	2250	Spinning at normal speed
Rear Tray Fourth fan	OK	2220	Spinning at normal speed
Rear Tray Bottom fan	OK	2280	Spinning at normal speed

## show chassis fan (T640 Router)

user@host&gt; show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3390	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5220	Spinning at normal speed
Rear Tray Second fan	OK	5220	Spinning at normal speed
Rear Tray Third fan	OK	5220	Spinning at normal speed
Rear Tray Fourth fan	OK	5220	Spinning at normal speed
Rear Tray Fifth fan	OK	5220	Spinning at normal speed
Rear Tray Sixth fan	OK	5220	Spinning at normal speed
Rear Tray Seventh fan	OK	5220	Spinning at normal speed
Rear Tray Bottom fan	OK	5220	Spinning at normal speed

## show chassis fan (T1600 Router)

user@host&gt; show chassis fan

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3450	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3390	Spinning at normal speed

Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	5190	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	5190	Spinning at normal speed
Rear Tray Sixth fan	OK	5190	Spinning at normal speed
Rear Tray Seventh fan	OK	5190	Spinning at normal speed
Rear Tray Bottom fan	OK	5190	Spinning at normal speed

### show chassis fan (T4000 Core Router)

```
user@host> show chassis fan
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	5190	Spinning at high speed
Top Left Middle fan	OK	5220	Spinning at high speed
Top Left Rear fan	OK	5190	Spinning at high speed
Top Right Front fan	OK	5160	Spinning at high speed
Top Right Middle fan	OK	5190	Spinning at high speed
Top Right Rear fan	OK	5160	Spinning at high speed
Bottom Left Front fan	OK	6030	Spinning at high speed
Bottom Left Middle fan	OK	6090	Spinning at high speed
Bottom Left Rear fan	OK	6090	Spinning at high speed
Bottom Right Front fan	OK	6030	Spinning at high speed
Bottom Right Middle fan	OK	6060	Spinning at high speed
Bottom Right Rear fan	OK	6060	Spinning at high speed
Rear Tray Top fan	OK	10000	Spinning at high speed
Rear Tray Second fan	OK	10000	Spinning at high speed
Rear Tray Third fan	OK	10000	Spinning at high speed
Rear Tray Fourth fan	OK	10000	Spinning at high speed
Rear Tray Fifth fan	OK	10000	Spinning at high speed
Rear Tray Sixth fan	OK	10000	Spinning at high speed
Rear Tray Seventh fan	OK	10000	Spinning at high speed
Rear Tray Bottom fan	OK	10000	Spinning at high speed

### show chassis fan (TX Matrix Router)

```
user@host> show chassis fan
scc-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3390	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3390	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3450	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3420	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray Top fan	OK	3420	Spinning at normal speed
Rear Tray Second fan	OK	5190	Spinning at normal speed
Rear Tray Third fan	OK	5190	Spinning at normal speed
Rear Tray Fourth fan	OK	5190	Spinning at normal speed
Rear Tray Fifth fan	OK	3420	Spinning at normal speed
Rear Tray Sixth fan	OK	3420	Spinning at normal speed

```

Rear Tray Seventh fan    OK      3420    Spinning at normal speed
Rear Tray Bottom fan     OK      3420    Spinning at normal speed

```

```
lcc2-re0:
```

```

-----
Item                Status  RPM    Measurement
Top Left Front fan   OK      3420    Spinning at normal speed
Top Left Middle fan  OK      3420    Spinning at normal speed
Top Left Rear fan    OK      3450    Spinning at normal speed
Top Right Front fan   OK      3420    Spinning at normal speed
Top Right Middle fan  OK      3450    Spinning at normal speed
Top Right Rear fan    OK      3360    Spinning at normal speed
Bottom Left Front fan OK      3420    Spinning at normal speed
Bottom Left Middle fan OK     3480    Spinning at normal speed
Bottom Left Rear fan  OK      3420    Spinning at normal speed
Bottom Right Front fan OK     3420    Spinning at normal speed
Bottom Right Middle fan OK     3390    Spinning at normal speed
Bottom Right Rear fan OK     3420    Spinning at normal speed
Rear Tray Top fan     OK      3420    Spinning at normal speed
Rear Tray Second fan  OK      3420    Spinning at normal speed
Rear Tray Third fan   OK      3420    Spinning at normal speed
Rear Tray Fourth fan  OK      3420    Spinning at normal speed
Rear Tray Fifth fan   OK      3420    Spinning at normal speed
Rear Tray Sixth fan   OK      3420    Spinning at normal speed
Rear Tray Seventh fan OK     3420    Spinning at normal speed
Rear Tray Bottom fan  OK     3420    Spinning at normal speed

```

#### show chassis fan (TX Matrix Plus Router)

```
user@host> show chassis fan
sfc0-re0:
```

```

-----
Item                Status  RPM    Measurement
Fan Tray 0 Fan 1     OK     4350    Spinning at normal speed
Fan Tray 0 Fan 2     OK     4380    Spinning at normal speed
Fan Tray 0 Fan 3     OK     4410    Spinning at normal speed
Fan Tray 0 Fan 4     OK     4380    Spinning at normal speed
Fan Tray 0 Fan 5     OK     4350    Spinning at normal speed
Fan Tray 0 Fan 6     OK     4380    Spinning at normal speed
Fan Tray 1 Fan 1     OK     4410    Spinning at normal speed
Fan Tray 1 Fan 2     OK     4380    Spinning at normal speed
Fan Tray 1 Fan 3     OK     4410    Spinning at normal speed
Fan Tray 1 Fan 4     OK     4380    Spinning at normal speed
Fan Tray 1 Fan 5     OK     4410    Spinning at normal speed
Fan Tray 1 Fan 6     OK     4410    Spinning at normal speed
Fan Tray 2 Fan 1     OK     4380    Spinning at normal speed
Fan Tray 2 Fan 2     OK     4380    Spinning at normal speed
Fan Tray 2 Fan 3     OK     4380    Spinning at normal speed
Fan Tray 2 Fan 4     OK     4410    Spinning at normal speed
Fan Tray 2 Fan 5     OK     4380    Spinning at normal speed
Fan Tray 2 Fan 6     OK     4410    Spinning at normal speed
Fan Tray 2 Fan 7     OK     4410    Spinning at normal speed
Fan Tray 2 Fan 8     OK     4380    Spinning at normal speed
Fan Tray 2 Fan 9     OK     4380    Spinning at normal speed
Fan Tray 3 Fan 1     OK     4350    Spinning at normal speed
Fan Tray 3 Fan 2     OK     4380    Spinning at normal speed
Fan Tray 3 Fan 3     OK     4410    Spinning at normal speed
Fan Tray 3 Fan 4     OK     4440    Spinning at normal speed
Fan Tray 3 Fan 5     OK     4380    Spinning at normal speed
Fan Tray 3 Fan 6     OK     4410    Spinning at normal speed
Fan Tray 3 Fan 7     OK     4410    Spinning at normal speed

```



Fan Tray 3 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4380	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4410	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4350	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4410	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4380	Spinning at normal speed
Fan Tray 5 Fan 9	OK	4410	Spinning at normal speed

```
lcc0-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3420	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3450	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3420	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3420	Spinning at normal speed
Bottom Left Rear fan	OK	3390	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3390	Spinning at normal speed
Rear Tray Top fan	OK	7050	Spinning at normal speed
Rear Tray Second fan	OK	7050	Spinning at normal speed
Rear Tray Third fan	OK	7050	Spinning at normal speed
Rear Tray Fourth fan	OK	7050	Spinning at normal speed
Rear Tray Fifth fan	OK	7050	Spinning at normal speed
Rear Tray Sixth fan	OK	7050	Spinning at normal speed
Rear Tray Seventh fan	OK	7050	Spinning at normal speed
Rear Tray Bottom fan	OK	7050	Spinning at normal speed

#### show chassis fan (TX Matrix Plus Router with 3D SIBs)

```
user@host> show chassis fan
sfc0-re0:
```

Item	Status	RPM	Measurement
Fan Tray 0 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 0 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 0 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 0 Fan 6	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 2	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 1 Fan 4	OK	4770	Spinning at normal speed

Fan Tray 1 Fan 5	OK	4770	Spinning at normal speed
Fan Tray 1 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 1	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 2	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 2 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 2 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 1	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 2	OK	4860	Spinning at normal speed
Fan Tray 3 Fan 3	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 6	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 3 Fan 8	OK	4800	Spinning at normal speed
Fan Tray 3 Fan 9	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 4	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 5	OK	4830	Spinning at normal speed
Fan Tray 4 Fan 6	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 7	OK	4800	Spinning at normal speed
Fan Tray 4 Fan 8	OK	4860	Spinning at normal speed
Fan Tray 4 Fan 9	OK	4770	Spinning at normal speed
Fan Tray 5 Fan 1	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 2	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 3	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 4	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 5	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 6	OK	4800	Spinning at normal speed
Fan Tray 5 Fan 7	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 8	OK	4830	Spinning at normal speed
Fan Tray 5 Fan 9	Check	2010	

1cc0-re0:

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3390	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed

Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

```
1cc2-re0:
```

Item	Status	RPM	Measurement
Top Left Front fan	OK	3420	Spinning at normal speed
Top Left Middle fan	OK	3390	Spinning at normal speed
Top Left Rear fan	OK	3420	Spinning at normal speed
Top Right Front fan	OK	3420	Spinning at normal speed
Top Right Middle fan	OK	3420	Spinning at normal speed
Top Right Rear fan	OK	3450	Spinning at normal speed
Bottom Left Front fan	OK	3420	Spinning at normal speed
Bottom Left Middle fan	OK	3390	Spinning at normal speed
Bottom Left Rear fan	OK	3420	Spinning at normal speed
Bottom Right Front fan	OK	3420	Spinning at normal speed
Bottom Right Middle fan	OK	3390	Spinning at normal speed
Bottom Right Rear fan	OK	3420	Spinning at normal speed
Rear Tray fan 1 (Top)	OK	7740	Spinning at normal speed
Rear Tray fan 2	OK	7740	Spinning at normal speed
Rear Tray fan 3	OK	7740	Spinning at normal speed
Rear Tray fan 4	OK	7740	Spinning at normal speed
Rear Tray fan 5	OK	7740	Spinning at normal speed
Rear Tray fan 6	OK	7740	Spinning at normal speed
Rear Tray fan 7	OK	7740	Spinning at normal speed
Rear Tray fan 8	OK	7740	Spinning at normal speed
Rear Tray fan 9	OK	7740	Spinning at normal speed
Rear Tray fan 10	OK	7740	Spinning at normal speed
Rear Tray fan 11	OK	7740	Spinning at normal speed
Rear Tray fan 12	OK	7740	Spinning at normal speed
Rear Tray fan 13	OK	7740	Spinning at normal speed
Rear Tray fan 14	OK	7740	Spinning at normal speed
Rear Tray fan 15	OK	7740	Spinning at normal speed
Rear Tray fan 16 (Bottom)	OK	7740	Spinning at normal speed

### show chassis fan (PTX5000 Packet Transport Router)

```
user@host> show chassis fan
user@host> show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	29%	2700 RPM
Fan Tray 0 Fan 2	OK	29%	2700 RPM
Fan Tray 0 Fan 3	OK	29%	2742 RPM
Fan Tray 0 Fan 4	OK	29%	2700 RPM
Fan Tray 0 Fan 5	OK	30%	2828 RPM
Fan Tray 0 Fan 6	OK	30%	2828 RPM
Fan Tray 0 Fan 7	OK	29%	2700 RPM
Fan Tray 0 Fan 8	OK	30%	2785 RPM
Fan Tray 0 Fan 9	OK	30%	2828 RPM
Fan Tray 0 Fan 10	OK	30%	2828 RPM
Fan Tray 0 Fan 11	OK	30%	2785 RPM
Fan Tray 0 Fan 12	OK	30%	2828 RPM
Fan Tray 0 Fan 13	OK	31%	2871 RPM
Fan Tray 0 Fan 14	OK	30%	2828 RPM

Fan Tray 1 Fan 1	OK	42%	3033 RPM
Fan Tray 1 Fan 2	OK	42%	3066 RPM
Fan Tray 1 Fan 3	OK	43%	3099 RPM
Fan Tray 1 Fan 4	OK	43%	3166 RPM
Fan Tray 1 Fan 5	OK	45%	3266 RPM
Fan Tray 1 Fan 6	OK	43%	3133 RPM
Fan Tray 2 Fan 1	OK	29%	2099 RPM
Fan Tray 2 Fan 2	OK	30%	2199 RPM
Fan Tray 2 Fan 3	OK	30%	2166 RPM
Fan Tray 2 Fan 4	OK	33%	2399 RPM
Fan Tray 2 Fan 5	OK	29%	2133 RPM
Fan Tray 2 Fan 6	OK	32%	2366 RPM

**show chassis fan (MX104 Router)**

```
user@host > show chassis fan
```

Item	Status	RPM	Measurement
Fan 1	OK	5640	Spinning at normal speed
Fan 2	OK	5640	Spinning at normal speed
Fan 3	OK	5760	Spinning at normal speed
Fan 4	OK	5640	Spinning at normal speed
Fan 5	OK	5640	Spinning at normal speed

**show chassis fan (MX2010 Router)**

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	38%	3480 RPM
Fan Tray 0 Fan 3	OK	37%	3360 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	38%	3480 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM
Fan Tray 1 Fan 1	OK	38%	3480 RPM
Fan Tray 1 Fan 2	OK	40%	3600 RPM
Fan Tray 1 Fan 3	OK	38%	3480 RPM
Fan Tray 1 Fan 4	OK	38%	3480 RPM
Fan Tray 1 Fan 5	OK	38%	3480 RPM
Fan Tray 1 Fan 6	OK	38%	3480 RPM
Fan Tray 2 Fan 1	OK	38%	3480 RPM
Fan Tray 2 Fan 2	OK	41%	3720 RPM
Fan Tray 2 Fan 3	OK	38%	3480 RPM
Fan Tray 2 Fan 4	OK	38%	3480 RPM
Fan Tray 2 Fan 5	OK	38%	3480 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	40%	3600 RPM
Fan Tray 3 Fan 3	OK	40%	3600 RPM
Fan Tray 3 Fan 4	OK	40%	3600 RPM
Fan Tray 3 Fan 5	OK	40%	3600 RPM
Fan Tray 3 Fan 6	OK	38%	3480 RPM

**show chassis fan (MX2020 Router)**

```
user@host > show chassis fan
```

Item	Status	% RPM	Measurement
Fan Tray 0 Fan 1	OK	37%	3360 RPM
Fan Tray 0 Fan 2	OK	37%	3360 RPM
Fan Tray 0 Fan 3	OK	36%	3240 RPM
Fan Tray 0 Fan 4	OK	37%	3360 RPM
Fan Tray 0 Fan 5	OK	37%	3360 RPM
Fan Tray 0 Fan 6	OK	37%	3360 RPM

Fan Tray 1 Fan 1	OK	37%	3360 RPM
Fan Tray 1 Fan 2	OK	37%	3360 RPM
Fan Tray 1 Fan 3	OK	37%	3360 RPM
Fan Tray 1 Fan 4	OK	37%	3360 RPM
Fan Tray 1 Fan 5	OK	37%	3360 RPM
Fan Tray 1 Fan 6	OK	36%	3240 RPM
Fan Tray 2 Fan 1	OK	37%	3360 RPM
Fan Tray 2 Fan 2	OK	37%	3360 RPM
Fan Tray 2 Fan 3	OK	37%	3360 RPM
Fan Tray 2 Fan 4	OK	37%	3360 RPM
Fan Tray 2 Fan 5	OK	37%	3360 RPM
Fan Tray 2 Fan 6	OK	38%	3480 RPM
Fan Tray 3 Fan 1	OK	38%	3480 RPM
Fan Tray 3 Fan 2	OK	38%	3480 RPM
Fan Tray 3 Fan 3	OK	38%	3480 RPM
Fan Tray 3 Fan 4	OK	37%	3360 RPM
Fan Tray 3 Fan 5	OK	37%	3360 RPM
Fan Tray 3 Fan 6	OK	37%	3360 RPM

#### show chassis fan (ACX4000 Router)

```

user@host > show chassis fan
  Item                Status  RPM    Measurement
  Fan 1               OK      4140   Spinning at normal speed
  Fan 2               OK      4200   Spinning at normal speed

```

#### show chassis fan (QFX5100 Switch)

```

user@switch > show chassis fan
  Item                Status  RPM    Measurement
  FPC 0 Tray 0 Fan 0  OK      6428   Spinning at normal speed
  FPC 0 Tray 0 Fan 1  OK      5515   Spinning at normal speed
  FPC 0 Tray 1 Fan 0  OK      6360   Spinning at normal speed
  FPC 0 Tray 1 Fan 1  OK      5532   Spinning at normal speed

```

## show chassis firmware

---

<b>List of Syntax</b>	<a href="#">Syntax on page 628</a> <a href="#">Syntax (TX Matrix Routers) on page 628</a> <a href="#">Syntax (TX Matrix Plus Routers) on page 628</a> <a href="#">Syntax (MX Series Routers) on page 628</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 628</a> <a href="#">Syntax (QFX Series) on page 628</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 628</a> <a href="#">Syntax (EX Series Switches) on page 628</a>
<b>Syntax</b>	show chassis firmware
<b>Syntax (TX Matrix Routers)</b>	show chassis firmware <fcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Routers)</b>	show chassis firmware <fcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis firmware <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis firmware
<b>Syntax (QFX Series)</b>	show chassis firmware interconnect-device <i>name</i> node-device <i>name</i>
<b>Syntax (ACX Series Universal Access Routers)</b>	show chassis firmware
<b>Syntax (EX Series Switches)</b>	show chassis firmware <detail>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.4 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced for EX8200 switches in Junos OS Release 10.2 for EX Series switches. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for ACX4000 Universal Access Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** On routers and switches, display the version levels of the firmware running on the System Control Board (SCB), Switching and Forwarding Module (SFM), System and Switch Board (SSB), Forwarding Engine Board (FEB), Flexible PIC Concentrators (FPCs), and Routing Engines. On a TX Matrix Plus router, display the version levels of the firmware running on the FPCs and the Switch Processor Mezzanine Board (SPMBs).

On EX2200, EX3200, and EX4200 switches, and the QFX Series, display the version levels of the firmware running on the switch. On an EX8208 switch, display the version levels of the firmware running on the Switch Fabric and Routing Engine (SRE) modules and on the line cards (shown as FPCs). On an EX8216 switch, display the version levels of the firmware running on the Routing Engine (RE) modules and on the line cards (shown as FPCs).

**Options** **none**—Display the version levels of the firmware running. For an EX4200 switch that is a member of a Virtual Chassis, display version levels for all members. For a TX Matrix router, display version levels for the firmware on the TX Matrix router and on all the T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, display version levels for the firmware on the TX Matrix Plus router and on all the routers connected to the TX Matrix Plus router.

**all-members**—(MX Series routers only) (Optional) Display the version levels of the firmware running for all members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems) (Optional) Display the version levels of the firmware running on the Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display version levels for the firmware on a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the version levels for the firmware on a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display the version levels of the firmware running for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the version levels of the firmware running for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device**—(QFabric systems only) (Optional) Display the version levels of the firmware running on the Node device.

**scc**—(TX Matrix router only) (Optional) Display version levels for the firmware on the TX Matrix router (switch-card chassis).

**sfc number**—(TX Matrix Plus router only) (Optional) Display version levels for the firmware on the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with **0**.

**detail**—(EX3200, EX3300, EX4200, and EX4500 standalone and Virtual Chassis member switches only) (Optional) Display version levels of the firmware running on the switch for its programmable hardware components.

**Required Privilege Level**

view

**Related Documentation**

- *Upgrading the HSM Firmware*

**List of Sample Output**

[show chassis firmware \(M10 Router\) on page 631](#)  
[show chassis firmware \(M20 Router\) on page 631](#)  
[show chassis firmware \(M40 Router\) on page 632](#)  
[show chassis firmware \(M120 Router\) on page 632](#)  
[show chassis firmware \(M160 Router\) on page 632](#)  
[show chassis firmware \(MX104 Router\) on page 632](#)  
[show chassis firmware \(MX240 Router\) on page 632](#)  
[show chassis firmware \(MX480 Router\) on page 633](#)  
[show chassis firmware \(MX960 Router\) on page 633](#)  
[show chassis firmware \(MX2010 Router\) on page 633](#)  
[show chassis firmware \(MX2020 Router\) on page 633](#)  
[show chassis firmware \(MX240, MX480, MX960 Router with Application Services Modular Line Card\) on page 634](#)  
[show chassis firmware \(EX4200 Switch\) on page 634](#)  
[show chassis firmware \(EX8200 Switch\) on page 634](#)  
[show chassis firmware \(EX9200 Switch\) on page 635](#)  
[show chassis firmware lcc \(TX Matrix Router\) on page 635](#)  
[show chassis firmware scc \(TX Matrix Router\) on page 635](#)  
[show chassis firmware \(TX Matrix Plus Router\) on page 635](#)  
[show chassis firmware lcc \(TX Matrix Plus Router\) on page 637](#)  
[show chassis firmware sfc \(TX Matrix Plus Router\) on page 637](#)  
[show chassis firmware \(QFX Series\) on page 637](#)  
[show chassis firmware interconnect-device \(QFabric System\) on page 638](#)  
[show chassis firmware \(ACX2000 Universal Access Router\) on page 638](#)  
[show chassis firmware detail \(EX3300 Switch\) on page 638](#)  
[show chassis firmware \(MX Routers with Media Services Blade \[MSB\]\) on page 638](#)

**Output Fields**

[Table 32 on page 631](#) lists the output fields for the **show chassis firmware** command. Output fields are listed in the approximate order in which they appear.



Table 32: show chassis firmware Output Fields

Field Name	Field Description
<b>Part</b>	(MX Series, MX2010, and MX2020 routers) Chassis part name.
<b>Type</b>	(MX Series, MX2010, and MX2020 routers) Type of firmware: On routers: <b>ROM</b> or <b>O/S</b> . On switches: <b>uboot</b> or <b>loader</b> .
<b>Version</b>	(MX Series, MX2010, and MX2020 routers) Version of firmware running on the chassis part.
<b>FPC</b>	( <i>detail</i> option only) Number of FPC. For a standalone switch, the value is 0. For a Virtual Chassis configuration, value in the range of 0-9; refers to the member ID assigned to the switch.
<b>AFEB</b>	(MX104 routers) Version of the compact Forwarding Engine Board.
<b>Boot</b>	( <i>detail</i> option only) Version of the SYSPLD.
<b>PoE</b>	( <i>detail</i> option only) Version of the PoE firmware.
<b>PFE-&lt;number&gt;</b>	( <i>detail</i> option only) Version of the PFE used in the switch.
<b>PHY-</b>	( <i>detail</i> option only) Version of the physical layer device (PHY) used in the switch.
<b>microcode</b>	( <i>detail</i> option only) Microcode of the physical layer devices (PHY) used in the switch.
<b>uboot</b>	( <i>detail</i> option only) Version of the u-boot used in the switch.
<b>loader</b>	( <i>detail</i> option only) Version of the loader used in the switch.

## Sample Output

### show chassis firmware (M10 Router)

```

user@host> show chassis firmware
Part          Type      Version
Forwarding engine board  ROM      Juniper ROM Monitor Version 4.1b2
                                O/S      Version 4.1I1 by tlim on 2000-04-24 11:27

```

### show chassis firmware (M20 Router)

```

user@host> show chassis firmware
Part          Type      Version
System switch board    ROM      Juniper ROM Monitor Version 3.4b26
                                O/S      Version 3.4I16 by smackie on 2000-02-29 2
FPC 1          ROM      Juniper ROM Monitor Version 3.0b1
                                O/S      Version 3.4I4 by smackie on 2000-02-25 21
FPC 2          ROM      Juniper ROM Monitor Version 3.0b1
                                O/S      Version 3.4I4 by smackie on 2000-02-25 21

```

**show chassis firmware (M40 Router)**

```

user@host> show chassis firmware
Part                Type      Version
System control board ROM       Juniper ROM Monitor Version 2.0i126Copyri
                  O/S       Version 2.0i1 by root on Thu Jul 23 00:51
FPC 5               ROM       Juniper ROM Monitor Version 2.0i49Copyrig
                  O/S       Version 2.0i1 by root on Thu Jul 23 00:59

```

**show chassis firmware (M120 Router)**

```

user@host> show chassis firmware
FPC 2               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FPC 3               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FPC 4               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:2
FEB 3               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:1
FEB 4               ROM       Juniper ROM Monitor Version 8.0b29
                  O/S       Version 8.2B1 by builder on 2006-10-18 16:1

```

**show chassis firmware (M160 Router)**

```

user@host> show chassis firmware
Part                Type      Version
SFM 0               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:50
SFM 1               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:50
FPC 0               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 1               ROM       Juniper ROM Monitor Version 4.0b2
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 2               ROM       Juniper ROM Monitor Version 4.0b3
                  O/S       Version 4.0I1 by tlim on 2000-02-29 11:56

```

**show chassis firmware (MX104 Router)**

```

user@host > show chassis firmware
Part                Type      Version
FPC 0               ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-
FPC 1               ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-
FPC 2               ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-
AFEB                ROM       Juniper ROM Monitor Version 13.1b24
                  O/S       Version 13.2-20130514.1 by builder on 2013-

```

**show chassis firmware (MX240 Router)**

```

user@host> show chassis firmware
Part                Type      Version
FPC 1               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20080103.0 by builder on 2008-0
FPC 2               ROM       Juniper ROM Monitor Version 8.3b1
                  O/S       Version 9.0-20080103.0 by builder on 2008-0

```

**show chassis firmware (MX480 Router)**

```

user@host> show chassis firmware
Part      Type      Version
FPC 1     ROM       Juniper ROM Monitor Version 8.3b1
           O/S       Version 9.0-20070916.3 by builder on 2007-0

```

**show chassis firmware (MX960 Router)**

```

user@host> show chassis firmware
Part      Type      Version
FPC 4     ROM       Juniper ROM Monitor Version 8.0b8
           O/S       Version 8.2I59 by artem on 2006-10-31 19:22
FPC 7     ROM       Juniper ROM Monitor Version 8.2b1
           O/S       Version 8.2-20061026.1 by builder on 2006-1

```

**show chassis firmware (MX2010 Router)**

```

user@host> show chassis firmware
Part      Type      Version
FPC 0     ROM       Juniper ROM Monitor Version 12.3b1
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 1     ROM       Juniper ROM Monitor Version 10.1b3
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 2     ROM       Juniper ROM Monitor Version 10.1b3
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 3     ROM       Juniper ROM Monitor Version 10.1b3
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 4     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 5     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 6     ROM       Juniper ROM Monitor Version 10.4b1
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 7     ROM       Juniper ROM Monitor Version 10.1b3
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 8     ROM       Juniper ROM Monitor Version 10.4b1
           O/S       Version 12.3-20121220.0 by builder on 2012-
FPC 9     ROM       Juniper ROM Monitor Version 10.4b1
           O/S       Version 12.3-20121220.0 by builder on 2012-
SPMB 0    ROM       Juniper ROM Monitor Version 12.1b1
           O/S       Version 12.3-20121220.0 by builder on 2012-
SPMB 1    ROM       Juniper ROM Monitor Version 12.1b1
           O/S       Version 12.3-20121220.0 by builder on 2012-

```

**show chassis firmware (MX2020 Router)**

```

user@host> show chassis firmware
Part      Type      Version
FPC 0     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20130415.0 by builder on 2013-
FPC 1     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20130415.0 by builder on 2013-
FPC 2     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20130415.0 by builder on 2013-
FPC 3     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20130415.0 by builder on 2013-
FPC 4     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20130415.0 by builder on 2013-
FPC 5     ROM       Juniper ROM Monitor Version 10.0b39
           O/S       Version 12.3-20130415.0 by builder on 2013-

```

FPC 6	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 7	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 8	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 9	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 10	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 11	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 12	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 13	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 14	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 15	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 16	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 17	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 18	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
FPC 19	ROM	Juniper ROM Monitor Version 10.0b39
	O/S	Version 12.3-20130415.0 by builder on 2013-
SPMB 0	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20130415.0 by builder on 2013-
SPMB 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.3-20130415.0 by builder on 2013-

#### show chassis firmware (MX240, MX480, MX960 Router with Application Services Modular Line Card)

```
user@host> show chassis firmware
```

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 12.1b1
	O/S	Version 12.2I21 by manish on 2012-06-19 17:

#### show chassis firmware (EX4200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 1	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1
FPC 2	uboot	U-Boot 1.1.6 (Feb 6 2008 - 11:27:42)
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.1

#### show chassis firmware (EX8200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 0	U-Boot	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 3	U-Boot	U-Boot 1.1.6 (Dec 4 2009 - 13:17:34) 3.1.0
	loader	FreeBSD/PowerPC U-Boot bootstrap loader 2.2

FPC 5	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 7	U-Boot loader	U-Boot 1.1.6 (Feb 6 2009 - 05:31:46) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 0	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 1	U-Boot loader	U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0 FreeBSD/PowerPC U-Boot bootstrap loader 2.2

#### show chassis firmware (EX9200 Switch)

```
user@switch> show chassis firmware
```

Part	Type	Version
FPC 2	ROM	Juniper ROM Monitor Version 11.4b2
	O/S	Version 14.1I20140312_0741_bavig by bavig o
FPC 3	ROM	Juniper ROM Monitor Version 10.4b1
	O/S	Version 14.1I20140312_0741_bavig by bavig o

#### show chassis firmware lcc (TX Matrix Router)

```
user@host> show chassis firmware lcc 0
```

lcc0-re0:

Part	Type	Version
FPC 1	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0
FPC 2	ROM	Juniper ROM Monitor Version 6.4b20
	O/S	Version 7.0-20040804.0 by builder on 2004-0
SPMB 0	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0

#### show chassis firmware scc (TX Matrix Router)

```
user@host> show chassis firmware scc
```

scc-re0:

Part	Type	Version
SPMB 0	ROM	Juniper ROM Monitor Version 6.4b18
	O/S	Version 7.0-20040804.0 by builder on 2004-0

#### show chassis firmware (TX Matrix Plus Router)

```
user@host> show chassis firmware
```

sfc0-re0:

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		
Global FPC 21		
Global FPC 22		
Global FPC 23		
Global FPC 24		
Global FPC 25		
Global FPC 26		
Global FPC 28		

```

Global FPC 29
Global FPC 31
SPMB 0          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1          ROM      Juniper ROM Monitor Version 9.5b1
                  O/S      Version 9.6-20090507.0 by builder on 2009-0

```

## lcc0-re1:

```

-----
Part      Type      Version
FPC 4     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0     ROM      Juniper ROM Monitor Version 9.5b1
           O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1     ROM      Juniper ROM Monitor Version 9.5b1
           O/S      Version 9.6-20090507.0 by builder on 2009-0

```

## lcc1-re1:

```

-----
Part      Type      Version
FPC 4     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0     ROM      Juniper ROM Monitor Version 9.5b1
           O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1     ROM      Juniper ROM Monitor Version 9.5b1
           O/S      Version 9.6-20090507.0 by builder on 2009-0

```

## lcc2-re1:

```

-----
Part      Type      Version
FPC 4     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 5     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 6     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 7     ROM      Juniper ROM Monitor Version 7.5b4
           O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 0     ROM      Juniper ROM Monitor Version 9.5b1
           O/S      Version 9.6-20090507.0 by builder on 2009-0
SPMB 1     ROM      Juniper ROM Monitor Version 9.5b1
           O/S      Version 9.6-20090507.0 by builder on 2009-0

```

## lcc3-re1:

```

-----
Part      Type      Version
FPC 0     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 1     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 2     ROM      Juniper ROM Monitor Version 9.0b2
           O/S      Version 9.6-20090507.0 by builder on 2009-0
FPC 4     ROM      Juniper ROM Monitor Version 7.5b4

```

	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 5	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

#### show chassis firmware lcc (TX Matrix Plus Router)

```
user@host> show chassis firmware lcc 0
lcc0-re1:
-----
```

Part	Type	Version
FPC 4	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 6	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
FPC 7	ROM	Juniper ROM Monitor Version 9.0b2
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

#### show chassis firmware sfc (TX Matrix Plus Router)

```
user@host> show chassis firmware sfc 0
sfc0-re0:
-----
```

Part	Type	Version
Global FPC 4		
Global FPC 6		
Global FPC 7		
Global FPC 12		
Global FPC 14		
Global FPC 15		
Global FPC 20		
Global FPC 21		
Global FPC 22		
Global FPC 23		
Global FPC 24		
Global FPC 25		
Global FPC 26		
Global FPC 28		
Global FPC 29		
Global FPC 31		
SPMB 0	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0
SPMB 1	ROM	Juniper ROM Monitor Version 9.5b1
	O/S	Version 9.6-20090507.0 by builder on 2009-0

#### show chassis firmware (QFX Series)

```
user@switch> show chassis firmware
Part                Type      Version
FPC 0
Routing Engine 0    U-Boot   U-Boot 1.1.6 (Sep 15 2010 - 02:11:11) 1.0.5
                    loader   FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

**show chassis firmware interconnect-device (QFabric System)**

```
user@switch> show chassis firmware interconnect-device interconnect1
Part                               Type      Version
Routing Engine 0                   U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
                                     loader     FreeBSD/MIPS U-Boot bootstrap loader 0.1
Routing Engine 1                   U-Boot    U-Boot 1.1.6 (May 10 2011 - 04:52:59) 1.1.1
                                     loader     FreeBSD/MIPS U-Boot bootstrap loader 0.1
```

**show chassis firmware (ACX2000 Universal Access Router)**

```
user@switch> show chassis firmware
Part      Type      Version
FPC       O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
FEB       O/S       Version 12.2I13 by jisjoy on 2012-05-29 06:
```

**show chassis firmware detail (EX3300 Switch)**

```
user@switch> show chassis firmware detail
FPC 0
  Boot SYSPLD                3
  PoE firmware                4.1.6
  PFE-0                       3
  PFE-1                       3
  PHY
    microcode                 0x514
  Boot Firmware
    uboot                     U-Boot 1.1.6 (Aug 21 2011 - 01:45:26) 1.0.0
    loader                     FreeBSD/arm U-Boot loader 1.0
```

**show chassis firmware (MX Routers with Media Services Blade [MSB])**

```
user@switch> show chassis firmware
Part      Type      Version
FPC 1     ROM       Juniper ROM Monitor Version 12.1b1
          O/S       Version 12.2I21 by manish on 2012-06-19 17:
```



## show chassis fpc

<b>List of Syntax</b>	<a href="#">Syntax on page 639</a> <a href="#">Syntax (EX Series Switches) on page 639</a> <a href="#">Syntax (T4000 Routers) on page 639</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 639</a> <a href="#">Syntax (MX Series Routers and EX Series switches) on page 639</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 639</a> <a href="#">Syntax (QFX Series) on page 639</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 639</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 639</a>
<b>Syntax</b>	<pre>show chassis fpc &lt;detail &lt;slot&gt;&gt;   &lt;pic-status &lt;slot&gt;&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show chassis fpc &lt;detail &lt;fpc-slot&gt;&gt;   &lt;pic-status &lt;fpc-slot&gt;&gt; &lt;fpc-slot&gt;</pre>
<b>Syntax (T4000 Routers)</b>	<pre>show chassis fpc &lt;detail &lt;fpc-slot&gt;&gt; &lt;pic-status &lt;fpc-slot&gt;&gt;</pre>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<pre>show chassis fpc &lt;detail &lt;fpc-slot&gt;&gt;   &lt;pic-status &lt;fpc-slot&gt;&gt; &lt;slot&gt;</pre>
<b>Syntax (MX Series Routers and EX Series switches)</b>	<pre>show chassis fpc &lt;detail &lt;slot&gt;&gt;   &lt;pic-status &lt;slot&gt;&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	<pre>show chassis fpc &lt;slot&gt; detail   &lt;detail &lt;slot&gt;&gt;   &lt;pic-status &lt;slot&gt;&gt; &lt;fpc-slot&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show chassis fpc &lt;detail&gt; &lt;interconnect-device name &lt;fpc-slot fpc-slot&gt;&gt; &lt;node-device name&gt;</pre>
<b>Syntax (PTX Series Packet Transport Routers)</b>	<pre>show chassis fpc &lt;detail &lt;fpc-slot&gt;&gt;   &lt;pic-status &lt;fpc-slot&gt;&gt; &lt;fpc-slot&gt;</pre>
<b>Syntax (ACX Series Universal Access Routers)</b>	<pre>show chassis fpc &lt;detail &lt;fpc-slot&gt;&gt;   &lt;pic-status &lt;fpc-slot&gt;&gt; &lt;fpc-slot&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for QFX Series.

Command introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.

Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.

Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.

**Description** Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

**Options** **none**—Display status information for all FPCs. On a TX Matrix router, display status information for all FPCs on the attached T640 routers in the routing matrix. On a TX Matrix Plus router, display status information for all FPCs on the attached routers in the routing matrix.



**NOTE:** In EX8200 switches, line cards initialize Packet Forwarding Engine during startup. If an error occurs during hardware initialization, the FPCs with bad hardware parts power down after transferring the debug information to the Routing Engine. The Routing Engine marks the FPC offline, logs the error in system log messages (/var/log/messages), and generates an alarm to inform the user.

See the following sample output:

```
user@host> show chassis fpc
```

Utilization (%)	Temp	CPU	Utilization (%)	Memory
Slot State	(C)	Total	Interrupt	DRAM (MB) Heap
Buffer				
0 Empty				
1 Empty				
2 Empty				
3 Empty				
4 Empty				
5 Offline				
6 Empty				
7 Online	26	4	0	1024 0
32				

The following sample output shows the alarm raised for the failed FPCs.

```
user@host > show chassis alarms
4 alarms currently active
```

Alarm time	Class	Description
2011-03-24 00:52:51 UTC	Major	FPC 5 Hard errors
2011-03-24 00:52:31 UTC	Major	Fan Tray Failure
2011-03-24 00:52:31 UTC	Major	Fan Tray Failure
2011-03-24 00:51:26 UTC	Minor	Loss of communication with Backup RE



**NOTE:** On T4000 routers, when you include the enhanced-mode statement at the [edit chassis network-services] hierarchy level and reboot the system, only the T4000 Type 5 FPCs present on the router become online while the remaining FPCs are offline, and FPC misconfiguration alarms are generated. The show chassis alarm command output displays FPC misconfiguration (FPC *fpc-slot* misconfig) as the reason for the generation the alarms.

The following sample output shows the FPC status after the enhanced-mode statement is configured on the T4000 router. The T4000 Type 5 FPC present in slot 5 becomes online while the remaining FPCs are offline.

```
user@host> show chassis fpc
```

	Temp	CPU Utilization (%)	Memory
Utilization (%)			
Slot State	(C)	Total	Interrupt
Buffer			
0 offline	---	FPC misconfiguration---	
1 offline	---	FPC misconfiguration---	
2 offline	---	FPC misconfiguration---	
3 Empty			
4 Empty			
5 Online	66	50	0
27			2816 29

The following sample output shows FPC misconfiguration alarms.

```
user@host > show chassis alarms
```

3 alarms currently active

Alarm time	Class	Description
2011-03-24 00:52:51 PST	Major	FPC 1 misconfig
2011-03-24 00:52:31 PST	Major	FPC 2 misconfig
2011-03-24 00:52:31 PST	Major	FPC 3 misconfig

**detail**—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see *fpc-slot* or *slot*).

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on all members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display status information for all FPCs on the Interconnect device.

***fpc-slot***—(Optional) FPC slot number:

- (TX Matrix and TX Matrix Plus router only)—On a TX Matrix router, if you specify the number of the T640 router (line-card chassis) by using the **lcc number** option (the recommended method), replace ***fpc-slot*** with a value from 0 through 7. Otherwise, replace ***fpc-slot*** with a value from 0 through 31. Likewise, on a TX Matrix Plus router, if you specify the number of the specified router (line-card chassis) by using the **lcc number** option (the recommended method), replace ***fpc-slot*** with

a value from 0 through 7. Otherwise, replace *fpc-slot* with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis fpc detail 1 lcc 1
user@host> show chassis fpc detail 9
```

- M120 router—Replace *fpc-slot* with a value from 0 through 5.
- MX80 router—Replace *fpc-slot* with a value from 0 through 1.
- MX104 router—Replace *fpc-slot* with a value from 0 through 2.
- MX240 router—Replace *fpc-slot* with a value from 0 through 2.
- MX480 router—Replace *fpc-slot* with a value from 0 through 5.
- MX-960 router—Replace *fpc-slot* with a value from 0 through 11.
- MX2010 router—Replace *fpc-slot-number* with a value from 0 through 9.
- MX2020 router—Replace *fpc-slot-number* with a value from 0 through 19.
- Other routers—Replace *fpc-slot* with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace *fpc-slot* with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace *fpc-slot* with a value from 0 through 9.
  - EX6210 switches—Replace *fpc-slot* with a value from 0 through 9.
  - EX8208 switches—Replace *fpc-slot* with a value from 0 through 7.
  - EX8216 switches—Replace *fpc-slot* with a value from 0 through 15.
- QFX Series:
  - QFX3500 switches—Replace *fpc-slot* with 0.
  - QFabric systems—Replace *fpc-slot* with 0 through 31 on the Interconnect device.
- PTX Series Packet Transport Routers:
  - PTX5000 Packet Transport Router—Replace *fpc-slot* with a value from 0 through 7.
- ACX Series Universal Access Routers:
  - ACX1000 and ACX2000 Universal Access Routers—Replace *fpc-slot* with 0.

**local**—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display status information for all FPCs on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display status information for each Node device. Each Node device is equivalent to an FPC.

**pic-status**—(Optional) Display status information for all PICs or for the PIC in the specified slot (see *fpc-slot*).



**NOTE:** On T1600 routers, Type 4 FPCs with ASICs based on the SL2.0 chipset do not support the 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (10x10GE [LAN/WAN] SFPP). If you issue the `show chassis fpc` command with the `pic-status` option, the CLI displays the string “Not Supported” for 10x10GE (LAN/WAN) SFPP PICs installed on such FPCs. The following is a sample output:

```
user@host> show chassis fpc pic-status
Slot 0  Online      E2-FPC Type 1
        PIC 0  Online      1x G/E SFP, 1000 BASE
        PIC 1  Online      Adaptive Services-II
        PIC 2  Online      1x G/E IQ, 1000 BASE
        PIC 3  Online      1x G/E IQ, 1000 BASE
Slot 1  Online      FPC Type 3-ES
        PIC 0  Present     UNUSED- Not Supported
Slot 2  Online      FPC Type 4-ES
        PIC 0  Offline     4x OC-192 SONET XFP
        PIC 1  Present     10x10GE(LAN/WAN) SFPP- Not Supported
<<<<<<
Slot 4  Offline     FPC Type 1-ES
Slot 5  Offline     FPC Type 2-ES
Slot 6  Online      E2-FPC Type 3
        PIC 0  Online      1x OC-192 SONET XFP
        PIC 1  Online      4x OC-48 SONET
        PIC 2  Online      4x OC-48 SONET
        PIC 3  Online      MultiServices 500
Slot 7  Online      FPC Type 4-ES
        PIC 0  Online      4x 10GE (LAN/WAN) XFP
        PIC 1  Online      4x 10GE (LAN/WAN) XFP
```

In addition, an entry is logged in the system log messages (/var/log/messages) that the PIC is not supported. The following is a sample message logged in the system log:

```
Apr  5 08:47:36 router1 chassisd[2770]: CHASSISD_UNSUPPORTED_PIC:
PIC 1 in FPC 2 (type 763, version 257) is not supported
```

If you see this issue, contact Juniper Networks Technical Assistance Center (JTAC) for a possible fix. For more information about this issue and a possible solution, see [PSN-2010-03-696](#).



**NOTE:** When there is a double-bit ECC error in a network processor's memory, the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP or Channelized E1/T1 Circuit Emulation MIC is switched to the offline state.

```
user@host> show chassis fpc pic-status
Slot 1   Online      MPC Type 2 3D Q
PIC 0    Offline     1xC0C12/4xC0C3 CH-CE- ECC error detected
```

**lcc number**—(TX Matrix router and TX Matrix Plus router only) (Optional) Line-card chassis number.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [request chassis fpc on page 384](#)
- *show chassis fpc-feb-connectivity*
- *show chassis fabric fpcs*
- *Configuring the Junos OS to Resynchronize FPC Sequence Numbers with Active FPCs when an FPC Comes Online*
- *MX960 Flexible PIC Concentrator Description*
- *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*
- *enhanced-mode*

**List of Sample Output**

[show chassis fpc \(EX6210 Switch\) on page 648](#)  
[show chassis fpc \(M10 Router\) on page 648](#)  
[show chassis fpc \(M20 Router\) on page 648](#)  
[show chassis fpc detail \(M Series Routers\) on page 648](#)  
[show chassis fpc detail \(MX80 Router\) on page 649](#)  
[show chassis fpc \(MX104 Router\) on page 649](#)  
[show chassis fpc detail \(MX104 Router\) on page 649](#)  
[show chassis fpc pic-status \(MX104 Router\) on page 650](#)

[show chassis fpc \(MX240 Router\) on page 650](#)  
[show chassis fpc \(EX Series Switch\) on page 650](#)  
[show chassis fpc detail \(EX9200 Switch\) on page 650](#)  
[show chassis fpc \(MX480 Router\) on page 650](#)  
[show chassis fpc \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 651](#)  
[show chassis fpc pic-status \(MX480 Router with 100-Gigabit Ethernet CFP\) on page 651](#)  
[show chassis fpc pic-status \(EX Series Switch\) on page 651](#)  
[show chassis fpc \(MX480 Router with MPC4E\) on page 651](#)  
[show chassis fpc detail \(MX480 Router with MPC4E\) on page 652](#)  
[show chassis fpc \(MX480 Router with MPC4E\) on page 652](#)  
[show chassis fpc detail \(MX480 Router with MPC4E\) on page 652](#)  
[show chassis fpc \(MX960 Router\) on page 653](#)  
[show chassis fpc \(MX960 Router with MPC5EQ\) on page 653](#)  
[show chassis fpc detail \(MX960 Router with MPC5EQ\) on page 653](#)  
[show chassis fpc pic-status \(MX960 Router with MPC5EQ\) on page 655](#)  
[show chassis fpc \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 656](#)  
[show chassis fpc \(MX240, MX480, MX960 with Application Services Modular Line Card\) on page 656](#)  
[show chassis fpc \(MX2010 Routers\) on page 656](#)  
[show chassis fpc \(MX2020 Routers\) on page 656](#)  
[show chassis fpc \(MX2020 Router with MPC4E\) on page 657](#)  
[show chassis fpc detail \(MX2020 Router with MPC4E\) on page 657](#)  
[show chassis fpc \(MX2020 Router with MPC5EQ and MPC6E\) on page 658](#)  
[show chassis fpc detail \(MX2020 Router with MPC5EQ and MPC6E\) on page 658](#)  
[show chassis fpc pic-status \(MX2020 Router with MPC5EQ and MPC6E\) on page 660](#)  
[show chassis fpc detail \(MX Series Routers\) on page 661](#)  
[show chassis fpc detail \(EX Series Switches\) on page 661](#)  
[show chassis fpc \(Hardware Not Supported\) on page 661](#)  
[show chassis fpc detail \(Hardware Not Supported\) on page 662](#)  
[show chassis fpc pic-status on page 662](#)  
[show chassis fpc pic-status \(M Series Routers\) on page 662](#)  
[show chassis fpc pic-status \(M120 Router\) on page 663](#)  
[show chassis fpc pic-status \(MX240, MX480, and MX960 Routers with Application Services Modular Line Card\) on page 663](#)  
[show chassis fpc lcc \(TX Matrix Router\) on page 663](#)  
[show chassis fpc pic-status \(TX Matrix Router\) on page 663](#)  
[show chassis fpc pic-status lcc \(TX Matrix Router\) on page 664](#)  
[show chassis fpc \(TX Matrix Plus Router\) on page 664](#)  
[show chassis fpc lcc \(TX Matrix Plus Router\) on page 665](#)  
[show chassis fpc detail \(TX Matrix Plus Router\) on page 665](#)  
[show chassis fpc pic-status \(TX Matrix Plus Router\) on page 667](#)  
[show chassis fpc \(T1600 Router\) on page 668](#)  
[show chassis fpc detail \(T1600 Router\) on page 668](#)  
[show chassis fpc <fpc-slot> \(EX Series Switch\) on page 669](#)  
[show chassis fpc slot \(T1600 Router\) on page 669](#)  
[show chassis fpc pic-status \(T1600 Router\) on page 669](#)  
[show chassis fpc \(T4000 Router\) on page 670](#)  
[show chassis fpc detail \(T4000 Router\) on page 670](#)

[show chassis fpc pic-status \(T4000 Router\) on page 671](#)  
[show chassis fpc \(QFX Series\) on page 671](#)  
[show chassis fpc detail \(QFX3500 Switches\) on page 671](#)  
[show chassis fpc pic-status \(QFX3500 Switches\) on page 671](#)  
[show chassis fpc interconnect-device \(QFabric System\) on page 671](#)  
[show chassis fpc interconnect-device \(QFabric System\) on page 672](#)  
[show chassis fpc interconnect-device detail \(QFabric System\) on page 672](#)  
[show chassis fpc pic-status interconnect-device \(QFabric System\) on page 672](#)  
[show chassis fpc pic-status node-device \(QFabric System\) on page 673](#)  
[show chassis fpc \(PTX5000 Packet Transport Router\) on page 673](#)  
[show chassis fpc detail \(PTX5000 Packet Transport Router\) on page 673](#)  
[show chassis fpc pic-status \(PTX5000 Packet Transport Router\) on page 674](#)  
[show chassis fpc \(ACX2000 Universal Access Router\) on page 674](#)  
[show chassis fpc 0 \(ACX2000 Universal Access Router\) on page 674](#)  
[show chassis fpc detail \(ACX2000 Universal Access Router\) on page 674](#)  
[show chassis fpc pic-status \(ACX2000 Universal Access Router\) on page 675](#)  
[show chassis FPC 1 \(MX Routers with Media Services Blade \[MSB\]\) on page 675](#)  
[show chassis FPC 1 detail \(MX Routers with Media Services Blade \[MSB\]\) on page 675](#)

**Output Fields** Table 33 on page 646 lists the output fields for the **show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

**Table 33: show chassis fpc Output Fields**

Field Name	Field Description	Level of Output
<b>Slot or Slot State</b>	Slot number and state. The state can be one of the following conditions: <ul style="list-style-type: none"> <li>• <b>Dead</b>—Held in reset because of errors.</li> <li>• <b>Diag</b>—Slot is being ignored while the FPC is running diagnostics.</li> <li>• <b>Dormant</b>—Held in reset.</li> <li>• <b>Empty</b>—No FPC is present.</li> <li>• <b>Offline</b>—(PTX Series Packet Transport Routers only) One of the following two states is displayed:               <ul style="list-style-type: none"> <li>• <b>FPC offlined due to unreachable destinations</b></li> <li>• <b>FPC Offlined due to degraded FPC action</b></li> </ul> </li> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Present</b>—FPC is detected by the chassis daemon but either is not supported by the current version of Junos OS or is inserted in the wrong slot. The output also states either <b>Hardware Not Supported</b> or <b>Hardware Not In Right Slot</b>. The FPC is coming up but not yet online.</li> <li>• <b>Probed</b>—Probe is complete; awaiting restart of the Packet Forwarding Engine.</li> <li>• <b>Probe-wait</b>—Waiting to be probed.</li> </ul>	all levels
<b>Logical slot</b>	Slot number.	all levels
<b>Temp (C) or Temperature</b>	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.	all levels all levels



Table 33: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Temperature (PTX Series)</b>	On PTX Series Packet Transport Routers, temperature details are provided in degrees Celsius and Fahrenheit. Output includes: <ul style="list-style-type: none"> <li>• Temperature (PMB)—Temperature of the air passing by the Processor Mezzanine Board (PMB) at the bottom of the FPC.</li> <li>• Temperature (Intake)—Temperature of the air flowing into the chassis.</li> <li>• Temperature (Exhaust)—Exhaust temperatures for multiple zones (Exhaust A and Exhaust B).</li> <li>• Temperature (TLn)—Temperature of the specified Lookup ASIC (TL) of the packet forwarding engine on the FPC.</li> <li>• Temperature (TQn)—Temperature of the specified Queuing and Memory Interface ASIC (TQ) of the packet forwarding engine on the FPC.</li> </ul>	<b>detail</b>
<b>Total CPU Utilization (%)</b>	Total percentage of CPU being used by the FPC's processor.	all levels
<b>Interrupt CPU Utilization (%)</b>	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.	none specified
<b>Memory DRAM (MB)</b>	Total DRAM, in megabytes, available to the FPC's processor.	none specified
<b>Heap Utilization (%)</b>	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).  <b>NOTE:</b> On MX Series routers and EX Series switches in a broadband edge environment, heap utilization levels higher than 70 percent can affect unified ISSU, router stability, or scaling capability.	none specified
<b>Buffer Utilization (%)</b>	Percentage of buffer space being used by the FPC's processor for buffering internal messages.	none specified
<b>Total CPU DRAM</b>	Amount of DRAM available to the FPC's CPU.	<b>detail</b>
<b>Total RLDRAM</b>	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.	<b>detail</b>
<b>Total DDR DRAM</b>	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.	<b>detail</b>
<b>Total SRAM</b>	Amount of static RAM (SRAM) used by the FPC's CPU.	<b>detail</b>
<b>Total SDRAM</b>	Total amount of memory used for storing packets and notifications.	<b>detail</b>
<b>I/O Manager ASICs information</b>	I/O Manager version number, manufacturer, and part number.	<b>detail</b>
<b>Start time</b>	Time when the Routing Engine detected that the FPC was running.	<b>detail</b>

Table 33: show chassis fpc Output Fields (*continued*)

Field Name	Field Description	Level of Output
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.	detail
PIC type	(pic-status output only) Type of PIC.	none specified

## Sample Output

### show chassis fpc (EX6210 Switch)

```

user@switch> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Empty					
1	Online	7	5 0	1024	0	32
2	Empty					
3	Empty					
4	Online	25	17 2	2048	0	30
5	Online	25	3 0	2048	0	24
6	Online	6	5 0	1024	0	32
7	Empty					
8	Empty					
9	Online	8	7 0	1024	0	32

### show chassis fpc (M10 Router)

```

user@host> show chassis fpc
FPC status:

```

Slot	State	Temp (C)
0	Online	27
1	Online	28

### show chassis fpc (M20 Router)

```

user@host> show chassis fpc
FPC status:

```

Slot	State	Temp (C)	CPU Utilization (%) Total Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
0	Empty	0	0 0	0	0	0
1	Online	38	0 0	8	0	4
2	Online	35	0 0	8	0	3
3	Empty	0	0 0	0	0	0

### show chassis fpc detail (M Series Routers)

```

user@host> show chassis fpc detail 1
Slot 1 information:
State Online
Temperature 48 degrees C
Total CPU DRAM 32 MB
Total SRAM 4 MB
Total SDRAM 256 MB
I/O Manager ASICs information Version 2.0, Foundry IBM, Part number 0
I/O Manager ASICs information Version 2.0, Foundry IBM, Part number 0

```

```

Start time          2000-02-08 02:18:49 UTC
Uptime              14 hours, 41 minutes, 41 seconds

```

### show chassis fpc detail (MX80 Router)

```

user@host> show chassis fpc detail
Slot 0 information:
  State              Online
  Temperature        47 degrees C / 116 degrees F
  Total CPU DRAM     1024 MB
  Total SRAM         331 MB
  Total SDRAM        1280 MB
  Start time         2010-02-08 12:25:33 PST
  Uptime             2 hours, 13 minutes, 19 seconds
Slot 1 information:
  State              Online
  Temperature        47 degrees C / 116 degrees F
  Total CPU DRAM     1024 MB
  Total SRAM         331 MB
  Total SDRAM        1280 MB
  Start time         2010-02-08 12:25:33 PST
  Uptime             2 hours, 13 minutes, 19 seconds

```

### show chassis fpc (MX104 Router)

```

user@host> show chassis fpc
Temp CPU Utilization (%) Memory Utilization (%)
Slot State (C) Total Interrupt DRAM (MB) Heap Buffer
0 Online 32 15 5 2048 22 13
1 Online 32 15 5 2048 22 13
2 Online 32 15 5 2048 22 13

```

### show chassis fpc detail (MX104 Router)

```

user@host> show chassis fpc detail
Slot 0 information:
  State              Online
  Temperature        32 (C)
  Total CPU DRAM     2048 MB
  Total SRAM         403 MB
  Total SDRAM        1316 MB
  Start time         2013-05-23 14:39:18 IST
  Uptime             1 hour, 20 minutes, 22 seconds
Slot 1 information:
  State              Online
  Temperature        32 (C)
  Total CPU DRAM     2048 MB
  Total SRAM         403 MB
  Total SDRAM        1316 MB
  Start time         2013-05-23 14:39:18 IST
  Uptime             1 hour, 20 minutes, 22 seconds
Slot 2 information:
  State              Online
  Temperature        32 (C)
  Total CPU DRAM     2048 MB
  Total SRAM         403 MB
  Total SDRAM        1316 MB
  Start time         2013-05-23 14:39:18 IST
  Uptime             1 hour, 20 minutes, 22 seconds

```

**show chassis fpc pic-status (MX104 Router)**

```

user@host> show chassis fpc pic-status
Slot 0   Online
Slot 1   Online
  PIC 0   Online      10x 1GE(LAN) -E SFP
  PIC 1   Online      10x 1GE(LAN) -E SFP
Slot 2   Online
  PIC 0   Online      4x 10GE(LAN) SFP+

```

**show chassis fpc (MX240 Router)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total	Interrupt	Heap Buffer
0	Empty				
1	Online	34	6	0	1024 18 30
2	Online	33	9	0	1024 24 30

**show chassis fpc (EX Series Switch)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total	Interrupt	Heap Buffer
0	Empty				
1	Online	41	13	0	2048 19 14
2	Online	42	12	0	2048 19 14

**show chassis fpc detail (EX9200 Switch)**

```

user@switch> show chassis fpc detail
Slot 2 information:
  State                               Online
  Temperature                         37
  Total CPU DRAM                      2048 MB
  Total RLDRAM                        331 MB
  Total DDR DRAM                      1536 MB
  Start time:                        2014-03-12 15:35:28 UTC
  Uptime:                            1 hour, 4 minutes, 29 seconds
  Max Power Consumption               239 Watts
Slot 3 information:
  State                               Online
  Temperature                         39
  Total CPU DRAM                      2048 MB
  Total RLDRAM                        1036 MB
  Total DDR DRAM                      6656 MB
  Start time:                        2014-03-12 15:00:18 UTC
  Uptime:                            1 hour, 39 minutes, 39 seconds
  Max Power Consumption               520 Watts

```

**show chassis fpc (MX480 Router)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total	Interrupt	Heap Buffer
0	Empty				
1	Online	36	9	0	1024 17 57
2	Empty				
3	Empty				
4	Empty				
5	Empty				

## show chassis fpc (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Online	33	4	0		2048	10	13
1	Online	36	7	0		2048	16	13
2	Online	29	6	0		1024	27	29
3	Online	33	0	0		0	0	0
4	Online	36	7	0		2048	19	13
5	Online	34	31	11		2048	14	13

## show chassis fpc pic-status (MX480 Router with 100-Gigabit Ethernet CFP)

```

user@host> show chassis fpc pic-status

```

Slot 1	Online	MPC Type 3
PIC 2	Online	1X100GE CFP
Slot 2	Online	DPCE 40x 1GE R EQ
PIC 0	Online	10x 1GE(LAN) EQ
PIC 1	Online	10x 1GE(LAN) EQ
PIC 2	Online	10x 1GE(LAN) EQ
PIC 3	Online	10x 1GE(LAN) EQ
Slot 3	Online	MPC Type 3
PIC 0	Online	1X100GE CFP
PIC 2	Online	1X100GE CFP
Slot 4	Online	MPC Type 3
PIC 0	Online	1X100GE CFP
PIC 2	Online	1X100GE CFP
Slot 5	Online	MPC Type 2 3D EQ
PIC 0	Online	2x 10GE XFP
PIC 1	Online	2x 10GE XFP
PIC 2	Online	10x 1GE(LAN) SFP
PIC 3	Online	10x 1GE(LAN) SFP

## show chassis fpc pic-status (EX Series Switch)

```

user@host> show chassis fpc pic-status

```

Slot 1	Online	EX9200 32x10G SFP
PIC 0	Online	8X10GE SFPP
PIC 1	Online	8X10GE SFPP
PIC 2	Online	8X10GE SFPP
PIC 3	Online	8X10GE SFPP
Slot 2	Online	EX9200 32x10G SFP
PIC 0	Online	8X10GE SFPP
PIC 1	Online	8X10GE SFPP
PIC 2	Online	8X10GE SFPP
PIC 3	Online	8X10GE SFPP

## show chassis fpc (MX480 Router with MPC4E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Interrupt	Utilization (%)	DRAM (MB)	Heap	Buffer
0	Empty							
1	Empty							
2	Online		38	7	0	2048	19	14
3	Online		39	8	0	2048	18	14
4	Online		39	7	0	2048	17	14
5	Empty							

**show chassis fpc detail (MX480 Router with MPC4E)**

```

user@host> show chassis fpc detail
Slot 2 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:06:57 PST
  Uptime:               17 hours, 41 minutes, 9 seconds
  Max Power Consumption 610 Watts
Slot 3 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:07:00 PST
  Uptime:               17 hours, 41 minutes, 6 seconds
  Max Power Consumption 610 Watts
Slot 4 information:
  State                Diagnostics
  Temperature           37
  Total CPU DRAM        0 MB
  Total RLDRAM          0 MB
  Total DDR DRAM        0 MB
  Max Power Consumption 520 Watts

```

**show chassis fpc (MX480 Router with MPC4E)**

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)	DRAM (MB)	Heap	Buffer
0	Empty						
1	Empty						
2	Online	38	7	0	2048	19	14
3	Online	39	8	0	2048	18	14
4	Online	39	7	0	2048	17	14
5	Empty						

**show chassis fpc detail (MX480 Router with MPC4E)**

```

user@host> show chassis fpc detail
Slot 2 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:06:57 PST
  Uptime:               17 hours, 41 minutes, 9 seconds
  Max Power Consumption 610 Watts
Slot 3 information:
  State                Online
  Temperature           38
  Total CPU DRAM        2048 MB
  Total RLDRAM          1036 MB
  Total DDR DRAM        11264 MB
  Start time:           2013-02-18 05:07:00 PST
  Uptime:               17 hours, 41 minutes, 6 seconds

```

```

Max Power Consumption      610 Watts
Slot 4 information:
State                      Diagnostics
Temperature                37
Total CPU DRAM             0 MB
Total RLD RAM              0 MB
Total DDR DRAM             0 MB
Max Power Consumption      520 Watts

```

#### show chassis fpc (MX960 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	Total	CPU Utilization (%) Interrupt	Memory DRAM (MB)	Heap	Utilization (%) Buffer
0	Empty						
1	Empty						
2	Empty						
3	Online	25	19	0	1024	15	57
4	Empty						
5	Online	26	27	0	1024	15	57
6	Empty						
7	Empty						
8	Empty						
9	Empty						
10	Empty						
11	Empty						

#### show chassis fpc (MX960 Router with MPC5EQ)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	Total	CPU Utilization (%) Interrupt	Memory DRAM (MB)	Heap	Utilization (%) Buffer
0	Online	38	16	0	3584	7	13
1	Online	31	15	0	2048	17	13
2	Empty						
3	Online	31	14	0	2048	20	13
4	Online	34	16	0	3584	7	13
5	Online	34	16	0	3584	7	13
6	Empty						
7	Online	32	9	0	2048	18	14
8	Online	36	19	0	3584	7	13
9	Online	31	9	0	2048	13	13
10	Online	35	14	0	3584	7	13
11	Online	33	11	0	2048	18	14

#### show chassis fpc detail (MX960 Router with MPC5EQ)

```

user@host> show chassis fpc detail
Slot 0 information:
State                      Online
Temperature                38
Total CPU DRAM             3584 MB
Total XR2                  291 MB
Total DDR DRAM             24960 MB
Start time:                2014-04-22 10:01:46 PDT
Uptime:                    1 hour, 23 minutes, 40 seconds
Max Power Consumption      607 Watts
Slot 1 information:
State                      Online
Temperature                31
Total CPU DRAM             2048 MB
Total RLD RAM              1036 MB

```

```
Total DDR DRAM                6656 MB
Start time:                    2014-04-22 10:01:50 PDT
Uptime:                        1 hour, 23 minutes, 36 seconds
Max Power Consumption          520 Watts
Slot 3 information:
State                          Online
Temperature                    31
Total CPU DRAM                 2048 MB
Total RLD RAM                  1324 MB
Total DDR DRAM                 5120 MB
Start time:                    2014-04-22 10:01:50 PDT
Uptime:                        1 hour, 23 minutes, 36 seconds
Max Power Consumption          440 Watts
Slot 4 information:
State                          Online
Temperature                    34
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:01:54 PDT
Uptime:                        1 hour, 23 minutes, 32 seconds
Max Power Consumption          607 Watts
Slot 5 information:
State                          Online
Temperature                    34
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:01:56 PDT
Uptime:                        1 hour, 23 minutes, 30 seconds
Max Power Consumption          607 Watts
Slot 7 information:
State                          Online
Temperature                    32
Total CPU DRAM                 2048 MB
Total RLD RAM                  1036 MB
Total DDR DRAM                 11264 MB
Start time:                    2014-04-22 10:02:02 PDT
Uptime:                        1 hour, 23 minutes, 24 seconds
Max Power Consumption          608 Watts
Slot 8 information:
State                          Online
Temperature                    36
Total CPU DRAM                 3584 MB
Total XR2                      291 MB
Total DDR DRAM                 24960 MB
Start time:                    2014-04-22 10:02:07 PDT
Uptime:                        1 hour, 23 minutes, 19 seconds
Max Power Consumption          607 Watts
Slot 9 information:
State                          Online
Temperature                    31
Total CPU DRAM                 2048 MB
Total RLD RAM                  734 MB
Total DDR DRAM                 3108 MB
Start time:                    2014-04-22 10:02:05 PDT
Uptime:                        1 hour, 23 minutes, 21 seconds
Max Power Consumption          368 Watts
Slot 10 information:
State                          Online
Temperature                    35
```



```

Total CPU DRAM          3584 MB
Total XR2                291 MB
Total DDR DRAM          24960 MB
Start time:              2014-04-22 10:02:11 PDT
Uptime:                  1 hour, 23 minutes, 15 seconds
Max Power Consumption    607 Watts
Slot 11 information:
State                    Online
Temperature              33
Total CPU DRAM          2048 MB
Total RLDRAM             1036 MB
Total DDR DRAM          11264 MB
Start time:              2014-04-22 10:02:16 PDT
Uptime:                  1 hour, 23 minutes, 10 seconds
Max Power Consumption    608 Watts

```

### show chassis fpc pic-status(MX960 Router with MPC5EQ)

```

user@host> show chassis fpc pic-status
Slot 0  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 1  Online      MPCE Type 3 3D
PIC 0   Online      10X10GE SFPP
PIC 2   Online      1X100GE CXP
Slot 3  Online      MPC 3D 16x 10GE
PIC 0   Online      4x 10GE(LAN) SFP+
PIC 1   Online      4x 10GE(LAN) SFP+
PIC 2   Online      4x 10GE(LAN) SFP+
PIC 3   Online      4x 10GE(LAN) SFP+
Slot 4  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 5  Online      MPC5E 3D Q 2CGE+4XGE
PIC 0   Online      2X10GE SFPP OTN
PIC 1   Online      1X100GE CFP2 OTN
PIC 2   Online      2X10GE SFPP OTN
PIC 3   Online      1X100GE CFP2 OTN
Slot 7  Online      MPC4E 3D 2CGE+8XGE
PIC 0   Online      4x10GE SFPP
PIC 1   Online      1X100GE CFP
PIC 2   Online      4x10GE SFPP
PIC 3   Online      1X100GE CFP
Slot 8  Online      MPC5E 3D Q 24XGE+6XLGE
PIC 0   Offline     12X10GE SFPP OTN
PIC 1   Offline     12X10GE SFPP OTN
PIC 2   Online      3X40GE QSFPP
PIC 3   Online      3X40GE QSFPP
Slot 9  Online      MPCE Type 2 3D P
PIC 0   Online      2x 10GE XFP
PIC 1   Online      2x 10GE XFP
Slot 10 Online      MPC5E 3D Q 24XGE+6XLGE
PIC 0   Online      12X10GE SFPP
PIC 1   Online      12X10GE SFPP
PIC 2   Offline     3X40GE QSFPP
PIC 3   Offline     3X40GE QSFPP

```

```

Slot 11 Online      MPC4E 3D 2CGE+8XGE
PIC 0 Online      4x10GE SFPP
PIC 1 Online      1X100GE CFP
PIC 2 Online      4x10GE SFPP
PIC 3 Online      1X100GE CFP

```

#### show chassis fpc (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis fpc 1
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total  Interrupt      DRAM (MB) Heap      Buffer
  1 Online           34      5          0      3072      5      13

```

#### show chassis fpc (MX240, MX480, MX960 with Application Services Modular Line Card)

```

user@host> show chassis fpc 1 detail
Slot 1 information:
State                               Online
Temperature                         34
Total CPU DRAM                     3072 MB
Total RLDRAM                       259 MB
Total DDR DRAM                     4864 MB
Start time:                        2012-06-19 10:51:43 PDT
Uptime:                            16 minutes, 48 seconds
Max Power Consumption              550 Watts

```

#### show chassis fpc (MX2010 Routers)

```

user@host> show chassis fpc
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total  Interrupt      DRAM (MB) Heap      Buffer
  0 Online           34      9          0      2048      18      13
  1 Online           32      9          0      2048      15      13
  2 Empty
  3 Empty
  4 Empty
  5 Empty
  6 Empty
  7 Empty
  8 Online           31     13          0      2048      11      13
  9 Online           33     10          0      2048      18      13

```

#### show chassis fpc (MX2020 Routers)

```

user@host> show chassis fpc
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total  Interrupt      DRAM (MB) Heap      Buffer
  0 Online          10     12          0      2048      18      13
  1 Online           8      9          0      2048      18      13
  2 Online           7      9          0      2048      18      13
  3 Online           8     10          0      2048      18      13
  4 Online           9     10          0      2048      18      13
  5 Online           8      9          0      2048      18      13
  6 Online           8     10          0      2048      18      13
  7 Online           9      9          0      2048      18      13
  8 Online           9     10          0      2048      18      13
  9 Online          10      9          0      2048      18      13
 10 Online          16      8          0      2048      18      13
 11 Online          11     10          0      2048      18      13
 12 Online          10     10          0      2048      18      13
 13 Online          11      9          0      2048      18      13

```

14	Online	12	10	0	2048	18	13
15	Online	13	9	0	2048	18	13
16	Online	13	9	0	2048	18	13
17	Online	12	9	0	2048	18	13
18	Online	12	8	0	2048	18	13
19	Online	14	10	0	2048	18	13

#### show chassis fpc (MX2020 Router with MPC4E)

```

user@host> show chassis fpc
      Temp CPU Utilization (%) Memory      Utilization (%)
Slot State      (C) Total Interrupt      DRAM (MB) Heap      Buffer
0  Online          33    12         2      2048    11      13
1  Empty
2  Empty
3  Empty
4  Empty
5  Empty
6  Empty
7  Empty
8  Empty
9  Online          31    10         0      2048    11      13
10 Online          32     7         0      2048    14      13
11 Empty
12 Empty
13 Empty
14 Online          28    12         0      2048    15      14
15 Empty
16 Empty
17 Empty
18 Empty
19 Online          38     8         0      2048    18      13

```

#### show chassis fpc detail (MX2020 Router with MPC4E)

```

user@host> show chassis fpc detail
Slot 0 information:
  State                Online
  Temperature          34
  Total CPU DRAM       2048 MB
  Total RLD RAM        806 MB
  Total DDR DRAM       2632 MB
  Start time:          2013-02-17 08:17:35 PST
  Uptime:              1 day, 14 hours, 50 minutes, 39 seconds
  Max Power Consumption 368 Watts
Slot 9 information:
  State                Online
  Temperature          32
  Total CPU DRAM       2048 MB
  Total RLD RAM        806 MB
  Total DDR DRAM       2632 MB
  Start time:          2013-02-17 08:17:43 PST
  Uptime:              1 day, 14 hours, 50 minutes, 31 seconds
  Max Power Consumption 368 Watts
Slot 10 information:
  State                Online
  Temperature          37
  Total CPU DRAM       2048 MB
  Total RLD RAM       1036 MB
  Total DDR DRAM       6656 MB
  Start time:          2013-02-17 08:17:54 PST

```

```

Uptime:                               1 day, 14 hours, 50 minutes, 20 seconds
Max Power Consumption                  520 Watts
Slot 14 information:
  State                               Online
  Temperature                         32
  Total CPU DRAM                      2048 MB
  Total RLDRAM                       1036 MB
  Total DDR DRAM                     11264 MB
  Start time:                        2013-02-17 08:18:01 PST
  Uptime:                            1 day, 14 hours, 50 minutes, 13 seconds
  Max Power Consumption               610 Watts
Slot 19 information:
  State                               Online
  Temperature                         38
  Total CPU DRAM                      2048 MB
  Total RLDRAM                       1324 MB
  Total DDR DRAM                     5120 MB
  Start time:                        2013-02-17 08:18:08 PST
  Uptime:                            1 day, 14 hours, 50 minutes, 6 seconds
  Max Power Consumption               440 Watts

```

#### show chassis fpc (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)	DRAM (MB)	Heap	Buffer
			Total	Interrupt			
0	Online	31	20	0	3584	7	13
1	Online	28	19	0	2048	17	13
2	Online	27	10	0	2048	18	14
3	Online	26	10	0	2048	13	13
4	Online	29	19	0	3584	7	13
5	Online	28	68	0	2048	20	13
6	Empty						
7	Empty						
8	Empty						
9	Online	36	19	0	3584	10	13
10	Online	37	26	0	3584	10	13
11	Empty						
12	Empty						
13	Empty						
14	Empty						
15	Empty						
16	Empty						
17	Online	28	43	0	3584	10	13
18	Online	29	19	0	3584	7	13
19	Online	31	19	0	3584	7	13

#### show chassis fpc detail (MX2020 Router with MPCEQ and MPC6E)

```

user@host> show chassis fpc detail
Slot 0 information:
  State                               Online
  Temperature                         31
  Total CPU DRAM                      3584 MB
  Total XR2                          291 MB
  Total DDR DRAM                     24960 MB
  Start time:                        2014-04-22 23:33:19 PDT
  Uptime:                            6 minutes, 24 seconds
  Max Power Consumption               607 Watts
Slot 1 information:

```

```

State                               Online
Temperature                         28
Total CPU DRAM                     2048 MB
Total RLD RAM                      1036 MB
Total DDR DRAM                     6656 MB
Start time:                        2014-04-22 23:33:24 PDT
Uptime:                            6 minutes, 19 seconds
Max Power Consumption              520 Watts
Slot 2 information:
State                               Online
Temperature                         27
Total CPU DRAM                     2048 MB
Total RLD RAM                      1036 MB
Total DDR DRAM                     11264 MB
Start time:                        2014-04-22 23:33:34 PDT
Uptime:                            6 minutes, 9 seconds
Max Power Consumption              608 Watts
Slot 3 information:
State                               Online
Temperature                         26
Total CPU DRAM                     2048 MB
Total RLD RAM                      734 MB
Total DDR DRAM                     3108 MB
Start time:                        2014-04-22 23:33:39 PDT
Uptime:                            6 minutes, 4 seconds
Max Power Consumption              368 Watts
Slot 4 information:
State                               Online
Temperature                         29
Total CPU DRAM                     3584 MB
Total XR2                          291 MB
Total DDR DRAM                     24960 MB
Start time:                        2014-04-22 23:33:51 PDT
Uptime:                            5 minutes, 52 seconds
Max Power Consumption              607 Watts
Slot 5 information:
State                               Online
Temperature                         28
Total CPU DRAM                     2048 MB
Total RLD RAM                      1324 MB
Total DDR DRAM                     5120 MB
Start time:                        2014-04-22 23:33:57 PDT
Uptime:                            5 minutes, 46 seconds
Max Power Consumption              440 Watts
Slot 9 information:
State                               Online
Temperature                         25
Total CPU DRAM                     3584 MB
Total XR2                          518 MB
Total DDR DRAM                     49920 MB
Start time:                        2014-04-22 23:31:20 PDT
Uptime:                            8 minutes, 23 seconds
Max Power Consumption              1130 Watts
Slot 10 information:
State                               Online
Temperature                         32
Total CPU DRAM                     3584 MB
Total XR2                          518 MB
Total DDR DRAM                     49920 MB
Start time:                        2014-04-22 23:31:25 PDT
Uptime:                            8 minutes, 18 seconds

```

```

Max Power Consumption          1130 Watts
Slot 17 information:
  State                        Online
  Temperature                  25
  Total CPU DRAM               3584 MB
  Total XR2                    518 MB
  Total DDR DRAM               49920 MB
  Start time:                  2014-04-22 23:31:29 PDT
  Uptime:                      8 minutes, 14 seconds
  Max Power Consumption        1130 Watts
Slot 18 information:
  State                        Online
  Temperature                  29
  Total CPU DRAM               3584 MB
  Total XR2                    291 MB
  Total DDR DRAM               24960 MB
  Start time:                  2014-04-22 23:34:11 PDT
  Uptime:                      5 minutes, 32 seconds
  Max Power Consumption        607 Watts
Slot 19 information:
  State                        Online
  Temperature                  32
  Total CPU DRAM               3584 MB
  Total XR2                    291 MB
  Total DDR DRAM               24960 MB
  Start time:                  2014-04-22 23:34:20 PDT
  Uptime:                      5 minutes, 23 seconds
  Max Power Consumption        607 Watts

```

#### show chassis fpc pic-status (MX2020 Router with MPC5EQ and MPC6E)

```

user@host> show chassis fpc pic-status
Slot 0  Online      MPC5E 3D Q 24XGE+6XLGE
  PIC 0  Online      12X10GE SFPP OTN
  PIC 1  Online      12X10GE SFPP OTN
  PIC 2  Offline     3X40GE QSFP
  PIC 3  Offline     3X40GE QSFP
Slot 1  Online      MPCE Type 3 3D
  PIC 0  Online      10X10GE SFPP
  PIC 2  Online      1X100GE CXP
Slot 2  Online      MPC4E 3D 2CGE+8XGE
  PIC 0  Online      4x10GE SFPP
  PIC 1  Online      1X100GE CFP
  PIC 2  Online      4x10GE SFPP
  PIC 3  Online      1X100GE CFP
Slot 3  Online      MPCE Type 2 3D P
  PIC 0  Online      2x 10GE XFP
  PIC 1  Online      2x 10GE XFP
Slot 4  Online      MPC5E 3D Q 2CGE+4XGE
  PIC 0  Online      2X10GE SFPP OTN
  PIC 1  Online      1X100GE CFP2 OTN
  PIC 2  Online      2X10GE SFPP OTN
  PIC 3  Online      1X100GE CFP2 OTN
Slot 5  Online      MPC 3D 16x 10GE
  PIC 0  Online      4x 10GE(LAN) SFP+
  PIC 1  Online      4x 10GE(LAN) SFP+
  PIC 2  Online      4x 10GE(LAN) SFP+
  PIC 3  Online      4x 10GE(LAN) SFP+
Slot 9  Online      MPC6E 3D
  PIC 0  Online      2X100GE CFP2 OTN
  PIC 1  Online      2X100GE CFP2 OTN

```

```

Slot 10 Online MPC6E 3D
PIC 0 Online 24X10GE SFPP OTN
PIC 1 Online 4X100GE CXP
Slot 17 Online MPC6E 3D
PIC 0 Online 24X10GE SFPP
PIC 1 Online 4X100GE CXP
Slot 18 Online MPC5E 3D Q 24XGE+6XLGE
PIC 0 Offline 12X10GE SFPP OTN
PIC 1 Offline 12X10GE SFPP OTN
PIC 2 Online 3X40GE QSFPP
PIC 3 Online 3X40GE QSFPP
Slot 19 Online MPC5E 3D Q 24XGE+6XLGE
PIC 0 Online 12X10GE SFPP OTN
PIC 1 Offline 12X10GE SFPP OTN
PIC 2 Offline 3X40GE QSFPP
PIC 3 Online 3X40GE QSFPP

```

#### show chassis fpc detail (MX Series Routers)

```

user@host> show chassis fpc detail 2
Slot 0 information:
State Online
Temperature 36 degrees C / 96 degrees F
Total CPU DRAM 1024 MB
Total RLDRAM 256 MB
Total DDR DRAM 4096 MB
Start time: 2009-08-11 21:20:30 PDT
Uptime: 2 hours, 8 minutes, 50 seconds
Max Power Consumption 335 Watts

```

#### show chassis fpc detail (EX Series Switches)

```

user@host> show chassis fpc detail 2
Slot 1 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:52 PDT
Uptime: 7 days, 9 hours, 47 minutes, 46 seconds
Max Power Consumption 610 Watts
Slot 2 information:
State Online
Temperature 41
Total CPU DRAM 2048 MB
Total RLDRAM 1036 MB
Total DDR DRAM 11264 MB
Start time: 2013-04-02 00:04:56 PDT
Uptime: 7 days, 9 hours, 47 minutes, 42 seconds
Max Power Consumption 610 Watts

```

#### show chassis fpc (Hardware Not Supported)

```

user@host> show chassis fpc
show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Interrupt	Memory DRAM (MB)	Utilization (%)	Heap	Buffer
0	Online	-----	CPU less FPC	-----				
1	Present	-----	Hardware Not In Right Slot	-----				
2	Online	0	0	0	0	0	0	0
3	Present	-----	Hardware Not Supported	-----				

```
4 Empty
5 Empty
6 Online          0          0          0          0          0
```

#### show chassis fpc detail (Hardware Not Supported)

```
user@host> show chassis fpc detail
Slot 0 information:
  State          Online
  Total CPU DRAM ----- CPU less FPC -----
  Start time     2006-07-07 03:21:00 UTC
  Uptime         27 minutes, 51 seconds
Slot 1 information:
  State          Present
  Reason         --- Hardware Not In Right Slot ---
Slot 2 information:
  State          Online
  Total CPU DRAM 32 MB
  Start time     2006-07-07 03:20:59 UTC
  Uptime         27 minutes, 52 seconds
Slot 3 information:
  State          Present
  Reason         --- Hardware Not Supported ---
  Total CPU DRAM 0 MB
Slot 6 information:
  State          Online
  Total CPU DRAM 32 MB
  Start time     2006-07-07 03:21:01 UTC
  Uptime         27 minutes, 50 seconds
```

#### show chassis fpc pic-status

```
user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1  1x OC-12 ATM, MM
  PIC 2  1x OC-12 ATM, MM
  PIC 3  1x OC-12 ATM, MM
Slot 1 Online
  PIC 0  1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0  1x OC-192 SONET, SMSR
```

#### show chassis fpc pic-status (M Series Routers)

```
user@host> show chassis fpc pic-status
Slot 1 Online      FPC Type 1
  PIC 0 Present    2x OC-3 ATM, MM- Hardware Error
  PIC 1 Online     4x OC-3 SONET, SMIR
Slot 2 Online      E-FPC Type 2
  PIC 0 Online     4x G/E, 1000 BASE-SX
  PIC 1 Online     2x G/E SFP, 1000 BASE
  PIC 3 Online     1x Tunnel
Slot 3 Online      E-FPC Type 1
  PIC 0 Online     1x G/E IQ, 1000 BASE
  PIC 2 Online     1x G/E SFP, 1000 BASE
Slot 4 Online      E-FPC Type 2
  PIC 0 Online     4x G/E SFP, 1000 BASE
  PIC 1 Online     4x G/E SFP, 1000 BASE
  PIC 2 Online     4x G/E SFP, 1000 BASE
  PIC 3 Online     4x G/E SFP, 1000 BASE
```



```
Slot 5   Online       FPC Type 2
...
```

#### show chassis fpc pic-status (M120 Router)

```
user@host> show chassis fpc pic-status
Slot 1   Online       M120 CFPC 10GE
  PIC 0   Online       1x 10GE(LAN/WAN) XFP
Slot 3   Online       M120 FPC Type 2 (proto)
  PIC 0   Online       2x G/E IQ, 1000 BASE
  PIC 1   Online       4x OC-3 SONET, SMIR
  PIC 2   Online       2x G/E IQ, 1000 BASE
  PIC 3   Online       8x 1GE(LAN), IQ2
Slot 4   Online       M120 FPC Type 3 (proto)
  PIC 0   Online       10x 1GE(LAN), 1000 BASE
Slot 5   Online       M120 FPC Type 1 (proto)
  PIC 0   Present      1x G/E, 1000 BASE-LX- Not Supported
  PIC 1   Online       1x CHOC3 IQ SONET, SMLR
  PIC 2   Online       4x CHDS3 IQ
  PIC 3   Online       1x G/E SFP, 1000 BASE
```

#### show chassis fpc pic-status (MX240, MX480, and MX960 Routers with Application Services Modular Line Card)

In the following output **Slot 1** and **Slot 5** are the Application Services Modular Carrier Cards (AS MCC), **PIC 0** is the Application Services Modular Storage Card (AS MSC), and **PIC 2** is the Application Services Modular Processing Card (AS MXC).

```
user@host> show chassis fpc pic-status
Slot 2   Online       MPC Type 1 3D Q
  Slot 1   Online       AS-MCC
  PIC 0   Online       AS-MSC
  PIC 2   Online       AS-MXC
Slot 4   Offline      MPC 3D 16x 10GE
Slot 5   Offline      AS-MCC
```

#### show chassis fpc lcc (TX Matrix Router)

```
user@host> show chassis fpc lcc 0
lcc0-re0:
-----
Slot State      Temp CPU      Utilization (%)  Memory  Utilization (%)
      (C) Total Interrupt      DRAM (MB)      Heap      Buffer
0 Empty
1 Online        27    2         0        256        8        44
2 Online        27    3         0        256       15        44
3 Empty
4 Empty
5 Empty
6 Empty
7 Empty
```

#### show chassis fpc pic-status (TX Matrix Router)

```
user@host> show chassis fpc pic-status
lcc0-re0:
-----
Slot 0   Online       FPC Type 3
  PIC 0   Online       1x OC-192 SM SR1
  PIC 1   Online       1x OC-192 SM SR2
  PIC 2   Online       1x OC-192 SM SR1
  PIC 3   Online       1x Tunnel
```

```

Slot 1  Online      FPC Type 2
PIC 0   Online      1x OC-48 SONET, SMSR
PIC 1   Online      1x OC-48 SONET, SMSR

```

```
lcc1-re0:
```

```
lcc2-re0:
```

```

Slot 1  Online      FPC Type 3
PIC 0   Online      1x OC-192 SM SR1
Slot 5  Online      FPC Type 2
PIC 0   Online      1x OC-48 SONET, SMSR
PIC 1   Online      2x G/E, 1000 BASE-LX
PIC 2   Online      2x G/E, 1000 BASE-LX
PIC 3   Online      1x OC-48 SONET, SMSR

```

```
lcc3-re0:
```

#### show chassis fpc pic-status lcc (TX Matrix Router)

```

user@host> show chassis fpc pic-status lcc 0
lcc0-re0:

```

```

Slot 0  Online      FPC Type 3
PIC 0   Online      1x OC-192 SM SR2
Slot 1  Online      FPC Type 2
PIC 0   Online      2x OC-12 ATM2 IQ, MM
PIC 1   Online      1x OC-48 SONET, SMSR
PIC 2   Online      1x OC-48 SONET, SMSR
PIC 3   Online      4x G/E, 1000 BASE-SX

```

#### show chassis fpc (TX Matrix Plus Router)

```

user@host> show chassis fpc
lcc0-re0:

```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Online	38	4	0	2048	3	24
2	Online	43	8	0	2048	6	24
3	Empty						
4	Online	43	6	0	2048	6	24
5	Empty						
6	Online	42	13	0	2048	6	24
7	Online	45	7	0	2048	3	24

```
lcc2-re0:
```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Online	42	10	0	2048	6	24
1	Empty						
2	Online	42	11	0	2048	6	24
3	Online	40	5	0	2048	3	24
4	Online	33	26	0	1024	8	49
5	Empty						
6	Online	43	8	0	2048	6	24
7	Online	46	6	0	2048	3	24

lcc3-re0:

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty						
1	Empty						
2	Online	39	30	0	2048	7	24
3	Empty						
4	Online	41	8	0	2048	6	24
5	Online	41	12	0	2048	6	24
6	Online	40	8	0	2048	6	24
7	Online	42	4	0	2048	3	24

**show chassis fpc lcc (TX Matrix Plus Router)**

user@host&gt; show chassis fpc lcc 0

lcc0-re0:

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Empty						
1	Online	38	4	0	2048	3	24
2	Online	43	8	0	2048	6	24
3	Empty						
4	Online	43	6	0	2048	6	24
5	Empty						
6	Online	42	14	0	2048	6	24
7	Online	45	6	0	2048	3	24

**show chassis fpc detail (TX Matrix Plus Router)**

user@host&gt; show chassis fpc details

lcc0-re0:

Slot 1 information:

```

State                               Online
Temperature                         38 degrees C / 100 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          64 MB
Total SDRAM                         1280 MB
Start time                          2010-10-04 20:06:22 PDT
Uptime                              1 hour, 32 minutes, 51 seconds

```

Slot 2 information:

```

State                               Online
Temperature                         43 degrees C / 109 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                          2010-10-04 20:06:37 PDT
Uptime                              1 hour, 32 minutes, 36 seconds

```

Slot 4 information:

```

State                               Online
Temperature                         43 degrees C / 109 degrees F
Total CPU DRAM                      2048 MB
Total SRAM                          128 MB
Total SDRAM                         2560 MB
Start time                          2010-10-04 20:06:40 PDT
Uptime                              1 hour, 32 minutes, 33 seconds

```

Slot 6 information:

State	Online
Temperature	42 degrees C / 107 degrees F
Total CPU DRAM	2048 MB
Total SRAM	128 MB
Total SDRAM	2560 MB
Start time	2010-10-04 20:06:42 PDT
Uptime	1 hour, 32 minutes, 31 seconds

## Slot 7 information:

State	Online
Temperature	45 degrees C / 113 degrees F
Total CPU DRAM	2048 MB
Total SRAM	64 MB
Total SDRAM	1280 MB
Start time	2010-10-04 20:06:43 PDT
Uptime	1 hour, 32 minutes, 30 seconds

lcc2-re0:  
-----

## Slot 0 information:

State	Online
Temperature	42 degrees C / 107 degrees F
Total CPU DRAM	2048 MB
Total SRAM	128 MB
Total SDRAM	2560 MB
Start time	2010-10-04 20:06:35 PDT
Uptime	1 hour, 32 minutes, 38 seconds

## Slot 2 information:

State	Online
Temperature	42 degrees C / 107 degrees F
Total CPU DRAM	2048 MB
Total SRAM	128 MB
Total SDRAM	2560 MB
Start time	2010-10-04 20:06:37 PDT
Uptime	1 hour, 32 minutes, 36 seconds

## Slot 3 information:

State	Online
Temperature	40 degrees C / 104 degrees F
Total CPU DRAM	2048 MB
Total SRAM	64 MB
Total SDRAM	1280 MB
Start time	2010-10-04 20:06:28 PDT
Uptime	1 hour, 32 minutes, 45 seconds

## Slot 4 information:

State	Online
Temperature	33 degrees C / 91 degrees F
Total CPU DRAM	1024 MB
Total SRAM	64 MB
Total SDRAM	1280 MB
Start time	2010-10-04 20:08:03 PDT
Uptime	1 hour, 31 minutes, 10 seconds

## Slot 6 information:

State	Online
Temperature	43 degrees C / 109 degrees F
Total CPU DRAM	2048 MB
Total SRAM	128 MB
Total SDRAM	2560 MB
Start time	2010-10-04 20:06:44 PDT
Uptime	1 hour, 32 minutes, 29 seconds

## Slot 7 information:

State	Online
Temperature	46 degrees C / 114 degrees F

```

Total CPU DRAM          2048 MB
Total SRAM               64 MB
Total SDRAM             1280 MB
Start time              2010-10-04 20:06:46 PDT
Uptime                  1 hour, 32 minutes, 27 seconds

```

lcc3-re0:

-----

Slot 2 information:

```

State                  Online
Temperature            38 degrees C / 100 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:31 PDT
Uptime                1 hour, 21 minutes, 42 seconds

```

Slot 4 information:

```

State                  Online
Temperature            41 degrees C / 105 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:34 PDT
Uptime                1 hour, 21 minutes, 39 seconds

```

Slot 5 information:

```

State                  Online
Temperature            41 degrees C / 105 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:36 PDT
Uptime                1 hour, 21 minutes, 37 seconds

```

Slot 6 information:

```

State                  Online
Temperature            40 degrees C / 104 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2010-10-04 20:17:39 PDT
Uptime                1 hour, 21 minutes, 34 seconds

```

Slot 7 information:

```

State                  Online
Temperature            42 degrees C / 107 degrees F
Total CPU DRAM        2048 MB
Total SRAM            64 MB
Total SDRAM           1280 MB
Start time            2010-10-04 20:17:41 PDT
Uptime                1 hour, 21 minutes, 32 seconds

```

### show chassis fpc pic-status (TX Matrix Plus Router)

```
user@host> show chassis fpc pic-status
```

lcc0-re0:

```

-----
Slot 1  Online      FPC Type 2-ES
PIC 0   Online      8x 1GE(LAN), IQ2
Slot 2  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP
Slot 4  Online      FPC Type 4-ES
PIC 0   Online      4x 10GE (LAN/WAN) XFP

```

```

Slot 6  Online      FPC Type 4-ES
      PIC 0  Online      4x 10GE (LAN/WAN) XFP
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
      PIC 0  Online      10x 1GE(LAN), 1000 BASE
      PIC 2  Online      1x OC-192 SM SR2
      PIC 3  Online      10x 1GE(LAN), 1000 BASE

```

lcc2-re0:

```

-----
Slot 0  Online      FPC Type 4-ES
      PIC 0  Online      4x 10GE (LAN/WAN) XFP
Slot 2  Online      FPC Type 4-ES
      PIC 0  Online      4x 10GE (LAN/WAN) XFP
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 3  Online      FPC Type 2-ES
      PIC 0  Online      8x 1GE(LAN), IQ2
Slot 4  Online      FPC Type 4
      PIC 0  Online      10x10GE(LAN/WAN) SFPP
Slot 6  Online      FPC Type 4-ES
      PIC 0  Online      4x OC-192 SONET XFP
Slot 7  Online      FPC Type 3-ES
      PIC 0  Online      10x 1GE(LAN), 1000 BASE
      PIC 1  Offline     1x 10GE(LAN/WAN) IQ2E
      PIC 2  Online      1x OC-192 SM SR2
      PIC 3  Online      1x Tunnel

```

lcc3-re0:

```

-----
Slot 2  Online      FPC Type 4-ES
      PIC 0  Online      10x10GE(LAN/WAN) SFPP
Slot 4  Online      FPC Type 4-ES
      PIC 0  Online      4x OC-192 SONET XFP
Slot 5  Online      FPC Type 4-ES
      PIC 0  Online      4x OC-192 SONET XFP
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 6  Online      FPC Type 4-ES
      PIC 1  Online      4x 10GE (LAN/WAN) XFP
Slot 7  Online      FPC Type 3-ES
      PIC 0  Online      10x 1GE(LAN), 1000 BASE
      PIC 1  Online      8x 1GE(TYPE3), IQ2E
      PIC 2  Online      4x OC-48 SONET

```

### show chassis fpc (T1600 Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)		Memory DRAM (MB)	Utilization (%)	
			Total	Interrupt		Heap	Buffer
0	Empty						
1	Empty						
2	Online	49	3	0	2048	3	24
3	Online	46	6	0	2048	6	24
4	Empty						
5	Online	46	5	0	2048	3	24
6	Empty						
7	Online	44	8	0	1024	7	49

### show chassis fpc detail (T1600 Router)

```

user@host> show chassis fpc detail

```

```

show chassis fpc detail
Slot 2 information:
  State                Online
  Temperature           49 degrees C / 120 degrees F
  Total CPU DRAM        2048 MB
  Total SRAM            64 MB
  Total SDRAM           1280 MB
  Start time            2010-10-04 21:12:52 PDT
  Uptime                32 minutes, 9 seconds
Slot 3 information:
  State                Online
  Temperature           47 degrees C / 116 degrees F
  Total CPU DRAM        2048 MB
  Total SRAM            128 MB
  Total SDRAM           2560 MB
  Start time            2010-10-04 21:13:06 PDT
  Uptime                31 minutes, 55 seconds
Slot 5 information:
  State                Online
  Temperature           46 degrees C / 114 degrees F
  Total CPU DRAM        2048 MB
  Total SRAM            64 MB
  Total SDRAM           1280 MB
  Start time            2010-10-04 21:12:56 PDT
  Uptime                32 minutes, 5 seconds
Slot 7 information:
  State                Online
  Temperature           44 degrees C / 111 degrees F
  Total CPU DRAM        1024 MB
  Total SRAM            64 MB
  Total SDRAM           1280 MB
  Start time            2010-10-04 21:14:34 PDT
  Uptime                30 minutes, 27 seconds

```

#### show chassis fpc <fpc-slot> (EX Series Switch)

```
user@host> show chassis fpc 2
```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
2	Online	40	12 0	2048 19	14

#### show chassis fpc slot (T1600 Router)

```
user@host> show chassis fpc slot 2
```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
2	Online	49	3 0	2048 3	24

#### show chassis fpc pic-status (T1600 Router)

```
user@host> show chassis fpc pic-status
```

```

Slot 2  Online  FPC Type 1-ES
PIC 0   Online  Load Type 1
PIC 1   Online  4x 1GE(LAN), IQ2E
PIC 3   Online  1x OC-12-3 SFP
Slot 3  Online  FPC Type 4-ES
PIC 0   Online  4x 10GE (LAN/WAN) XFP
PIC 1   Online  4x OC-192 SONET XFP

```

```

Slot 5  Online      FPC Type 2-ES
PIC 0   Online      Load Type 2
PIC 1   Online      8x 1GE(LAN), IQ2E
PIC 2   Online      8x 1GE(LAN), IQ2E
PIC 3   Online      1x OC-48-12-3 SFP
Slot 7  Online      FPC Type 4
PIC 0   Online      4x 10GE (LAN/WAN) XFP

```

### show chassis fpc (T4000 Router)

```
user@host> show chassis fpc
```

```

regress@stymphalian# run show chassis fpc

```

Slot	State	Temp (C)	CPU Total	Utilization (%) Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Utilization (%) Buffer
0	Online	48	15	0	2816	21	27
1	Empty						
2	Empty						
3	Online	51	15	0	2816	21	27
4	Empty						
5	Online	39	8	0	2048	6	23
6	Online	49	15	0	2816	21	27
7	Empty						

### show chassis fpc detail (T4000 Router)

```
user@host> show chassis fpc detail
```

```
Slot 0 information:
```

```

State                Online
Temperature           48 degrees C / 118 degrees F
Total CPU DRAM        2816 MB
Total SRAM            1554 MB
Total SDRAM           10752 MB
Start time            2012-02-09 22:56:25 PST
Uptime                2 hours, 40 minutes, 52 seconds

```

```
Slot 3 information:
```

```

State                Online
Temperature           51 degrees C / 123 degrees F
Total CPU DRAM        2816 MB
Total SRAM            1554 MB
Total SDRAM           10752 MB
Start time            2012-02-09 22:56:22 PST
Uptime                2 hours, 40 minutes, 55 seconds

```

```
Slot 5 information:
```

```

State                Online
Temperature           39 degrees C / 102 degrees F
Total CPU DRAM        2048 MB
Total SRAM            128 MB
Total SDRAM           2560 MB
Start time            2012-02-09 22:51:27 PST
Uptime                2 hours, 45 minutes, 50 seconds

```

```
Slot 6 information:
```

```

State                Online
Temperature           49 degrees C / 120 degrees F
Total CPU DRAM        2816 MB
Total SRAM            1554 MB
Total SDRAM           10752 MB
Start time            2012-02-09 22:56:29 PST
Uptime                2 hours, 40 minutes, 48 seconds

```



**show chassis fpc pic-status (T4000 Router)**

```

user@host> show chassis fpc pic-status
Slot 0  Online      FPC Type 5-3D
        PIC 0  Online  12x10GE (LAN/WAN) SFPP
        PIC 1  Online  12x10GE (LAN/WAN) SFPP
Slot 3  Online      FPC Type 5-3D
        PIC 0  Online  1x100GE
        PIC 1  Online  12x10GE (LAN/WAN) SFPP
Slot 5  Online      FPC Type 4-ES
        PIC 0  Online  100GE
        PIC 1  Online  100GE CFP
Slot 6  Online      FPC Type 5-3D
        PIC 0  Online  12x10GE (LAN/WAN) SFPP
        PIC 1  Online  12x10GE (LAN/WAN) SFPP

```

**show chassis fpc (QFX Series)**

```

user@switch> show chassis fpc
Temp CPU Utilization (%) Memory      Utilization (%)
Slot State              (C) Total Interrupt    DRAM (MB) Heap      Buffer
0 Online                26      2          0        2820      0        49

```

**show chassis fpc detail (QFX3500 Switches)**

```

user@switch> show chassis fpc detail
Slot 0 information:
State                               Online
Temperature                         28 degrees C / 82 degrees F
Total CPU DRAM                      2820 MB
Total SRAM                          0 MB
Total SDRAM                         0 MB
Start time                         2010-09-20 01:34:13 PDT
Uptime                             3 days, 3 hours, 31 minutes, 48 seconds

```

**show chassis fpc pic-status (QFX3500 Switches)**

```

user@switch> show chassis fpc pic-status
Slot 0  Online      QFX 48x10G 4x40G Switch
        PIC 0  Online  48x 10G-SFP+
        PIC 1  Online  15x 10G-SFP+

```

**show chassis fpc interconnect-device (QFabric System)**

```

user@switch> show chassis fpc interconnect-device interconnect1
FPC status:
Temp
Slot State      (C)
0 Online        0
1 Online        0
2 Online        0
3 Online        0
4 Online        0
5 Online        0
6 Online        0
7 Online        0
8 Online        0
9 Online        0
10 Online       0
11 Online       0
12 Online       0

```

```

13 Online      0
14 Online      0
15 Online      0

```

### show chassis fpc interconnect-device (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1 3
FPC status:

Slot State      Temp
          (C)
  3 Online        0

```

### show chassis fpc interconnect-device detail (QFabric System)

```

user@switch> show chassis fpc interconnect-device interconnect1 3 detail
Slot 3 information:
State Online
Temperature 0 degrees C / 32 degrees F
Start time 2011-08-18 10:45:04 PDT
Uptime 1 minute, 49 seconds

```

### show chassis fpc pic-status interconnect-device (QFabric System)

```

user@switch> show chassis fpc pic-status interconnect-device interconnect1
Slot 0 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 1 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 2 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 3 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 4 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 5 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 6 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 7 Online QFX 16-port QSFP+ Front Card
  PIC 0 Online 16x 40G-QSFP+
  PIC 1 Online 16x 40G-GE
Slot 8 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE
Slot 9 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE
Slot 10 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE
Slot 11 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE
Slot 12 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE
Slot 13 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE
Slot 14 Online QFX Fabric Rear Card
  PIC 0 Online 16x 40G-GE

```

```

Slot 15 Online      QFX Fabric Rear Card
PIC 0  Online      16x 40G-GE

```

### show chassis fpc pic-status node-device (QFabric System)

```

user@switch> show chassis fpc pic-status node-device node1
Slot node1 Online      QFX 48x10G 4x40G Switch
PIC 0  Online      48x 10G-SFP+
PIC 1  Online      4x 40G-QSFP+

```

### show chassis fpc (PTX5000 Packet Transport Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory DRAM (MB)	Utilization (%)
			Total Interrupt	Heap	Buffer
0	Empty				
1	Empty				
2	Online	50	6	0	2816
3	Empty				
4	Empty				
5	Online	48	9	0	2816
6	Empty				
7	Online	49	8	0	2816

### show chassis fpc detail (PTX5000 Packet Transport Router)

```

user@host> show chassis fpc detail
Slot 2 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 35 degrees C / 95 degrees F (Intake)
Temperature 50 degrees C / 122 degrees F (Exhaust A)
Temperature 54 degrees C / 129 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 61 degrees C / 141 degrees F (TL1)
Temperature 58 degrees C / 136 degrees F (TQ1)
Temperature 57 degrees C / 134 degrees F (TL2)
Temperature 58 degrees C / 136 degrees F (TQ2)
Temperature 62 degrees C / 143 degrees F (TL3)
Temperature 61 degrees C / 141 degrees F (TQ3)
Total CPU DRAM 2816 MB
Total SRAM 0 MB
Total SDRAM 0 MB
Start time 2012-01-12 12:05:42 PST
Uptime 3 hours, 14 minutes, 7 seconds
Slot 5 information:
State Online
Temperature 35 degrees C / 95 degrees F (PMB)
Temperature 34 degrees C / 93 degrees F (Intake)
Temperature 48 degrees C / 118 degrees F (Exhaust A)
Temperature 53 degrees C / 127 degrees F (Exhaust B)
Temperature 54 degrees C / 129 degrees F (TL0)
Temperature 52 degrees C / 125 degrees F (TQ0)
Temperature 69 degrees C / 156 degrees F (TL1)
Temperature 56 degrees C / 132 degrees F (TQ1)
Temperature 54 degrees C / 129 degrees F (TL2)
Temperature 56 degrees C / 132 degrees F (TQ2)
Temperature 59 degrees C / 138 degrees F (TL3)
Temperature 60 degrees C / 140 degrees F (TQ3)
Total CPU DRAM 2816 MB

```

```

Total SRAM                0 MB
Total SDRAM                0 MB
Start time                2012-01-12 12:05:43 PST
Uptime                    3 hours, 14 minutes, 6 seconds
Slot 7 information:
State                     Online
Temperature               35 degrees C / 95 degrees F (PMB)
Temperature               33 degrees C / 91 degrees F (Intake)
Temperature               50 degrees C / 122 degrees F (Exhaust A)
Temperature               55 degrees C / 131 degrees F (Exhaust B)
Temperature               56 degrees C / 132 degrees F (TL0)
Temperature               56 degrees C / 132 degrees F (TQ0)
Temperature               61 degrees C / 141 degrees F (TL1)
Temperature               57 degrees C / 134 degrees F (TQ1)
Temperature               55 degrees C / 131 degrees F (TL2)
Temperature               59 degrees C / 138 degrees F (TQ2)
Temperature               62 degrees C / 143 degrees F (TL3)
Temperature               62 degrees C / 143 degrees F (TQ3)
Total CPU DRAM            2816 MB
Total SRAM                0 MB
Total SDRAM                0 MB
Start time                2012-01-12 12:05:44 PST
Uptime                    3 hours, 14 minutes, 5 seconds

```

#### show chassis fpc pic-status (PTX5000 Packet Transport Router)

```

user@host> show chassis fpc pic-status
Slot 2  Online      FPC
PIC 0   Online      24x 10GE(LAN) SFP+
PIC 1   Online      24x 10GE(LAN) SFP+
Slot 5  Online      FPC
PIC 0   Online      24x 10GE(LAN) SFP+
PIC 1   Online      2x 40GE CFP
Slot 7  Online      FPC
PIC 0   Online      24x 10GE(LAN) SFP+
PIC 1   Online      2x 40GE CFP

```

#### show chassis fpc (ACX2000 Universal Access Router)

```

user@host> show chassis fpc

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			Total Interrupt	DRAM (MB) Heap Buffer
0	Online	61	17 6	512 21 37

#### show chassis fpc 0 (ACX2000 Universal Access Router)

```

user@host> show chassis fpc 0

```

Slot	State	Temp (C)	CPU Utilization (%)	Memory Utilization (%)
			Total Interrupt	DRAM (MB) Heap Buffer
0	Online	61	17 6	512 21 37

#### show chassis fpc detail (ACX2000 Universal Access Router)

```

user@host> show chassis fpc detail
Slot 0 information:
State                     Online
Temperature               61 degrees C / 141 degrees F
Total CPU DRAM            512 MB
Start time                2012-05-29 02:52:06 PDT
Uptime                    27 minutes, 17 seconds

```

**show chassis fpc pic-status (ACX2000 Universal Access Router)**

```

user@host> show chassis fpc pic-status
Slot 0  Online
  PIC 0  Online      16x CHE1T1, RJ48
  PIC 1  Online      8x 1GE(LAN) RJ45
  PIC 2  Online      2x 1GE(LAN) SFP
  PIC 3  Online      2x 10GE(LAN) SFP+

```

**show chassis FPC 1 (MX Routers with Media Services Blade [MSB])**

```

user@switch> show chassis fpc 1

```

Slot	State	Temp (C)	CPU Utilization (%) Total	Interrupt	Memory DRAM (MB)	Utilization (%) Heap	Buffer
1	Online	34	5	0	3072	5	13

**show chassis FPC 1 detail (MX Routers with Media Services Blade [MSB])**

```

user@switch> show chassis fpc 1 detail
Slot 1 information:
  State                               Online
  Temperature                         34
  Total CPU DRAM                      3072 MB
  Total RLDRAM                        259 MB
  Total DDR DRAM                      4864 MB
  Start time:                        2012-06-19 10:51:43 PDT
  Uptime:                            16 minutes, 48 seconds
  Max Power Consumption               550 Watts

```

## show chassis hardware

---

<b>List of Syntax</b>	<a href="#">Syntax on page 676</a> <a href="#">Syntax (EX Series) on page 676</a> <a href="#">Syntax (T4000 Router) on page 676</a> <a href="#">Syntax (TX Matrix Router) on page 676</a> <a href="#">Syntax (TX Matrix Plus Router) on page 676</a> <a href="#">Syntax (MX Series Routers) on page 676</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 676</a> <a href="#">Syntax (QFX Series) on page 677</a> <a href="#">Syntax (PTX Series Packet Transport Routers) on page 677</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 677</a>
<b>Syntax</b>	show chassis hardware <detail   extensive> <clei-models> <models>
<b>Syntax (EX Series)</b>	show chassis hardware <clei-models> <detail   extensive> <models>
<b>Syntax (T4000 Router)</b>	show chassis hardware <clei-models> <detail   extensive> <models>
<b>Syntax (TX Matrix Router)</b>	show chassis hardware <clei-models> <detail   extensive> <models> <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis hardware <clei-models> <detail   extensive> <models> <lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Routers)</b>	show chassis hardware <detail   extensive> <clei-models> <models> <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis hardware <clei-models> <detail   extensive> <models>

<b>Syntax (QFX Series)</b>	<pre>show chassis hardware &lt;detail   extensive&gt; &lt;clei-models&gt; &lt;interconnect-device <i>name</i>&gt; &lt;node-device <i>name</i>&gt; &lt;models&gt;</pre>
<b>Syntax (PTX Series Packet Transport Routers)</b>	<pre>show chassis hardware &lt;detail   extensive&gt; &lt;clei-models&gt; &lt;models&gt;</pre>
<b>Syntax (ACX Series Universal Access Routers)</b>	<pre>show chassis hardware &lt;detail   extensive&gt; &lt;clei-models&gt; &lt;models&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>models</b> option introduced in Junos OS Release 8.2.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.</p>
<b>Description</b>	<p>Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.</p> <p>In the EX Series switch command output, FPC refers to the following:</p> <ul style="list-style-type: none"> <li>On EX2200 switches, EX3200 switches, EX4200 standalone switches, and EX4500 switches—Refers to the switch; FPC <i>number</i> is always 0.</li> <li>On EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC <i>number</i> equals the member ID, from 0 through 9.</li> <li>On EX8208 and EX8216 switches—Refers to a line card; FPC <i>number</i> equals the slot number for the line card.</li> </ul> <p>On QFX3500 and QFX5100 standalone switches, both the FPC and FPC <i>number</i> are always 0.</p> <p>On T4000 Type 5 FPCs, there are no <b>top temperature sensor</b> or <b>bottom temperature sensor</b> parameters. Instead, <b>fan intake temperature sensor</b> and <b>fan exhaust temperature sensors</b> parameters are displayed.</p> <p>Starting from Junos OS Release 11.4, the output of the <b>show chassis hardware models</b> operational mode command displays the enhanced midplanes FRU model numbers (CHAS-BP3-MX240-S, CHAS-BP3-MX480-S or CHAS-BP3-MX960-S) based on the</p>

router. Prior to release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through the Junos OS Release 13.3, but can be supported on all Junos OS releases.

Starting with Junos OS Release 14.1, the output of the **show chassis hardware detail | extensive | clei-models | models** operational mode command displays the new DC power supply module (PSM) and power distribution unit (PDU) that are added to provide power to the high-density FPC (FPC2-PTX-P1A) and other components in a PTX5000 Packet Transport Router.

**Options** **none**—Display information about hardware. For a TX Matrix router, display information about the TX Matrix router and its attached T640 routers. For a TX Matrix Plus router, display information about the TX Matrix Plus router and its attached routers.

**clei-models**—(Optional) Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs).

**detail**—(Optional) Include RAM and disk information in output.

**extensive**—(Optional) Display ID EEPROM information.

**all-members**—(MX Series routers only) (Optional) Display hardware-specific information for all the members of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display hardware information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display hardware information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display hardware-specific information for the local Virtual Chassis members.

**member *member-id***—(MX Series routers and EX Series switches) (Optional) Display hardware-specific information for the specified member of the Virtual Chassis configuration. Replace *member-id* variable with a value 0 or 1.



**models**—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

**node-device *name***—(QFabric systems only) (Optional) Display hardware-specific information for the Node device.

**scc**—(TX Matrix router only) (Optional) Display hardware information for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus router only) (Optional) Display hardware information for the TX Matrix Plus router (switch-fabric chassis). Replace *number* variable with 0.

**Additional Information** The **show chassis hardware detail** command now displays DIMM information for the following Routing Engines:

**Table 34: Routing Engines Displaying DIMM Information**

Routing Engines	Routers
RE-S-1800x2 and RE-S-1800x4	MX240, MX480, and MX960 routers
RE-A-1800x2	M120 and M320 routers

In Junos OS Release 11.4 and later, the output for the **show chassis hardware models** operational mode command for MX Series routers display the enhanced midplanes FRU model numbers—CHAS-BP3-MX240-S, CHAS-BP3-MX480-S, or CHAS-BP3-MX960-S—based on the router. In releases before Junos OS Release 11.4, the FRU model numbers are left blank when the router has enhanced midplanes. Note that the enhanced midplanes are introduced through Junos OS Release 13.3, but can be supported on all Junos OS releases.

**Required Privilege Level** view

**Related Documentation**

- *show chassis power*

**List of Sample Output**

- [show chassis hardware \(EX8216 Switch\) on page 685](#)
- [show chassis hardware clei-models \(EX8216 Switch\) on page 686](#)
- [show chassis hardware clei-models \(T1600 Router\) on page 687](#)
- [show chassis hardware detail \(EX4200 Switch\) on page 688](#)
- [show chassis hardware \(EX4300 Switch\) on page 688](#)
- [show chassis hardware models \(EX4500 Switch\) on page 688](#)
- [show chassis hardware detail \(EX9200 Switch\) on page 688](#)
- [show chassis hardware \(J6350 Router\) on page 689](#)
- [show chassis hardware \(J6300 Router\) on page 689](#)
- [show chassis hardware \(M7i Router\) on page 690](#)
- [show chassis hardware \(M10 Router\) on page 690](#)
- [show chassis hardware models \(M10 Router\) on page 691](#)
- [show chassis hardware \(M20 Router\) on page 691](#)
- [show chassis hardware models \(M20 Router\) on page 692](#)

[show chassis hardware \(M40 Router\) on page 692](#)  
[show chassis hardware \(M40e Router\) on page 693](#)  
[show chassis hardware \(M120 Router\) on page 693](#)  
[show chassis hardware detail \(M120 Router\) on page 694](#)  
[show chassis hardware models \(M120 Router\) on page 695](#)  
[show chassis hardware \(M160 Router\) on page 696](#)  
[show chassis hardware models \(M160 Router\) on page 696](#)  
[show chassis hardware detail \(M160 Router\) on page 697](#)  
[show chassis hardware \(M320 Router\) on page 698](#)  
[show chassis hardware models \(M320 Router\) on page 699](#)  
[show chassis hardware \(MX5 Router\) on page 700](#)  
[show chassis hardware \(MX10 Router\) on page 700](#)  
[show chassis hardware \(MX40 Router\) on page 701](#)  
[show chassis hardware \(Fixed MX80 Router\) on page 701](#)  
[show chassis hardware \(Modular MX80 Router\) on page 702](#)  
[show chassis hardware \(MX104 Router\) on page 702](#)  
[show chassis hardware detail \(MX104 Router\) on page 703](#)  
[show chassis hardware extensive \(MX104 Router\) on page 704](#)  
[show chassis hardware models \(MX104 Router\) on page 707](#)  
[show chassis hardware clei-models \(MX104 Router\) on page 707](#)  
[show chassis hardware \(MX240 Router\) on page 707](#)  
[show chassis hardware detail \(MX 240 Router with Routing Engine Displaying DIMM information\) on page 708](#)  
[show chassis hardware \(MX240 Router with Enhanced MX SCB\) on page 708](#)  
[show chassis hardware \(MX480 Router\) on page 709](#)  
[show chassis hardware \(MX480 Router with Enhanced MX SCB\) on page 710](#)  
[show chassis hardware \(MX480 Routers with MPC5E and built-in OTN PIC\) on page 710](#)  
[show chassis hardware detail \(MX480 Routers with MPC5E and built-in OTN PIC\) on page 711](#)  
[show chassis hardware extensive \(MX480 Routers with MPC5E and built-in OTN PIC\) on page 713](#)  
[show chassis hardware \(MX960 Router\) on page 716](#)  
[show chassis hardware \(MX960 Router with Bidirectional Optics\) on page 716](#)  
[show chassis hardware \(MX960 Router with Enhanced MX SCB\) on page 717](#)  
[show chassis hardware models \(MX960 Router with Enhanced MX SCB\) on page 719](#)  
[show chassis hardware \(MX960 Router with MPC5EQ\) on page 719](#)  
[show chassis hardware detail \(MX960 Router\) on page 722](#)  
[show chassis hardware detail \(MX960 Router with MPC5EQ\) on page 723](#)  
[show chassis hardware extensive \(MX960 Router with MPC5EQ\) on page 726](#)  
[show chassis hardware models \(MX960 Router with MPC5EQ\) on page 734](#)  
[show chassis hardware clei-models \(MX960 Router with MPC5EQ\) on page 735](#)  
[show chassis hardware \(MX2010 Router\) on page 735](#)  
[show chassis hardware detail \(MX2010 Router\) on page 738](#)  
[show chassis hardware extensive \(MX2010 Router\) on page 742](#)  
[show chassis hardware models \(MX2010 Router\) on page 748](#)  
[show chassis hardware clei-models \(MX2010 Routers\) on page 748](#)  
[show chassis hardware \(MX2010 Routers with MPC6E and OTN MIC\) on page 749](#)  
[show chassis hardware detail \(MX2010 Routers with MPC6E and OTN MIC\) on page 751](#)

[show chassis hardware extensive \(MX2010 Routers with MPC6E and OTN MIC\) on page 753](#)

[show chassis hardware \(MX2020 Router\) on page 758](#)

[show chassis hardware detail \(MX2020 Router\) on page 766](#)

[show chassis hardware models \(MX2020 Router\) on page 775](#)

[show chassis hardware clei-models \(MX2020 Router\) on page 776](#)

[show chassis hardware \(MX2020 Router with MPC5EQ and MPC6E\) on page 778](#)

[show chassis hardware detail \(MX2020 Router with MPC5EQ and MPC6E\) on page 782](#)

[show chassis hardware extensive \(MX2020 Router with MPC5EQ and MPC6E\) on page 784](#)

[show chassis hardware models \(MX2020 Routers with MPC5EQ and MPC6E\) on page 789](#)

[show chassis hardware clei-models \(MX2020 Router with MPC5EQ and MPC6E\) on page 791](#)

[show chassis hardware \(MX Series routers with ATM MIC\) on page 792](#)

[show chassis hardware \(MX240, MX480, MX960 routers with Application Services Modular Line Card\) on page 792](#)

[show chassis hardware extensive \(MX240, MX480, MX960 routers with Application Services Modular Line Card\) on page 793](#)

[show chassis hardware \(MX480 Router with MPC4E\) on page 794](#)

[show chassis hardware \(MX2020 Router with MPC4E\) on page 794](#)

[show chassis hardware \(MX5, MX10, MX40, MX80, MX240, MX480, and MX960 routers with Enhanced 20-port Gigabit Ethernet MIC\) on page 796](#)

[show chassis hardware models \(MX5, MX10, MX40, MX80, MX240, MX480, and MX960 routers with Enhanced 20-port Gigabit Ethernet MIC\) on page 797](#)

[show chassis hardware \(T320 Router\) on page 797](#)

[show chassis hardware \(T640 Router\) on page 798](#)

[show chassis hardware models \(T640 Router\) on page 799](#)

[show chassis hardware extensive \(T640 Router\) on page 799](#)

[show chassis hardware \(T4000 Router\) on page 800](#)

[show chassis hardware \(T4000 Router with 16 GB line card chassis \(LCC\) Routing Engine\) on page 802](#)

[show chassis hardware \(T4000 Router with LSR FPC\) on page 803](#)

[show chassis hardware clei-models \(T4000 Router\) on page 803](#)

[show chassis hardware detail \(T4000 Router\) on page 803](#)

[show chassis hardware models \(T4000 Router\) on page 805](#)

[show chassis hardware lcc \(TX Matrix Router\) on page 806](#)

[show chassis hardware scc \(TX Matrix Router\) on page 807](#)

[show chassis hardware \(T1600 Router\) on page 807](#)

[show chassis hardware \(TX Matrix Plus Router\) on page 809](#)

[show chassis hardware sfc \(TX Matrix Plus Router\) on page 814](#)

[show chassis hardware extensive \(TX Matrix Plus Router\) on page 816](#)

[show chassis hardware clei-models \(TX Matrix Plus Router\) on page 817](#)

[show chassis hardware detail \(TX Matrix Plus Router\) on page 819](#)

[show chassis hardware models \(TX Matrix Plus Router\) on page 821](#)

[show chassis hardware \(TX Matrix Plus router with 3D SIBs\) on page 824](#)

[show chassis hardware clei-models \(TX Matrix Plus router with 3D SIBs\) on page 827](#)

[show chassis hardware detail \(TX Matrix Plus router with 3D SIBs\) on page 831](#)

[show chassis hardware lcc \(TX Matrix Plus router with 3D SIBs\) on page 834](#)

[show chassis hardware sfc \(TX Matrix Plus router with 3D SIBs\) on page 835](#)  
[show chassis hardware \(16-Port 10-Gigabit Ethernet MPC with SFP+ Optics \[MX Series Routers\]\) on page 837](#)  
[show chassis hardware \(MPC3E \[MX Series Routers\]\) on page 837](#)  
[show chassis hardware \(QFX3500 Switches\) on page 838](#)  
[show chassis hardware detail \(QFX3500 Switches\) on page 839](#)  
[show chassis hardware models \(QFX3500 Switches\) on page 840](#)  
[show chassis hardware clei-models \(QFX3500 Switches\) on page 840](#)  
[show chassis hardware clei-models \(QFX5100 Switches\) on page 840](#)  
[show chassis hardware interconnect-device \(QFabric Systems\) on page 840](#)  
[show chassis hardware node-device \(QFabric Systems\) on page 841](#)  
[show chassis hardware \(PTX5000 Packet Transport Router\) on page 841](#)  
[show chassis hardware \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 842](#)  
[show chassis hardware clei-models \(PTX5000 Packet Transport Router\) on page 843](#)  
[show chassis hardware clei-models \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 843](#)  
[show chassis hardware detail \(PTX5000 Packet Transport Router\) on page 843](#)  
[show chassis hardware detail \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 845](#)  
[show chassis hardware models \(PTX5000 Packet Transport Router\) on page 845](#)  
[show chassis hardware models \(PTX5000 Packet Transport Router with FPC2-PTX-P1A\) on page 846](#)  
[show chassis hardware extensive \(PTX5000 Packet Transport Router\) on page 846](#)  
[show chassis hardware \(MX Routers with Media Services Blade \[MSB\]\) on page 847](#)  
[show chassis hardware extensive \(MX Routers with Media Services Blade \[MSB\]\) on page 847](#)  
[show chassis hardware \(QFX3500 Switch running Enhanced Layer 2 Software\) on page 849](#)  
[show chassis hardware \(QFX5100 Switch running Enhanced Layer 2 Software\) on page 849](#)

**Output Fields**    [Table 35 on page 683](#) lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 35: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
<b>Item</b>	<p>Chassis component:</p> <ul style="list-style-type: none"> <li>(EX Series switches)—Information about the chassis, Routing Engine (SRE and Routing Engine modules in EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for EX8200 switches. See <i>EX Series Switches Hardware and CLI Terminology Mapping</i>.</li> <li>(MX Series routers and EX Series switches)—Information about the backplane, Routing Engine, Power Entry Modules (PEMs), and fan trays. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs), Modular Port Concentrators (MPCs) and associated Modular Interface Cards (MICs), or Dense Port Concentrators (DPCs). MX80 routers have a single Routing Engine and a built-in Packet Forwarding Engine that attaches directly to MICs. The Packet Forwarding Engine has two “pseudo” FPCs (FPC 0 and FPC1). MX80 routers also have a Forwarding Engine Board (FEB). MX104 routers have a built-in Packet forwarding Engine and a Forwarding Engine Board (FEB). The Packet Forwarding Engine of the MX104 router has three “pseudo” FPCs (FPC0, FPC1, and FPC2).</li> <li>(M Series routers, except for the M320 router)—Information about the backplane; power supplies; fan trays; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); SCB, SSB, SFM, or FEB; MCS and PCG (for the M160 router only); each FPC and PIC; and each fan, blower, and impeller.</li> <li>(M120, M320, and T Series routers)—Information about the backplane, power supplies, fan trays, midplane, FPM (craft interface), CIP, PEM, SCG, CB, FPC, PIC, SFP, SPMB, and SIB.</li> <li>(QFX Series)—Information about the chassis, Pseudo CB, Routing Engine, power supplies, fan trays, Interconnect devices, and Node devices. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs).</li> <li>(PTX Series)—Information about the chassis, midplane, craft interface (FPM), power distribution units (PDUs) and Power Supply Modules (PSMs), Centralized Clock Generators (CCGs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Flexible PIC Concentrators (FPCs), PICs, Switch Interface Boards (SIBs), and fan trays (vertical and horizontal).</li> <li>(MX2010 and MX2020 routers)—Information about the chassis, midplane, craft interface (FPM), power midplane (PMP), Power Supply Modules (PSMs), Power Distribution Modules (PDMs), Routing Engines, Control Boards (CBs) and Switch Processor Mezzanine Boards (SPMBs), Switch Fabric Boards (SFBs), Flexible PIC Concentrators (FPCs), PICs, adapter cards (ADCs) and fan trays.</li> </ul>	All levels
<b>Version</b>	Revision level of the chassis component.	All levels
<b>Part number</b>	Part number of the chassis component.	All levels
<b>Serial number</b>	Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact Juniper Networks Customer Support about the router or switch chassis.	All levels

Table 35: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Assb ID or Assembly ID</b>	( <b>extensive</b> keyword only) Identification number that describes the FRU hardware.	<b>extensive</b>
<b>Assembly Version</b>	( <b>extensive</b> keyword only) Version number of the FRU hardware.	<b>extensive</b>
<b>Assembly Flags</b>	( <b>extensive</b> keyword only) Flags.	<b>extensive</b>
<b>FRU model number</b>	( <b>clei-models</b> , <b>extensive</b> , and <b>models</b> keyword only) Model number of the FRU hardware component.	none specified
<b>CLEI code</b>	( <b>clei-models</b> and <b>extensive</b> keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
<b>EEPROM Version</b>	ID EEPROM version used by the hardware component: <b>0x00</b> (version 0), <b>0x01</b> (version 1), or <b>0x02</b> (version 2).	<b>extensive</b>
<b>Description</b>	<p>Brief description of the hardware item:</p> <ul style="list-style-type: none"> <li>Type of power supply.</li> <li>Type of PIC. If the PIC type is not supported on the current software release, the output states <b>Hardware Not Supported</b>.</li> <li>Type of FPC: <b>FPC Type 1</b>, <b>FPC Type 2</b>, <b>FPC Type 3</b>, <b>FPC Type 4</b>, or <b>FPC TypeOC192</b>.</li> </ul> <p>On EX Series switches, a brief description of the FPC.</p> <p>On the J Series routers, the FPC type corresponds to the Physical Interface Module (PIM). The following list shows the PIM abbreviation in the output and the corresponding PIM name.</p> <ul style="list-style-type: none"> <li><b>2x FE</b>—Either two built-in Fast Ethernet interfaces (fixed PIM) or dual-port Fast Ethernet PIM</li> <li><b>4x FE</b>—4-port Fast Ethernet ePIM</li> <li><b>1x GE Copper</b>—Copper Gigabit Ethernet ePIM (one 10-Mbps, 100-Mbps, or 1000-Mbps port)</li> <li><b>1x GE SFP</b>—SFP Gigabit Ethernet ePIM (one fiber port)</li> <li><b>4x GE Base PIC</b>—Four built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM)</li> <li><b>2x Serial</b>—Dual-port serial PIM</li> <li><b>2x T1</b>—Dual-port T1 PIM</li> <li><b>2x E1</b>—Dual-port E1 PIM</li> <li><b>2x CTIE1</b>—Dual-port channelized T1/E1 PIM</li> <li><b>1x T3</b>—T3 PIM (one port)</li> <li><b>1x E3</b>—E3 PIM (one port)</li> <li><b>4x BRI S/T</b>—4-port ISDN BRI S/T PIM</li> <li><b>4x BRI U</b>—4-port ISDN BRI U PIM</li> <li><b>1x ADSL Annex A</b>—ADSL 2/2+ Annex A PIM (one port, for POTS)</li> <li><b>1x ADSL Annex B</b>—ADSL 2/2+ Annex B PIM (one port, for ISDN)</li> </ul>	All levels

Table 35: show chassis hardware Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> <li>• <b>2xSHDSL (ATM)</b>—G SHDSL PIM (2-port two-wire module or 1-port four-wire module)</li> <li>• <b>1x TGM550</b>—TGM550 Telephony Gateway Module (Avaya VoIP gateway module with one console port, two analog <b>LINE</b> ports, and two analog <b>TRUNK</b> ports)</li> <li>• <b>1x DS1 TIM510</b>—TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with one E1 or T1 trunk termination port and ISDN PRI backup)</li> <li>• <b>4x FXS, 4x FXO, TIM514</b>—TIM514 Analog Telephony Interface Module (Avaya VoIP media module with four analog <b>LINE</b> ports and four analog <b>TRUNK</b> ports)</li> <li>• <b>4x BRI TIM521</b>—TIM521 BRI Telephony Interface Module (Avaya VoIP media module with four ISDN BRI ports)</li> <li>• <b>Crypto Accelerator Module</b>—For enhanced performance of cryptographic algorithms used in IP Security (IPsec) services</li> <li>• <b>MPC M 16x10GE</b>—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on EX Series switches.)</li> <li>• For hosts, the Routing Engine type.</li> <li>• For small form-factor pluggable transceiver (SFP) modules, the type of fiber: <b>LX</b>, <b>SX</b>, <b>LH</b>, or <b>T</b>.</li> <li>• LCD description for EX Series switches (except EX2200 switches).</li> <li>• <b>MPC2</b>—1-port MPC2 that supports two separate slots for MICs.</li> <li>• <b>MPC3E</b>—1-port MPC3E that supports two separate slots for MICs (MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP) on MX960, MX480, and MX240 routers. The MPC3E maps one MIC to one PIC (1 MIC, 1 PIC), which differs from the mapping of legacy MPCs.</li> <li>• 100GBASE-LR4, pluggable CFP optics</li> <li>• Supports the Enhanced MX Switch Control Board with fabric redundancy and existing SCBs without fabric redundancy.</li> <li>• Interoperates with existing MX Series line cards, including Flexible Port Concentrators (FPC), Dense Port Concentrators (DPCs), and Modular Port Concentrators (MPCs).</li> <li>• <b>MPC4E</b>—Fixed configuration MPC4E that is available in two flavors: MPC4E-3D-32XGE-SFP and MPC4E-3D-2CGE-8XGE on MX2020, MX960, MX480, and MX240 routers.</li> <li>• LCD description for MX Series routers</li> </ul>	

## Sample Output

### show chassis hardware (EX8216 Switch)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       REV 06   710-016845   CY0109220035  EX8216
Midplane      REV 06   710-016845   BA0909120112  EX8216-MP
CB 0          REV 22   710-020771   AX0109197723  EX8216-RE320
CB 1          REV 22   710-020771   AX0109197726  EX8216-RE320
  Routing Engine 1    BUILTIN      BUILTIN        RE-EX8216
FPC 3         REV 19   710-020683   BC0109083125  EX8200-48F

```



CPU	REV 13	710-020598	BF0109144549	EX8200-CPU
FPC 4	REV 17	710-020683	BC0108500127	EX8200-48F
CPU	REV 10	710-020598	BF0108460510	EX8200-CPU
PIC 0		BUILTIN	BUILTIN	48x 100 Base-QFX/1000
Base-X				
Xcvr 1	REV 01	740-011613	PE70V89	SFP-SX
Xcvr 11	REV 01	740-011613	PE70YCE	SFP-SX
Xcvr 12	REV 01	740-011613	PE70VSH	SFP-SX
Xcvr 13	REV 01	740-011613	E08C02063	SFP-SX
Xcvr 14	REV 01	740-011613	PE70VKU	SFP-SX
Xcvr 15	REV 01	740-011613	E08E03372	SFP-SX
Xcvr 21	REV 01	740-011613	PE70VAD	SFP-SX
Xcvr 22	REV 01	740-011613	E08E01228	SFP-SX
Xcvr 23	REV 01	740-011613	PE70VSL	SFP-SX
Xcvr 24	REV 01	740-011613	E08E03409	SFP-SX
Xcvr 25	REV 01	740-011613	PE70VL4	SFP-SX
Xcvr 26	REV 01	740-011613	PDQ4L2Z	SFP-SX
Xcvr 27	REV 01	740-011613	PE70WFK	SFP-SX
Xcvr 28	REV 01	740-011782	PBD2B5U	SFP-SX
Xcvr 29	REV 01	740-011613	PE70UQX	SFP-SX
Xcvr 30	REV 01	740-011613	PE70VL5	SFP-SX
Xcvr 31	REV 01	740-011613	PE70V0F	SFP-SX
Xcvr 32	REV 01	740-011613	E08C02052	SFP-SX
Xcvr 33	REV 01	740-011613	E08C02197	SFP-SX
Xcvr 34	REV 01	740-011613	PE70V0L	SFP-SX
Xcvr 35	REV 01	740-011613	E08E03390	SFP-SX
Xcvr 36	REV 01	740-011613	PDQ4VL9	SFP-SX
Xcvr 37	REV 01	740-011613	E08E03370	SFP-SX
Xcvr 38	REV 01	740-011613	E08E03362	SFP-SX
Xcvr 39	REV 01	740-011613	E08C02065	SFP-SX
Xcvr 40	REV 01	740-011613	E08E03405	SFP-SX
Xcvr 41	REV 01	740-011613	E08E03411	SFP-SX
Xcvr 43	REV 01	740-011613	E08C02171	SFP-SX
Xcvr 45	REV 01	740-011613	E08E03410	SFP-SX
FPC 13	REV 16	710-016837	BB0109051344	EX8200-8XS
CPU				
SIB 0	REV 10	710-021613	AY0109166244	EX8216-SF320
SIB 1	REV 10	710-021613	AY0109166357	EX8216-SF320
SIB 2	REV 10	710-021613	AY0109166362	EX8216-SF320
SIB 3	REV 10	710-021613	AY0109166338	EX8216-SF320
SIB 4	REV 10	710-021613	AY0109166350	EX8216-SF320
SIB 5	REV 10	710-021613	AY0109166365	EX8216-SF320
SIB 6	REV 10	710-021613	AY0109166361	EX8216-SF320
SIB 7	REV 10	710-021613	AY0109166399	EX8216-SF320
PSU 0	REV 17	740-021466	BG0709170003	EX8200-AC2K
PSU 1	REV 17	740-021466	BG0709170004	EX8200-AC2K
PSU 2	REV 17	740-021466	BG0709170020	EX8200-AC2K
PSU 3	REV 17	740-021466	BG0709170017	EX8200-AC2K
PSU 4	REV 17	740-021466	BG0709170008	EX8200-AC2K
PSU 5	REV 17	740-021466	BG0709170018	EX8200-AC2K
Top Fan Tray				
FTC 0	REV 4	760-022620	CX1209140212	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140212	EX8216-FT
Bottom Fan Tray				
FTC 0	REV 4	760-022620	CX1209140211	EX8216-FT
FTC 1	REV 4	760-022620	CX1209140211	EX8216-FT
LCD 0	REV 04	710-025742	CE0109186919	EX8200 LCD

### show chassis hardware clei-models (EX8216 Switch)

```
user@host> show chassis hardware clei-models
```



## Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 08	710-016845		
PSU 0	REV 05	740-023002	COUPAEAEAA	EX8200-PWR-AC3KR
PSU 1	REV 05	740-023002	COUPAEAEAA	EX8200-PWR-AC3KR
PSU 2	REV 05	740-023002	COUPAEAEAA	EX8200-PWR-AC3KR
PSU 3	REV 05	740-023002	COUPAEAEAA	EX8200-PWR-AC3KR
PSU 4	REV 05	740-023002	COUPAEAEAA	EX8200-PWR-AC3KR
PSU 5	REV 05	740-023002	COUPAEAEAA	EX8200-PWR-AC3KR
Top Fan Tray				
Bottom Fan Tray				

## show chassis hardware clei-models (T1600 Router)

user@host&gt; show chassis hardware clei-models

## Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-005608		CHAS-BP-T640-S
FPM Display	REV 05	710-002897		CRAFT-T640-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 07	740-017906	IPUPAC7KTA	PWR-T1600-3-80-DC-S
PEM 1	Rev 18	740-002595		PWR-T-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 08	740-014082		RE-A-2000-4096-S
Routing Engine 1	REV 07	740-014082		RE-A-2000-4096-S
CB 0	REV 05	710-007655		CB-T-S
CB 1	REV 03	710-017707		CB-T-S
FPC 0	REV 07	710-013558		T640-FPC2-E2
PIC 0	REV 01	750-010618		PB-4GE-SFP
PIC 1	REV 06	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 14	750-001901		PB-40C12-SON-SMIR
PIC 3	REV 07	750-001900		PB-10C48-SON-SMSR
FPC 1	REV 06	710-013553		T640-FPC1-E2
PIC 0	REV 08	750-001072		P-1GE-SX
PIC 1	REV 10	750-012266		PB-4GE-TYPE1-SFP-IQ2
PIC 2	REV 22	750-005634		PB-1CHOC12SMIR-QPP
FPC 2				
PIC 0	REV 16	750-007141		PC-10GE-SFP
PIC 1	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 05	750-004695		PC-TUNNEL
PIC 3	REV 17	750-009553		PC-40C48-SON-SFP
FPC 3	REV 01	710-010154		T640-FPC3-E
PIC 0	REV 07	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 25	750-007141		PC-10GE-SFP
PIC 2	REV 17	750-009553		PC-40C48-SON-SFP
PIC 3	REV 32	750-003700		PC-10C192-SON-VSR
FPC 4	REV 16	710-013037		T1600-FPC4-ES
PIC 1	REV 06	750-034781		PD-1CE-CFP
FPC 5	REV 02	710-013037		T1600-FPC4-ES
PIC 0	REV 16	750-012518		PD-40C192-SON-XFP
PIC 1	REV 01	750-010850		PD-10C768-SON-SR
FPC 6	REV 14	710-013037		T1600-FPC4-ES
PIC 0	REV 11	750-017405		PD-4XGE-XFP
PIC 1	REV 13	750-017405		PD-4XGE-XFP
FPC 7	REV 09	710-007529		T640-FPC3
PIC 0	REV 10	750-012793		PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 01	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 3	REV 15	750-009450		PC-10C192-SON-SR2
SIB 0	REV 07	710-013074		SIB-I-T1600-S
SIB 1	REV 07	710-013074		SIB-I-T1600-S

SIB 2	REV 07	710-013074	SIB-I-T1600-S
SIB 3	REV 07	710-013074	SIB-I-T1600-S
SIB 4	REV 07	710-013074	SIB-I-T1600-S
Fan Tray 0			FANTRAY-T-S
Fan Tray 1			FANTRAY-T-S
Fan Tray 2			FAN-REAR-TX-T640-S

**show chassis hardware detail (EX4200 Switch)**

```
user@host> show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			BM0208327733	EX4200-24T
Routing Engine 0	REV 11	750-021256	BM0208327733	EX4200-24T, 8 POE
Routing Engine 0			BM0208327733	EX4200-24T, 8 POE
FPC 0	REV 11	750-021256	BM0208327733	EX4200-24T, 8 POE
CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0		BUILTIN	BUILTIN	24x 10/100/1000 Base-T
PIC 1	REV 03B	711-021270	AR0208162285	4x GE SFP
BRD	REV 08	711-021264	AK0208328289	EX4200-24T, 8 POE
Power Supply 0	REV 03	740-020957	AT0508346354	PS 320W AC
Fan Tray				Fan Tray

**show chassis hardware (EX4300 Switch)**

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			PD3713160055	EX4300-48P
Routing Engine 0	REV 04	650-044930	PD3713160055	EX4300-48P
FPC 0	REV 04	650-044930	PD3713160055	EX4300-48P
CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0	REV 04	BUILTIN	BUILTIN	48x 10/100/1000 Base-T
PIC 1	REV 04	BUILTIN	BUILTIN	4x 40GE
Power Supply 0	REV 01	740-046871	1EDA3090026	JPSU-1100-AC-AFO-A
Fan Tray 0 (AFO)				Fan Module, Airflow Out
Fan Tray 1 (AFO)				Fan Module, Airflow Out

**show chassis hardware models (EX4500 Switch)**

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Routing Engine 0	REV 01	750-035700	GG0210271867	EX4500-40F-FB-C
FPC 0	REV 01	750-035700	GG0210271867	EX4500-40F-FB-C
PIC 0		BUILTIN	BUILTIN	EX4500-40F-FB-C
Power Supply 1	REV 01	740-029654	H884FS00JC09	EX4500-PWR1-AC-FB

**show chassis hardware detail (EX9200 Switch)**

```
user@switch> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN111DA44RFB	EX9208
Midplane	REV 05	710-017414	TS2912	EX9208-BP
FPM Board	REV 02	710-017254	XN1804	Front Panel Display
PEM 0	Rev 01	740-022697	QCS0906C033	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 01	740-022697	QCS0906C095	PS 1.2-1.7kW; 100-240V

```

AC in
Routing Engine 0 REV 08 740-031116 9009122883 RE-S-EX9200-1800X4
CB 0 REV 16 750-031391 CAAW4391 EX9200-SCBEF
PC 0 REV 07 750-049612 CABJ9312 EX9200 40x1G Copper
CPU REV 04 711-038484 CABH8268 MPCE PMB 2G
MIC 0 REV 02 750-049607 CABT9623 40x 1GE RJ45
PIC 0 BUILTIN BUILTIN 10x 1GE RJ45
PIC 1 BUILTIN BUILTIN 10x 1GE RJ45
PIC 2 BUILTIN BUILTIN 10x 1GE RJ45
PIC 3 BUILTIN BUILTIN 10x 1GE RJ45
FPC 1 REV 10 710-013699 CAAN3529 EX9200-40x1G-SFP
CPU REV 04 711-038484 CAAL7608 MPCE PMB 2G
MIC 0 REV 26 750-028392 CAAS5151 20x 1GE SFP
PIC 0 BUILTIN BUILTIN 10x 1GE SFP
PIC 1 BUILTIN BUILTIN 10x 1GE SFP
MIC 1 REV 26 750-028392 CAAC8006 20x 1GE SFP
PIC 2 BUILTIN BUILTIN 10x 1GE SFP
Xcvr 8 REV 01 740-011613 E08L03674 SFP-SX
Xcvr 9 REV 01 740-011613 E08M00243 SFP-SX
PIC 3 BUILTIN BUILTIN 10x 1GE SFP
FPC 3 REV 10 710-013699 CAAR5261 EX9200-40x1G-SFP
CPU REV 04 711-038484 CAAS2118 MPCE PMB 2G
MIC 0 REV 26 750-028392 CAAS5067 20x 1GE SFP
PIC 0 BUILTIN BUILTIN 10x 1GE SFP
Xcvr 2 REV 01 740-031851 PNA7L8U SFP-SX
Xcvr 3 REV 02 740-011613 AM0943SEKGZ SFP-SX
Xcvr 4 REV 02 740-011613 AM0943SEJZ9 SFP-SX
PIC 1 BUILTIN BUILTIN 10x 1GE SFP
MIC 1 REV 26 750-028392 CAAS5132 20x 1GE SFP
PIC 2 BUILTIN BUILTIN 10x 1GE SFP
Xcvr 4 REV 01 740-011613 E08D02625 SFP-SX
Xcvr 9 REV 02 740-011613 PJH4RD9 SFP-SX
PIC 3 BUILTIN BUILTIN 10x 1GE SFP
Xcvr 0 REV 01 740-011613 AM0813S8YME SFP-SX
Fan Tray Left Fan Tray

```

### show chassis hardware (J6350 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1090E07ADB   JSR6350
Midplane      REV 03   710-014593   NP1265
System IO     REV 01   710-016210   NN9950        JX350 System IO
Crypto Module                               Crypto Acceleration
Routing Engine REV 08   710-015273   NM6509        RE-J6350-3400
ad0          248 MB  256MB  CKS          00102006C24A00000039 Compact
Flash
FPC 0  FPC
PIC 0  4x GE Base PIC
FPC 1      REV 06   750-010355   AI07030023    FPC
PIC 0  2x T1
FPC 3      REV 06   750-011148   AJ06520151    FPC
PIC 0  2x E1
FPC 6      REV 06   750-013492   NC4170        FPC
PIC 0  4x FE
Power Supply 0

```

### show chassis hardware (J6300 Router)

```

user@host> show chassis hardware

```

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN000164AB	J6300
Midplane	REV 02.04	710-010001	CORE99570	
System IO	REV 02.00	710-010003	CORE100848	System IO board
Routing Engine	RevX2.6	750-010006	IWGS40735390	RE-J.3
FPC 0				FPC
PIC 0				2x FE
FPC 1	RevX2.0	750-011380	N3960005	FPC
PIC 0				1xADSL pic Annex A
FPC 2	RevX2.0	750-011380	N3960002	FPC
PIC 0				1xADSL pic Annex B
FPC 3	REV 03	750-010354	N0780028	FPC
PIC 0				1x T3

## show chassis hardware (M7i Router)

user@host&gt; show chassis hardware

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			31959	M7i
Midplane	REV 02	710-008761	CA0209	M7i Midplane
Power Supply 0	Rev 04	740-008537	PD10272	AC Power Supply
Routing Engine	REV 01	740-008846	1000396803	RE-5.0
CFEB	REV 02	750-009492	CA0166	Internet Processor IIv1
FPC 0				E-FPC
PIC 0	REV 04	750-003163	HJ6416	1x G/E, 1000 BASE-SX
PIC 1	REV 04	750-003163	HJ6423	1x G/E, 1000 BASE-SX
PIC 2	REV 04	750-003163	HJ6421	1x G/E, 1000 BASE-SX
PIC 3	REV 02	750-003163	HJ0425	1x G/E, 1000 BASE-SX
FPC 1				E-FPC
PIC 2	REV 01	750-009487	HM2275	ASP - Integrated
PIC 3	REV 01	750-009098	CA0142	2x F/E, 100 BASE-TX

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			B1157	M7i
Midplane	REV 05	710-008761	DM0840	M7i Midplane
Power Supply 0	Rev 08	740-008537	TE53755	AC Power Supply
Routing Engine	REV 07	740-011202	1000736567	RE-850
CFEB	REV 09	750-010463	DK6952	Internet Processor II
FPC 0				E-FPC
PIC 0	REV 12	750-012838	DL7993	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011614	PD94TDJ	SFP-LX10
Xcvr 1	REV 01	740-011615	PAD5EER	UNSUPPORTED
Xcvr 2	REV 01	740-011614	PD94THU	SFP-LX10
Xcvr 3		NON-JNPR	PDC2E7A	SFP-LX10
PIC 1	REV 03	750-023116	JT0203	4x CHSTM1 SDH CE SFP
Xcvr 0	REV 01	740-012434	AGT063832PS	SFP-SR
Xcvr 1	REV 01	740-012434	AGT063832LY	SFP-SR
Xcvr 3	REV 01	740-016064	C06J19018	SFP-LR
PIC 2	REV 15	750-014895	DM5757	MultiServices 100
PIC 3	REV 01	750-025390	JW9448	12x T1/E1 CE
FPC 1				E-FPC
PIC 2		BUILTIN	BUILTIN	1x Tunnel
PIC 3	REV 09	750-009099	DM0899	1x G/E, 1000 BASE
Xcvr 0	REV 01	740-012434	AGT07150HGJ	UNSUPPORTED
Fan Tray				Rear Fan Tray

## show chassis hardware (M10 Router)

user@host&gt; show chassis hardware

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			1122	M10
Midplane	REV 1.1	710-001950	S/N AC6626	
Power supply A	Rev 01	740-002497	S/N LC36095	AC
Power supply B	Rev 01	740-002497	S/N LC36100	AC
Display	REV 1.2	710-001995	S/N AC6656	
Host			18000005dfb3fb01	teknor
FEB	REV 01	710-001948	S/N AC6632	Internet Processor II
FPC 0				
PIC 0	REV 08	750-001072	S/N AB2485	1x G/E, 1000 BASE-SX
PIC 1	REV 01	750-000613	S/N AA1048	1x OC-12 SONET, SMIR
FPC 1				
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

## show chassis hardware models (M10 Router)

user@host&gt; show chassis hardware models

## Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-008920		CHAS-MP-M10i-S
Power Supply 0	Rev 06	740-008537		PWR-M10i-M7i-AC-S
Power Supply 1	Rev 06	740-008537		PWR-M10i-M7i-AC-S
HCM 0	REV 03	710-010580		HCM-M10i-S
HCM 1	REV 03	710-010580		HCM-M10i-S
Routing Engine 0	REV 09	740-009459		RE-400-256-S
CFEB 0	REV 05	750-010465		FEB-M10i-M7i-S
FPC 0				
PIC 0	REV 10	750-002971		PE-40C3-SON-MM
PIC 1	REV 11	750-002992		PE-4FE-TX
PIC 2	REV 03	750-002977		PE-20C3-ATM-MM
PIC 3	REV 08	750-005724		PE-20C3-ATM2-MM
FPC 1				
PIC 2	REV 12	750-008425		PE-AS
PIC 3	REV 13	750-005636		PE-4CHDS3-QPP
Fan Tray 0				FANTRAY-M10I-S
Fan Tray 1				FANTRAY-M10I-S

## show chassis hardware (M20 Router)

user@host&gt; show chassis hardware

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			20033	M20
Backplane	REV 07	710-001517	S/N AA7940	
Power supply B	Rev 01	740-001465	S/N 000001	AC
Display	REV 02	710-001519	S/N AA9704	
Host 0			98000004f8f27501	teknor
SSB slot 0	REV 01	710-001951	S/N AD5905	Internet Processor II
SSRAM bank 0	REV 01	710-001385	S00480	2 MB
SSRAM bank 1	REV 01	710-001385	S00490	2 MB
SSRAM bank 2	REV 01	710-001385	S001:?	2 MB
SSRAM bank 3	REV 01	710-001385	S00483	2 MB
SSB slot 1	N/A	N/A	N/A	Backup
FPC 1	REV 01	710-001292	S/N AB7528	
SSRAM	REV 01	710-000077	S/N 304209	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 000603	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 000414	64 MB
PIC 0	REV 03	750-000612	S/N AB8433	2x OC-3 ATM, MM
PIC 1	REV 01	750-000616	S/N AA1168	1x OC-12 ATM, MM

PIC 2	REV 01	750-000613	S/N AA1008	1x OC-12 SONET, SMIR
PIC 3	REV 01	750-002501	S/N AD5810	4x E3
FPC 2	REV 01	710-001292	S/N AC0119	
SSRAM	REV 01	710-000077	S/N 503241	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 306835	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 306832	64 MB
Fan Tray 0				Front Upper Fan Tray
Fan Tray 1				Front Middle Fan Tray
Fan Tray 2				Front Bottom Fan Tray
Fan Tray 3				Rear Fan Tray

### show chassis hardware models (M20 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Backplane	REV 03	710-002334		CHAS-MP-M20-S
Power Supply A	REV 06	740-001465		PWR-M20-AC-S
Display	REV 04	710-001519		CRAFT-M20-S
Routing Engine 0	REV 06	740-003239		RE-333-768-S
Routing Engine 1	REV 06	740-003239		RE-333-768-S
SSB 0	REV 02	710-001951		SSB-E-M20
SSB 1	N/A	N/A		
FPC 0	REV 03	710-003308		FPC-E
PIC 0	REV 08	750-002303		P-4FE-TX
PIC 1	REV 07	750-004745		P-2MCDS3
PIC 2	REV 03	750-002965		PE-4CHDS3
FPC 1	REV 03	710-003308		FPC-E
PIC 0	REV 03	750-002914		P-2OC3-ATM-MM
Fan Tray 0				FANTRAY-F-M20-S
Fan Tray 1				FANTRAY-F-M20-S
Fan Tray 2				FANTRAY-F-M20-S
Fan Tray 3				FANTRAY-R-M20-S

### show chassis hardware (M40 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Backplane	REV 02	710-000073	S/N AA0053	
Power supply A	Rev 2	740-000235	S/N 000042	DC
Maxicab	REV X1	710-000229	S/N AA0139	
Minicab	REV X1	710-000482	S/N AA0201	
Display	REV 06	710-000150	S/N AA0905	
Host				cpv5000
SCB	REV X1	710-000075	S/N AA0158	Internet Processor I
SSRAM bank 0	REV 02	710-000077	S/N AA2267	1 MB
SSRAM bank 1	REV 02	710-000077	S/N AA2270	1 MB
SSRAM bank 2	REV 02	710-000077	S/N AA2269	1 MB
SSRAM bank 3	REV 02	710-000077	S/N AA2268	1 MB
FPC 0	REV 01	710-000175	S/N AA0048	
SSRAM	REV 01	710-000077	S/N AA2333	1 MB
SDRAM bank 0	REV 01	710-000099	S/N AA2332	64 MB
SDRAM bank 1	REV X1	710-000099	S/N AA2337	64 MB
PIC 0	REV 04	750-000613	S/N aa0343	1x OC-12 SONET, SMIR
PIC 1	REV 04	750-000613	S/N AA0379	1x OC-12 SONET, SMIR
PIC 2	REV 04	750-000613	S/N AA0377	1x OC-12 SONET, SMIR
PIC 3	REV 04	750-000613	S/N AA0378	1x Tunnel
FPC 2	REV 01	710-000175	S/N AA0042	
SSRAM	REV 02	710-000077	S/N AA2288	1 MB
SDRAM bank 0	REV 01	710-000099	S/N AA2331	64 MB

SDRAM bank 1	REV 01	710-000099	S/N AA2330	64 MB
PIC 0	REV X1	750-000603	S/N AA0143	4x OC-3 SONET, SMIR
PIC 1	REV X1	750-000615	S/N AA0149	4x OC-3 SONET, MM
PIC 2	REV X1	750-000611	S/N AA0148	4x OC-3 SONET, MM
PIC 3	REV 04	750-000613	S/N AA0330	1x OC-12 SONET, SMIR
FPC 4	REV 01	710-000175	S/N AA0050	
SSRAM	REV 01	710-000077	S/N AA2327	1 MB
SDRAM bank 0	REV 01	710-000099	S/N AA2329	64 MB
SDRAM bank 1	REV 01	710-000099	S/N AA2328	64 MB
PIC 0	REV 04	750-000613	S/N AA0320	1x OC-12 SONET, SMIR
PIC 2	REV 05	750-000616	S/N AA1341	1x OC-12 ATM, MM
PIC 3	REV 08	750-001072	S/N AB2462	1x G/E, 1000 BASE-SX
FPC 5	REV 10	710-000175	S/N AA7663	
SSRAM	REV 01	710-000077	S/N 501590	1 MB
SDRAM bank 0	REV 01	710-000099	S/N 300949	64 MB
SDRAM bank 1	REV 01	710-000099	S/N 300868	64 MB
PIC 1	REV 01	750-001323	S/N AB1670	1x Tunnel

### show chassis hardware (M40e Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				m40e
Midplane	REV 01	710-005071	AX3671	
FPM CMB	REV 03	710-001642	AR9074	
FPM Display	REV 03	710-001647	AR7331	
CIP	REV 04	710-002649	BB4449	
PEM 0	Rev 01	740-003787	MC12364	Power Entry Module
PEM 1	Rev 01	740-003787	MC12383	Power Entry Module
PCG 0	REV 07	710-001568	AG1332	
PCG 1	REV 07	710-001568	AR3789	
Host 0			3e000007c8176601	Present
MCS 0	REV 11	710-001226	AN5813	
SFM 0 SPP	REV 07	710-001228	AG4676	
SFM 0 SPR	REV 05	710-002189	AE4735	Internet Processor II
SFM 1 SPP	REV 07	710-001228	AP1347	
SFM 1 SPR	REV 05	710-002189	BE0063	Internet Processor II
FPC 0	REV 01	710-011725	BE0669	M40e-EP-FPC Type 1
CPU	REV 01	710-004600	BD9504	
PIC 0	REV 03	750-003737	AY3991	4x G/E, 1000 BASE-SX
FPC 1	REV 01	710-005197	BD9842	M40e-FPC Type 2
CPU	REV 01	710-004600	BB4869	
PIC 0	REV 07	750-001900	AR8278	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005197	BD9824	M40e-FPC Type 2
CPU	REV 01	710-004600	BD9531	
PIC 0	REV 03	750-003737	AY3986	4x G/E, 1000 BASE-SX
FPC 4	REV 02	710-005078	BE0664	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9559	
PIC 0	REV 03	750-001894	AG7963	1x G/E, 1000 BASE-SX
PIC 2	REV 01	750-002575	AF2472	4x OC-3 SONET, SMIR
FPC 6	REV 02	710-005078	BE0652	M40e-FPC Type 1
CPU	REV 01	710-004600	BD9607	
PIC 0	REV 02	750-002911	AN2286	4x F/E, 100 BASE-TX
PIC 2	REV 01	750-002577	AP6345	4x OC-3 SONET, MM

### show chassis hardware (M120 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
------	---------	-------------	---------------	-------------

Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module
PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000
CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board
FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4Q0R9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

### show chassis hardware detail (M120 Router)

```
user@host> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN000054AC	M120
Midplane	REV 01	710-013667	RB4170	M120 Midplane
FPM Board	REV 02	710-011407	CJ9186	M120 FPM Board
FPM Display	REV 02	710-011405	CJ9173	M120 FPM Display
FPM CIP	REV 02	710-011410	CJ9221	M120 FPM CIP
PEM 0	Rev 05	740-011936	RM28320	AC Power Entry Module



PEM 1	Rev 05	740-011936	RM28321	AC Power Entry Module
Routing Engine 0	REV 03	740-014080	1000642883	RE-A-1000
ad0 248 MB		SILICONSYSTEMS INC 256M 126CT505S0763SC00110		Compact Flash
ad2 38154 MB		HTE541040G9SA00	MPBBT0X2HS2E3M	Hard Disk
CB 0	REV 03	710-011403	CM8346	M120 Control Board
CB 1	REV 06	710-011403	CP6728	M120 Control Board
FPC 1	REV 02	710-015908	CP6925	M120 CFPC 10GE
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	62E204N00007	XFP-10G-LR
FPC 3	REV 03	710-011393	CJ9234	M120 FPC Type 2
PIC 0	REV 16	750-008155	NB5229	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F15JB	SFP-SX
Xcvr 1	REV 01	740-007326	P4QOR9G	SFP-SX
PIC 1	REV 09	750-007745	CG4360	4x OC-3 SONET, SMIR
PIC 2	REV 16	750-008155	ND7787	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	P9F12AS	SFP-SX
Xcvr 1	REV 01	740-011613	P9F1ALU	SFP-SX
PIC 3	REV 07	750-011800	JW1284	8x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	P9F1AM6	SFP-SX
Xcvr 6	REV 01	740-011613	P9F16NN	SFP-SX
Xcvr 7	REV 01	740-011782	P8C29Y7	SFP-SX
Board B	REV 02	710-011395	CN3754	M120 FPC Mezz
FPC 4	REV 02	710-011398	CP6741	M120 FPC Type 3
PIC 0	REV 16	750-007141	NB2855	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	P922A1F	SFP-SX
Xcvr 1	REV 01	740-011782	P922A16	SFP-SX
Xcvr 2	REV 01	740-011782	P922A0U	SFP-SX
Xcvr 3	REV 01	740-011782	P9229UZ	SFP-SX
Xcvr 4	REV 01	740-009029	P11JXWP	SFP-LX
Xcvr 6	REV 01	740-011613	P9F1ALW	SFP-SX
FPC 5	REV 01	710-011388	CJ9088	M120 FPC Type 1
PIC 0	*** Hardware Not Supported ***			
PIC 1	REV 05	750-012052	NB0410	1x CHOC3 IQ SONET, SMLR
PIC 2	REV 01	750-013167	CM3824	4x CHDS3 IQ
PIC 3	REV 01	750-010240	CB5366	1x G/E SFP, 1000 BASE
Board B	REV 01	710-011390	CJ9103	M120 FPC Mezz Board
FEB 3	REV 04	710-011663	CP6673	M120 FEB
FEB 4	REV 04	710-011663	CJ9368	M120 FEB
FEB 5	REV 04	710-011663	CJ9386	M120 FEB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Top Fan Tray
Fan Tray 3				Rear Bottom Fan Tray

### show chassis hardware models (M120 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 01    710-013667
FPM CIP       REV 02    710-011410
PEM 0         Rev 05    740-011936
PEM 1         Rev 05    740-011936
Routing Engine 0 REV 03    740-014080
CB 0          REV 03    710-011403
CB 1          REV 06    710-011403
FPC 1         REV 02    710-015908
FPC 3
PIC 0         REV 16    750-008155

```

CRAFT-M120-S  
 PWR-M120-AC-S  
 PWR-M120-AC-S  
 RE-A-1000-2048-S  
 CB-M120-S  
 CB-M120-S  
 M120-cFPC-1XGE-XFP  
 PB-2GE-SFP-QPP

PIC 1	REV 09	750-007745	PC-40C3-SON-SMIR
PIC 2	REV 16	750-008155	PB-2GE-SFP-QPP
PIC 3	REV 07	750-011800	PB-8GE-TYPE2-SFP-IQ2
FPC 4			
PIC 0	REV 16	750-007141	PC-10GE-SFP
FPC 5			
PIC 1	REV 05	750-012052	PB-1CHOC3-SMIR-QPP
PIC 2	REV 01	750-013167	PE-4CHDS3-QPP
PIC 3	REV 01	750-010240	PB-1GE-SFP
Fan Tray 0			FFANTRAY-M120-S
Fan Tray 1			FFANTRAY-M120-S
Fan Tray 2			RFANTRAY-M120-S
Fan Tray 3			RFANTRAY-M120-S

### show chassis hardware (M160 Router)

```
user@host> show chassis hardware
```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
CPU	REV 03	710-001217	S/N AB3329	
PIC 0	REV 01			1x OC-192 SM SR-2
Fan Tray 0				Rear Bottom Blower
Fan Tray 1				Rear Top Blower
Fan Tray 2				Front Top Blower
Fan Tray 3				Front Fan Tray

### show chassis hardware models (M160 Router)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-009120		CHAS-BP-M320-S
FPM Display	REV 02	710-009351		CRAFT-M320-S
CIP	REV 03	710-005926		CIP-M320-S
PEM 2	Rev X4	740-009148		PWR-M-DC-S
PEM 3	Rev X4	740-009148		PWR-M-DC-S
Routing Engine 0	REV 02	740-008883		RE-1600-2048-S

Routing Engine 1	REV 02	740-008883	RE-1600-2048-S
FPC 0	REV 02	710-010419	M320-FPC1
PIC 0	REV 01	750-001323	P-TUNNEL
PIC 1	REV 02	750-002987	PE-10C12-SON-SMIR
PIC 2	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 04	750-001896	PB-10C12-SON-SMIR
FPC 1	REV 02	710-010419	M320-FPC1
PIC 0	REV 04	750-001894	PB-1GE-SX
PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 03	750-001894	PB-1GE-SX
FPC 2	REV 02	710-010419	M320-FPC1
PIC 0	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634	PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634	PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634	PB-1CHOC12SMIR-QPP
FPC 3			
PIC 0	REV 03	750-001895	PB-10C12-SON-MM
PIC 1	REV 04	750-001894	PB-1GE-SX
PIC 3	REV 04	750-003141	PB-1GE-SX-B
FPC 4	REV 02	710-010419	M320-FPC1
FPC 5	REV 02	710-010419	M320-FPC1
FPC 6	REV 02	710-010419	M320-FPC1
FPC 7			
PIC 0	REV 15	750-001901	PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900	PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900	PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737	PB-4GE-SX
SIB 0	REV 03	710-009184	SIB-M-S
SIB 1	REV 03	710-009184	SIB-M-S
SIB 2	REV 03	710-009184	SIB-M-S
SIB 3	REV 03	710-009184	SIB-M-S
Fan Tray 0			FFANTRAY-M320-S
Fan Tray 1			FFANTRAY-M320-S
Fan Tray 2			RFANTRAY-M320-S

### show chassis hardware detail (M160 Router)

```

user@host> show chassis hardware detail
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 306456	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 306474	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 306388	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 306392	1 MB
SFM 1 SPP	REV 04	710-001228	S/N AA2860	

SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
SSRAM bank 0	REV 01	710-000077	S/N 302917	1 MB
SSRAM bank 1	REV 01	710-000077	S/N 302662	1 MB
SSRAM bank 2	REV 01	710-000077	S/N 302593	1 MB
SSRAM bank 3	REV 01	710-000077	S/N 100160	1 MB
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
SSRAM	REV 01	710-000077	S/N 302836	1 MB
SDRAM 0	REV 01	710-001196	S00141	32 MB
SDRAM 1	REV 01	710-001196	S0010;	32 MB
SSRAM	REV 01	710-000077	S/N 302633	1 MB
SDRAM 0	REV 01	710-001196	S00143	32 MB
SDRAM 1	REV 01	710-001196	S00115	32 MB
SSRAM	REV 01	710-000077	S/N 302952	1 MB
SDRAM 0	REV 01	710-001196	S00135	32 MB
SDRAM 1	REV 01	710-001196	S001=3	32 MB
SSRAM	REV 01	710-000077	S/N 302892	1 MB
SDRAM 0	REV 01	710-001196	S000?6	32 MB
SDRAM 1	REV 01	710-001196	S001=5	32 MB
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
SSRAM	REV 01	710-000077	S/N 306340	1 MB
SDRAM 0	REV 01	710-001196	S00012	32 MB
SDRAM 1	REV 01	710-001196	S0001?	32 MB
SSRAM	REV 01	710-000077	S/N 306454	1 MB
SDRAM 0	REV 01	710-001196	S00028	32 MB
SDRAM 1	REV 01	710-001196	S0002?	32 MB
SSRAM	REV 01	710-000077	S/N 306492	1 MB
SDRAM 0	REV 01	710-001196	S00015	32 MB
SDRAM 1	REV 01	710-001196	S00031	32 MB
SSRAM	REV 01	710-000077	S/N 306363	1 MB
SDRAM 0	REV 01	710-001196	S00013	32 MB
SDRAM 1	REV 01	710-001196	S00032	32 MB
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX
FPC 2				FPC Type OC192
... SSRAM	REV 01	710-000077	S/N 306466	1 MB

### show chassis hardware (M320 Router)

user@host> show chassis hardware

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			67245	M320
Midplane	REV 05	710-009120	RB1202	M320 Midplane
FPM GBUS	REV 04	710-005928	HZ5697	M320 Board
FPM Display	REV 05	710-009351	HR1464	M320 FPM Display
CIP	REV 04	710-005926	HT8672	M320 CIP
PEM 0	Rev 05	740-009148	QK34208	DC Power Entry Module
PEM 1	Rev 05	740-009148	QK34262	DC Power Entry Module
PEM 2	Rev 05	740-009148	QF10449	DC Power Entry Module
PEM 3	Rev 05	740-009148	QJ18257	DC Power Entry Module
Routing Engine 0	REV 06	740-008883	P11123901185	RE-4.0
CB 0	REV 07	710-009115	JB2382	M320 Control Board
FPC 0	REV 02	710-005017	CD9926	M320 FPC Type 2
CPU	REV 01	710-011659	CJ6940	M320 PCA SCPU
PIC 0	REV 07	750-001900	AT1594	1x OC-48 SONET, SMSR
PIC 1	REV 03	750-001850	HS2746	1x Tunnel

PIC 2	REV 05	750-010618	JE7117	4x G/E SFP, 1000 BASE
PIC 3	REV 06	750-001900	HE6083	1x OC-48 SONET, SMSR
FPC 2	REV 02	710-005017	CH0319	M320 FPC Type 1
CPU	REV 01	710-011659	CJ6942	M320 PCA SCPU
PIC 0	REV 05	750-003034	BD8705	4x OC-3 SONET, SMIR
FPC 5	REV 02	710-005017	CD9938	M320 FPC Type 2
CPU				
FPC 7	REV 02	710-005017	CD9934	M320 FPC Type 2
CPU				
SIB 0	REV 09	710-009184	JA6540	M320 SIB
SIB 1	REV 09	710-009184	HV9511	M320 SIB
SIB 2	REV 09	710-009184	HW2057	M320 SIB
SIB 3	REV 09	710-009184	JA6687	M320 SIB
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

### show chassis hardware models (M320 Router)

```

user@host> show chassis hardware models
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-009120		CHAS-BP-M320-S
FPM Display	REV 02	710-009351		CRAFT-M320-S
CIP	REV 03	710-005926		CIP-M320-S
PEM 2	Rev X4	740-009148		PWR-M-DC-S
PEM 3	Rev X4	740-009148		PWR-M-DC-S
Routing Engine 0	REV 02	740-008883		RE-1600-2048-S
Routing Engine 1	REV 02	740-008883		RE-1600-2048-S
FPC 0	REV 02	710-010419		M320-FPC1
PIC 0	REV 01	750-001323		P-TUNNEL
PIC 1	REV 02	750-002987		PE-10C12-SON-SMIR
PIC 2	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-001896		PB-10C12-SON-SMIR
FPC 1	REV 02	710-010419		M320-FPC1
PIC 0	REV 04	750-001894		PB-1GE-SX
PIC 1	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 03	750-001894		PB-1GE-SX
FPC 2	REV 02	710-010419		M320-FPC1
PIC 0	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 1	REV 10	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 07	750-005634		PB-1CHOC12SMIR-QPP
PIC 3	REV 07	750-005634		PB-1CHOC12SMIR-QPP
FPC 3				
PIC 0	REV 03	750-001895		PB-10C12-SON-MM
PIC 1	REV 04	750-001894		PB-1GE-SX
PIC 3	REV 04	750-003141		PB-1GE-SX-B
FPC 4	REV 02	710-010419		M320-FPC1
FPC 5	REV 02	710-010419		M320-FPC1
FPC 6	REV 02	710-010419		M320-FPC1
FPC 7				
PIC 0	REV 15	750-001901		PB-40C12-SON-SMIR
PIC 1	REV 06	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 07	750-001900		PB-10C48-SON-SMSR
PIC 3	REV 05	750-003737		PB-4GE-SX
SIB 0	REV 03	710-009184		SIB-M-S
SIB 1	REV 03	710-009184		SIB-M-S
SIB 2	REV 03	710-009184		SIB-M-S

SIB 3	REV 03	710-009184	SIB-M-S
Fan Tray 0			FFANTRAY-M320-S
Fan Tray 1			FFANTRAY-M320-S
Fan Tray 2			RFANTRAY-M320-S

**show chassis hardware (MX5 Router)**

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			E1368	MX5-T
Midplane	REV 01	711-038215	YF5288	MX5-T
PEM 0	Rev 04	740-028288	VA01215	AC Power Entry Module
PEM 1	Rev 04	740-028288	VA01218	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9136	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX9820	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1045SUAQ3	SFP-SX
Xcvr 1	REV 01	740-031851	AM1045SUAPA	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUAN7	SFP-SX
Xcvr 3	REV 01	740-031851	AM1045SU91Q	SFP-SX
Xcvr 4	REV 01	740-031851	AM1045SUDDR	SFP-SX
Xcvr 9	REV 01	740-011613	AM0848SB6A1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AM1045SUANO	SFP-SX
Xcvr 1	REV 01	740-011613	AS0812S0719	SFP-SX
Xcvr 2	REV 01	740-011613	AM0821SA121	SFP-SX
Xcvr 3	REV 01	740-011613	PF21K21	SFP-SX
Xcvr 4	REV 01	740-011613	AM0848SB69Z	SFP-SX
Xcvr 5	REV 01	740-011782	P9P0XV3	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8WJN	SFP-SX
Xcvr 7	REV 01	740-011613	PAM3G9Q	SFP-SX
Xcvr 8	REV 01	740-011613	AM0848SB4A6	SFP-SX
Xcvr 9	REV 01	740-011782	P9MOU37	SFP-SX
MIC 1	REV 20	750-028380	ZG2657	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Fan Tray				Fan Tray

**show chassis hardware (MX10 Router)**

```
user@host> show chassis hardware
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			E1372	MX10-T
Midplane	REV 01	711-038211	YF5285	MX10-T
PEM 0	Rev 04	740-028288	VB01678	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9053	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP

```

FPC 1          BUILTIN      BUILTIN      MPC BUILTIN
  MIC 0        REV 24      750-028392  YX9436      3D 20x 1GE(LAN) SFP
    PIC 0      BUILTIN      BUILTIN      10x 1GE(LAN) SFP
      Xcvr 0    REV 01      740-031851  AM1107SUFQW SFP-SX
    PIC 1      BUILTIN      BUILTIN      10x 1GE(LAN) SFP
Fan Tray                               Fan Tray

```

### show chassis hardware (MX40 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			E1367	MX40-T
Midplane	REV 01	711-038211	YF5284	MX40-T
PEM 0	Rev 04	740-028288	VB01680	AC Power Entry Module
PEM 1	Rev 04	740-028288	VB01700	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZA9048	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
Xcvr 0	REV 01	740-014279	M7067UPP	XFP-10G-LR
Xcvr 1		NON-JNPR	K9J02UN	XFP-10G-LR
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 24	750-028392	YX3504	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	AM0812S8WTE	SFP-SX
Xcvr 1	REV 01	740-011613	PFA6KV2	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUDDM	SFP-SX
Xcvr 3	REV 01	740-011613	PD63C7M	SFP-SX
Xcvr 4	REV 01	740-011613	PD63DJY	SFP-SX
Xcvr 5	REV 02	740-011613	AA0950STLL9	SFP-SX
Xcvr 6	REV 01	740-011782	PAR1YHC	SFP-SX
Xcvr 7	REV 01	740-011782	P9P0XXL	SFP-SX
Xcvr 8	REV 01	740-011613	PD63D95	SFP-SX
Xcvr 9	REV 01	740-031851	AM1045SU9B8	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	PF21L3Z	SFP-SX
Xcvr 1	REV 01	740-031851	AM1045SU7M9	SFP-SX
Xcvr 2	REV 01	740-031851	AM1045SUAPT	SFP-SX
Xcvr 3	REV 01	740-011613	PFF2BZH	SFP-SX
Xcvr 4	REV 01	740-031851	AM1045SUDDN	SFP-SX
Xcvr 5	REV 01	740-031851	AM1039S00ZR	SFP-SX
Xcvr 6	REV 01	740-031851	AM1045SUD6Y	SFP-SX
Xcvr 8	REV 01	740-011613	PFM1QBS	SFP-SX
Xcvr 9	REV 01	740-011613	PFF2E25	SFP-SX
MIC 1	REV 01	750-021130	KG4391	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-011571	C645XJ04G	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0		NON-JNPR	CA49BK0AE	XFP-10G-SR
Fan Tray				Fan Tray

### show chassis hardware (Fixed MX80 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis				MX80-48T

Midplane	REV 01	711-031603	KF9250	MX80-48T
Routing Engine		BUILTIN	BUILTIN	Routing Engine
FEB 0		BUILTIN	BUILTIN	Forwarding Engine Board
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
Xcvr 0		NON-JNPR	M6439D41	XFP-10G-LR
Xcvr 1	REV 01	740-014279	6XE931N00202	XFP-10G-LR
Xcvr 2	REV 01	740-014289	C715XU05F	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C650XU0EP	XFP-10G-SR
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 01	711-029399	JR6981	12x 1GE(LAN) RJ45
PIC 0		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 1		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
MIC 1	REV 01	BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 2		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
PIC 3		BUILTIN	BUILTIN	12x 1GE(LAN) RJ45
Fan Tray				Fan Tray

### show chassis hardware (Modular MX80 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				MX80
Midplane	REV 02	711-031594	JR7084	MX80
PEM 0	Rev 01	740-028288	000018	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
FEB 0		BUILTIN	BUILTIN	Forwarding Engine Board
QXM 0	REV 05	711-028408	JR7041	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 02	750-028380	JR6598	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M86365	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M71094	XFP-10G-SR
MIC 1	REV 02	750-028380	JG8548	3D 2x 10GE XFP
PIC 2		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	T08L86302	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	C810XU0BA	XFP-10G-SR
Fan Tray				Fan Tray

### show chassis hardware (MX104 Router)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			G3503	MX104
Midplane	REV 28	750-044219	CAAX5741	MX104
PEM 0	REV 03	740-045933	1H072500016	AC Power Entry Module
PEM 1	REV 03	740-045932	1H073050017	DC Power Entry Module
Routing Engine 0	REV 20	750-044228	CAAY7935	RE-MX-104
Routing Engine 1	REV 13	750-044228	CAAM6380	RE-MX-104
AFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				



```

FPC 0          BUILTIN      BUILTIN      MPC BUILTIN
FPC 1          BUILTIN      BUILTIN      MPC BUILTIN
  MIC 0        REV 15      750-036132  CAAF7948    2x0C12/8x0C3 CC-CE
    PIC 0      BUILTIN      BUILTIN      2x0C12/8x0C3 CC-CE
      Xcvr 0    REV 01      740-011615  PCQ0U2J     SFP-IR
      Xcvr 1    REV 01      740-016068  PJJL7A6G    SFP-SR
      Xcvr 2    REV 01      740-016068  PJJL7A5J    SFP-SR
      Xcvr 3    REV 01      740-016065  PJJN5HPZ    SFP-SR
      Xcvr 4    REV 01      740-029122  PKB38TL     SFP-LR
      Xcvr 5    REV 01      740-011787  P6A107G     SFP-LR
      Xcvr 6    REV 01      740-029122  PKB38TR     SFP-LR
      Xcvr 7    REV 01      740-011787  PBKONK3     SFP-LR
    MIC 1
  FPC 2          BUILTIN      BUILTIN      MPC BUILTIN
  MIC 0          BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
    PIC 0      BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
      Xcvr 0    REV 01      740-031980  B10F00465   SFP+-10G-SR
      Xcvr 1    REV 01      740-031980  B10F00461   SFP+-10G-SR
      Xcvr 2    REV 01      740-031980  B10G01545   SFP+-10G-SR
      Xcvr 3    REV 01      740-031980  B10G01385   SFP+-10G-SR
  Fan Tray 0    REV 02      711-049570  CAAX6538    Fan Tray

```

#### show chassis hardware detail (MX104 Router)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               G3503         MX104
Midplane      REV 28    750-044219  CAAX5741      MX104
PEM 0         REV 03    740-045933  1H072500016  AC Power Entry Module
PEM 1         REV 03    740-045932  1H073050017  DC Power Entry Module
Routing Engine 0 REV 20    750-044228  CAAY7935     RE-MX-104
  da0 7836 MB ATP IG eUSB SSD          Nand Flash 0
  usb0 (addr 1) EHCI root hub 0      Freescale     uhub0
  usb0 (addr 2) USB2513Bi 9491       SMSC          uhub1
  usb0 (addr 3) ATP IG eUSB SSD 44801 ATP Electronics umass0
Routing Engine 1 REV 13    750-044228  CAAM6380     RE-MX-104
  da0 7836 MB ATP IG eUSB SSD          Nand Flash 0
AFEB 0          BUILTIN      BUILTIN      Forwarding Engine
Processor
FPC 0          BUILTIN      BUILTIN      MPC BUILTIN
FPC 1          BUILTIN      BUILTIN      MPC BUILTIN
  MIC 0        REV 15      750-036132  CAAF7948    2x0C12/8x0C3 CC-CE
    PIC 0      BUILTIN      BUILTIN      2x0C12/8x0C3 CC-CE
      Xcvr 0    REV 01      740-011615  PCQ0U2J     SFP-IR
      Xcvr 1    REV 01      740-016068  PJJL7A6G    SFP-SR
      Xcvr 2    REV 01      740-016068  PJJL7A5J    SFP-SR
      Xcvr 3    REV 01      740-016065  PJJN5HPZ    SFP-SR
      Xcvr 4    REV 01      740-029122  PKB38TL     SFP-LR
      Xcvr 5    REV 01      740-011787  P6A107G     SFP-LR
      Xcvr 6    REV 01      740-029122  PKB38TR     SFP-LR
      Xcvr 7    REV 01      740-011787  PBKONK3     SFP-LR
    MIC 1
  FPC 2          BUILTIN      BUILTIN      MPC BUILTIN
  MIC 0          BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
    PIC 0      BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
      Xcvr 0    REV 01      740-031980  B10F00465   SFP+-10G-SR
      Xcvr 1    REV 01      740-031980  B10F00461   SFP+-10G-SR
      Xcvr 2    REV 01      740-031980  B10G01545   SFP+-10G-SR
      Xcvr 3    REV 01      740-031980  B10G01385   SFP+-10G-SR
  Fan Tray 0    REV 02      711-049570  CAAX6538    Fan Tray

```

## show chassis hardware extensive (MX104 Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x02
S/N:          G3503
Assembly ID:  0x0560          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: MX104
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 60 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 47 33 35 30 33 00 00 00 00 00 00 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 28    750-044219    CAAX5741      MX104
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          750-044219      S/N:          CAAX5741
Assembly ID:  0x0560          Assembly Version: 01.28
Date:         03-27-2013      Assembly Flags: 0x00
Version:      REV 28          CLEI Code:    PROTOXCLEI
ID: MX104      FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ad 01 08 00 b0 a8 6e a7 f8 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 60 01 1c 52 45 56 20 32 38 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 31 39 00 00
Address 0x20: 53 2f 4e 20 43 41 41 58 35 37 34 31 00 1b 03 07
Address 0x30: dd ff ff ff ad 01 08 00 b0 a8 6e a7 f8 00 ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 47 33 35 30 33 00 00 00 00 00 00 00
PEM 0          REV 03    740-045933    1H072500016    AC Power Entry Module
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-045933      S/N:          1H072500016
Assembly ID:  0x0475          Assembly Version: 00.03
Date:         12-14-2012      Assembly Flags: 0x00
Version:      REV 03          CLEI Code:    IPUPAJ9KAA
ID: AC Power Entry Module      FRU Model Number: PWR-AMX1100-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 02 02 00 ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 75 00 03 52 45 56 20 30 33 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 39 33 33 00 00
Address 0x20: 31 48 30 37 32 35 30 30 30 31 36 00 00 0e 0c 07
Address 0x30: dc 30 43 ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 02 02 00 ff 01 49 50 55 50 41 4a 39 4b 41 41 50
Address 0x50: 57 52 2d 41 4d 58 31 31 30 30 2d 41 43 2d 53 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 70 ff ff ff ff ff ff ff ff ff ff ff ff
PEM 1          REV 03    740-045932    1H073050017    DC Power Entry Module
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-045932      S/N:          1H073050017

```

```

Assembly ID: 0x0476      Assembly Version: 00.03
Date: 01-30-2013      Assembly Flags: 0x00
Version: REV 03      CLEI Code: IPUPAJ8KAA
ID: DC Power Entry Module      FRU Model Number: PWR-AMX1100-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 02 02 00 ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 76 00 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 39 33 32 00 00
  Address 0x20: 31 48 30 37 33 30 35 30 30 31 37 00 00 1e 01 07
  Address 0x30: dd 30 44 ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 02 02 00 ff 01 49 50 55 50 41 4a 38 4b 41 41 50
  Address 0x50: 57 52 2d 41 4d 58 31 31 30 30 2d 44 43 2d 53 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff
  Address 0x70: ff ff ff 72 ff ff ff ff ff ff ff ff ff ff ff
Routing Engine 0 REV 20 750-044228 CAAY7935 RE-MX-104
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 750-044228      S/N: CAAY7935
Assembly ID: 0x0b81      Assembly Version: 01.20
Date: 03-18-2013      Assembly Flags: 0x00
Version: REV 20      CLEI Code: PROTOXCLEI
ID: RE-MX-104      FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
  Address 0x00: ad 01 00 08 b0 a8 6e a6 fc 10 ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 81 01 14 52 45 56 20 32 30 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 32 38 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 59 37 39 33 35 00 12 03 07
  Address 0x30: dd ff ff ff ad 01 00 08 b0 a8 6e a6 fc 10 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff
  Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff
da0 7836 MB ATP IG eUSB SSD Nand Flash 0
usb0 (addr 1) EHCI root hub 0 Freescale uhub0
usb0 (addr 2) USB2513Bi 9491 SMSC uhub1
usb0 (addr 3) ATP IG eUSB SSD 44801 ATP Electronics umass0
Routing Engine 1 REV 13 750-044228 CAAM6380 RE-MX-104
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 750-044228      S/N: CAAM6380
Assembly ID: 0x0b81      Assembly Version: 01.13
Date: 09-17-2012      Assembly Flags: 0x00
Version: REV 13      CLEI Code: PROTOXCLEI
ID: RE-MX-104      FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
  Address 0x00: ad 01 00 08 64 87 88 27 08 18 ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 fe 0b 81 01 0d 52 45 56 20 31 33 00 00
  Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 32 32 38 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 4d 36 33 38 30 00 11 09 07
  Address 0x30: dc ff ff ff ad 01 00 08 64 87 88 27 08 18 ff ff
  Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
  Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
da0 7836 MB ATP IG eUSB SSD Nand Flash 0
AFEB 0 BUILTIN BUILTIN Forwarding Engine
Processor
FPC 0 BUILTIN BUILTIN MPC BUILTIN
FPC 1 BUILTIN BUILTIN MPC BUILTIN
MIC 0 REV 15 750-036132 CAAF7948 2xOC12/8xOC3 CC-CE

```

```

Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 750-036132        S/N: CAAF7948
Assembly ID: 0x0a1a     Assembly Version: 01.15
Date: 07-03-2012       Assembly Flags: 0x00
Version: REV 15         CLEI Code: IP9IAM2DAA
ID: 2x0C12/8x0C3 CC-CE FRU Model Number: MIC-3D-80C3-20C12-ATM

Board Information Record:
Address 0x00: 12 01 05 03 05 ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 1a 01 0f 52 45 56 20 31 35 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 36 31 33 32 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 37 39 34 38 00 03 07 07
Address 0x30: dc ff ff ff 12 01 05 03 05 ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 39 49 41 4d 32 44 41 41 4d
Address 0x50: 49 43 2d 33 44 2d 38 4f 43 33 2d 32 4f 43 31 32
Address 0x60: 2d 41 54 4d 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff e3 c0 02 a3 9c 00 00 00 00 0a 60 00 00
PIC 0      BUILTIN      BUILTIN      2x0C12/8x0C3 CC-CE
  Xcvr 0    REV 01      740-011615    PCQOU2J      SFP-IR
  Xcvr 1    REV 01      740-016068    P3L7A6G      SFP-SR
  Xcvr 2    REV 01      740-016068    P3L7A5J      SFP-SR
  Xcvr 3    REV 01      740-016065    P3N5HPZ      SFP-SR
  Xcvr 4    REV 01      740-029122    PKB38TL      SFP-LR
  Xcvr 5    REV 01      740-011787    P6A107G      SFP-LR
  Xcvr 6    REV 01      740-029122    PKB38TR      SFP-LR
  Xcvr 7    REV 01      740-011787    PBKONK3      SFP-LR
MIC 1
FPC 2      BUILTIN      BUILTIN      MPC BUILTIN
MIC 0      BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N: BUILTIN           S/N: BUILTIN
Assembly ID: 0x0a60     Assembly Version: 00.00
Date: 00-00-0000       Assembly Flags: 0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 60 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 a5 04 7f b0 02 ff 0a 1a 01 0f
PIC 0      BUILTIN      BUILTIN      4x 10GE(LAN) SFP+
  Xcvr 0    REV 01      740-031980    B10F00465    SFP+-10G-SR
  Xcvr 1    REV 01      740-031980    B10F00461    SFP+-10G-SR
  Xcvr 2    REV 01      740-031980    B10G01545    SFP+-10G-SR
  Xcvr 3    REV 01      740-031980    B10G01385    SFP+-10G-SR
Fan Tray 0 REV 02      711-049570    CAAX6538      Fan Tray
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 711-049570        S/N: CAAX6538
Assembly ID: 0x0b82     Assembly Version: 01.02
Date: 03-01-2013       Assembly Flags: 0x00
Version: REV 02         CLEI Code: PROTOXCLEI
ID: Fan Tray           FRU Model Number: PROTO-ASSEMBLY

Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 82 01 02 52 45 56 20 30 32 00 00

```

```

Address 0x10: 00 00 00 00 37 31 31 2d 30 34 39 35 37 30 00 00
Address 0x20: 53 2f 4e 20 43 41 41 58 36 35 33 38 00 01 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

```

#### show chassis hardware models (MX104 Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item                Version  Part number  Serial number  FRU model number
Midplane            REV 20   750-044219   CAAS5849       PROTO-ASSEMBLY
PEM 0               REV 01   740-045932   1H072400065
Routing Engine 0    REV 16   750-044228   CAAR5915       PROTO-ASSEMBLY
AFEB 0              BUILTIN BUILTIN
FPC 0               BUILTIN BUILTIN
FPC 1               BUILTIN BUILTIN
  MIC 0             REV 01   750-046905   CAAK7103       MIC-3D-20GE-SFP-EH
FPC 2               BUILTIN BUILTIN
Fan Tray            REV 02   711-049570   CAAX6538       PROTO-ASSEMBLY

```

#### show chassis hardware clei-models (MX104 Router)

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code      FRU model number
Midplane            REV 20   750-044219   PROTOXCLEI     PROTO-ASSEMBLY
PEM 0               REV 01   740-045932
Routing Engine 0    REV 16   750-044228   PROTOXCLEI     PROTO-ASSEMBLY
AFEB 0              BUILTIN
FPC 0               BUILTIN
FPC 1               BUILTIN
  MIC 0             REV 01   750-046905   PROTOXCLEI     MIC-3D-20GE-SFP-EH
FPC 2               BUILTIN
Fan Tray            REV 02   711-049570   CAAX6538       PROTO-ASSEMBLY

```

#### show chassis hardware (MX240 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              REV 01   710-021041   JN10C7F7EAFC  MX240
Midplane             REV 01   710-017254   TR1502         MX240 Backplane
FPM Board            REV 01   710-017254   KD4017         Front Panel Display
PEM 0                Rev 02   740-017330   000332         PS 1.2-1.7kW; 100-240V
AC in
PEM 1                Rev 02   740-017330   000226         PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0     REV 06   740-013063   1000703522     RE-S-2000
Routing Engine 1     REV 06   740-015113   1000687625     RE-S-1300
CB 0                 REV 07   710-013385   KC9057         MX SCB
CB 1                 REV 05   710-013385   JY4760         MX SCB
FPC 1                REV 01   750-021679   KC7340         DPCE 40x 1GE R
  CPU                 REV 06   710-013713   KD4078         DPC PMB
  PIC 0               BUILTIN BUILTIN         10x 1GE(LAN)
    Xcvr 0            REV 01   740-011613   P9F18ME        SFP-SX
  PIC 1               BUILTIN BUILTIN         10x 1GE(LAN)
  PIC 2               BUILTIN BUILTIN         10x 1GE(LAN)
  PIC 3               BUILTIN BUILTIN         10x 1GE(LAN)
FPC 2                REV 04   710-016669   JS4529         DPCE 40x 1GE R EQ

```

CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

#### show chassis hardware detail (MX 240 Router with Routing Engine Displaying DIMM information)

```
user@host> show chassis hardware detail
```

Item	Version	Part number	Serial number	Description
Chassis			JN11279B4AFC	MX240 Backplane
Midplane	REV 07	760-021404	TS2474	MX240 Backplane
FPM Board	REV 03	760-021392	XC2643	Front Panel Display
PEM 0	Rev 03	740-017343	QCS0908A068	DC Power Entry Module
Routing Engine 0	REV 01	740-031117	AARCH00	RE-S-1800x4
ad0 3764 MB	STEC M2+	CF 9.0.2	STIM2Q3209239145303	Removable Compact Flash
ad1 28626 MB	WDC SSD-F0030S-5000		C933Z036237215548S00	Compact Flash
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 1	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 2	VL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
DIMM 3	SL31B5263E-F8S DIE REV-0 PCB REV-0			MFR ID-ce80
CB 0	REV 03	710-021523	XD7225	MX SCB
Fan Tray 0	REV 01	710-021113	WZ4986	MX240 Fan Tray

#### show chassis hardware (MX240 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7EAFC	MX240
Midplane	REV 01	710-021041	TR1502	MX240 Backplane
FPM Board	REV 01	710-017254	KD4017	Front Panel Display
PEM 0	Rev 02	740-017330	000332	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	000226	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 06	740-013063	1000703522	RE-S-2000
Routing Engine 1	REV 06	740-015113	1000687625	RE-S-1300
CB 0	REV 02	710-031391	YE8494	Enhanced MX SCB

CB 1	REV 05	710-031391	YOP5764	Enhanced MX SCB
FPC 1	REV 01	750-021679	KC7340	DPCE 40x 1GE R
CPU	REV 06	710-013713	KD4078	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18ME	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
FPC 2	REV 04	710-016669	JS4529	DPCE 40x 1GE R EQ
CPU	REV 06	710-013713	KB3969	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y79	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XU8	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YG6	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3XUG	SFP-SX
Xcvr 4	REV 01	740-011613	PBG3XTJ	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3ZUM	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3Y5H	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3UZT	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3US1	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3YG7	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3XZ9	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3XTY	SFP-SX
Xcvr 3	REV 01	740-011613	PBG3UZG	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 01	740-011613	PBG3Y8W	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3YVX	SFP-SX
Xcvr 2	REV 01	740-011613	PBG3YB3	SFP-SX
Xcvr 3	REV 01	740-011613	PBG43VQ	SFP-SX
Fan Tray 0	REV 01	710-021113	JS4642	MX240 Fan Tray

#### show chassis hardware (MX480 Router)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN10C7F7FAFB	MX480
Midplane	REV 04	710-017414	TR2071	MX480 Midplane
FPM Board	REV 02	710-017254	KB8459	Front Panel Display
PEM 0	Rev 02	740-017330	QCS07519029	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 02	740-017330	QCS07519041	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 02	740-017330	QCS07519097	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 07	740-013063	1000733381	RE-S-2000
Routing Engine 1	REV 07	740-013063	1000733540	RE-S-2000
CB 0	REV 07	710-013385	KA8022	MX SCB
CB 1	REV 07	710-013385	KA8303	MX SCB
FPC 0	REV 09	750-020452	KA8660	DPCE 40x 1GE X EQ
CPU	REV 06	710-013713	KA8185	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Fan Tray				Left Fan Tray

**show chassis hardware (MX480 Router with Enhanced MX SCB)**

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis           REV 04   710-017414   JN10C7F7FAFB   MX480
Midplane          REV 02   710-017254   TR2071         MX480 Midplane
FPM Board         Rev 02   740-017330   KB8459         Front Panel Display
PEM 0             Rev 02   740-017330   QCS07519029    PS 1.2-1.7kW; 100-240V
AC in
PEM 1             Rev 02   740-017330   QCS07519041    PS 1.2-1.7kW; 100-240V
AC in
PEM 2             Rev 02   740-017330   QCS07519097    PS 1.2-1.7kW; 100-240V
AC in
Routing Engine 0  REV 07   740-013063   1000733381     RE-S-2000
Routing Engine 1  REV 07   740-013063   1000733540     RE-S-2000
CB 0              REV 07   710-013385   KA8022         Enhanced MX SCB
CB 1              REV 07   710-013385   KA8303         Enhanced MX SCB
FPC 0             REV 09   750-020452   KA8660         DPCE 40x 1GE X EQ
CPU              REV 06   710-013713   KA8185         DPC PMB
PIC 0             BUILTIN BUILTIN       10x 1GE(LAN) EQ
PIC 1             BUILTIN BUILTIN       10x 1GE(LAN) EQ
PIC 2             BUILTIN BUILTIN       10x 1GE(LAN) EQ
PIC 3             BUILTIN BUILTIN       10x 1GE(LAN) EQ
Fan Tray
Left Fan Tray

```

**show chassis hardware (MX480 Routers with MPC5E and built-in OTN PIC)**

```

user@host> show chassis hardware
Hardware inventory:
Item              Version  Part number  Serial number  Description
Chassis           REV 05   710-017414   JN11C0338AFB   MX480
Midplane          REV 02   710-017254   ABAB8430       MX480 Midplane
FPM Board         Rev 05   740-029970   ZS8005         Front Panel Display
PEM 0             Rev 05   740-029970   QCS1024U089    PS 1.4-2.52kW; 90-264V
AC in
PEM 1             Rev 10   740-029970   QCS1314U0FJ    PS 1.4-2.52kW; 90-264V
AC in
PEM 2             Rev 07   740-029970   QCS1121U076    PS 1.4-2.52kW; 90-264V
AC in
Routing Engine 0  REV 05   740-031116   9009092471     RE-S-1800x4
Routing Engine 1  REV 05   740-031116   9009097958     RE-S-1800x4
CB 0              REV 16   750-031391   CAAX0789       Enhanced MX SCB
CB 1              REV 16   750-031391   CAAX0856       Enhanced MX SCB
FPC 0             REV 32   750-028467   ABBP1782       MPC 3D 16x 10GE
CPU              REV 10   711-029089   ABBP5410       AMPC PMB
PIC 0             BUILTIN BUILTIN       4x 10GE(LAN) SFP+
Xcvr 0           REV 01   740-021308   983152A00038   SFP+-10G-SR
Xcvr 1           REV 01   740-031980   B11F00211      SFP+-10G-SR
Xcvr 2           REV 01   740-031980   AQ72LPB        SFP+-10G-SR
Xcvr 3           REV 01   740-031980   AHNOWR5        SFP+-10G-SR
PIC 1             BUILTIN BUILTIN       4x 10GE(LAN) SFP+
Xcvr 0           REV 01   740-031980   B11J03627      SFP+-10G-SR
Xcvr 1           REV 01   740-031980   B11F00300      SFP+-10G-SR
Xcvr 2           REV 01   740-021308   AQ42WSS        SFP+-10G-SR
Xcvr 3           REV 01   740-021308   AQ43HGC        SFP+-10G-SR
PIC 2             BUILTIN BUILTIN       4x 10GE(LAN) SFP+
Xcvr 0           REV 01   740-021308   ANAONDO        SFP+-10G-SR
Xcvr 1           REV 01   740-021308   ANAONGF        SFP+-10G-SR
Xcvr 2           REV 01   740-021308   ANAONG9        SFP+-10G-SR
Xcvr 3           REV 01   740-021308   ANAOMP9        SFP+-10G-SR

```



PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQA06CG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	19T511100493	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	APR040J	SFP+-10G-SR
FPC 1	REV 26	750-046005	CACN1894	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACN8698	RMPD PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	163363A03046	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ40JS8	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	153363A00593	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ40JUJ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQC0B53	CFP2-100G-LR4-D
FPC 2	REV 26	750-046005	CACN1891	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACN8694	RMPD PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0		NON-JNPR	URA012A	SFP+-10G-LR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	J13F47042	CFP2-100G-LR4-D
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	AJC0BM3	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	11T511100917	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQK07SU	CFP2-100G-LR4-D
FPC 3	REV 03	750-045372	CAAD9425	MPCE Type 3 3D
CPU	REV 08	711-035209	CAAD9094	HMPD PMB 2G
MIC 0	REV 14	750-033196	CAAW9204	1X100GE CXP
PIC 0		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC034	CFP2-100G-SR10
MIC 1	REV 19	750-033199	CAAJ1814	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 4	REV 21.0.11	750-045715	CAAY3568	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7430	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	AP406NG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AR41NLP	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11D05630	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 12	750-049136	CACM6678	MPC5E 24XGE OTN Mezz
FPC 5	REV 11	750-045372	CABK7539	MPCE Type 3 3D
CPU	REV 08	711-035209	CABJ2466	HMPD PMB 2G
MIC 0	REV 19	750-033199	CAAJ9719	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	UP1020P	CFP-100G-SR10
MIC 1	REV 07	750-033196	YZ0797	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XC42FC022	CFP2-100G-SR10
Fan Tray				Enhanced Left Fan Tray

#### show chassis hardware detail (MX480 Routers with MPC5E and built-in OTN PIC)

```
user@host> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11C0338AFB	MX480
Midplane	REV 05	710-017414	ABAB8430	MX480 Midplane
FPM Board	REV 02	710-017254	ZS8005	Front Panel Display

PEM 0	Rev 05	740-029970	QCS1024U089	PS 1.4-2.52kW; 90-264V
AC in				
PEM 1	Rev 10	740-029970	QCS1314U0FJ	PS 1.4-2.52kW; 90-264V
AC in				
PEM 2	Rev 07	740-029970	QCS1121U076	PS 1.4-2.52kW; 90-264V
AC in				
Routing Engine 0	REV 05	740-031116	9009092471	RE-S-1800x4
ad0 3896 MB	VRFCF14096DIHK1		VM4096MB 6862	Compact Flash
ad1 30533 MB	UGB94ARF32H0S3-KC		UNIGEN-478612-001127	Disk 1
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
DIMM 1	SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
DIMM 2	SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
DIMM 3	SGU04G72H1BB2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80	
Routing Engine 1	REV 05	740-031116	9009097958	RE-S-1800x4
ad0 3896 MB	VRFCF14096DIHK1		VM4096MB 6145	Compact Flash
ad1 30533 MB	UGB94ARF32H0S3-KC		UNIGEN-499551-000273	Disk 1
CB 0	REV 16	750-031391	CAAX0789	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAX0856	Enhanced MX SCB
FPC 0	REV 32	750-028467	ABBP1782	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBP5410	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	983152A00038	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00211	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AQ72LPB	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AHNRW5	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11J03627	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00300	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ42WSS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HGC	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	ANAONDO	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANAONGF	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	ANAONG9	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	ANAOMP9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQA06CG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	19T511100493	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	APR040J	SFP+-10G-SR
FPC 1	REV 26	750-046005	CACN1894	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACN8698	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	163363A03046	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ40JS8	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	153363A00593	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ40JUJ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQC0B53	CFP2-100G-LR4-D
FPC 2	REV 26	750-046005	CACN1891	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACN8694	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0		NON-JNPR	URA012A	SFP+-10G-LR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	J13F47042	CFP2-100G-LR4-D
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	AJCOBM3	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	11T511100917	SFP+-10G-SR

PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQK07SU	CFP2-100G-LR4-D
FPC 3	REV 03	750-045372	CAAD9425	MPCE Type 3 3D
CPU	REV 08	711-035209	CAAD9094	HMPCE PMB 2G
MIC 0	REV 14	750-033196	CAAW9204	1X100GE CXP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-046563	XD16FC034	CFP2-100G-SR10
MIC 1	REV 19	750-033199	CAAJ1814	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 4	REV 21.0.11	750-045715	CAAY3568	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7430	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	AP406NG	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AR41NLP	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11D05630	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 12	750-049136	CACM6678	MPC5E 24XGE OTN Mezz
FPC 5	REV 11	750-045372	CABK7539	MPCE Type 3 3D
CPU	REV 08	711-035209	CABJ2466	HMPCE PMB 2G
MIC 0	REV 19	750-033199	CAAJ9719	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	UP1020P	CFP-100G-SR10
MIC 1	REV 07	750-033196	YZ0797	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XC42FC022	CFP2-100G-SR10
Fan Tray				Enhanced Left Fan Tray

### show chassis hardware extensive (MX480 Routers with MPC5E and built-in OTN PIC)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11C0338AFB  MX480
Jedec Code:   0x7fb0                  EEPROM Version: 0x02
S/N:          JN11C0338AFB
Assembly ID:  0x01fe                  Assembly Version: 00.00
Date:         00-00-0000              Assembly Flags:  0x02
ID: MX480
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 01 fe 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 43 30 33 33 38 41 46 42 02 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane     REV 05   710-017414  ABAB8430      MX480 Midplane
Jedec Code:   0x7fb0                  EEPROM Version: 0x01
P/N:         710-017414              S/N:          ABAB8430
Assembly ID:  0x01fe                  Assembly Version: 01.05
Date:         12-13-2011             Assembly Flags: 0x00
Version:      REV 05
ID: MX480 Midplane                  FRU Model Number: CHAS-BP-MX480-S
Board Information Record:
Address 0x00: ad 01 08 00 00 23 9c fc 98 00 ff ff ff ff ff ff

```

```

I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 fe 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 34 31 34 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 38 34 33 30 00 0d 0c 07
Address 0x30: db ff ff ff ad 01 08 00 00 23 9c fc 98 00 ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 34 38 30 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board          REV 02    710-017254    ZS8005          Front Panel Display
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 710-017254          S/N: ZS8005
Assembly ID: 0x01ff          Assembly Version: 01.02
Date: 11-21-2011          Assembly Flags: 0x00
Version: REV 02
ID: Front Panel Display          FRU Model Number: CRAFT-MX480-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 01 ff 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 30 2d 30 31 37 32 35 34 00 00
Address 0x20: 53 2f 4e 20 5a 53 38 30 30 35 00 00 00 15 0b 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
Address 0x50: 52 41 46 54 2d 4d 58 34 38 30 2d 53 00 00 00 00
Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PEM 0              Rev 05    740-029970    QCS1024U089    PS 1.4-2.52kW; 90-264V
AC in
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 740-029970          S/N: QCS1024U089
Assembly ID: 0x0432          Assembly Version: 01.05
Date: 06-17-2010          Assembly Flags: 0x00
Version: Rev 05
ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number: PWR-MX480-2520-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 32 01 05 52 65 76 20 30 35 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
Address 0x20: 51 43 53 31 30 32 34 55 30 38 39 00 00 11 06 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1              Rev 10    740-029970    QCS1314U0FJ    PS 1.4-2.52kW; 90-264V
AC in
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 740-029970          S/N: QCS1314U0FJ
Assembly ID: 0x0432          Assembly Version: 01.10
Date: 04-04-2013          Assembly Flags: 0x00
Version: Rev 10
ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number: PWR-MX480-2520-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 32 01 0a 52 65 76 20 31 30 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
Address 0x20: 51 43 53 31 33 31 34 55 30 46 4a 00 00 04 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 2          Rev 07   740-029970   QCS1121U076   PS 1.4-2.52kW; 90-264V
AC in
Jedec Code:    0x7fb0          EEPROM Version: 0x01
P/N:           740-029970      S/N:           QCS1121U076
Assembly ID:   0x0432          Assembly Version: 01.07
Date:          05-23-2011      Assembly Flags: 0x00
Version:       Rev 07
ID: PS 1.4-2.52kW; 90-264V AC in FRU Model Number: PWR-MX480-2520-AC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 32 01 07 52 65 76 20 30 37 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 32 39 39 37 30 00 00
Address 0x20: 51 43 53 31 31 32 31 55 30 37 36 00 00 17 05 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
Address 0x50: 57 52 2d 4d 58 34 38 30 2d 32 35 32 30 2d 41 43
Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 05   740-031116   9009092471   RE-S-1800x4
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-031116      S/N:           9009092471
Assembly ID:   0x09c0          Assembly Version: 01.05
Date:          11-01-2011      Assembly Flags: 0x00
Version:       REV 05          CLEI Code:     COUCALDBAA
ID: RE-S-1800x4              FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 43 41 2d 34 32 46 42 23 23 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
Address 0x20: 39 30 30 39 30 39 32 34 37 31 00 00 00 01 0b 07
Address 0x30: db ff ff ff 54 32 30 32 37 43 41 2d 34 32 46 42
Address 0x40: 23 23 23 00 01 43 4f 55 43 41 4c 44 42 41 41 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4b ff ff ff ff ff ff ff ff ff ff ff ff
ad0   3896 MB   VRFCF14096DIHK1   VM4096MB 6862   Compact Flash
ad1   30533 MB  UGB94ARF32H0S3-KC   UNIGEN-478612-001127 Disk 1
usb0 (addr 1)  EHCI root hub 0     Intel          uhub0
usb0 (addr 2)  product 0x0020 32    vendor 0x8087  uhub1
DIMM 0         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3         SGU04G72H1BB2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 05   740-031116   9009097958   RE-S-1800x4
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-031116      S/N:           9009097958
Assembly ID:   0x09c0          Assembly Version: 01.05
Date:          02-06-2012      Assembly Flags: 0x00
Version:       REV 05          CLEI Code:     COUCALDBAA
ID: RE-S-1800x4              FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 43 41 2d 34 32 46 42 23 23 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00

```

```

Address 0x20: 39 30 30 39 30 39 37 39 35 38 00 00 00 06 02 07
Address 0x30: dc ff ff ff 54 32 30 32 37 43 41 2d 34 32 46 42
Address 0x40: 23 23 23 00 01 43 4f 55 43 41 4c 44 42 41 41 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 4b ff ff ff ff ff ff ff ff ff ff ff ff
ad0   3896 MB  VRFCF14096DIHK1      VM4096MB 6145      Compact Flash
ad1   30533 MB UGB94ARF32H0S3-KC    UNIGEN-499551-000273 Disk 1

```

...

### show chassis hardware (MX960 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Midplane         REV 01    710-013698   AA6082         MX960 Midplane
PIM              Rev 01    740-013110   000008         Power Inlet Module
PEM 2
PEM 3            Rev 01    740-013682   000038         PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 00    740-015113   1000617944     RE-S-1300
CB 0             REV 05    710-013725   JK6947         MX960 Test SCB
FPC 4            REV 01    710-013305   JM7617         MX960 Test DPC
CPU
PIC 0            BUILTIN   BUILTIN       1x 10GE(LAN/WAN)
PIC 1            BUILTIN   BUILTIN       10x 1GE
FPC 7            REV 01    710-013305   JL9634         MX960 Test DPC
CPU
PIC 0            BUILTIN   BUILTIN       1x 10GE(LAN/WAN)
Xcvr 0           NON-JNPR   MYBG65I82C    XFP-10G-SR
PIC 1            BUILTIN   BUILTIN       10x 1GE
Xcvr 1           REV 01    740-011782   P7N0368        SFP-SX
Xcvr 4           REV 01    740-011782   P8J1W27        SFP-SX
Xcvr 6           REV 01    740-011782   P8J1VSD        SFP-SX
Xcvr 9           REV 01    740-011782   P8J1W25        SFP-SX
Fan Tray 0
Fan Tray 1

```

### show chassis hardware (MX960 Router with Bidirectional Optics)

```

user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Midplane         REV 03    710-013698   TR0234         MX960 Backplane
FPM Board        REV 03    710-014974   JA0878         Front Panel Display
PDM              Rev 03    740-013110   QCS11135028    Power Distribution Module
PEM 0            Rev 03    740-013682   QCS11154036    PS 1.7kW; 200-240VAC in
PEM 1            Rev 03    740-013682   QCS11154010    PS 1.7kW; 200-240VAC in
PEM 2            Rev 03    740-013682   QCS11154022    PS 1.7kW; 200-240VAC in
Routing Engine 0 REV 06    740-013063   1000691458     RE-S-2000
CB 0             REV 07    710-013385   KA2190         MX SCB
CB 1             REV 07    710-013385   KA0837         MX SCB
FPC 3            REV 02    750-018122   KB3890         DPCE 40x 1GE R
CPU
FPC 4            REV 01    750-018122   KB3889         DPCE 40x 1GE R
CPU              REV 06    710-013713   KB3976         DPC PMB
PIC 0            BUILTIN   BUILTIN       10x 1GE(LAN)
Xcvr 1           REV 01    740-020426   4910549        SFP-1000BASE-BX40-D
Xcvr 2           REV 01    740-020426   4910551        SFP-1000BASE-BX40-D

```

Xcvr 5	REV 01	740-021340	77E245N00006	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-020425	4882821	SFP-1000BASE-BX40-U
Xcvr 8	REV 01	740-020425	4882820	SFP-1000BASE-BX40-U
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020465	77E555N00894	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020465	75E467X00818	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020465	75E467X00573	SFP-1000BASE-BX10-D
Xcvr 3	REV 01	740-020465	4888227	SFP-1000BASE-BX10-D
Xcvr 4	REV 01	740-020465	4888241	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021340	77E245N00005	SFP-1000BASE-BX10-U
Xcvr 6	REV 01	740-021340	76E245X00487	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021341	5255889	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255887	SFP-1000BASE-BX10-U
Xcvr 9	REV 01	740-021340	77E245N00004	SFP-1000BASE-BX10-U
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-020424	5007582	SFP-1000BASE-BX10-D
Xcvr 1	REV 01	740-020424	4888187	SFP-1000BASE-BX10-D
Xcvr 2	REV 01	740-020424	4656500	SFP-1000BASE-BX10-D
Xcvr 5	REV 01	740-021341	5255886	SFP-1000BASE-BX10-U
Xcvr 7	REV 01	740-021340	77E245N00003	SFP-1000BASE-BX10-U
Xcvr 8	REV 01	740-021341	5255888	SFP-1000BASE-BX10-U
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-017726	74S184H30341	SFP-EX
Xcvr 1	REV 01	740-017726	4814061	SFP-EX
Xcvr 5	REV 01	740-017726	6ZS184H31108	SFP-EX
Xcvr 9	REV 01	740-021340	76E245X00486	SFP-1000BASE-BX10-U
Fan Tray 0				
Fan Tray 1	REV 03	740-014971	TP0850	Fan Tray

#### show chassis hardware (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1096805AFA	MX960
Midplane	REV 03	710-013698	TR0183	MX960 Backplane
Fan Extender	REV 02	710-018051	JY5227	Extended Cable Manager
FPM Board	REV 03	710-014974	JZ6876	Front Panel Display
PDM	Rev 03	740-013110	QCS11035023	Power Distribution Module
PEM 1	Rev 03	740-013682	QCS1109400L	PS 1.7kW; 200-240VAC in
PEM 2	Rev 03	740-013682	QCS11094015	PS 1.7kW; 200-240VAC in
PEM 3	Rev 03	740-013682	QCS11094012	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000
CB 0	REV 11	750-031391	YZ6072	Enhanced MX SCB
CB 1	REV 11	750-031391	YZ6068	Enhanced MX SCB
CB 2	REV 11	750-031391	YZ6081	Enhanced MX SCB
FPC 0	REV 01	750-018122	KA5576	DPCE 40x 1GE R
CPU	REV 06	710-013713	KB3961	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	P9F18GF	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TL9	SFP-SX
Xcvr 7	REV 01	740-011782	P9P0XXH	SFP-SX
Xcvr 9	REV 01	740-011782	P9M0TN1	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PAJ4UHC	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011613	PFF2CD0	SFP-SX
Xcvr 1	REV 01	740-011613	PBG3ZUT	SFP-SX
Xcvr 2	REV 01	740-011613	PFF2DDV	SFP-SX
Xcvr 5	REV 01	740-011613	P8E2SST	SFP-SX

Xcvr 9	REV 01	740-011782	PB8329N	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-026192	1U0201084503342	SFP-100BASE-BX10-U
Xcvr 1	REV 01	740-026193	1U1201084503313	SFP-100BASE-BX10-D
Xcvr 2	REV 01	740-011613	PAJ4Y5B	SFP-SX
Xcvr 6	REV 01	740-011782	P9M0U3M	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0TLA	SFP-SX
FPC 1	REV 16	750-031089	YL0719	MPC Type 2 3D
CPU	REV 06	711-030884	YL1463	MPC PMB 2G
MIC 0	REV 07	750-028387	JR6500	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	733019A00154	XFP-10G-LR
Xcvr 1	REV 02	740-014289	T09F55034	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014279	913019B00791	XFP-10G-LR
Xcvr 1	REV 01	740-014289	98S803A90384	XFP-10G-SR
MIC 1	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014279	T10B36134	XFP-10G-LR
Xcvr 1	REV 01	740-014289	T07M86354	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 2	REV 08	710-014219	JY9654	DPCE 4x 10GE R
CPU	REV 06	710-013713	JZ6549	DPC PMB
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 1		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 2		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
Xcvr 0	REV 03	740-011571	C931BK028	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
FPC 3	REV 10	750-024199	XJ6692	MX FPC Type 3
CPU	REV 03	710-022351	XF5182	DPC PMB
PIC 0	REV 17	750-009553	RJ2945	4x 0C-48 SONET
Xcvr 1	REV 01	740-011785	PCP3YLL	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMRY	SFP-SR
PIC 1	REV 32	750-003700	DP2113	1x 0C-192 12xMM VSR
FPC 5	REV 25	750-028467	YM8256	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YL3029	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 1	REV 01	740-031980	AHNOX1Z	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
FPC 7	REV 02	750-031092	JR6658	MPC Type 1 3D Q
CPU	REV 01	711-030884	JZ9038	MPC PMB 2G
MIC 0	REV 08	750-028392	JZ8737	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PBE2C6Y	SFP-SX
Xcvr 2		NON-JNPR	U8105N8	SFP-SX
Xcvr 4	REV 01	740-011613	PFM18EF	SFP-SX
Xcvr 7	REV 01	740-011613	PFF2AM8	SFP-SX
Xcvr 8	REV 01	740-011613	PFF2CT6	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011782	PB82VHH	SFP-SX
Xcvr 1	REV 01	740-011613	PFF2CSW	SFP-SX
Xcvr 9	REV 01	740-011613	PFF2BY0	SFP-SX
QXM 0	REV 04	711-028408	JR6372	MPC QXM
FPC 8	REV 05	750-024387	JW9754	MX FPC Type 2
CPU	REV 03	710-022351	KF1651	DPC PMB
PIC 0	REV 08	750-014730	DM3664	4x 0C-3 1x 0C-12 SFP
Xcvr 0	REV 01	740-016065	81S290N00077	SFP-SR
Xcvr 1		NON-JNPR	2191844	SFP-SR
Xcvr 2	REV 01	740-011618	PD81EE5	SFP-IR



PIC 1	REV 08	750-014637	DM3671	4x OC-12-3 SFP
Xcvr 0	REV 01	740-011785	PCK3UNK	SFP-SR
Xcvr 3	REV 01	740-011785	PDSOMPZ	SFP-SR
FPC 10	REV 04	710-013699	JY4654	DPCE 40x 1GE R
CPU	REV 05	710-013713	JS9717	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 5	REV 01	740-011782	PAR1L72	SFP-SX
Xcvr 6	REV 01	740-011782	P8N1YQ4	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN)
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN)
Xcvr 0	REV 01	740-011782	P8Q2AVL	SFP-SX
Xcvr 5	REV 01	740-011782	PAR1L7B	SFP-SX
Xcvr 6	REV 01	740-011782	PAR1L2J	SFP-SX
Xcvr 8	REV 01	740-011782	P8N1YMY	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN)
Fan Tray 0	REV 03	740-014971	TP0567	Fan Tray
Fan Tray 1	REV 03	740-014971	TP0702	Fan Tray

### show chassis hardware models (MX960 Router with Enhanced MX SCB)

```
user@host> show chassis hardware models
```

Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-013698	TR0183	CHAS-BP-MX960-S
Fan Extender	REV 02	710-018051	JY5227	ECM-MX960
FPM Board	REV 03	710-014974	JZ6876	CRAFT-MX960-S
Routing Engine 0	REV 06	740-013063	1000687969	RE-S-2000-4096-S
Routing Engine 1	REV 06	740-013063	1000687955	RE-S-2000-4096-S
CB 0	REV 11	750-031391	YZ6072	SCBE-MX-S
CB 1	REV 11	750-031391	YZ6068	SCBE-MX-S
CB 2	REV 11	750-031391	YZ6081	SCBE-MX-S
FPC 0	REV 01	750-018122	KA5576	DPCE-R-40GE-SFP
FPC 1	REV 16	750-031089	YL0719	MX-MPC2-3D
MIC 0	REV 07	750-028387	JR6500	MIC-3D-4XGE-XFP
MIC 1	REV 24	750-028387	YJ3950	MIC-3D-4XGE-XFP
FPC 2	REV 08	710-014219	JY9654	DPCE-R-4XGE-XFP
FPC 3	REV 10	750-024199	XJ6692	MX-FPC3
PIC 0	REV 17	750-009553	RJ2945	PC-40C48-SON-SFP
PIC 1	REV 32	750-003700	DP2113	PC-10C192-SON-VSR
FPC 5	REV 25	750-028467	YM8256	MPC-3D-16XGE-SFP
FPC 7	REV 02	750-031092	JR6658	MX-MPC1-3D-Q
MIC 0	REV 08	750-028392	JZ8737	MIC-3D-20GE-SFP
FPC 8	REV 05	750-024387	JW9754	MX-FPC2
PIC 0	REV 08	750-014730	DM3664	PB-40C3-10C12-SON2-SFP
PIC 1	REV 08	750-014637	DM3671	PB-40C3-40C12-SON-SFP
FPC 10	REV 04	710-013699	JY4654	DPC-R-40GE-SFP
Fan Tray 0	REV 03	740-014971	TP0567	FFANTRAY-MX960-S
Fan Tray 1	REV 03	740-014971	TP0702	FFANTRAY-MX960-S

### show chassis hardware (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1214852AFA	MX960
Midplane	REV 01	710-030012	ACAX3674	MX960 Backplane
FPM Board	REV 03	710-014974	CAAZ9326	Front Panel Display
PDM	Rev 03	740-013110	QCS17025017	Power Distribution Module
PEM 0	Rev 10	740-027760	QCS1702N062	PS 4.1kW; 200-240V AC
in				
PEM 1	Rev 04	740-027760	QCS1422N02C	PS 4.1kW; 200-240V AC

in				
PEM 2	Rev 09	740-027760	QCS1614N01X	PS 4.1kW; 200-240V AC
in				
Routing Engine 0	REV 08	740-031116	9009131803	RE-S-1800x4
Routing Engine 1	REV 08	740-031116	9009124913	RE-S-1800x4
CB 0	REV 18	750-031391	CABF0579	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAZ2471	Enhanced MX SCB
CB 2	REV 16	750-031391	CAAW9595	Enhanced MX SCB
FPC 0	REV 18	750-046005	CACE6574	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8908	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA0DYT	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOMS7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03Z	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	ANA0NAJ	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOMRQ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-049775	J13K72993	CFP2-100G-LR4
FPC 1	REV 11	750-045372	CABK8154	MPCE Type 3 3D
CPU	REV 08	711-035209	CABE7370	HMPC PMB 2G
MIC 0	REV 07	750-033307	CABD5255	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-021308	AQ50319	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ5035V	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502XJ	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HHR	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQ502YA	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ502EU	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ502HR	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ502A6	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ43H8M	SFP+-10G-SR
MIC 1	REV 14	750-033196	CAAP1398	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC064	CFP-100G-SR10
FPC 3	REV 35	750-028467	CAAT9156	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAV4645	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HZ1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43HZC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HD2	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502HN	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HGF	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501RZ	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5029V	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ501X9	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ502ZN	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43H86	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502ZY	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502PZ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ503E6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502XN	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11F00213	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ50336	SFP+-10G-SR
FPC 4	REV 18	750-046005	CACE6568	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8900	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN

Xcvr 0	REV 01	740-021308	AQA095A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0M1E	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000F	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0LYC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LYB	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-048813	XD32FE00Z	CFP2-100G-SR10
FPC 5	REV 18	750-046005	CACE6577	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8902	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0MXE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LVY	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03T	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0LW1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LW3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000J	CFP2-100G-SR10
FPC 7	REV 09	750-037355	CAAF0937	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAD8004	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	ANA0MM3	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X000C163	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQG0MS6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0MRX	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQG0M6Y	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQG0LZM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00499	CFP-100G-SR10
FPC 8	REV 39	750-045715	CACD1903	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1815	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QC480289	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QC480274	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130190	QSFP+-40G-SR4
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130197	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130180	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130199	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0415	MPC5E 24XGE OTN Mezz
FPC 9	REV 05	750-044444	CAAY9801	MPCE Type 2 3D P
CPU	REV 04	711-038484	CAAW3673	MPCE PMB 2G
MIC 0	REV 28	750-028387	CAAX1071	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T12L92342	XFP-10G-SR
Xcvr 1		NON-JNPR	T12L92303	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	CC07BK02X	XFP-10G-SR
QXM 0	REV 06	711-028408	CAAW4883	MPC QXM
QXM 1	REV 06	711-028408	CAAW4603	MPC QXM
FPC 10	REV 21.0.11	750-045715	CAAY3541	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7426	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-031980	AHK01AP	SFP+-10G-SR

Xcvr 1	REV 01	740-021308	AQ502ZU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP41BLS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQA08YA	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQA0K26	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA06S3	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQA06AS	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQA053N	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0E97	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA0GS4	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA0JVA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-021308	AQA057A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANA0MLS	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQA093A	SFP+-10G-SR
Xcvr 3	REV 01	740-021309	943153A00075	SFP+-10G-LR
Xcvr 4	REV 01	740-021308	AQA077B	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQA0JSC	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA0735	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ5028N	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AP40VN5	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0K0J	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA07AP	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA08YB	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 07	750-045717	CAAX3123	MPC5E 24XGE Mezz
FPC 11	REV 17	750-037355	CAAT3986	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAR3972	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQA0DSE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501Y3	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ501XU	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ5036Y	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00247	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	ALQ1DKF	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403YA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP40TY0	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14G0	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00095	CFP-100G-SR10
Fan Tray 0	REV 08	740-031521	ACAF4219	Enhanced Fan Tray
Fan Tray 1	REV 08	740-031521	ACAF4225	Enhanced Fan Tray

### show chassis hardware detail (MX960 Router)

```
user@host> show chassis hardware detail
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis				MX960
Midplane	REV 01	710-013698	AA6082	MX960 Midplane
PIM	Rev 01	740-013110	000008	Power Inlet Module
PEM 2				
PEM 3	Rev 01	740-013682	000038	PS 1.7kW; 200-240VAC in
Routing Engine 0	REV 00	740-015113	1000617944	RE-S-1300
ad0	245 MB	SanDisk SDCFB-256	111419E1805T1141	Compact Flash
ad2	38154 MB	FUJITSU MHT2040BH	NR0WT5925N77	Hard Disk
CB 0	REV 05	710-013725	JK6947	MX960 Test SCB
FPC 4	REV 01	710-013305	JM7617	MX960 Test DPC
CPU				

PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
PIC 1		BUILTIN	BUILTIN	10x 1GE
FPC 7	REV 01	710-013305	JL9634	MX960 Test DPC
CPU				
PIC 0		BUILTIN	BUILTIN	1x 10GE(LAN/WAN)
Xcvr 0		NON-JNPR	MYBG65I82C	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	10x 1GE
Xcvr 1	REV 01	740-011782	P7N0368	SFP-SX
Xcvr 4	REV 01	740-011782	P8J1W27	SFP-SX
Xcvr 6	REV 01	740-011782	P8J1VSD	SFP-SX
Xcvr 9	REV 01	740-011782	P8J1W25	SFP-SX
Fan Tray 0				
Fan Tray 1				

### show chassis hardware detail (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware detail
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1214852AFA	MX960
Midplane	REV 01	710-030012	ACAX3674	MX960 Backplane
FPM Board	REV 03	710-014974	CAAZ9326	Front Panel Display
PDM	Rev 03	740-013110	QCS17025017	Power Distribution Module
PEM 0	Rev 10	740-027760	QCS1702N062	PS 4.1kW; 200-240V AC
in				
PEM 1	Rev 04	740-027760	QCS1422N02C	PS 4.1kW; 200-240V AC
in				
PEM 2	Rev 09	740-027760	QCS1614N01X	PS 4.1kW; 200-240V AC
in				
Routing Engine 0	REV 08	740-031116	9009131803	RE-S-1800x4
ad0 3831 MB	UGB30SFA4000T1		SFA4000T1 000016CD	Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000061346	Disk 1
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	VL31B5263F-F8SD DIE	REV-0 PCB	REV-0	MFR ID-ce80
DIMM 1	VL31B5263F-F8SD DIE	REV-0 PCB	REV-0	MFR ID-ce80
DIMM 2	VL31B5263F-F8SD DIE	REV-0 PCB	REV-0	MFR ID-ce80
DIMM 3	VL31B5263F-F8SD DIE	REV-0 PCB	REV-0	MFR ID-ce80
Routing Engine 1	REV 08	740-031116	9009124913	RE-S-1800x4
ad0 3831 MB	UGB30SFA4000T1		SFA4000T1 0000106D	Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000052402	Disk 1
CB 0	REV 18	750-031391	CABF0579	Enhanced MX SCB
CB 1	REV 16	750-031391	CAAZ2471	Enhanced MX SCB
CB 2	REV 16	750-031391	CAAW9595	Enhanced MX SCB
FPC 0	REV 18	750-046005	CACE6574	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8908	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA0DYT	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0MS7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03Z	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	ANA0NAJ	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0MRQ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-049775	J13K72993	CFP2-100G-LR4
FPC 1	REV 11	750-045372	CABK8154	MPCE Type 3 3D
CPU	REV 08	711-035209	CABE7370	HMPC PMB 2G
MIC 0	REV 07	750-033307	CABD5255	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-021308	AQ50319	SFP+-10G-SR

Xcvr 1	REV 01	740-021308	AQ5035V	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502XJ	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43HHR	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQ502YA	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ502EU	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ502HR	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ502A6	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ43H8M	SFP+-10G-SR
MIC 1	REV 14	750-033196	CAAP1398	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC064	CFP2-100G-SR10
FPC 3	REV 35	750-028467	CAAT9156	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAV4645	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HZ1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43HZC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HD2	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502HN	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43HGF	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501RZ	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5029V	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ501X9	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ502ZN	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43H86	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502ZY	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502PZ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ503E6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502XN	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11F00213	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ50336	SFP+-10G-SR
FPC 4	REV 18	750-046005	CACE6568	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8900	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA095A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0M1E	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000F	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0LYC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LYB	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-048813	XD32FE00Z	CFP2-100G-SR10
FPC 5	REV 18	750-046005	CACE6577	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACG8902	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0MXE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LVY	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03T	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQG0LW1	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQG0LW3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000J	CFP2-100G-SR10
FPC 7	REV 09	750-037355	CAAF0937	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAD8004	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	ANA0MM3	SFP+-10G-SR

PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X000C163	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQGOMS6	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOMRX	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQGOM6Y	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQGOLZM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00499	CFP-100G-SR10
FPC 8	REV 39	750-045715	CACD1903	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1815	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QC480289	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QC480274	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130190	QSFP+-40G-SR4
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130197	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130180	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130199	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0415	MPC5E 24XGE OTN Mezz
FPC 9	REV 05	750-044444	CAAY9801	MPCE Type 2 3D P
CPU	REV 04	711-038484	CAAW3673	MPCE PMB 2G
MIC 0	REV 28	750-028387	CAAX1071	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T12L92342	XFP-10G-SR
Xcvr 1		NON-JNPR	T12L92303	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	CC07BK02X	XFP-10G-SR
QXM 0	REV 06	711-028408	CAAW4883	MPC QXM
QXM 1	REV 06	711-028408	CAAW4603	MPC QXM
FPC 10	REV 21.0.11	750-045715	CAAY3541	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 07	711-045719	CAAW7426	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-031980	AHK01AP	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502ZU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP41BLS	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQA08YA	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQA0K26	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA06S3	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQA06AS	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQA053N	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0E97	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA0GS4	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA0JVA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP
Xcvr 0	REV 01	740-021308	AQA057A	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	ANAOMLS	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQA093A	SFP+-10G-SR
Xcvr 3	REV 01	740-021309	943153A00075	SFP+-10G-LR
Xcvr 4	REV 01	740-021308	AQA077B	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQA0JSC	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQA0735	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ5028N	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AP40VN5	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQA0K0J	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQA07AP	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQA08YB	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP

WAN MEZZ	REV 07	750-045717	CAAX3123	MPC5E 24XGE Mezz
FPC 11	REV 17	750-037355	CAAT3986	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAR3972	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	AQA0DSE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501Y3	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ501XU	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ5036Y	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00247	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	ALQ1DKF	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403YA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AP40TY0	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14G0	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00095	CFP-100G-SR10
Fan Tray 0	REV 08	740-031521	ACAF4219	Enhanced Fan Tray
Fan Tray 1	REV 08	740-031521	ACAF4225	Enhanced Fan Tray

### show chassis hardware extensive (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware extensive
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1214852AFA	MX960
Jedec Code:	0x7fb0	EEPROM Version:	0x02	
		S/N:	JN1214852AFA	
Assembly ID:	0x0512	Assembly Version:	00.00	
Date:	00-00-0000	Assembly Flags:	0x00	
ID:	MX960			
Board Information Record:				
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00				
I2C Hex Data:				
Address 0x00: 7f b0 02 ff 05 12 00 00 00 00 00 00 00 00 00 00				
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00				
Address 0x20: 4a 4e 31 32 31 34 38 35 32 41 46 41 00 00 00 00				
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00				
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00				
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00				
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00				
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00				
Midplane	REV 01	710-030012	ACAX3674	MX960 Backplane
Jedec Code:	0x7fb0	EEPROM Version:	0x02	
P/N:	710-030012	S/N:	ACAX3674	
Assembly ID:	0x01df	Assembly Version:	01.01	
Date:	01-19-2013	Assembly Flags:	0x00	
Version:	REV 01	CLEI Code:	COM8T00CRB	
ID:	MX960 Backplane	FRU Model Number:	CHAS-BP-MX960-S	
Board Information Record:				
Address 0x00: ad 01 08 00 54 e0 32 bc 68 00 ff ff ff ff ff ff				
I2C Hex Data:				
Address 0x00: 7f b0 02 ff 01 df 01 01 52 45 56 20 30 31 00 00				
Address 0x10: 00 00 00 00 37 31 30 2d 30 33 30 30 31 32 00 00				
Address 0x20: 53 2f 4e 20 41 43 41 58 33 36 37 34 00 13 01 07				
Address 0x30: dd ff ff ff ad 01 08 00 54 e0 32 bc 68 00 ff ff				
Address 0x40: ff ff ff ff 01 43 4f 4d 38 54 30 30 43 52 42 43				
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 39 36 30 2d 53 00 00				
Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff				
Address 0x70: ff ff ff aa ff ff ff ff ff ff ff ff ff ff ff				
FPM Board	REV 03	710-014974	CAAZ9326	Front Panel Display



```

Jedec Code: 0x7fb0      EEPROM Version: 0x01
P/N: 710-014974      S/N: CAAZ9326
Assembly ID: 0x01e6    Assembly Version: 01.03
Date: 12-31-2012      Assembly Flags: 0x00
Version: REV 03
ID: Front Panel Display      FRU Model Number: CRAFT-MX960-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 01 e6 01 03 52 45 56 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 31 34 39 37 34 00 00
  Address 0x20: 53 2f 4e 20 43 41 41 5a 39 33 32 36 00 1f 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43
  Address 0x50: 52 41 46 54 2d 4d 58 39 36 30 2d 53 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PDM      Rev 03 740-013110 QCS17025017      Power Distribution Module
Jedec Code: 0x7fb0      EEPROM Version: 0x01
P/N: 740-013110      S/N: QCS17025017
Assembly ID: 0x0416    Assembly Version: 01.03
Date: 01-10-2013      Assembly Flags: 0x00
Version: Rev 03
ID: Power Distribution Module
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 16 01 03 52 65 76 20 30 33 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 31 33 31 31 30 00 00
  Address 0x20: 51 43 53 31 37 30 32 35 30 31 37 00 00 0a 01 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 0      Rev 10 740-027760 QCS1702N062      PS 4.1kW; 200-240V AC
in
Jedec Code: 0x7fb0      EEPROM Version: 0x01
P/N: 740-027760      S/N: QCS1702N062
Assembly ID: 0x0430    Assembly Version: 01.10
Date: 01-15-2013      Assembly Flags: 0x00
Version: Rev 10
ID: PS 4.1kW; 200-240V AC in      FRU Model Number: PWR-MX960-4100-AC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 30 01 0a 52 65 76 20 31 30 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
  Address 0x20: 51 43 53 31 37 30 32 4e 30 36 32 00 00 0f 01 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 1      Rev 04 740-027760 QCS1422N02C      PS 4.1kW; 200-240V AC
in
Jedec Code: 0x7fb0      EEPROM Version: 0x01
P/N: 740-027760      S/N: QCS1422N02C
Assembly ID: 0x0430    Assembly Version: 01.04
Date: 06-04-2010      Assembly Flags: 0x00
Version: Rev 04

```

```

ID: PS 4.1kW; 200-240V AC in    FRU Model Number: PWR-MX960-4100-AC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 30 01 04 52 65 76 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
  Address 0x20: 51 43 53 31 34 32 32 4e 30 32 43 00 00 04 06 07
  Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
PEM 2          Rev 09    740-027760    QCS1614N01X    PS 4.1kW; 200-240V AC
in
  Jedec Code: 0x7fb0          EEPROM Version: 0x01
  P/N: 740-027760          S/N: QCS1614N01X
Assembly ID: 0x0430          Assembly Version: 01.09
  Date: 04-07-2012          Assembly Flags: 0x00
  Version: Rev 09
ID: PS 4.1kW; 200-240V AC in    FRU Model Number: PWR-MX960-4100-AC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 30 01 09 52 65 76 20 30 39 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 32 37 37 36 30 00 00
  Address 0x20: 51 43 53 31 36 31 34 4e 30 31 58 00 00 07 04 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00 50
  Address 0x50: 57 52 2d 4d 58 39 36 30 2d 34 31 30 30 2d 41 43
  Address 0x60: 2d 53 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 08    740-031116    9009131803    RE-S-1800x4
  Jedec Code: 0x7fb0          EEPROM Version: 0x02
  P/N: 740-031116          S/N: 9009131803
  Assembly ID: 0x09c0          Assembly Version: 01.08
  Date: 03-04-2013          Assembly Flags: 0x00
  Version: REV 08          CLEI Code: COUCASKBAA
ID: RE-S-1800x4          FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
  Address 0x00: 54 32 30 32 37 44 42 2d 34 34 47 42 23 42 23 00
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
  Address 0x20: 39 30 30 39 31 33 31 38 30 33 00 00 00 04 03 07
  Address 0x30: dd ff ff ff 54 32 30 32 37 44 42 2d 34 34 47 42
  Address 0x40: 23 42 23 00 01 43 4f 55 43 41 53 4b 42 41 41 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 59 ff ff ff ff ff ff ff ff ff ff ff ff
ad0    3831 MB    UGB30SFA4000T1    SFA4000T1 000016CD Compact Flash
ad1    30533 MB   UGB94BPH32H0S1-KCI    11000061346    Disk 1
usb0 (addr 1) EHCI root hub 0    Intel    uhub0
usb0 (addr 2) product 0x0020 32    vendor 0x8087    uhub1
DIMM 0    VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
DIMM 1    VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
DIMM 2    VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
DIMM 3    VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
Routing Engine 1 REV 08    740-031116    9009124913    RE-S-1800x4
  Jedec Code: 0x7fb0          EEPROM Version: 0x02
  P/N: 740-031116          S/N: 9009124913
  Assembly ID: 0x09c0          Assembly Version: 01.08

```

```

Date:          01-09-2013      Assembly Flags:    0x00
Version:       REV 08         CLEI Code:       COUCASKBAA
ID: RE-S-1800x4      FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
Address 0x00: 54 32 30 32 37 44 42 2d 34 34 47 42 23 42 23 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 c0 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 31 31 31 36 00 00
Address 0x20: 39 30 30 39 31 32 34 39 31 33 00 00 00 09 01 07
Address 0x30: dd ff ff ff 54 32 30 32 37 44 42 2d 34 34 47 42
Address 0x40: 23 42 23 00 01 43 4f 55 43 41 53 4b 42 41 41 52
Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 59 ff ff ff ff ff ff ff ff ff ff ff ff
ad0   3831 MB   UGB30SFA4000T1      SFA4000T1 0000106D Compact Flash
ad1   30533 MB  UGB94BPH32H0S1-KCI  11000052402      Disk 1
CB 0          REV 18   750-031391  CABF0579      Enhanced MX SCB
Jedec Code:   0x7fb0      EEPROM Version: 0x02
P/N:          750-031391  S/N:          CABF0579
Assembly ID:  0x09b0      Assembly Version: 01.18
Date:         04-15-2013  Assembly Flags: 0x00
Version:      REV 18      CLEI Code:     COUCASRBAA
ID: Enhanced MX SCB      FRU Model Number: SCBE-MX-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 b0 01 12 52 45 56 20 31 38 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
Address 0x20: 53 2f 4e 20 43 41 42 46 30 35 37 39 00 0f 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 43 41 53 52 42 41 41 53
Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 43 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 7d ff ff ff ff ff ff ff ff ff ff ff ff
CB 1          REV 16   750-031391  CAAZ2471      Enhanced MX SCB
Jedec Code:   0x7fb0      EEPROM Version: 0x02
P/N:          750-031391  S/N:          CAAZ2471
Assembly ID:  0x09b0      Assembly Version: 01.16
Date:         03-09-2013  Assembly Flags: 0x00
Version:      REV 16      CLEI Code:     COUCARCBAB
ID: Enhanced MX SCB      FRU Model Number: SCBE-MX-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 b0 01 10 52 45 56 20 31 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00
Address 0x20: 53 2f 4e 20 43 41 41 5a 32 34 37 31 00 09 03 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 43 41 52 43 42 41 42 53
Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 6d ff ff ff ff ff ff ff ff ff ff ff ff
CB 2          REV 16   750-031391  CAAW9595      Enhanced MX SCB
Jedec Code:   0x7fb0      EEPROM Version: 0x02
P/N:          750-031391  S/N:          CAAW9595
Assembly ID:  0x09b0      Assembly Version: 01.16
Date:         02-01-2013  Assembly Flags: 0x00
Version:      REV 16      CLEI Code:     COUCARCBAB
ID: Enhanced MX SCB      FRU Model Number: SCBE-MX-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

## I2C Hex Data:

Address 0x00: 7f b0 02 ff 09 b0 01 10 52 45 56 20 31 36 00 00  
 Address 0x10: 00 00 00 00 37 35 30 2d 30 33 31 33 39 31 00 00  
 Address 0x20: 53 2f 4e 20 43 41 41 57 39 35 39 35 00 01 02 07  
 Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
 Address 0x40: ff ff ff ff 01 43 4f 55 43 41 52 43 42 41 42 53  
 Address 0x50: 43 42 45 2d 4d 58 2d 53 00 00 00 00 00 00 00 00  
 Address 0x60: 00 00 00 00 00 00 42 00 00 ff ff ff ff ff ff ff  
 Address 0x70: ff ff ff 6d ff ff ff ff ff ff ff ff ff ff ff ff

FPC 0 REV 18 750-046005 CACE6574 MPC5E 3D Q 2CGE+4XGE

Jedec Code: 0x7fb0 EEPROM Version: 0x02  
 P/N: 750-046005 S/N: CACE6574  
 Assembly ID: 0x0b8c Assembly Version: 01.18  
 Date: 11-20-2013 Assembly Flags: 0x00  
 Version: REV 18 CLEI Code: PROTOXCLEI  
 ID: MPC5E 3D Q 2CGE+4XGE FRU Model Number: PROTO-ASSEMBLY

## Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

## I2C Hex Data:

Address 0x00: 7f b0 02 ff 0b 8c 01 12 52 45 56 20 31 38 00 00  
 Address 0x10: 00 00 00 00 37 35 30 2d 30 34 36 30 30 35 00 00  
 Address 0x20: 53 2f 4e 20 43 41 43 45 36 35 37 34 00 14 0b 07  
 Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
 Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50  
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00  
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff  
 Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

CPU REV 09 711-045719 CACG8908 RMPC PMB

Jedec Code: 0x7fb0 EEPROM Version: 0x02  
 P/N: 711-045719 S/N: CACG8908  
 Assembly ID: 0x0b85 Assembly Version: 01.09  
 Date: 11-13-2013 Assembly Flags: 0x00  
 Version: REV 09

ID: RMPC PMB

## Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

## I2C Hex Data:

Address 0x00: 7f b0 02 ff 0b 85 01 09 52 45 56 20 30 39 00 00  
 Address 0x10: 00 00 00 00 37 31 31 2d 30 34 35 37 31 39 00 00  
 Address 0x20: 53 2f 4e 20 43 41 43 47 38 39 30 38 00 0d 0b 07  
 Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
 Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 50  
 Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00  
 Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff  
 Address 0x70: ff ff ff c2 00 00 00 00 00 00 00 00 00 00 00 00

PIC 0 BUILTIN BUILTIN 2X10GE SFPP OTN

Jedec Code: 0x0000 EEPROM Version: 0x00  
 P/N: BUILTIN S/N: BUILTIN  
 Assembly ID: 0x0a90 Assembly Version: 00.00  
 Date: 00-00-0000 Assembly Flags: 0x00

ID: 2X10GE SFPP OTN

## Board Information Record:

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

## I2C Hex Data:

Address 0x00: 00 00 00 00 0a 90 00 00 00 00 00 00 00 00 00 00  
 Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20  
 Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00  
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

```

Address 0x70: 00 00 00 00 c0 02 ae dc 00 00 00 00 0a 6e 00 00
Xcvr 0      REV 01  740-021308  AQA0DYT      SFP+-10G-SR
  Xcvr 1      REV 01  740-021308  AQGOMS7      SFP+-10G-SR
  PIC 1      BUILTIN  BUILTIN      1X100GE CFP2 OTN
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:      BUILTIN      S/N:      BUILTIN
Assembly ID: 0x0a6e      Assembly Version: 00.00
Date:      00-00-0000      Assembly Flags: 0x00
ID: 1X100GE CFP2 OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 6e 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 03 f3 8c 31 5c e7 80 00 00 00 02
  Xcvr 0      REV 01  740-046563  XD16FC03Z      CFP2-100G-SR10
  PIC 2      BUILTIN  BUILTIN      2X10GE SFPP OTN
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:      BUILTIN      S/N:      BUILTIN
Assembly ID: 0x0a90      Assembly Version: 00.00
Date:      00-00-0000      Assembly Flags: 0x00
ID: 2X10GE SFPP OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 90 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 03 f5 6c 31 5c db 40 00 00 00 02
  Xcvr 0      REV 01  740-021308  ANA0NAJ      SFP+-10G-SR
  Xcvr 1      REV 01  740-021308  AQGOMRQ      SFP+-10G-SR
  PIC 3      BUILTIN  BUILTIN      1X100GE CFP2 OTN
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:      BUILTIN      S/N:      BUILTIN
Assembly ID: 0x0a6e      Assembly Version: 00.00
Date:      00-00-0000      Assembly Flags: 0x00
ID: 1X100GE CFP2 OTN
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 6e 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 03 ed ec 31 5c e2 e8 00 00 00 02
Xcvr 0      REV 01  740-049775  J13K72993      CFP2-100G-LR4
FPC 1      REV 11  750-045372  CABK8154      MPCE Type 3 3D
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N:      750-045372      S/N:      CABK8154

```

```

Assembly ID: 0x09db          Assembly Version: 04.11
Date: 05-18-2013           Assembly Flags: 0x00
Version: REV 11            CLEI Code: COUIBBNBAA
ID: MPCE Type 3 3D         FRU Model Number: MX-MPC3E-3D
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 09 db 04 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 35 33 37 32 00 00
Address 0x20: 53 2f 4e 20 43 41 42 4b 38 31 35 34 00 12 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4e 42 41 41 4d
Address 0x50: 58 2d 4d 50 43 33 45 2d 33 44 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 44 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff cf ff ff ff ff ff ff ff ff ff ff ff ff
CPU REV 08 711-035209 CABE7370 HMPC PMB 2G
Jedec Code: 0x7fb0          EEPROM Version: 0x01
P/N: 711-035209            S/N: CABE7370
Assembly ID: 0x0b04         Assembly Version: 01.08
Date: 05-08-2013           Assembly Flags: 0x00
Version: REV 08
ID: HMPC PMB 2G
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
Address 0x20: 53 2f 4e 20 43 41 42 45 37 33 37 30 00 08 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0 REV 07 750-033307 CABD5255 10X10GE SFPP
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 750-033307            S/N: CABD5255
Assembly ID: 0x0a2a         Assembly Version: 02.07
Date: 04-25-2013           Assembly Flags: 0x00
Version: REV 07            CLEI Code: COUIBBJBAA
ID: 10X10GE SFPP          FRU Model Number: MIC3-3D-10XGE-SFPP
Board Information Record:
Address 0x00: 34 01 03 03 05 ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 fe 0a 2a 02 07 52 45 56 20 30 37 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 33 33 30 37 00 00
Address 0x20: 53 2f 4e 20 43 41 42 44 35 32 35 35 00 19 04 07
Address 0x30: dd ff ff ff 34 01 03 03 05 ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4a 42 41 41 4d
Address 0x50: 49 43 33 2d 33 44 2d 31 30 58 47 45 2d 53 46 50
Address 0x60: 50 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 82 c0 03 f0 bc 57 79 83 80 00 00 00 02
PIC 0 BUILTIN BUILTIN 10X10GE SFPP
Xcvr 0 REV 01 740-021308 AQ50319 SFP+-10G-SR
Xcvr 1 REV 01 740-021308 AQ5035V SFP+-10G-SR
Xcvr 2 REV 01 740-021308 AQ502XJ SFP+-10G-SR
Xcvr 3 REV 01 740-021308 AQ43HHR SFP+-10G-SR
Xcvr 4 REV 01 740-021308 AQ502YA SFP+-10G-SR
Xcvr 5 REV 01 740-021308 AQ502EU SFP+-10G-SR
Xcvr 6 REV 01 740-021308 AQ502HR SFP+-10G-SR
Xcvr 7 REV 01 740-021308 AQ502A6 SFP+-10G-SR
Xcvr 8 REV 01 740-021308 AQ43H8M SFP+-10G-SR

```

```

MIC 1          REV 14    750-033196    CAAP1398          1X100GE CXP
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-033196      S/N:           CAAP1398
Assembly ID:   0x0a29          Assembly Version: 03.14
Date:          10-27-2012      Assembly Flags: 0x00
Version:       REV 14          CLEI Code:     COUIBBKBAA
ID: 1X100GE CXP                FRU Model Number: MIC3-3D-1X100GE-CXP
Board Information Record:
Address 0x00: 34 01 07 07 08 ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 fe 0a 29 03 0e 52 45 56 20 31 34 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 33 31 39 36 00 00
Address 0x20: 53 2f 4e 20 43 41 41 50 31 33 39 38 00 1b 0a 07
Address 0x30: dc ff ff ff 34 01 07 07 08 ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 43 4f 55 49 42 42 4b 42 41 41 4d
Address 0x50: 49 43 33 2d 33 44 2d 31 58 31 30 30 47 45 2d 43
Address 0x60: 58 50 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 96 c0 03 ef cc 57 79 85 08 00 00 00 02
PIC 2          BUILTIN      BUILTIN          1X100GE CXP
Xcvr 0         REV 01      740-046563    XD16FC064          CFP2-100G-SR10
FPC 3          REV 35      750-028467    CAAT9156          MPC 3D 16x 10GE
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           750-028467      S/N:           CAAT9156
Assembly ID:   0x0997          Assembly Version: 01.35
Date:          12-17-2012      Assembly Flags: 0x00
Version:       REV 35
ID: MPC 3D 16x 10GE            FRU Model Number: MPC-3D-16XGE-SFPP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 97 01 23 52 45 56 20 33 35 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 32 38 34 36 37 00 00
Address 0x20: 53 2f 4e 20 43 41 41 54 39 31 35 36 00 11 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 4d
Address 0x50: 50 43 2d 33 44 2d 31 36 58 47 45 2d 53 46 50 50
Address 0x60: 00 00 00 00 00 00 00 00 ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CPU            REV 11      711-029089    CAAV4645          AMPC PMB
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-029089      S/N:           CAAV4645
Assembly ID:   0x0998          Assembly Version: 01.11
Date:          12-13-2012      Assembly Flags: 0x00
Version:       REV 11
ID: AMPC PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 09 98 01 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 32 39 30 38 39 00 00
Address 0x20: 53 2f 4e 20 43 41 41 56 34 36 34 35 00 0d 0c 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN          4x 10GE(LAN) SFP+
Jedec Code:    0x0000          EEPROM Version:    0x00
P/N:           BUILTIN          S/N:           BUILTIN
Assembly ID:   0x02fe          Assembly Version: 00.00
Date:          00-00-0000      Assembly Flags: 0x00

```

```

ID: 4x 10GE(LAN) SFP+
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 02 fe 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 6b 94 00 00 00 00 02 fe 00 00
Xcvr 0      REV 01 740-021308 AQ43HZ1      SFP+-10G-SR
Xcvr 1      REV 01 740-021308 AQ43HZC      SFP+-10G-SR
Xcvr 2      REV 01 740-021308 AQ43HD2      SFP+-10G-SR
Xcvr 3      REV 01 740-021308 AQ502HN      SFP+-10G-SR
PIC 1      BUILTIN BUILTIN      4x 10GE(LAN) SFP+
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:      BUILTIN      S/N:      BUILTIN
Assembly ID: 0x02fe      Assembly Version: 00.00
Date:      00-00-0000      Assembly Flags: 0x00
ID: 4x 10GE(LAN) SFP+
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 02 fe 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 25 73 3a 20
Address 0x20: 42 55 49 4c 54 49 4e 00 25 73 3a 20 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ac 0c 00 00 00 00 02 fe 00 00
Xcvr 0      REV 01 740-021308 AQ43HGF      SFP+-10G-SR
Xcvr 1      REV 01 740-021308 AQ501RZ      SFP+-10G-SR
Xcvr 2      REV 01 740-021308 AQ5029V      SFP+-10G-SR
Xcvr 3      REV 01 740-021308 AQ501X9      SFP+-10G-SR
PIC 2      BUILTIN BUILTIN      4x 10GE(LAN) SFP+
Jedec Code: 0x0000      EEPROM Version: 0x00
P/N:      BUILTIN      S/N:      BUILTIN
Assembly ID: 0x02fe      Assembly Version: 00.00
Date:      00-00-0000      Assembly Flags: 0x00
.....

```

### show chassis hardware models (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 01	710-030012	ACAX3674	CHAS-BP-MX960-S
FPM Board	REV 03	710-014974	CAAZ9326	CRAFT-MX960-S
PEM 0	Rev 10	740-027760	QCS1702N062	PWR-MX960-4100-AC-S
PEM 1	Rev 04	740-027760	QCS1422N02C	PWR-MX960-4100-AC-S
PEM 2	Rev 09	740-027760	QCS1614N01X	PWR-MX960-4100-AC-S
Routing Engine 0	REV 08	740-031116	9009131803	RE-S-1800X4-16G-S
Routing Engine 1	REV 08	740-031116	9009124913	RE-S-1800X4-16G-S
CB 0	REV 18	750-031391	CABF0579	SCBE-MX-S
CB 1	REV 16	750-031391	CAAZ2471	SCBE-MX-S
CB 2	REV 16	750-031391	CAAW9595	SCBE-MX-S
FPC 0	REV 18	750-046005	CACE6574	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	CABK8154	MX-MPC3E-3D



MIC 0	REV 07	750-033307	CABD5255	MIC3-3D-10XGE-SFPP
MIC 1	REV 14	750-033196	CAAP1398	MIC3-3D-1X100GE-CXP
FPC 3	REV 35	750-028467	CAAT9156	MPC-3D-16XGE-SFPP
FPC 4	REV 18	750-046005	CACE6568	PROTO-ASSEMBLY
FPC 5	REV 18	750-046005	CACE6577	PROTO-ASSEMBLY
FPC 7	REV 09	750-037355	CAAF0937	MPC4E-2CGE-8XGE
FPC 8	REV 39	750-045715	CACD1903	PROTO-ASSEMBLY
FPC 9	REV 05	750-044444	CAAY9801	MX-MPC2E-3D-P
MIC 0	REV 28	750-028387	CAAX1071	MIC-3D-4XGE-XFP
FPC 10	REV 21.0.11	750-045715	CAAY3541	PROTO-ASSEMBLY
FPC 11	REV 17	750-037355	CAAT3986	MPC4E-3D-2CGE-8XGE
Fan Tray 0	REV 08	740-031521	ACAF4219	FFANTRAY-MX960-HC-S
Fan Tray 1	REV 08	740-031521	ACAF4225	FFANTRAY-MX960-HC-S

### show chassis hardware clei-models (MX960 Router with MPC5EQ)

```
user@host> show chassis hardware clei-models
```

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-030012	COM8T00CRB	CHAS-BP-MX960-S
FPM Board	REV 03	710-014974		CRAFT-MX960-S
PEM 0	Rev 10	740-027760		PWR-MX960-4100-AC-S
PEM 1	Rev 04	740-027760		PWR-MX960-4100-AC-S
PEM 2	Rev 09	740-027760		PWR-MX960-4100-AC-S
Routing Engine 0	REV 08	740-031116	COUCASKBAA	RE-S-1800X4-16G-S
Routing Engine 1	REV 08	740-031116	COUCASKBAA	RE-S-1800X4-16G-S
CB 0	REV 18	750-031391	COUCASRBAA	SCBE-MX-S
CB 1	REV 16	750-031391	COUCARCBAB	SCBE-MX-S
CB 2	REV 16	750-031391	COUCARCBAB	SCBE-MX-S
FPC 0	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	COUIBBNBAA	MX-MPC3E-3D
MIC 0	REV 07	750-033307	COUIBBJBAA	MIC3-3D-10XGE-SFPP
MIC 1	REV 14	750-033196	COUIBBKBAA	MIC3-3D-1X100GE-CXP
FPC 3	REV 35	750-028467		MPC-3D-16XGE-SFPP
FPC 4	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 5	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 7	REV 09	750-037355	PROTOXCLEI	MPC4E-2CGE-8XGE
FPC 8	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 9	REV 05	750-044444	COUIBBGBAA	MX-MPC2E-3D-P
MIC 0	REV 28	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 10	REV 21.0.11	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 11	REV 17	750-037355	IPU3A4DHAA	MPC4E-3D-2CGE-8XGE
Fan Tray 0	REV 08	740-031521		FFANTRAY-MX960-HC-S
Fan Tray 1	REV 08	740-031521		FFANTRAY-MX960-HC-S

### show chassis hardware (MX2010 Router)

```
user@host > show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11E3217AFK	MX2010
Midplane	REV 01	750-044636	ABAB8506	Lower Backplane
Midplane 1	REV 01	711-044557	ZY8296	Upper Backplane
PMP	REV 03	711-032426	ACAJ1388	Power Midplane
FPM Board	REV 06	711-032349	ZX8744	Front Panel Display
PSM 4	REV 0C	740-033727	VK00254	DC 52V Power Supply
Module				
PSM 5	REV 0B	740-033727	VG00015	DC 52V Power Supply
Module				
PSM 6	REV 0B	740-033727	VH00097	DC 52V Power Supply
Module				

PSM 7 Module	REV 0C	740-033727	VJ00151	DC 52V Power Supply
PSM 8 Module	REV 0C	740-033727	VJ00149	DC 52V Power Supply
PDM 0	REV 0B	740-038109	WA00008	DC Power Dist Module
PDM 1	REV 0B	740-038109	WA00014	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009094134	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009094141	RE-S-1800x4
CB 0	REV 08	750-040257	CAAB3491	Control Board
CB 1	REV 08	750-040257	CAAB3489	Control Board
SPMB 0	REV 02	711-041855	CAA6135	PMB Board
SPMB 1	REV 02	711-041855	CAA6137	PMB Board
SFB 0	REV 06	711-032385	ZV1828	Switch Fabric Board
SFB 1	REV 07	711-032385	ZZ2568	Switch Fabric Board
SFB 2	REV 07	711-032385	ZZ2563	Switch Fabric Board
SFB 3	REV 07	711-032385	ZZ2564	Switch Fabric Board
SFB 4	REV 07	711-032385	ZZ2580	Switch Fabric Board
SFB 5	REV 07	711-032385	ZZ2579	Switch Fabric Board
SFB 6	REV 07	711-032385	CAAB4882	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4898	Switch Fabric Board
FPC 0	REV 33	750-028467	CAAB1919	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAB7174	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH02RE	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH038C	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH0390	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMG0SUA	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0579	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0SGP	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH04SV	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04X3	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH0135	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02NC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02XB	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH02PN	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMH057Y	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMG0JHE	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AMH02HT	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AMH04V4	SFP+-10G-SR
FPC 1	REV 21	750-033205	ZG5027	MPC Type 3
CPU	REV 04	711-035209	YT4780	HMPC PMB 2G
MIC 0	REV 03	750-033307	ZV6299	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	083363A00410	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	083363A00334	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	113363A00125	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	083363A00953	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AHR013D	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ40JUR	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JKL	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ30ECK	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100864	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511100868	SFP+-10G-SR
MIC 1	REV 03	750-033307	ZV6268	10X10GE SFPP
PIC 2		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-031980	AJC0JML	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ403PC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ10N25	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	AJ40JF4	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JSJ	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJ403V7	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJ40JN3	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJ40JSU	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	19T511100468	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	19T511101363	SFP+-10G-SR
FPC 8	REV 22	750-031089	ZT9746	MPC Type 2 3D
CPU	REV 06	711-030884	ZS1271	MPC PMB 2G
MIC 0	REV 26	750-028392	ABBS1150	3D 20x 1GE(LAN) SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	PLG023C	SFP-SX
Xcvr 1	REV 01	740-031851	PLG09C6	SFP-SX
Xcvr 2	REV 02	740-011613	AM0950SF9L7	SFP-SX
Xcvr 3	REV 02	740-011613	AM1001SFN1H	SFP-SX
Xcvr 4	REV 02	740-011613	AM1001SFM9D	SFP-SX
Xcvr 5	REV 02	740-011613	AM1001SFLTJ	SFP-SX
Xcvr 6	REV 01	740-031851	AC1108S03L9	SFP-SX
Xcvr 7	REV 01	740-031851	AC1102S00NC	SFP-SX
Xcvr 8	REV 01	740-031851	AC1102S00MX	SFP-SX
Xcvr 9	REV 01	740-031851	AC1102S0085	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-031851	AC1102S00KU	SFP-SX
Xcvr 1	REV 01	740-031851	AC1102S00NG	SFP-SX
Xcvr 2	REV 01	740-031851	AC1102S00K3	SFP-SX
Xcvr 3	REV 01	740-031851	AC1102S008R	SFP-SX
Xcvr 4	REV 01	740-031851	AM1107SUFVJ	SFP-SX
Xcvr 5	REV 01	740-031851	AC1108S03LG	SFP-SX
MIC 1	REV 26	750-028387	ABBR9582	3D 4x 10GE XFP
PIC 2		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T10A91703	XFP-10G-SR
Xcvr 1		NON-JNPR	T09L42604	XFP-10G-SR
PIC 3		BUILTIN	BUILTIN	2x 10GE XFP
FPC 9	REV 11	750-036284	ZL3591	MPC 3D 16x 10GE EM
CPU	REV 10	711-029089	ZL0513	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101825	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101821	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101682	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ13R6	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101828	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101716	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALP0TR1	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101741	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101829	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ14E3	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	1YT517101826	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	1YT517101817	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	1YT517101735	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	ALQ159A	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2073	Adapter Card
ADC 1	REV 01	750-043596	ZV4117	Adapter Card
ADC 8	REV 01	750-043596	ZV4107	Adapter Card
ADC 9	REV 02	750-043596	ZW1555	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0015	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0019	172mm FanTray - 6 Fans

Fan Tray 2	REV 2A	760-046960	ACAY0020	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0021	172mm FanTray - 6 Fans

**show chassis hardware detail (MX2010 Router)**

```

user@host > show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11E233DAFK  MX2010
Midplane      REV 26   750-044636   ABAB9357      Lower Backplane
Midplane 1    REV 01   711-044557   ABAB8643      Upper Backplane
PMP           REV 04   711-032426   ACAJ1677      Power Midplane
FPM Board     REV 08   760-044634   ABBV9726      Front Panel Display
PSM 0         REV 01   740-045050   1E02224000P   DC 52V Power Supply
Module
PSM 1         REV 01   740-045050   1E02224000M   DC 52V Power Supply
Module
PSM 2         REV 01   740-045050   1E022240010   DC 52V Power Supply
Module
PSM 3         REV 01   740-045050   1E02224000G   DC 52V Power Supply
Module
PSM 4         REV 01   740-045050   1E022240013   DC 52V Power Supply
Module
PSM 5         REV 01   740-045050   1E022240007   DC 52V Power Supply
Module
PSM 6         REV 01   740-045050   1E02224001C   DC 52V Power Supply
Module
PSM 7         REV 01   740-045050   1E02224001D   DC 52V Power Supply
Module
PSM 8         REV 01   740-045050   1E02224001B   DC 52V Power Supply
Module
PDM 0         REV 01   740-045234   1E262250067   DC Power Dist Module
Routing Engine 0 REV 02   740-041821   9009099704    RE-S-1800x4
  ad0  3831 MB  UGB30SFA4000T1  SFA4000T1 00000651 Compact Flash
  ad1  30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
  usb0 (addr 1) EHCI root hub 0 Intel uhub0
  usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
  DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
  DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02   740-041821   9009099706    RE-S-1800x4
  ad0  3998 MB Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
  ad1  30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0          REV 13   750-040257   CAAF8436      Control Board
CB 1          REV 13   750-040257   CAAF8434      Control Board
SPMB 0        REV 02   711-041855   ABBV3825      PMB Board
SPMB 1        REV 02   711-041855   ABBV3833      PMB Board
SFB 0         REV 05   711-044466   ABBX5682      Switch Fabric Board
SFB 1         REV 05   711-044466   ABBX5676      Switch Fabric Board
SFB 2         REV 05   711-044466   ABBX5665      Switch Fabric Board
SFB 3         REV 05   711-044466   ABBX5699      Switch Fabric Board
SFB 4         REV 05   711-044466   ABBX5603      Switch Fabric Board
SFB 5         REV 05   711-044466   ABBX5587      Switch Fabric Board
SFB 6         REV 05   711-044466   ABBX5607      Switch Fabric Board
SFB 7         REV 05   711-044466   ABBX5669      Switch Fabric Board
FPC 0         REV 09   750-037355   CAAF0924      MPC Type 4-2
CPU           REV 08   711-035209   CAAB9842      HMPC PMB 2G
PIC 0         BUILTIN BUILTIN      4x10GE SFPP
  Xcvr 0       REV 01   740-021308   19T511101656 SFP+-10G-SR
  Xcvr 1       REV 01   740-031980   AMA04RU      SFP+-10G-SR

```

Xcvr 2	REV 01	740-031980	193363A00558	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M00202	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00328	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	AMA088W	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10L04211	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	19T511101602	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10L04151	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00332	CFP-100G-SR10
FPC 1	REV 18	750-033205	ZE0128	MPC Type 3
CPU	REV 06	711-035209	ZG5431	HMPC PMB 2G
MIC 0	REV 15	750-033199	ZP6435	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	J11E46118	CFP-100G-LR4
MIC 1	REV 15	750-033199	ZP6442	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	UMN03T4	CFP-100G-LR4
FPC 2	REV 16	750-037358	CAAL1001	MPC Type 4-1
CPU	REV 08	711-035209	CAAK7927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00589	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00028	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00376	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00016	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00499	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00039	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E01239	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B10M00075	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00014	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA0638	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00063	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AMA0629	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00053	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00344	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00046	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062M	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00080	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00580	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00064	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	093363A01494	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00020	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	123363A00047	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00072	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01033	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00022	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00013	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01028	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	973152A00079	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01018	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	973152A00025	SFP+-10G-SR
FPC 3	REV 33	750-028467	CAAF5400	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7626	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00066	SFP+-10G-SR

Xcvr 1	REV 01	740-021308	973152A00021	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00062	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00027	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00065	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00069	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00026	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00003	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00035	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00004	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00049	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00055	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00010	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	973152A00001	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	973152A00073	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	973152A00012	SFP+-10G-SR
FPC 4	REV 21	750-033205	ZG5028	MPC Type 3
CPU	REV 05	711-035209	YX3911	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2036	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220708	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220735	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2028	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB220727	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB220715	QSFP+-40G-SR4
FPC 5	REV 11	750-037358	CAAE2196	MPC Type 4-1
CPU	REV 08	711-035209	CAAD9074	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062S	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA062P	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA052R	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA0632	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00564	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00229	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00363	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00278	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04CC	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A001W	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA04N2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA062U	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00491	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	183363A01511	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00565	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00405	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA07QX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AMA06MS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00318	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	193363A00402	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00174	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00388	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00377	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00234	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA062T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00550	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00364	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	AMA0630	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	193363A00509	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	193363A00459	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	113363A00191	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00352	SFP+-10G-SR
FPC 6	REV 33	750-028467	CAAF5552	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAH7601	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AD0927A0036	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AD0927A003M	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0927A003G	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0927A0031	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	193363A00331	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00325	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00417	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A02509	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75140	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11A04356	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01952	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01914	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	T09K75157	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	T09K75194	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01926	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01936	SFP+-10G-SR
FPC 7	REV 16	750-037358	CAAL1012	MPC Type 4-1
CPU	REV 08	711-035209	CAAJ3851	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	AMA04NK	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11F00260	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11E02192	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04CP	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJ40JJK	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11F00238	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B10M00275	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	193363A00211	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	B11D05577	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11G00586	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AMA08B7	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AMA04Q0	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11D05840	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11E00467	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11E00029	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	19T511101712	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00568	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00166	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10M00212	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11D05823	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	03DZ06A01005	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	03DZ06A01003	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	03DZ06A01009	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01004	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	8X10GE SFPP
Xcvr 0	REV 01	740-021308	03DZ06A01017	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	03DZ06A01016	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	03DZ06A01024	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	03DZ06A01008	SFP+-10G-SR

Xcvr 4	REV 01	740-030658	AD0946A02UH	SFP+-10G-USR
Xcvr 5	REV 01	740-021308	T09J67913	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AD0837ES09G	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	03DZ06A01015	SFP+-10G-SR
FPC 8	REV 03	750-045372	CAAD3111	MPC Type 3
CPU	REV 08	711-035209	CAAD8033	HMPC PMB 2G
MIC 0	REV 03	750-036233	ZL2032	2X40GE QSFP
PIC 0		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB230273	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB230254	QSFP+-40G-SR4
MIC 1	REV 03	750-036233	ZL2021	2X40GE QSFP
PIC 2		BUILTIN	BUILTIN	2X40GE QSFP
Xcvr 0	REV 01	740-032986	QB390962	QSFP+-40G-SR4
Xcvr 1	REV 01	740-032986	QB390960	QSFP+-40G-SR4
FPC 9	REV 09	750-037355	CAAF1531	MPC Type 4-2
CPU	REV 08	711-035209	CAAB9927	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	193363A00525	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	193363A00504	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	193363A00368	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJ40JSS	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	123363A00042	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B10M00023	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ802EM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11E02348	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
ADC 0	REV 13	750-043596	ABBX5532	Adapter Card
ADC 1	REV 13	750-043596	ABBX5550	Adapter Card
ADC 2	REV 13	750-043596	ABBX5571	Adapter Card
ADC 3	REV 13	750-043596	ABBX5568	Adapter Card
ADC 4	REV 13	750-043596	ABBX5556	Adapter Card
ADC 5	REV 13	750-043596	ABBX5553	Adapter Card
ADC 6	REV 13	750-043596	ABBX5541	Adapter Card
ADC 7	REV 13	750-043596	ABBX5578	Adapter Card
ADC 8	REV 13	750-043596	ABBX5560	Adapter Card
ADC 9	REV 07	750-043596	ABBV7188	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0127	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0068	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0072	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0070	172mm FanTray - 6 Fans

### show chassis hardware extensive (MX2010 Router)

```
user@host > show chassis hardware extensive
```

```
Hardware inventory:
```

```

Item              Version  Part number  Serial number  Description
Chassis
Jedec Code:       0x7fb0          EEPROM Version: 0x02
                                   S/N:           JN11E233DAFK
Assembly ID:      0x0557          Assembly Version: 00.00
Date:             00-00-0000      Assembly Flags:  0x00
ID: MX2010
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 45 32 33 33 44 41 46 4b 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00

```



```

Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane          REV 26    750-044636    ABAB9357          Lower Backplane
Jedec Code:      0x7fb0          EEPROM Version:    0x02
P/N:             750-044636          S/N:             ABAB9357
Assembly ID:     0x0b66          Assembly Version: 01.26
Date:            08-28-2012        Assembly Flags:   0x00
Version:         REV 26          CLEI Code:        PROTOXCLEI
ID: Lower Backplane          FRU Model Number: PROTO-ASSEMBLY
Board Information Record:
Address 0x00: ad 01 08 00 2c 21 72 70 a0 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 66 01 1a 52 45 56 20 32 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 35 37 00 1c 08 07
Address 0x30: dc ff ff ff ad 01 08 00 2c 21 72 70 a0 00 ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1        REV 01    711-044557    ABAB8643          Upper Backplane
Jedec Code:      0x7fb0          EEPROM Version:    0x01
P/N:             711-044557          S/N:             ABAB8643
Assembly ID:     0x0b65          Assembly Version: 01.01
Date:            07-27-2012        Assembly Flags:   0x00
Version:         REV 01
ID: Upper Backplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 65 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 38 36 34 33 00 1b 07 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP               REV 04    711-032426    ACAJ1677          Power Midplane
Jedec Code:      0x7fb0          EEPROM Version:    0x01
P/N:             711-032426          S/N:             ACAJ1677
Assembly ID:     0x045d          Assembly Version: 01.04
Date:            07-20-2012        Assembly Flags:   0x00
Version:         REV 04
ID: Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 31 36 37 37 00 14 07 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board         REV 08    760-044634    ABBV9726          Front Panel Display
Jedec Code:      0x7fb0          EEPROM Version:    0x02
P/N:             760-044634          S/N:             ABBV9726

```

```

Assembly ID: 0x0b64      Assembly Version: 01.08
Date: 09-10-2012      Assembly Flags: 0x00
Version: REV 08      CLEI Code: IPMYA4EJRA
ID: Front Panel Display  FRU Model Number: MX2010-CRAFT-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 64 01 08 52 45 56 20 30 38 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 42 56 39 37 32 36 00 0a 09 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
  Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0      REV 01  740-045050  1E02224000P  DC 52V Power Supply
Module
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-045050      S/N: 1E02224000P
Assembly ID: 0x0478      Assembly Version: 01.01
Date: 12-06-2012      Assembly Flags: 0x00
Version: REV 01      CLEI Code: XXXXXXXXXX
ID: DC 52V Power Supply Module  FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
  Address 0x20: 31 45 30 32 32 32 34 30 30 30 50 00 00 06 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
  Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1      REV 01  740-045050  1E02224000M  DC 52V Power Supply
Module
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-045050      S/N: 1E02224000M
Assembly ID: 0x0478      Assembly Version: 01.01
Date: 12-06-2012      Assembly Flags: 0x00
Version: REV 01      CLEI Code: XXXXXXXXXX
ID: DC 52V Power Supply Module  FRU Model Number: MX2000-PSM-HC-DC-S-A
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 30 35 30 00 00
  Address 0x20: 31 45 30 32 32 32 34 30 30 30 4d 00 00 06 0c 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 58 58 58 58 58 58 58 58 58 58 4d
  Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 48 43 2d 44 43 2d
  Address 0x60: 53 2d 41 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 4a 00 00 00 00 00 00 00 00 00 00 00 00
...
PDM 0      REV 01  740-045234  1E262250067  DC Power Dist Module
Jedec Code: 0x7fb0      EEPROM Version: 0x02
P/N: 740-045234      S/N: 1E262250067
Assembly ID: 0x047b      Assembly Version: 01.01
Date: 06-28-2012      Assembly Flags: 0x00
Version: REV 01      CLEI Code: IPUPAJSKAA
ID: DC Power Dist Module  FRU Model Number: MX2000-PDM-DC-S-A

```

```

Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 7b 01 01 52 45 56 20 30 31 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00
  Address 0x20: 31 45 32 36 32 32 35 30 30 36 37 00 00 1c 06 07
  Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d
  Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 2d 41
  Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 89 00 00 00 00 00 00 00 00 00 00 00 00
Routing Engine 0 REV 02 740-041821 9009099704 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099704
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 03-15-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
  Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
  Address 0x20: 39 30 30 39 30 39 39 37 30 34 00 00 00 0f 03 07
  Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
  Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3831 MB UGB30SFA4000T1 SFA4000T1 00000651 Compact Flash
ad1 30533 MB UGB94BPH32H0S1-KCI 11000019592 Disk 1
usb0 (addr 1) EHCI root hub 0 Intel uhub0
usb0 (addr 2) product 0x0020 32 vendor 0x8087 uhub1
DIMM 0 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 1 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 2 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
DIMM 3 SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80
Routing Engine 1 REV 02 740-041821 9009099706 RE-S-1800x4
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 740-041821 S/N: 9009099706
Assembly ID: 0x09c0 Assembly Version: 01.02
Date: 02-23-2012 Assembly Flags: 0x00
Version: REV 02
ID: RE-S-1800x4 FRU Model Number: RE-S-1800X4-16G-S
Board Information Record:
  Address 0x00: 54 32 30 32 37 44 41 2d 34 34 47 42 23 41 23 00
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 09 c0 01 02 52 45 56 20 30 32 00 00
  Address 0x10: 00 00 00 00 37 34 30 2d 30 34 31 38 32 31 00 00
  Address 0x20: 39 30 30 39 30 39 39 37 30 36 00 00 00 17 02 07
  Address 0x30: dc ff ff ff 54 32 30 32 37 44 41 2d 34 34 47 42
  Address 0x40: 23 41 23 00 01 00 00 00 00 00 00 00 00 00 00 52
  Address 0x50: 45 2d 53 2d 31 38 30 30 58 34 2d 31 36 47 2d 53
  Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 8c ff ff ff ff ff ff ff ff ff ff ff ff
ad0 3998 MB Virtium - TuffDrive VCF P1T0200262860208 114 Compact Flash
ad1 30533 MB UGB94ARF32H0S3-KC UNIGEN-499551-000404 Disk 1
CB 0 REV 13 750-040257 CAAF8436 Control Board
Jedec Code: 0x7fb0 EEPROM Version: 0x02
P/N: 750-040257 S/N: CAAF8436
Assembly ID: 0x0b26 Assembly Version: 01.13

```

```

Date:          08-29-2012      Assembly Flags:    0x00
Version:       REV 13          CLEI Code:       PROTOXCLEI
ID: Control Board              FRU Model Number:  PROTO-ASSEMBLY

```

## Board Information Record:

```
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

## I2C Hex Data:

```

Address 0x00: 7f b0 02 ff 0b 26 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 35 37 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 38 34 33 36 00 1d 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 ff ff ff ff ff ff ff ff ff ff ff ff

```

...

```
SPMB 0          REV 02    711-041855    ABBV3825          PMB Board
```

```

Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-041855      S/N:              ABBV3825
Assembly ID:   0x0b29          Assembly Version:  01.02
Date:          08-14-2012      Assembly Flags:    0x00
Version:       REV 02
ID: PMB Board

```

## Board Information Record:

```
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

## I2C Hex Data:

```

Address 0x00: 7f b0 01 ff 0b 29 01 02 52 45 56 20 30 32 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 31 38 35 35 00 00
Address 0x20: 53 2f 4e 20 41 42 42 56 33 38 32 35 00 0e 08 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00

```

...

```
SFB 0          REV 05    711-044466    ABBX5682          Switch Fabric Board
```

```

Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           711-044466      S/N:              ABBX5682
Assembly ID:   0x0b25          Assembly Version:  01.05
Date:          09-07-2012      Assembly Flags:    0x00
Version:       REV 05          CLEI Code:       PROTOXCLEI
ID: Switch Fabric Board        FRU Model Number:  PROTO-ASSEMBLY

```

## Board Information Record:

```
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
```

## I2C Hex Data:

```

Address 0x00: 7f b0 02 ff 0b 25 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 34 36 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 36 38 32 00 07 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 50
Address 0x50: 52 4f 54 4f 2d 41 53 53 45 4d 42 4c 59 00 00 00
Address 0x60: 00 00 00 00 00 00 41 30 30 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c2 00 00 00 01 00 00 00 00 00 00 48 00

```

...

```
FPC 0          REV 09    750-037355    CAAF0924          MPC Type 4-2
```

```

Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037355      S/N:              CAAF0924
Assembly ID:   0x0b4e          Assembly Version:  01.09
Date:          05-21-2012      Assembly Flags:    0x00
Version:       REV 09          CLEI Code:       PROTOXCLEI
ID: MPC Type 4-2              FRU Model Number:  MPC4E-2CGE-8XGE

```

## Board Information Record:

```

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 4e 01 09 52 45 56 20 30 39 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 33 35 35 00 00
Address 0x20: 53 2f 4e 20 43 41 41 46 30 39 32 34 00 15 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 4d
Address 0x50: 50 43 34 45 2d 32 43 47 45 2d 38 58 47 45 00 00
Address 0x60: 00 00 00 00 00 00 30 39 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff c6 ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 08    711-035209    CAAB9842          HMPC PMB 2G
Jedec Code:  0x7fb0          EEPROM Version:  0x01
P/N:         711-035209          S/N:          CAAB9842
Assembly ID: 0x0b04          Assembly Version: 01.08
Date:        05-17-2012          Assembly Flags: 0x00
Version:     REV 08
ID: HMPC PMB 2G
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0b 04 01 08 52 45 56 20 30 38 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 35 32 30 39 00 00
Address 0x20: 53 2f 4e 20 43 41 41 42 39 38 34 32 00 11 05 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
PIC 0          BUILTIN          BUILTIN          4x10GE SFPP
Jedec Code:  0x0000          EEPROM Version:  0x00
P/N:         BUILTIN          S/N:          BUILTIN
Assembly ID: 0x0a53          Assembly Version: 00.00
Date:        00-00-0000          Assembly Flags:  0x00
ID: 4x10GE SFPP
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 00 00 00 00 0a 53 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 42 55 49 4c 54 49 4e 00 4d 58 43 00
Address 0x20: 42 55 49 4c 54 49 4e 00 4d 58 43 00 00 00 00 00
Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 c0 02 ae 64 00 00 00 00 0a 52 00 00
Xcvr 0      REV 01    740-021308    19T511101656      SFP+-10G-SR
Xcvr 1      REV 01    740-031980    AMA04RU           SFP+-10G-SR
Xcvr 2      REV 01    740-031980    193363A00558      SFP+-10G-SR
Xcvr 3      REV 01    740-031980    B10M00202         SFP+-10G-SR
...
ADC 0      REV 13    750-043596    ABBX5532          Adapter Card
Jedec Code: 0x7fb0          EEPROM Version:  0x02
P/N:        750-043596          S/N:          ABBX5532
Assembly ID: 0x0b3d          Assembly Version: 01.13
Date:       09-12-2012          Assembly Flags: 0x00
Version:    REV 13          CLEI Code:     IPUCBA8CAA
ID: Adapter Card          FRU Model Number: MX2000-LC-ADAPTER
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 3d 01 0d 52 45 56 20 31 33 00 00

```

```

Address 0x10: 00 00 00 00 37 35 30 2d 30 34 33 35 39 36 00 00
Address 0x20: 53 2f 4e 20 41 42 42 58 35 35 33 32 00 0c 09 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 43 42 41 38 43 41 41 4d
Address 0x50: 58 32 30 30 30 2d 4c 43 2d 41 44 41 50 54 45 52
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff
Address 0x70: ff ff ff 3a 00 00 00 00 00 00 00 00 00 00 00
...

```

### show chassis hardware models (MX2010 Router)

```

user@host > show chassis hardware models
Hardware inventory:
Item                Version  Part number  Serial number  FRU model number
FPM Board           REV 06   711-032349   ZX8744         711-032349
PSM 4               REV 0C   740-033727   VK00254       000000000000000000000000
PSM 5               REV 0B   740-033727   VG00015       000000000000000000000000
PSM 6               REV 0B   740-033727   VH00097       000000000000000000000000
PSM 7               REV 0C   740-033727   VJ00151       000000000000000000000000
PSM 8               REV 0C   740-033727   VJ00149       000000000000000000000000
PDM 0               REV 0B   740-038109   WA00008
PDM 1               REV 0B   740-038109   WA00014
Routing Engine 0    REV 02   740-041821   9009094134    RE-S-1800X4-16G-S
Routing Engine 1    REV 02   740-041821   9009094141    RE-S-1800X4-16G-S
CB 0                REV 08   750-040257   CAAB3491      750-040257
CB 1                REV 08   750-040257   CAAB3489      750-040257
SFB 0               REV 06   711-032385   ZV1828        711-032385
SFB 1               REV 07   711-032385   ZZ2568        711-032385
SFB 2               REV 07   711-032385   ZZ2563        711-032385
SFB 3               REV 07   711-032385   ZZ2564        711-032385
SFB 4               REV 07   711-032385   ZZ2580        711-032385
SFB 5               REV 07   711-032385   ZZ2579        711-032385
SFB 6               REV 07   711-032385   CAAB4882      711-044170
SFB 7               REV 07   711-032385   CAAB4898      711-044170
FPC 0               REV 33   750-028467   CAAB1919      MPC-3D-16XGE-SFPP
FPC 1               REV 21   750-033205   ZG5027        MX-MPC3-3D
    MIC 0            REV 03   750-033307   ZV6299        MIC3-3D-10XGE-SFPP
    MIC 1            REV 03   750-033307   ZV6268        MIC3-3D-10XGE-SFPP
FPC 8               REV 22   750-031089   ZT9746        MX-MPC2-3D
    MIC 0            REV 26   750-028392   ABBS1150      MIC-3D-20GE-SFP
    MIC 1            REV 26   750-028387   ABBR9582      MIC-3D-4XGE-XFP
FPC 9               REV 11   750-036284   ZL3591        MPCE-3D-16XGE-SFPP
ADC 0               REV 05   750-043596   CAAC2073      750-043596
ADC 1               REV 01   750-043596   ZV4117        750-043596
ADC 8               REV 01   750-043596   ZV4107        750-043596
ADC 9               REV 02   750-043596   ZW1555        750-043596
Fan Tray 0          REV 2A   760-046960   ACAY0015
Fan Tray 1          REV 2A   760-046960   ACAY0019
Fan Tray 2          REV 2A   760-046960   ACAY0020
Fan Tray 3          REV 2A   760-046960   ACAY0021

```

### show chassis hardware clei-models (MX2010 Routers)

```

user@host > show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code      FRU model number
FPM Board           REV 06   711-032349   PROTOXCLEI     711-032349
PSM 4               REV 0C   740-033727   0000000000     000000000000000000000000
PSM 5               REV 0B   740-033727   0000000000     000000000000000000000000
PSM 6               REV 0B   740-033727   0000000000     000000000000000000000000
PSM 7               REV 0C   740-033727   0000000000     000000000000000000000000

```

PSM 8	REV 0C	740-033727	0000000000	000000000000000000000000
PDM 0	REV 0B	740-038109		
PDM 1	REV 0B	740-038109		
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 08	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 06	711-032385	PROTOXCLEI	711-032385
SFB 1	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 2	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 3	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 4	REV 07	711-032385	PROTOXCLEI	711-032385
SFB 5	REV 07	711-032385	PROTOXCLEI	711-0323856
SFB 6	REV 07	711-032385	PROTOXCLEI	711-044170
SFB 7	REV 07	711-032385	PROTOXCLEI	711-044170
FPC 0	REV 33	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 21	750-033205		MX-MPC3-3D
MIC 0	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
MIC 1	REV 03	750-033307	PROTOXCLEI	MIC3-3D-10XGE-SFPP
FPC 8	REV 22	750-031089	COUIBAYBAA	MX-MPC2-3D
MIC 0	REV 26	750-028392	COUIA15BAA	MIC-3D-20GE-SFP
MIC 1	REV 26	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 9	REV 11	750-036284	CMUIACGBAA	MPCE-3D-16XGE-SFPP
ADC 0	REV 05	750-043596	PROTOXCLEI	750-043596
ADC 1	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 8	REV 01	750-043596	PROTOXCLEI	750-043596
ADC 9	REV 02	750-043596	PROTOXCLEI	750-043596
Fan Tray 0	REV 2A	760-046960		
Fan Tray 1	REV 2A	760-046960		
Fan Tray 2	REV 2A	760-046960		
Fan Tray 3	REV 2A	760-046960		

#### show chassis hardware (MX2010 Routers with MPC6E and OTN MIC)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11C9AFEAFK	MX2010
Midplane	REV 35	750-044636	ABAB9188	Lower Backplane
Midplane 1	REV 02	711-044557	ABAB8729	Upper Backplane
PMP	REV 04	711-032426	ACAJ2432	Power Midplane
Front Panel Board	REV 09	760-044634	ABCA4314	Front Panel Display
PSM 0	REV 01	740-050037	1EDB321015C	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB321015J	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB32000K8	DC 52V Power Supply
Module				
PSM 3	REV 01	740-050037	1EDB32101JW	DC 52V Power Supply
Module				
PSM 4	REV 01	740-050037	1EDB321015G	DC 52V Power Supply
Module				
PSM 5	REV 01	740-050037	1EDB32101HH	DC 52V Power Supply
Module				
PSM 6	REV 01	740-050037	1EDB32101HD	DC 52V Power Supply
Module				
PSM 7	REV 01	740-050037	1EDB321015F	DC 52V Power Supply
Module				
PSM 8	REV 01	740-050037	1EDB321015B	DC 52V Power Supply
Module				
PDM 0	REV 03	740-045234	1EFA3220433	DC Power Dist Module

PDM 1	REV 03	740-045234	1EFA3220425	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009115685	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009099711	RE-S-1800x4
CB 0	REV 23	750-040257	CABE8395	Control Board
CB 1	REV 12	750-040257	CAAD9499	Control Board
SPMB 0	REV 02	711-041855	ABCG8426	PMB Board
SPMB 1	REV 02	711-041855	ABBS1481	PMB Board
SFB 0	REV 06	711-044466	ABCD5013	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCD5160	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCD5175	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCD4938	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCD4944	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCD4968	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCD5267	Switch Fabric Board
SFB 7	REV 06	711-044466	ABCD4997	Switch Fabric Board
FPC 0	REV 59	750-044130	ABCT7676	MPC6E 3D
CPU	REV 10	711-045719	ABCK8527	RMPK PMB
XLM 0	REV 13	711-046638	ABCT7810	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7811	MPC6E XL
FPC 2	REV 27	750-033205	ZL6014	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9068	HMPK PMB 2G
MIC 0	REV 14	750-033196	CAAW9214	1X100GE CXP
PIC 0		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XC49FC030	CFP2-100G-SR10
MIC 1	REV 18	750-033199	CAAC3231	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 3	REV 59	750-044130	ABCT7682	MPC6E 3D
CPU	REV 10	711-045719	ABCK8531	RMPK PMB
XLM 0	REV 13	711-046638	ABCT7818	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7819	MPC6E XL
FPC 4	REV 33	750-044130	ABBY9278	MPC6E 3D
CPU	REV 09	711-045719	ABBY8677	RMPK PMB
XLM 0	REV 06.2.00	711-046638	ABBY8844	MPC6E XL
XLM 1	REV 06.2.00	711-046638	ABBY8830	MPC6E XL
FPC 5	REV 59	750-044130	ABCT7675	MPC6E 3D
CPU	REV 10	711-045719	ABCK8526	RMPK PMB
XLM 0	REV 13	711-046638	ABCT7808	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7809	MPC6E XL
FPC 6	REV 30	750-028467	ZM4986	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6541	AMPK PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43GAC	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	ALM0A6D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AQFORB3	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	153363A00333	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AN10KYE	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	APK04YM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AQFOH44	SFP+-10G-SR
FPC 8	REV 38	750-031090	CABF7313	MPC Type 2 3D EQ
CPU	REV 08	711-030884	CABE6727	MPC PMB 2G
MIC 0	REV 18	750-028380	YK8253	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 03	740-014289	AD1148M00TP	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
QXM 0	REV 06	711-028408	CABC5614	MPC QXM
QXM 1	REV 06	711-028408	CABC5550	MPC QXM
FPC 9	REV 39	750-044130	ABCK1652	MPC6E 3D
CPU	REV 09	711-045719	ABCK1655	RMPK PMB



MIC 0	REV 09	750-049457	ABCP1230	2X100GE CFP2 OTN
PIC 0		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	37300222WP0002	CFP2-100G-LR4-D
Xcvr 1		NON-JNPR	FD46F001Y	CFP2-100G-SR10
MIC 1	REV 07	750-049457	ABCV6662	2X100GE CFP2 OTN
PIC 1		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQD0014	CFP2-100G-LR4-D
Xcvr 1		NON-JNPR	J13J68335	CFP2-100G-LR4-D
XLM 0	REV 07.2.00	711-046638	ABCK5491	MPC6E XL
XLM 1	REV 07.2.00	711-046638	ABCK5475	MPC6E XL
ADC 1	REV 17	750-043596	ABCG9023	Adapter Card
ADC 2	REV 01	750-043596	ZV4079	Adapter Card
ADC 6	REV 17	750-043596	ABCG8866	Adapter Card
ADC 8	REV 17	750-043596	ABCA8993	Adapter Card
Fan Tray 0	REV 06	760-046960	ACAY0354	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0831	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0892	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0839	172mm FanTray - 6 Fans

### show chassis hardware detail (MX2010 Routers with MPC6E and OTN MIC)

```

user@host> show chassis hardware detail
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis
Midplane            REV 35   750-044636   ABAB9188      Lower Backplane
Midplane 1          REV 02   711-044557   ABAB8729      Upper Backplane
PMP                  REV 04   711-032426   ACAJ2432      Power Midplane
FPM Board            REV 09   760-044634   ABCA4314      Front Panel Display
PSM 0                REV 01   740-050037   1EDB321015C   DC 52V Power Supply
Module
PSM 1                REV 01   740-050037   1EDB321015J   DC 52V Power Supply
Module
PSM 2                REV 01   740-050037   1EDB32000K8    DC 52V Power Supply
Module
PSM 3                REV 01   740-050037   1EDB32101JW    DC 52V Power Supply
Module
PSM 4                REV 01   740-050037   1EDB321015G    DC 52V Power Supply
Module
PSM 5                REV 01   740-050037   1EDB32101HH    DC 52V Power Supply
Module
PSM 6                REV 01   740-050037   1EDB32101HD    DC 52V Power Supply
Module
PSM 7                REV 01   740-050037   1EDB321015F    DC 52V Power Supply
Module
PSM 8                REV 01   740-050037   1EDB321015B    DC 52V Power Supply
Module
PDM 0                REV 03   740-045234   1EFA3220433    DC Power Dist Module
PDM 1                REV 03   740-045234   1EFA3220425    DC Power Dist Module
Routing Engine 0     REV 02   740-041821   9009115685     RE-S-1800x4
  ad0    3998 MB   Virtium - TuffDrive VCF P1T0200274310822 191 Compact Flash
  ad1    30533 MB UGB94BPH32H0S1-KCI 11000043190      Disk 1
  usb0 (addr 1)    EHCI root hub 0      Intel          uhub0
  usb0 (addr 2)    product 0x0020 32    vendor 0x8087  uhub1
  DIMM 0          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
  DIMM 1          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
  DIMM 2          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
  DIMM 3          VL31B5263F-F8SD DIE REV-0 PCB REV-0    MFR ID-ce80
Routing Engine 1     REV 02   740-041821   9009099711     RE-S-1800x4
  ad0    3998 MB   Virtium - TuffDrive VCF P1T0200262860208 30 Compact Flash
  ad1    30533 MB   UGB94ARF32H0S3-KC   UNIGEN-499551-000146 Disk 1

```

CB 0	REV 23	750-040257	CABE8395	Control Board
CB 1	REV 12	750-040257	CAAD9499	Control Board
SPMB 0	REV 02	711-041855	ABCG8426	PMB Board
SPMB 1	REV 02	711-041855	ABBS1481	PMB Board
SFB 0	REV 06	711-044466	ABCD5013	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCD5160	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCD5175	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCD4938	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCD4944	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCD4968	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCD5267	Switch Fabric Board
SFB 7	REV 06	711-044466	ABCD4997	Switch Fabric Board
FPC 0	REV 59	750-044130	ABCT7676	MPC6E 3D
CPU	REV 10	711-045719	ABCK8527	RMPD PMB
XLM 0	REV 13	711-046638	ABCT7810	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7811	MPC6E XL
FPC 2	REV 27	750-033205	ZL6014	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9068	HMPD PMB 2G
MIC 0	REV 14	750-033196	CAAW9214	1X100GE CXP
PIC 0		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-046563	XC49FC030	CFP2-100G-SR10
MIC 1	REV 18	750-033199	CAAC3231	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
FPC 3	REV 59	750-044130	ABCT7682	MPC6E 3D
CPU	REV 10	711-045719	ABCK8531	RMPD PMB
XLM 0	REV 13	711-046638	ABCT7818	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7819	MPC6E XL
FPC 4	REV 33	750-044130	ABBY9278	MPC6E 3D
CPU	REV 09	711-045719	ABBY8677	RMPD PMB
XLM 0	REV 06.2.00	711-046638	ABBY8844	MPC6E XL
XLM 1	REV 06.2.00	711-046638	ABBY8830	MPC6E XL
FPC 5	REV 59	750-044130	ABCT7675	MPC6E 3D
CPU	REV 10	711-045719	ABCK8526	RMPD PMB
XLM 0	REV 13	711-046638	ABCT7808	MPC6E XL
XLM 1	REV 13	711-046638	ABCT7809	MPC6E XL
FPC 6	REV 30	750-028467	ZM4986	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6541	AMPD PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ43GAC	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	ALM0A6D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AQFORB3	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	153363A00333	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AN10KYE	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	APK04YM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AQFOH44	SFP+-10G-SR
FPC 8	REV 38	750-031090	CABF7313	MPC Type 2 3D EQ
CPU	REV 08	711-030884	CABE6727	MPC PMB 2G
MIC 0	REV 18	750-028380	YK8253	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 03	740-014289	AD1148M00TP	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
QXM 0	REV 06	711-028408	CABC5614	MPC QXM
QXM 1	REV 06	711-028408	CABC5550	MPC QXM
FPC 9	REV 39	750-044130	ABCK1652	MPC6E 3D
CPU	REV 09	711-045719	ABCK1655	RMPD PMB
MIC 0	REV 09	750-049457	ABCP1230	2X100GE CFP2 OTN
PIC 0		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	37300222WP0002	CFP2-100G-LR4-D

Xcvr 1		NON-JNPR	FD46F001Y	CFP2-100G-SR10
MIC 1	REV 07	750-049457	ABCV6662	2X100GE CFP2 OTN
PIC 1		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	UQD0014	CFP2-100G-LR4-D
Xcvr 1		NON-JNPR	J13J68335	CFP2-100G-LR4-D
XLM 0	REV 07.2.00	711-046638	ABCK5491	MPC6E XL
XLM 1	REV 07.2.00	711-046638	ABCK5475	MPC6E XL
ADC 1	REV 17	750-043596	ABCG9023	Adapter Card
ADC 2	REV 01	750-043596	ZV4079	Adapter Card
ADC 6	REV 17	750-043596	ABCG8866	Adapter Card
ADC 8	REV 17	750-043596	ABCA8993	Adapter Card
Fan Tray 0	REV 06	760-046960	ACAY0354	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0831	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0892	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0839	172mm FanTray - 6 Fans

### show chassis hardware extensive (MX2010 Routers with MPC6E and OTN MIC)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x02
S/N:          JN11C9AFEAFK
Assembly ID:  0x0557          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: MX2010
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 31 43 39 41 46 45 41 46 4b 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 35      750-044636  ABAB9188      Lower Backplane
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          750-044636      S/N:          ABAB9188
Assembly ID:  0x0b66          Assembly Version: 01.35
Date:         06-21-2013      Assembly Flags: 0x00
Version:      REV 35          CLEI Code:    IPMU810ARA
ID: Lower Backplane          FRU Model Number: CHAS-BP-MX2010-S
Board Information Record:
Address 0x00: ad 01 08 00 3c 8a b0 38 68 00 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 66 01 23 52 45 56 20 33 35 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 34 36 33 36 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 31 38 38 00 15 06 07
Address 0x30: dd ff ff ff ad 01 08 00 3c 8a b0 38 68 00 ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 55 38 31 30 41 52 41 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 31 30 2d 53 00
Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff f8 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1    REV 02      711-044557  ABAB8729      Upper Backplane
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          711-044557      S/N:          ABAB8729
Assembly ID:  0x0b65          Assembly Version: 01.02
Date:         03-21-2013      Assembly Flags: 0x00

```

```

Version:      REV 02
ID: Upper Backplane
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 0b 65 01 02 52 45 56 20 32 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 34 34 35 35 37 00 00
  Address 0x20: 53 2f 4e 20 41 42 41 42 38 37 32 39 00 15 03 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP          REV 04    711-032426    ACAJ2432          Power Midplane
Jedec Code:  0x7fb0          EEPROM Version:  0x01
P/N:         711-032426      S/N:         ACAJ2432
Assembly ID: 0x045d          Assembly Version: 01.04
Date:        03-28-2013      Assembly Flags: 0x00
Version:     REV 04
ID: Power Midplane
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 34 00 00
  Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
  Address 0x20: 53 2f 4e 20 41 43 41 4a 32 34 33 32 00 1c 03 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 00
  Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board    REV 09    760-044634    ABCA4314          Front Panel Display
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         760-044634      S/N:         ABCA4314
Assembly ID: 0x0b64          Assembly Version: 01.09
Date:        03-28-2013      Assembly Flags: 0x00
Version:     REV 09          CLEI Code:      IPMYA4EJRA
ID: Front Panel Display      FRU Model Number: MX2010-CRAFT-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 0b 64 01 09 52 45 56 20 39 00 00
  Address 0x10: 00 00 00 00 37 36 30 2d 30 34 34 36 33 34 00 00
  Address 0x20: 53 2f 4e 20 41 42 43 41 34 33 31 34 00 1c 03 07
  Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
  Address 0x40: ff ff ff ff 01 49 50 4d 59 41 34 45 4a 52 41 4d
  Address 0x50: 58 32 30 31 30 2d 43 52 41 46 54 2d 53 00 00 00
  Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff ff
  Address 0x70: ff ff ff 93 ff ff ff ff ff ff ff ff ff ff ff ff ff
PSM 0        REV 01    740-050037    1EDB321015C      DC 52V Power Supply
Module
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         740-050037      S/N:         1EDB321015C
Assembly ID: 0x0478          Assembly Version: 01.01
Date:        05-28-2013      Assembly Flags: 0x00
Version:     REV 01          CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
  Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 31 00 00

```

```

Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 35 43 00 00 1c 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01  740-050037  1EDB321015J  DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:          740-050037      S/N:          1EDB321015J
Assembly ID:   0x0478         Assembly Version: 01.01
Date:          05-28-2013     Assembly Flags: 0x00
Version:       REV 01         CLEI Code:     IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 35 4a 00 00 1c 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2          REV 01  740-050037  1EDB32000K8  DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:          740-050037      S/N:          1EDB32000K8
Assembly ID:   0x0478         Assembly Version: 01.01
Date:          05-23-2013     Assembly Flags: 0x00
Version:       REV 01         CLEI Code:     IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 30 30 30 4b 38 00 00 17 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3          REV 01  740-050037  1EDB32101JW  DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:          740-050037      S/N:          1EDB32101JW
Assembly ID:   0x0478         Assembly Version: 01.01
Date:          05-30-2013     Assembly Flags: 0x00
Version:       REV 01         CLEI Code:     IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 4a 57 00 00 1e 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d

```

```

Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 4          REV 01  740-050037  1EDB321015G      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:  0x02
P/N:           740-050037      S/N:             1EDB321015G
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-28-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 35 47 00 00 1c 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 5          REV 01  740-050037  1EDB32101HH      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:  0x02
P/N:           740-050037      S/N:             1EDB32101HH
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-30-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 48 48 00 00 1e 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 6          REV 01  740-050037  1EDB32101HD      DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:  0x02
P/N:           740-050037      S/N:             1EDB32101HD
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-30-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 31 30 31 48 44 00 00 1e 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7          REV 01  740-050037  1EDB321015F      DC 52V Power Supply

```

## Module

Jedec Code: 0x7fb0                      EEPROM Version: 0x02  
P/N: 740-050037                      S/N: 1EDB321015F  
Assembly ID: 0x0478                      Assembly Version: 01.01  
Date: 05-28-2013                      Assembly Flags: 0x00  
Version: REV 01                      CLEI Code: IPUPAKRKAA  
ID: DC 52V Power Supply Module      FRU Model Number: MX2000-PSM-DC-S

## Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
I2C Hex Data:  
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00  
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00  
Address 0x20: 31 45 44 42 33 32 31 30 31 35 46 00 00 1c 05 07  
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d  
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00  
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff  
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00

PSM 8                      REV 01      740-050037      1EDB321015B                      DC 52V Power Supply

## Module

Jedec Code: 0x7fb0                      EEPROM Version: 0x02  
P/N: 740-050037                      S/N: 1EDB321015B  
Assembly ID: 0x0478                      Assembly Version: 01.01  
Date: 05-28-2013                      Assembly Flags: 0x00  
Version: REV 01                      CLEI Code: IPUPAKRKAA  
ID: DC 52V Power Supply Module      FRU Model Number: MX2000-PSM-DC-S

## Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
I2C Hex Data:  
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00  
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00  
Address 0x20: 31 45 44 42 33 32 31 30 31 35 42 00 00 1c 05 07  
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d  
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00  
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff  
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00

PDM 0                      REV 03      740-045234      1EFA3220433                      DC Power Dist Module

Jedec Code: 0x7fb0                      EEPROM Version: 0x02  
P/N: 740-045234                      S/N: 1EFA3220433  
Assembly ID: 0x047b                      Assembly Version: 01.03  
Date: 05-30-2013                      Assembly Flags: 0x00  
Version: REV 03                      CLEI Code: IPUPAJSKAA  
ID: DC Power Dist Module              FRU Model Number: MX2000-PDM-DC-S

## Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
I2C Hex Data:  
Address 0x00: 7f b0 02 ff 04 7b 01 03 52 45 56 20 30 33 00 00  
Address 0x10: 00 00 00 00 37 34 30 2d 30 34 35 32 33 34 00 00  
Address 0x20: 31 45 46 41 33 32 32 30 34 33 33 00 00 1e 05 07  
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4a 53 4b 41 41 4d  
Address 0x50: 58 32 30 30 30 2d 50 44 4d 2d 44 43 2d 53 00 00  
Address 0x60: 00 00 00 00 00 00 31 30 33 ff ff ff ff ff ff ff  
Address 0x70: ff ff ff 1d 00 00 00 00 00 00 00 00 00 00 00 00

PDM 1                      REV 03      740-045234      1EFA3220425                      DC Power Dist Module

Jedec Code: 0x7fb0                      EEPROM Version: 0x02  
P/N: 740-045234                      S/N: 1EFA3220425  
Assembly ID: 0x047b                      Assembly Version: 01.03  
Date: 05-30-2013                      Assembly Flags: 0x00  
Version: REV 03                      CLEI Code: IPUPAJSKAA

```

ID: DC Power Dist Module          FRU Model Number:  MX2000-PDM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
..

```

### show chassis hardware (MX2020 Router)

```
user@host > show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane	REV 27	750-040240	ABAB9384	Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ1579	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ1524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E022240056	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240054	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 4	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224006W	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 01	740-045050	1E02224005B	DC 52V Power Supply
Module				
PSM 10	REV 01	740-045050	1E02224005A	DC 52V Power Supply
Module				
PSM 11	REV 01	740-045050	1E022240052	DC 52V Power Supply
Module				
PSM 12	REV 01	740-045050	1E022240051	DC 52V Power Supply
Module				
PSM 13	REV 01	740-045050	1E022240058	DC 52V Power Supply
Module				
PSM 14	REV 01	740-045050	1E02224004L	DC 52V Power Supply
Module				
PSM 15	REV 01	740-045050	1E02224005M	DC 52V Power Supply
Module				
PSM 16	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				
PSM 17	REV 01	740-045050	1E02224005Z	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E012150028	DC Power Dist Module
PDM 3	REV 01	740-045234	1E012150045	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089704	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009094138	RE-S-1800x4
CB 0	REV 14	750-040257	CAAF8430	Control Board
CB 1	REV 08	750-040257	CAAB3482	Control Board
SPMB 0	REV 01	711-041855	ZS2290	PMB Board
SPMB 1	REV 02	711-041855	CAAA6141	PMB Board
SFB 0	REV 03	711-044466	ABBV6789	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBX5666	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX5678	Switch Fabric Board



SFB 3	REV 05	711-044466	ABBX5687	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBX5609	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBX5675	Switch Fabric Board
SFB 6	REV 03	711-044466	ABBV6805	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBX5701	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 2	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR

PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLS		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9		SFP+-10G-SR
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0		SFP+-10G-SR
FPC 3	REV 30	750-028467	ABBN0302		MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495		AMPC PMB
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752		SFP+-10G-USR
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819		SFP+-10G-USR
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333		SFP+-10G-USR
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01784		SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP		SFP+-10G-SR
FPC 4	REV 30	750-028467	ABBN0308		MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1095		AMPC PMB
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743		SFP+-10G-USR
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057		SFP+-10G-USR
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789		SFP+-10G-USR
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250		SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847		SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787		SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803		SFP+-10G-USR
FPC 5	REV 30	750-028467	ABBN0316		MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082		AMPC PMB
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523		SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848		SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865		SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540		SFP+-10G-SR

PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 6	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 7	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 8	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB

PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR	
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR	
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR	
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR	
FPC 9	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE	
CPU	REV 10	711-029089	ABBK6556	AMPC PMB	
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR	
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR	
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR	
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR	
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR	
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR	
FPC 10	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE	
CPU	REV 10	711-029089	ABBK6542	AMPC PMB	
PIC 0			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR	
PIC 1			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR	
PIC 2			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR	
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR	
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR	
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR	
PIC 3			BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR	

Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LF2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
FPC 11	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 12	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 14	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR

Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 17	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02638	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03058	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 19	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR

Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5538	Adapter Card
ADC 11	REV 13	750-043596	ABBX5566	Adapter Card
ADC 12	REV 13	750-043596	ABBX5542	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 2A	760-046960	ACAY0030	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0039	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0033	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0062	172mm FanTray - 6 Fans

### show chassis hardware detail (MX2020 Router)

```
user@host> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11E2227AFJ	MX2020
Midplane				Lower Power Midplane
Midplane 1	REV 04	711-032386	ABAB9386	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ1821	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ1524	Lower Power Midplane
FPM Board	REV 06	760-040242	ABBT8837	Front Panel Display
PSM 0	REV 01	740-045050	1E02224006G	DC 52V Power Supply
Module				
PSM 1	REV 01	740-045050	1E022240053	DC 52V Power Supply
Module				
PSM 2	REV 01	740-045050	1E02224004K	DC 52V Power Supply
Module				
PSM 3	REV 01	740-045050	1E022240056	DC 52V Power Supply



Module				
PSM 4	REV 01	740-045050	1E022240054	DC 52V Power Supply
Module				
PSM 5	REV 01	740-045050	1E02224005H	DC 52V Power Supply
Module				
PSM 6	REV 01	740-045050	1E02224006S	DC 52V Power Supply
Module				
PSM 7	REV 01	740-045050	1E02224005M	DC 52V Power Supply
Module				
PSM 8	REV 01	740-045050	1E022240062	DC 52V Power Supply
Module				
PSM 9	REV 03	740-045050	1EDB2350095	DC 52V Power Supply
Module				
PSM 10	REV 03	740-045050	1EDB235009L	DC 52V Power Supply
Module				
PSM 11	REV 03	740-045050	1EDB2350092	DC 52V Power Supply
Module				
PSM 12	REV 03	740-045050	1EDB23500AT	DC 52V Power Supply
Module				
PSM 13	REV 03	740-045050	1EDB2350094	DC 52V Power Supply
Module				
PSM 15	REV 03	740-045050	1EDB235008X	DC 52V Power Supply
Module				
PDM 0	REV 01	740-045234	1E012150033	DC Power Dist Module
PDM 1	REV 01	740-045234	1E012150027	DC Power Dist Module
PDM 2	REV 01	740-045234	1E262250072	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800x4
ad0	3998 MB	Virtium - TuffDisk	VCF3 20110825A021D0000064	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-499551-000347	Disk 1
usb0 (addr 1)		EHCI root hub 0	Intel	uhub0
usb0 (addr 2)		product 0x0020 32	vendor 0x8087	uhub1
DIMM 0		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 1		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 2		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
DIMM 3		SGU04G72H1BD2SA-BB DIE	REV-52 PCB REV-54	MFR ID-ce80
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800x4
ad0	3831 MB	UGB30SFA4000T1	SFA4000T1 00000113	Compact Flash
ad1	30533 MB	UGB94ARF32H0S3-KC	UNIGEN-478612-001044	Disk 1
CB 0	REV 08	750-040257	CAAB3482	Control Board
CB 1	REV 04	750-040257	ZT2864	Control Board
SPMB 0	REV 02	711-041855	CAA6141	PMB Board
SPMB 1	REV 01	711-041855	ZS2275	PMB Board
SFB 0	REV 05	711-044466	ABBT2161	Switch Fabric Board
SFB 1	REV 05	711-044466	ABBT2159	Switch Fabric Board
SFB 2	REV 05	711-044466	ABBX3718	Switch Fabric Board
SFB 3	REV 05	711-044466	ABBT2152	Switch Fabric Board
SFB 4	REV 05	711-044466	ABBT2160	Switch Fabric Board
SFB 5	REV 05	711-044466	ABBT2145	Switch Fabric Board
SFB 6	REV 05	711-044466	ABBT2150	Switch Fabric Board
SFB 7	REV 05	711-044466	ABBT2163	Switch Fabric Board
FPC 0	REV 30	750-028467	ABBN0284	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0507	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00990	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04357	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01327	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04375	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02760	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02904	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E03963	SFP+-10G-USR

Xcvr 3	REV 01	740-030658	B11E00756	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04418	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01077	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01128	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01253	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01140	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01626	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01075	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01177	SFP+-10G-USR
FPC 1	REV 30	750-028467	ABBN0308	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1095	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04305	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01147	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01195	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01743	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01892	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02880	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00725	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01057	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02816	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11C04501	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02764	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00789	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01250	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00787	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E03803	SFP+-10G-USR
FPC 2	REV 30	750-028467	ABBN0316	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ1082	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00523	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01848	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01865	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00540	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00422	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00428	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00423	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01855	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K01847	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00526	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K00529	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00525	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00425	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00530	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01851	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00528	SFP+-10G-SR
FPC 3	REV 32	750-028467	ABBN6832	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6534	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MB4	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FQ6	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N1F	SFP+-10G-SR

Xcvr 3	REV 01	740-031980	AK80NLQ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80KDR	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FGJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N5G	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KD8	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LET	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80N1X	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRF	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL2	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N3D	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MRB	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LEQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LER	SFP+-10G-SR
FPC 4	REV 32	750-028467	ABBN6811	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7288	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NK8	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LJG	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LBU	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N21	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEU	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NL6	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LES	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEN	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80ME0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LMG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM1	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MG7	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KF9	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NRQ	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLE	SFP+-10G-SR
FPC 5	REV 32	750-028467	ABBN6791	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7289	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00424	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01849	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01862	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K01852	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP
Xcvr 0	REV 01	740-031980	B11K00427	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K00430	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01854	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00426	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	B11K00429	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01864	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01850	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11K00522	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E01144	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00985	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00796	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	B11K01866	SFP+-10G-SR
FPC 6	REV 30	750-028467	ABBM4592	MPC 3D 16x 10GE

CPU	REV 10	711-029089	ABBN0465	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01435	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01052	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01328	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01254	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02738	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02881	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01624	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00889	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02883	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00681	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04306	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02813	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01801	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02753	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01156	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04324	SFP+-10G-USR
FPC 7	REV 32	750-028467	ABBN6810	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7237	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03058	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02082	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01674	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02638	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03048	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02729	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02566	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02567	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02878	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02739	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01959	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02660	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02731	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02588	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02673	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02654	SFP+-10G-SR
FPC 8	REV 30	750-028467	ABBM4739	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0487	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02569	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02886	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03082	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	133363A00297	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02726	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03050	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02884	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03076	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02581	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02873	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02582	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03083	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031981	UL70BU6	SFP+-10G-LR
Xcvr 1	REV 01	740-031981	UL50QC6	SFP+-10G-LR
Xcvr 2	REV 01	740-031981	UL708N6	SFP+-10G-LR
Xcvr 3	REV 01	740-031981	UL603KK	SFP+-10G-LR
FPC 9	REV 32	750-028467	ABBN6827	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6508	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A01688	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A01724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01773	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02593	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A03061	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A03056	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02669	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03070	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02572	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02697	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02585	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03052	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02591	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02649	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02577	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02698	SFP+-10G-SR
FPC 10	REV 30	750-028467	ABBN0302	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0495	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01581	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01176	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01251	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02752	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00786	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01020	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01023	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02819	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02812	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11D04437	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01279	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01333	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00978	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E01018	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01784	SFP+-10G-USR
Xcvr 3	REV 01	740-031980	AK80NKP	SFP+-10G-SR
FPC 11	REV 32	750-028467	ABBN6790	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6515	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LZM	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCC	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCM	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE0	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021310	C10F99155	SFP+-10G-LRM
Xcvr 1	REV 01	740-021310	C10F99049	SFP+-10G-LRM
Xcvr 2	REV 01	740-021310	C10F99128	SFP+-10G-LRM
Xcvr 3	REV 01	740-021310	C10F99169	SFP+-10G-LRM
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	AK80LF3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02597	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A03060	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03057	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEX	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80FEU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FNM	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AJQQQ5G	SFP+-10G-SR
FPC 12	REV 30	750-028467	ZM5111	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZP6607	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LJA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MFZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKL	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KF4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FBJ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MM2	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LJV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NXV	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1H	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLS	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FL5	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL9	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG2	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80KDU	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MG1	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80MM0	SFP+-10G-SR
FPC 13	REV 30	750-028467	ABBN0208	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABB11084	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04745	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01570	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E04388	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01439	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04739	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01869	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01675	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01901	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01346	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11F01288	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01824	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E04312	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E02811	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03847	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01495	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11F01265	SFP+-10G-USR
FPC 14	REV 23	750-028467	YN2977	MPC 3D 16x 10GE
CPU	REV 10	711-029089	YP1856	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00875	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00851	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00772	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00882	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	183363A00735	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00169	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00726	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00077	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00168	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00676	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00732	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00091	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	183363A00725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00642	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	183363A00871	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	183363A00853	SFP+-10G-SR
FPC 15	REV 32	750-028467	ABBN6798	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6556	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	9ZDZ06A00055	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	183363A00239	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AD0915E003K	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AD0915E003A	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80MRC	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NL5	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKN	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N3U	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N1T	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJ808DJ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NG4	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FND	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80FKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLT	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NKR	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LKM	SFP+-10G-SR
FPC 16	REV 30	750-028467	ABBN0270	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBJ0966	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NL1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NXW	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KD2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80FMD	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NKQ	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MGH	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80N38	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NL7	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80M5J	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NKD	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80KCY	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LHK	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LEL	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MBE	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80NLG	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LFH	SFP+-10G-SR
FPC 17	REV 32	750-028467	ABBN6796	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN7259	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+

Xcvr 0	REV 01	740-031980	B11K01856	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11K01853	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11K01863	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02863	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02668	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02881	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A01671	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02627	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02725	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02692	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02730	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A03081	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	163363A02736	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	163363A02568	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	163363A02747	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	163363A02579	SFP+-10G-SR
FPC 18	REV 30	750-028467	ABBN0281	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBN0526	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11F01326	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E03973	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E00950	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E00674	SFP+-10G-USR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E00775	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E04461	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E01074	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E02821	SFP+-10G-USR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04501	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E00757	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11F01623	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01022	SFP+-10G-USR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-030658	B11E04359	SFP+-10G-USR
Xcvr 1	REV 01	740-030658	B11E02751	SFP+-10G-USR
Xcvr 2	REV 01	740-030658	B11E02736	SFP+-10G-USR
Xcvr 3	REV 01	740-030658	B11E01178	SFP+-10G-USR
FPC 19	REV 32	750-028467	ABBN6813	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ABBK6542	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NA3	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80NLF	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80MRH	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80KE4	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	973152A00030	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80L9H	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80ME8	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80NLR	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80NG1	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80MCA	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80LFC	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80LEM	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80N9X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AK80LAC	SFP+-10G-SR



Xcvr 2	REV 01	740-031980	AK80LF2	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AK80N8T	SFP+-10G-SR
ADC 0	REV 13	750-043596	ABBX5561	Adapter Card
ADC 1	REV 13	750-043596	ABBX5546	Adapter Card
ADC 2	REV 13	750-043596	ABBX5535	Adapter Card
ADC 3	REV 13	750-043596	ABBX5552	Adapter Card
ADC 4	REV 13	750-043596	ABBX5581	Adapter Card
ADC 5	REV 13	750-043596	ABBX5545	Adapter Card
ADC 6	REV 13	750-043596	ABBX5554	Adapter Card
ADC 7	REV 07	750-043596	ABBV7194	Adapter Card
ADC 8	REV 07	750-043596	ABBV7251	Adapter Card
ADC 9	REV 07	750-043596	ABBV7202	Adapter Card
ADC 10	REV 13	750-043596	ABBX5579	Adapter Card
ADC 11	REV 13	750-043596	ABBX5548	Adapter Card
ADC 12	REV 13	750-043596	ABBX5575	Adapter Card
ADC 13	REV 13	750-043596	ABBX5539	Adapter Card
ADC 14	REV 13	750-043596	ABBX5555	Adapter Card
ADC 15	REV 13	750-043596	ABBX5557	Adapter Card
ADC 16	REV 13	750-043596	ABBX5536	Adapter Card
ADC 17	REV 13	750-043596	ABBX5559	Adapter Card
ADC 18	REV 13	750-043596	ABBX5537	Adapter Card
ADC 19	REV 11	750-043596	ABBW5685	Adapter Card
Fan Tray 0	REV 04	760-046960	ACAY0090	172mm FanTray - 6 Fans
Fan Tray 1	REV 04	760-046960	ACAY0088	172mm FanTray - 6 Fans
Fan Tray 2	REV 04	760-046960	ACAY0089	172mm FanTray - 6 Fans
Fan Tray 3	REV 04	760-046960	ACAY0108	172mm FanTray - 6 Fans

#### show chassis hardware models (MX2020 Router)

```
user@host > show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 27	750-040240	ABAB9384	750-040240
FPM Board	REV 06	760-040242	ABBT8837	760-040242
PSM 0	REV 01	740-045050	1E02224006G	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	1E022240053	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	1E02224004K	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	1E022240056	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	1E022240054	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	1E02224005H	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	1E02224006S	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	1E02224005M	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	1E022240062	MX2000-PSM-HC-DC-S-A
PSM 9	REV 03	740-045050	1EDB2350095	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	1EDB235009L	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	1EDB2350092	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	1EDB23500AT	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	1EDB2350094	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	1EDB235008X	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234	1E012150033	
PDM 1	REV 01	740-045234	1E012150027	
PDM 2	REV 01	740-045234	1E262250072	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821	9009094138	RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821	9009089709	RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	CAAB3482	750-040257
CB 1	REV 04	750-040257	ZT2864	750-040257
SFB 0	REV 05	711-044466	ABBT2161	MX2000-SFB-S
SFB 1	REV 05	711-044466	ABBT2159	MX2000-SFB-S
SFB 2	REV 05	711-044466	ABBX3718	MX2000-SFB-S
SFB 4	REV 05	711-044466	ABBT2160	MX2000-SFB-S
SFB 5	REV 05	711-044466	ABBT2145	MX2000-SFB-S

SFB 7	REV 05	711-044466	ABBT2163	MX2000-SFB-S
FPC 0	REV 30	750-028467	ABBN0284	MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467	ABBN0308	MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467	ABBN0316	MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467	ABBN6832	MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467	ABBN6811	MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467	ABBN6791	MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467	ABBM4592	MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467	ABBN6810	MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467	ABBM4739	MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467	ABBN6827	MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467	ABBN0302	MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467	ABBN6790	MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467	ZM5111	MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467	ABBN0208	MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467	YN2977	MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467	ABBN6798	MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467	ABBN0270	MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467	ABBN6796	MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467	ABBN0281	MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467	ABBN6813	MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	ABBX5561	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	ABBX5546	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	ABBX5535	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	ABBX5552	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	ABBX5581	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	ABBX5545	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	ABBX5554	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	ABBV7194	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	ABBV7251	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	ABBV7202	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	ABBX5579	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	ABBX5575	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	ABBX5539	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	ABBX5555	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	ABBX5557	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	ABBX5536	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	ABBX5559	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	ABBX5537	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	ABBW5685	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960	ACAY0090	
Fan Tray 1	REV 04	760-046960	ACAY0088	
Fan Tray 2	REV 04	760-046960	ACAY0089	
Fan Tray 3	REV 04	760-046960	ACAY0108	

**show chassis hardware clei-models (MX2020 Router)**

user@ host &gt; show chassis hardware clei-models

Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 27	750-040240	PROTOXCLEI	750-040240
FPM Board	REV 06	760-040242	PROTOXCLEI	760-040242
PSM 0	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 1	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 2	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 3	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 4	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 5	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 6	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 7	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A
PSM 8	REV 01	740-045050	IPUPAJMKAA	MX2000-PSM-HC-DC-S-A

PSM 9	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 10	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 11	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 12	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 13	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PSM 15	REV 03	740-045050	IPUPAJMKAA	MX2000-PSM-DC-S-A
PDM 0	REV 01	740-045234		
PDM 1	REV 01	740-045234		
PDM 2	REV 01	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S-A
Routing Engine 0	REV 02	740-041821		RE-S-1800X4-16G-S
Routing Engine 1	REV 02	740-041821		RE-S-1800X4-16G-S
CB 0	REV 08	750-040257	PROTOXCLEI	750-040257
CB 1	REV 04	750-040257	PROTOXCLEI	750-040257
SFB 0	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 1	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 2	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 4	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 5	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 7	REV 05	711-044466	IPUCBA6CAA	MX2000-SFB-S
FPC 0	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 1	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 2	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 3	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 4	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 5	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 6	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 7	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 8	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 9	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 10	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 11	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 12	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 13	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 14	REV 23	750-028467		MPC-3D-16XGE-SFPP
FPC 15	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 16	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 17	REV 32	750-028467		MPC-3D-16XGE-SFPP
FPC 18	REV 30	750-028467		MPC-3D-16XGE-SFPP
FPC 19	REV 32	750-028467		MPC-3D-16XGE-SFPP
ADC 0	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 1	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 2	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 3	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 4	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 5	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 6	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 7	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 8	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 9	REV 07	750-043596	PROTOXCLEI	MX2000-LC-ADAPTER
ADC 10	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 12	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 13	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 14	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 15	REV 13	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 16	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 17	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 18	REV 13	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
ADC 19	REV 11	750-043596	PROTOXCLEI	PROTO-ASSEMBLY
Fan Tray 0	REV 04	760-046960		
Fan Tray 1	REV 04	760-046960		

```

Fan Tray 2      REV 04    760-046960
Fan Tray 3      REV 04    760-046960

```

### show chassis hardware (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN120BADBAFJ	MX2020
Midplane	REV 51	750-040240	ABAB9243	Lower Backplane
Midplane 1	REV 04	711-032386	ABAB9399	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ2541	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ2194	Lower Power Midplane
FPM Board	REV 13	760-040242	ABCA8835	Front Panel Display
PSM 0	REV 01	740-050037	1EDB32403L5	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB32403L3	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB32403KM	DC 52V Power Supply
Module				
PSM 3	REV 01	740-050037	1EDB3130079	DC 52V Power Supply
Module				
PSM 4	REV 01	740-050037	1EDB3130077	DC 52V Power Supply
Module				
PSM 5	REV 01	740-050037	1EDB3130020	DC 52V Power Supply
Module				
PSM 6	REV 01	740-050037	1EDB313009S	DC 52V Power Supply
Module				
PSM 7	REV 01	740-050037	1EDB313008E	DC 52V Power Supply
Module				
PSM 8	REV 01	740-050037	1EDB3130063	DC 52V Power Supply
Module				
PSM 12	REV 01	740-050037	1EDB3130026	DC 52V Power Supply
Module				
PSM 13	REV 01	740-050037	1EDB3130074	DC 52V Power Supply
Module				
PSM 14	REV 01	740-050037	1EDB313009D	DC 52V Power Supply
Module				
PSM 15	REV 01	740-050037	1EDB3130024	DC 52V Power Supply
Module				
PSM 16	REV 01	740-050037	1EDB3130054	DC 52V Power Supply
Module				
PSM 17	REV 01	740-050037	1EDB3130080	DC 52V Power Supply
Module				
PDM 0	REV 03	740-045234	1EGA3170144	DC Power Dist Module
PDM 1	REV 03	740-045234	1EGA3170158	DC Power Dist Module
PDM 2	REV 03	740-045234	1EGA3170182	DC Power Dist Module
PDM 3	REV 03	740-045234	1EGA3170207	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009112112	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009112087	RE-S-1800x4
CB 0	REV 23	750-040257	CABA2295	Control Board
CB 1	REV 23	750-040257	CABE8379	Control Board
SPMB 0	REV 02	711-041855	ABCE8851	PMB Board
SPMB 1	REV 02	711-041855	ABCE8839	PMB Board
SFB 0	REV 06	711-044466	ABCD5001	Switch Fabric Board
SFB 1	REV 06	711-044466	ABCD5034	Switch Fabric Board
SFB 2	REV 06	711-044466	ABCH3899	Switch Fabric Board
SFB 3	REV 06	711-044466	ABCD5020	Switch Fabric Board
SFB 4	REV 06	711-044466	ABCD4975	Switch Fabric Board
SFB 5	REV 06	711-044466	ABCH3881	Switch Fabric Board
SFB 6	REV 06	711-044466	ABCD5026	Switch Fabric Board

SFB 7	REV 06	711-044466	ABCD5032	Switch Fabric Board
FPC 0	REV 39	750-045715	CACD1902	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACB1933	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	B11F00361	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	19T511101854	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	19T511100377	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	ANT0878	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	19T511100398	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ4363J	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	19T511101377	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	ANT072M	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AG90C7N	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AM30M09	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10E01016	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-031980	B10L04151	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	19T511101379	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5036J	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AG90C4M	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	19T511101104	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ502ZM	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AN10KY2	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ43G41	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ41F04	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AMS16N3	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AMH04Y3	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	ANA093E	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
WAN MEZZ	REV 09	750-049136	CABN0410	MPC5E 24XGE OTN Mezz
FPC 1	REV 11	750-045372	CABK8112	MPCE Type 3 3D
CPU	REV 08	711-035209	CABJ6621	HMPD PMB 2G
MIC 0	REV 07	750-033307	CAAZ2897	10X10GE SFPP
PIC 0		BUILTIN	BUILTIN	10X10GE SFPP
Xcvr 0	REV 01	740-021308	AQ501VK	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501YC	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HJF	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43H8D	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	19T511100370	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	153363A00763	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	APH2LXB	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AMCOLVV	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11F00230	SFP+-10G-SR
MIC 1	REV 14	750-033196	CAAP1390	1X100GE CXP
PIC 2		BUILTIN	BUILTIN	1X100GE CXP
Xcvr 0	REV 01	740-032166	XB11F000M	CFP2-100G-SR10
FPC 2	REV 17	750-037355	CAAS5826	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAR3986	HMPD PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	T09F43722	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	ALPOKXF	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502FG	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502T7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00571	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-031980	AJ71KEH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11E01355	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11F00249	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP

FPC 3	REV 05	750-044444	CAAY9920	MPCE Type 2 3D P
CPU	REV 04	711-038484	CAAW3639	MPCE PMB 2G
MIC 0	REV 28	750-028387	CAAX1083	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	CC07BK05B	XFP-10G-SR
Xcvr 1	REV 01	740-011571	C728XJ00U	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T12L92339	XFP-10G-SR
QXM 0	REV 06	711-028408	CAAW4915	MPC QXM
QXM 1	REV 06	711-028408	CAAW4894	MPC QXM
FPC 4	REV 18	750-046005	CACH5661	MPC5E 3D Q 2CGE+4XGE
CPU	REV 09	711-045719	CACF2880	RMPC PMB
PIC 0		BUILTIN	BUILTIN	2X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-046563	XD16FC03Y	CFP2-100G-SR10
PIC 2		BUILTIN	BUILTIN	2X10GE SFPP OTN
PIC 3		BUILTIN	BUILTIN	1X100GE CFP2 OTN
Xcvr 0	REV 01	740-049775	J13K72997	CFP2-100G-LR4-D
FPC 5	REV 35	750-028467	CAAR2623	MPC 3D 16x 10GE
CPU	REV 11	711-029089	CAAR0491	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ5027T	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ502J0	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ5027S	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ501Y7	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ501YB	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ503EB	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ43HJH	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ43J0Y	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ50352	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ501X6	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQ502NV	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502ZJ	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AQ502H4	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQ43HJK	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJ30CU7	SFP+-10G-SR
FPC 9	REV 30	750-044130	ABCF5773	MPC6E 3D
CPU	REV 09	711-045719	ABCF1270	RMPC PMB
MIC 0	REV 05	750-049457	ABCD7829	2X100GE CFP2 OTN
PIC 0		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0		NON-JNPR	FE13F000K	CFP2-100G-SR10
Xcvr 1	REV 01	740-048813	XD32FE017	CFP2-100G-LR-D
MIC 1	REV 07	750-049457	ABCK2812	2X100GE CFP2 OTN
PIC 1		BUILTIN	BUILTIN	2X100GE CFP2 OTN
Xcvr 0	REV 01	740-048813	XD32FE018	CFP2-100G-SR10
Xcvr 1		NON-JNPR	FE13F000E	CFP2-100G-LR4-D
XLM 0	REV 05.2.00	711-046638	ABCF5915	MPC6E XL
XLM 1	REV 05.2.00	711-046638	ABCF5916	MPC6E XL
FPC 10	REV 36	750-044130	ABCS8602	MPC6E 3D
CPU	REV 09	711-045719	ABCS8779	RMPC PMB
MIC 0	REV 06	750-049979	ABCK2656	24X10GE SFPP OTN
PIC 0		BUILTIN	BUILTIN	24X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQ43J08	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQE1Y2E	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQE1UW4	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQE1MQF	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQGOMN1	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQE1L9M	SFP+-10G-SR

Xcvr 6	REV 01	740-021308	AQGOMPD	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQE1Y2B	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQGOLT5	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQD2ET4	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQGOMPC	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQGOM63	SFP+-10G-SR
Xcvr 12	REV 01	740-021308	AQGOLT1	SFP+-10G-SR
Xcvr 13	REV 01	740-021308	AQGOM4L	SFP+-10G-SR
Xcvr 14	REV 01	740-021308	AQGOLS7	SFP+-10G-SR
Xcvr 15	REV 01	740-021308	AQE1MQB	SFP+-10G-SR
Xcvr 16	REV 01	740-021308	AQGOLZP	SFP+-10G-SR
Xcvr 17	REV 01	740-021308	AQE1LU9	SFP+-10G-SR
Xcvr 18	REV 01	740-021308	AQGOMRZ	SFP+-10G-SR
Xcvr 19	REV 01	740-021308	AQE1MQ9	SFP+-10G-SR
Xcvr 20	REV 01	740-021308	AQGOLRX	SFP+-10G-SR
Xcvr 21	REV 01	740-021308	AQE1UWD	SFP+-10G-SR
Xcvr 22	REV 01	740-021308	AQGOLT4	SFP+-10G-SR
Xcvr 23	REV 01	740-021308	AQE1MQL	SFP+-10G-SR
MIC 1	REV 12	750-050008	ABCK5372	4X100GE CXP
PIC 1		BUILTIN	BUILTIN	4X100GE CXP
Xcvr 3	REV 01	740-046563	XD16FC02Z	CFP2-100G-SR10
XLM 0	REV 07.2.00	711-046638	ABCK3481	MPC6E XL
XLM 1	REV 07.2.00	711-046638	ABCK4725	MPC6E XL
FPC 17	REV 28	750-044130	ABBZ3873	MPC6E 3D
CPU	REV 08	711-045719	ABBZ3770	RMPD PMB
MIC 0	REV 11	750-046535	ABCC7731	24X10GE SFPP
PIC 0		BUILTIN	BUILTIN	24X10GE SFPP
Xcvr 1	REV 01	740-021308	APK0543	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B10G01119	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQ502SX	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQ43H84	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQ501TB	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQ502JZ	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQ502SC	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQ502JW	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQ502RM	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AHK013B	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQGOMRT	SFP+-10G-SR
Xcvr 13	REV 01	740-031980	AMC0JTC	SFP+-10G-SR
Xcvr 14	REV 01	740-021308	ANAMQ0	SFP+-10G-SR
Xcvr 15	REV 01	740-021308	AQ502GS	SFP+-10G-SR
Xcvr 16	REV 01	740-021308	AQGOM0J	SFP+-10G-SR
Xcvr 17	REV 01	740-021308	AQGOMUR	SFP+-10G-SR
Xcvr 18	REV 01	740-021308	AQGOMRR	SFP+-10G-SR
Xcvr 19	REV 01	740-021308	AQGOM0F	SFP+-10G-SR
Xcvr 20	REV 01	740-021308	AQ50312	SFP+-10G-SR
Xcvr 21	REV 01	740-021308	AQ5032U	SFP+-10G-SR
Xcvr 22	REV 01	740-021308	APE17B5	SFP+-10G-SR
Xcvr 23	REV 01	740-021309	91D104A00011	SFP+-10G-LR
MIC 1	REV 03	750-050008	ABCC4522	4X100GE CXP
PIC 1		BUILTIN	BUILTIN	4X100GE CXP
Xcvr 0	REV 01	740-046563	XD16FC02U	CFP2-100G-SR10
Xcvr 1	REV 01	740-046563	XC42FC03K	CFP2-100G-SR10
Xcvr 2	REV 01	740-046563	XC42FC01Z	CFP2-100G-SR10
Xcvr 3	REV 01	740-046563	XC42FC02U	CFP2-100G-SR10
XLM 0	REV 04.2.00	711-046638	ABBZ3779	MPC6E XL
XLM 1	REV 04.2.00	711-046638	ABBZ3780	MPC6E XL
FPC 18	REV 39	750-045715	CACD1910	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1817	RMPD PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN

PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130194	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130193	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130196	QSFP+-40G-SR4
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130191	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130198	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130192	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0411	MPC5E 24XGE OTN Mezz
FPC 19	REV 39	750-045715	CACD1908	MPC5E 3D Q 24XGE+6XLGE
CPU	REV 09	711-045719	CACD1820	RMPC PMB
PIC 0		BUILTIN	BUILTIN	12X10GE SFPP OTN
Xcvr 0	REV 01	740-021308	AQA0EXJ	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AQGOM6D	SFP+-10G-SR
Xcvr 2	REV 01	740-021308	AQGOLW7	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	AQA0JKB	SFP+-10G-SR
Xcvr 4	REV 01	740-021308	AQGOMTM	SFP+-10G-SR
Xcvr 5	REV 01	740-021308	AQA07NE	SFP+-10G-SR
Xcvr 6	REV 01	740-021308	AQGOM41	SFP+-10G-SR
Xcvr 7	REV 01	740-021308	AQGOMU7	SFP+-10G-SR
Xcvr 8	REV 01	740-021308	AQGOMUG	SFP+-10G-SR
Xcvr 9	REV 01	740-021308	AQGOMMX	SFP+-10G-SR
Xcvr 10	REV 01	740-021308	AQGOM5K	SFP+-10G-SR
Xcvr 11	REV 01	740-021308	AQGOLVZ	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	12X10GE SFPP OTN
PIC 2		BUILTIN	BUILTIN	3X40GE QSFPP
PIC 3		BUILTIN	BUILTIN	3X40GE QSFPP
Xcvr 0	REV 01	740-046565	QD130242	QSFP+-40G-SR4
Xcvr 1	REV 01	740-046565	QD130245	QSFP+-40G-SR4
Xcvr 2	REV 01	740-046565	QD130613	QSFP+-40G-SR4
WAN MEZZ	REV 09	750-049136	CABN0418	MPC5E 24XGE OTN Mezz
ADC 0	REV 17	750-043596	ABCD5378	Adapter Card
ADC 1	REV 17	750-043596	ABCD5465	Adapter Card
ADC 2	REV 17	750-043596	ABCD5431	Adapter Card
ADC 3	REV 17	750-043596	ABCD5356	Adapter Card
ADC 4	REV 02	750-043596	ZW1545	Adapter Card
ADC 5	REV 17	750-043596	ABCD5517	Adapter Card
ADC 18	REV 17	750-043596	ABCD5535	Adapter Card
ADC 19	REV 01	750-043596	ZV4127	Adapter Card
Fan Tray 0	REV 06	760-046960	ACAY0791	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0788	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0755	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0441	172mm FanTray - 6 Fans

## show chassis hardware detail (MX2020 Router with MPC5EQ and MPC6E)

user@host&gt;show chassis hardware detail

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN120BADBAFJ	MX2020
Midplane	REV 51	750-040240	ABAB9243	Lower Backplane
Midplane 1	REV 04	711-032386	ABAB9399	Upper Backplane
PMP 1	REV 05	711-032428	ACAJ2541	Upper Power Midplane
PMP 0	REV 04	711-032426	ACAJ2194	Lower Power Midplane
FPM Board	REV 13	760-040242	ABCA8835	Front Panel Display
PSM 0	REV 01	740-050037	1EDB32403L5	DC 52V Power Supply
Module				
PSM 1	REV 01	740-050037	1EDB32403L3	DC 52V Power Supply
Module				
PSM 2	REV 01	740-050037	1EDB32403KM	DC 52V Power Supply
Module				



PSM 3 Module	REV 01	740-050037	1EDB3130079	DC 52V Power Supply
PSM 4 Module	REV 01	740-050037	1EDB3130077	DC 52V Power Supply
PSM 5 Module	REV 01	740-050037	1EDB3130020	DC 52V Power Supply
PSM 6 Module	REV 01	740-050037	1EDB313009S	DC 52V Power Supply
PSM 7 Module	REV 01	740-050037	1EDB313008E	DC 52V Power Supply
PSM 8 Module	REV 01	740-050037	1EDB3130063	DC 52V Power Supply
PSM 12 Module	REV 01	740-050037	1EDB3130026	DC 52V Power Supply
PSM 13 Module	REV 01	740-050037	1EDB3130074	DC 52V Power Supply
PSM 14 Module	REV 01	740-050037	1EDB313009D	DC 52V Power Supply
PSM 15 Module	REV 01	740-050037	1EDB3130024	DC 52V Power Supply
PSM 16 Module	REV 01	740-050037	1EDB3130054	DC 52V Power Supply
PSM 17 Module	REV 01	740-050037	1EDB3130080	DC 52V Power Supply
PDM 0	REV 03	740-045234	1EGA3170144	DC Power Dist Module
PDM 1	REV 03	740-045234	1EGA3170158	DC Power Dist Module
PDM 2	REV 03	740-045234	1EGA3170182	DC Power Dist Module
PDM 3	REV 03	740-045234	1EGA3170207	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009112112	RE-S-1800x4
ad0 3998 MB	Virtium - TuffDrive		VCF P1T0200274310822	113 Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000031656	Disk 1
usb0 (addr 1)	EHCI root hub 0		Intel	uhub0
usb0 (addr 2)	product 0x0020 32		vendor 0x8087	uhub1
DIMM 0	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
DIMM 1	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
DIMM 2	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
DIMM 3	SGU04G72H1BD2SA-BB DIE REV-52 PCB REV-54 MFR ID-ce80			
Routing Engine 1	REV 02	740-041821	9009112087	RE-S-1800x4
ad0 3998 MB	Virtium - TuffDrive		VCF P1T0200274310822	366 Compact Flash
ad1 30533 MB	UGB94BPH32H0S1-KCI		11000039979	Disk 1
CB 0	REV 23	750-040257	CABA2295	Control Board
CB 1	REV 23	750-040257	CABE8379	Control Board
SPMB 0				
SPMB 1				
FPC 0 CPU	REV 39	750-045715	CACD1902	MPC5E 3D Q 24XGE+6XLGE
FPC 1 CPU	REV 11	750-045372	CABK8112	MPCE Type 3 3D
FPC 2 CPU	REV 17	750-037355	CAAS5826	MPC4E 3D 2CGE+8XGE
FPC 3 CPU	REV 05	750-044444	CAAY9920	MPCE Type 2 3D P
FPC 4 CPU	REV 18	750-046005	CACH5661	MPC5E 3D Q 2CGE+4XGE
FPC 5 CPU	REV 35	750-028467	CAAR2623	MPC 3D 16x 10GE
FPC 9 CPU	REV 30	750-044130	ABCF5773	MPC6E 3D
FPC 10 CPU	REV 36	750-044130	ABCS8602	MPC6E 3D
FPC 17	REV 28	750-044130	ABBZ3873	MPC6E 3D

CPU				
FPC 18	REV 39	750-045715	CACD1910	MPC5E 3D Q 24XGE+6XLGE
CPU				
FPC 19	REV 39	750-045715	CACD1908	MPC5E 3D Q 24XGE+6XLGE
CPU				
Fan Tray 0	REV 06	760-046960	ACAY0791	172mm FanTray - 6 Fans
Fan Tray 1	REV 06	760-046960	ACAY0788	172mm FanTray - 6 Fans
Fan Tray 2	REV 06	760-046960	ACAY0755	172mm FanTray - 6 Fans
Fan Tray 3	REV 06	760-046960	ACAY0441	172mm FanTray - 6 Fans

### show chassis hardware extensive (MX2020 Router with MPC5EQ and MPC6E)

```

Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Jedec Code:   0x7fb0          EEPROM Version: 0x02
S/N:          JN120BADBAFJ
Assembly ID:  0x0557          Assembly Version: 00.00
Date:         00-00-0000      Assembly Flags:  0x00
ID: MX2020
Board Information Record:
Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
Address 0x00: 7f b0 02 ff 05 57 00 00 00 00 00 00 00 00 00 00
Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x20: 4a 4e 31 32 30 42 41 44 42 41 46 4a 00 00 00 00
Address 0x30: 00 00 00 ff 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 51    750-040240  ABAB9243      Lower Backplane
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          750-040240      S/N:          ABAB9243
Assembly ID:  0x0b22          Assembly Version: 01.51
Date:         05-30-2013      Assembly Flags: 0x00
Version:      REV 51          CLEI Code:    IPMU710ARA
ID: Lower Backplane          FRU Model Number: CHAS-BP-MX2020-S
Board Information Record:
Address 0x00: ad 01 10 00 4c 96 14 72 30 08 ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 22 01 33 52 45 56 20 35 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 34 30 32 34 30 00 00
Address 0x20: 53 2f 4e 20 41 42 41 42 39 32 34 33 00 1e 05 07
Address 0x30: dd ff ff ff ad 01 10 00 4c 96 14 72 30 08 ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 55 37 31 30 41 52 41 43
Address 0x50: 48 41 53 2d 42 50 2d 4d 58 32 30 32 30 2d 53 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff d3 ff ff ff ff ff ff ff ff ff ff ff ff
Midplane 1    REV 04    711-032386  ABAB9399      Upper Backplane
Jedec Code:   0x7fb0          EEPROM Version: 0x01
P/N:          711-032386      S/N:          ABAB9399
Assembly ID:  0x0b23          Assembly Version: 01.04
Date:         10-22-2012      Assembly Flags: 0x00
Version:      REV 04
ID: Upper Backplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 fe 0b 23 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 33 38 36 00 00

```

```

Address 0x20: 53 2f 4e 20 41 42 41 42 39 33 39 39 00 16 0a 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP 1          REV 05    711-032428    ACAJ2541          Upper Power Midplane
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-032428      S/N:           ACAJ2541
Assembly ID:   0x045c          Assembly Version: 01.05
Date:          04-26-2013      Assembly Flags: 0x00
Version:       REV 05
ID: Upper Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5c 01 05 52 45 56 20 30 35 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 38 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 32 35 34 31 00 1a 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
PMP 0          REV 04    711-032426    ACAJ2194          Lower Power Midplane
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           711-032426      S/N:           ACAJ2194
Assembly ID:   0x045d          Assembly Version: 01.04
Date:          01-29-2013      Assembly Flags: 0x00
Version:       REV 04
ID: Lower Power Midplane
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 04 5d 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 32 34 32 36 00 00
Address 0x20: 53 2f 4e 20 41 43 41 4a 32 31 39 34 00 1d 01 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM Board      REV 13    760-040242    ABCA8835          Front Panel Display
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           760-040242      S/N:           ABCA8835
Assembly ID:   0x0b24          Assembly Version: 01.13
Date:          04-13-2013      Assembly Flags: 0x00
Version:       REV 13          CLEI Code:       IPMYAE5JRA
ID: Front Panel Display        FRU Model Number: MX2020-CRAFT-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 24 01 0d 52 45 56 20 31 33 00 00
Address 0x10: 00 00 00 00 37 36 30 2d 30 34 30 32 34 32 00 00
Address 0x20: 53 2f 4e 20 41 42 43 41 38 38 33 35 00 0d 04 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 4d 59 41 45 35 4a 52 41 4d
Address 0x50: 58 32 30 32 30 2d 43 52 41 46 54 2d 53 00 00 00
Address 0x60: 00 00 00 00 00 00 41 00 00 ff ff ff ff ff ff
Address 0x70: ff ff ff 95 ff ff ff ff ff ff ff ff ff ff ff
PSM 0          REV 01    740-050037    1EDB32403L5      DC 52V Power Supply

```

```

Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB32403L5
Assembly ID: 0x0478        Assembly Version: 01.01
Date: 06-21-2013          Assembly Flags: 0x00
Version: REV 01            CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 34 30 33 4c 35 00 00 15 06 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 1          REV 01 740-050037 1EDB32403L3 DC 52V Power Supply
Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB32403L3
Assembly ID: 0x0478        Assembly Version: 01.01
Date: 06-21-2013          Assembly Flags: 0x00
Version: REV 01            CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 34 30 33 4c 33 00 00 15 06 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 2          REV 01 740-050037 1EDB32403KM DC 52V Power Supply
Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB32403KM
Assembly ID: 0x0478        Assembly Version: 01.01
Date: 06-21-2013          Assembly Flags: 0x00
Version: REV 01            CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 32 34 30 33 4b 4d 00 00 15 06 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 3          REV 01 740-050037 1EDB3130079 DC 52V Power Supply
Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB3130079
Assembly ID: 0x0478        Assembly Version: 01.01

```

```

Date:          05-16-2013      Assembly Flags:   0x00
Version:       REV 01          CLEI Code:        IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 37 39 00 00 10 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 4          REV 01 740-050037 1EDB3130077 DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-050037      S/N:           1EDB3130077
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-17-2013      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 37 37 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 5          REV 01 740-050037 1EDB3130020 DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-050037      S/N:           1EDB3130020
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-16-2013      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 32 30 00 00 10 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 6          REV 01 740-050037 1EDB313009S DC 52V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version: 0x02
P/N:           740-050037      S/N:           1EDB313009S
Assembly ID:   0x0478          Assembly Version: 01.01
Date:          05-17-2013      Assembly Flags: 0x00
Version:       REV 01          CLEI Code:     IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:

```

```

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 39 53 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 7          REV 01  740-050037  1EDB313008E      DC 52V Power Supply
Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB313008E
Assembly ID: 0x0478        Assembly Version: 01.01
Date: 05-17-2013          Assembly Flags: 0x00
Version: REV 01           CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 38 45 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 8          REV 01  740-050037  1EDB3130063      DC 52V Power Supply
Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB3130063
Assembly ID: 0x0478        Assembly Version: 01.01
Date: 05-17-2013          Assembly Flags: 0x00
Version: REV 01           CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 36 33 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 12         REV 01  740-050037  1EDB3130026      DC 52V Power Supply
Module
Jedec Code: 0x7fb0          EEPROM Version: 0x02
P/N: 740-050037          S/N: 1EDB3130026
Assembly ID: 0x0478        Assembly Version: 01.01
Date: 05-16-2013          Assembly Flags: 0x00
Version: REV 01           CLEI Code: IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00

```

```

Address 0x20: 31 45 44 42 33 31 33 30 30 32 36 00 00 10 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00
PSM 13          REV 01   740-050037   1EDB3130074       DC 52V Power Supply
Module
Jedec Code:    0x7fb0           EEPROM Version: 0x02
P/N:           740-050037       S/N:            1EDB3130074
Assembly ID:   0x0478           Assembly Version: 01.01
Date:          05-17-2013       Assembly Flags:  0x00
Version:       REV 01           CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 37 34 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 14          REV 01   740-050037   1EDB313009D       DC 52V Power Supply
Module
Jedec Code:    0x7fb0           EEPROM Version: 0x02
P/N:           740-050037       S/N:            1EDB313009D
Assembly ID:   0x0478           Assembly Version: 01.01
Date:          05-17-2013       Assembly Flags:  0x00
Version:       REV 01           CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 35 30 30 33 37 00 00
Address 0x20: 31 45 44 42 33 31 33 30 30 39 44 00 00 11 05 07
Address 0x30: dd ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 49 50 55 50 41 4b 52 4b 41 41 4d
Address 0x50: 58 32 30 30 30 2d 50 53 4d 2d 44 43 2d 53 00 00
Address 0x60: 00 00 00 00 00 00 31 30 31 ff ff ff ff ff ff
Address 0x70: ff ff ff 2a 00 00 00 00 00 00 00 00 00 00 00 00
PSM 15          REV 01   740-050037   1EDB3130024       DC 52V Power Supply
Module
Jedec Code:    0x7fb0           EEPROM Version: 0x02
P/N:           740-050037       S/N:            1EDB3130024
Assembly ID:   0x0478           Assembly Version: 01.01
Date:          05-16-2013       Assembly Flags:  0x00
Version:       REV 01           CLEI Code:      IPUPAKRKAA
ID: DC 52V Power Supply Module FRU Model Number: MX2000-PSM-DC-S
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 78 01 01 52 45 56 20 30 31 00 00
...

```

**show chassis hardware models (MX2020 Routers with MPC5EQ and MPC6E)**

```
user@host> show chassis hardware models
```

## Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 51	750-040240	ABAB9243	CHAS-BP-MX2020-S
FPM Board	REV 13	760-040242	ABCA8835	MX2020-CRAFT-S
PSM 0	REV 01	740-050037	1EDB32403L5	MX2000-PSM-DC-S
PSM 1	REV 01	740-050037	1EDB32403L3	MX2000-PSM-DC-S
PSM 2	REV 01	740-050037	1EDB32403KM	MX2000-PSM-DC-S
PSM 3	REV 01	740-050037	1EDB3130079	MX2000-PSM-DC-S
PSM 4	REV 01	740-050037	1EDB3130077	MX2000-PSM-DC-S
PSM 5	REV 01	740-050037	1EDB3130020	MX2000-PSM-DC-S
PSM 6	REV 01	740-050037	1EDB313009S	MX2000-PSM-DC-S
PSM 7	REV 01	740-050037	1EDB313008E	MX2000-PSM-DC-S
PSM 8	REV 01	740-050037	1EDB3130063	MX2000-PSM-DC-S
PSM 12	REV 01	740-050037	1EDB3130026	MX2000-PSM-DC-S
PSM 13	REV 01	740-050037	1EDB3130074	MX2000-PSM-DC-S
PSM 14	REV 01	740-050037	1EDB313009D	MX2000-PSM-DC-S
PSM 15	REV 01	740-050037	1EDB3130024	MX2000-PSM-DC-S
PSM 16	REV 01	740-050037	1EDB3130054	MX2000-PSM-DC-S
PSM 17	REV 01	740-050037	1EDB3130080	MX2000-PSM-DC-S
PDM 0	REV 03	740-045234	1EGA3170144	MX2000-PDM-DC-S
PDM 1	REV 03	740-045234	1EGA3170158	MX2000-PDM-DC-S
PDM 2	REV 03	740-045234	1EGA3170182	MX2000-PDM-DC-S
PDM 3	REV 03	740-045234	1EGA3170207	MX2000-PDM-DC-S
Routing Engine 0	REV 02	740-041821	9009112112	RE-MX2000-1800X4-S
Routing Engine 1	REV 02	740-041821	9009112087	RE-MX2000-1800X4-S
CB 0	REV 23	750-040257	CABA2295	RE-MX2000-1800X4-S
CB 1	REV 23	750-040257	CABE8379	RE-MX2000-1800X4-S
SFB 0	REV 06	711-044466	ABCD5001	MX2000-SFB-S
SFB 1	REV 06	711-044466	ABCD5034	MX2000-SFB-S
SFB 2	REV 06	711-044466	ABCH3899	MX2000-SFB-S
SFB 3	REV 06	711-044466	ABCD5020	MX2000-SFB-S
SFB 4	REV 06	711-044466	ABCD4975	MX2000-SFB-S
SFB 5	REV 06	711-044466	ABCH3881	MX2000-SFB-S
SFB 6	REV 06	711-044466	ABCD5026	MX2000-SFB-S
SFB 7	REV 06	711-044466	ABCD5032	MX2000-SFB-S
FPC 0	REV 39	750-045715	CACD1902	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	CABK8112	MX-MPC3E-3D
FPC 2	REV 17	750-037355	CAAS5826	MPC4E-3D-2CGE-8XGE
FPC 3	REV 05	750-044444	CAAY9920	MX-MPC2E-3D-P
FPC 4	REV 18	750-046005	CACH5661	PROTO-ASSEMBLY
FPC 5	REV 35	750-028467	CAAR2623	MPC-3D-16XGE-SFPP
FPC 9	REV 30	750-044130	ABCF5773	PROTO-ASSEMBLY
FPC 10	REV 36	750-044130	ABCS8602	PROTO-ASSEMBLY
FPC 17	REV 28	750-044130	ABBZ3873	PROTO-ASSEMBLY
FPC 18	REV 39	750-045715	CACD1910	PROTO-ASSEMBLY
FPC 19	REV 39	750-045715	CACD1908	PROTO-ASSEMBLY
ADC 0	REV 17	750-043596	ABCD5378	MX2000-LC-ADAPTER
ADC 1	REV 17	750-043596	ABCD5465	MX2000-LC-ADAPTER
ADC 2	REV 17	750-043596	ABCD5431	MX2000-LC-ADAPTER
ADC 3	REV 17	750-043596	ABCD5356	MX2000-LC-ADAPTER
ADC 4	REV 02	750-043596	ZW1545	750-043596
ADC 5	REV 17	750-043596	ABCD5517	MX2000-LC-ADAPTER
ADC 18	REV 17	750-043596	ABCD5535	MX2000-LC-ADAPTER
ADC 19	REV 01	750-043596	ZV4127	750-043596
Fan Tray 0	REV 06	760-046960	ACAY0791	MX2000-FANTRAY-S
Fan Tray 1	REV 06	760-046960	ACAY0788	MX2000-FANTRAY-S
Fan Tray 2	REV 06	760-046960	ACAY0755	MX2000-FANTRAY-S
Fan Tray 3	REV 06	760-046960	ACAY0441	MX2000-FANTRAY-S



## show chassis hardware clei-models (MX2020 Router with MPC5EQ and MPC6E)

```
user@host> show chassis hardware clei-models
```

```
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 51	750-040240	IPMU710ARA	CHAS-BP-MX2020-S
FPM Board	REV 13	760-040242	IPMYAE5JRA	MX2020-CRAFT-S
PSM 0	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 1	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 2	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 3	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 4	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 5	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 6	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 7	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 8	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 12	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 13	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 14	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 15	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 16	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PSM 17	REV 01	740-050037	IPUPAKRKAA	MX2000-PSM-DC-S
PDM 0	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
PDM 1	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
PDM 2	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
PDM 3	REV 03	740-045234	IPUPAJSKAA	MX2000-PDM-DC-S
CB 0	REV 23	750-040257	IPUCBA7CTA	RE-MX2000-1800X4-S
CB 1	REV 23	750-040257	IPUCBA7CTA	RE-MX2000-1800X4-S
SFB 0	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 1	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 2	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 3	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 4	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 5	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 6	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
SFB 7	REV 06	711-044466	IPUCBA6CAA	MX2000-SFB-S
FPC 0	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 1	REV 11	750-045372	COUIBBNBAA	MX-MPC3E-3D
FPC 2	REV 17	750-037355	IPU3A4DHAA	MPC4E-3D-2CGE-8XGE
FPC 3	REV 05	750-044444	COUIBBGBAA	MX-MPC2E-3D-P
MIC 0	REV 28	750-028387	COUIA16BAA	MIC-3D-4XGE-XFP
FPC 4	REV 18	750-046005	PROTOXCLEI	PROTO-ASSEMBLY
FPC 5	REV 35	750-028467		MPC-3D-16XGE-SFPP
FPC 9	REV 30	750-044130	PROTOXCLEI	PROTO-ASSEMBLY
MIC 0	REV 05	750-049457	PROTOXCLEI	PROTO-ASSEMBLY
FPC 10	REV 36	750-044130	PROTOXCLEI	PROTO-ASSEMBLY
MIC 0	REV 06	750-049979	PROTOXCLEI	PROTO-ASSEMBLY
MIC 1	REV 12	750-050008	PROTOXCLEI	PROTO-ASSEMBLY
FPC 17	REV 28	750-044130	PROTOXCLEI	PROTO-ASSEMBLY
MIC 1	REV 03	750-050008	PROTOXCLEI	PROTO-ASSEMBLY
FPC 18	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
FPC 19	REV 39	750-045715	PROTOXCLEI	PROTO-ASSEMBLY
ADC 0	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 1	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 2	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 3	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 4	REV 02	750-043596	PROTOXCLEI	750-043596
ADC 5	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 18	REV 17	750-043596	IPUCBA8CAA	MX2000-LC-ADAPTER
ADC 19	REV 01	750-043596	PROTOXCLEI	750-043596
Fan Tray 0	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S

Fan Tray 1	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S
Fan Tray 2	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S
Fan Tray 3	REV 06	760-046960	IPUCBA5CAA	MX2000-FANTRAY-S

### show chassis hardware (MX Series routers with ATM MIC)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN115736EAFc	MX240
Midplane	REV 07	760-021404	ABAA5038	MX240 Backplane
FPM Board	REV 03	760-021392	ABBA2758	Front Panel Display
PEM 0	Rev 01	740-022697	QCS0937C07K	PS 1.2-1.7kW; 100-240V
AC in				
PEM 1	Rev 01	740-022697	QCS0939C04X	PS 1.2-1.7kW; 100-240V
AC in				
PEM 2	Rev 01	740-022697	QCS0937C06B	PS 1.2-1.7kW; 100-240V
AC in				
PEM 3	Rev 01	740-022697	QCS0937C07U	PS 1.2-1.7kW; 100-240V
AC in				
Routing Engine 0	REV 12	740-013063	9009042291	RE-S-2000
Routing Engine 1	REV 12	740-013063	9009042266	RE-S-2000
CB 0	REV 06	710-021523	ABBC1435	MX SCB
CB 1	REV 06	710-021523	ABBC1497	MX SCB
FPC 2	REV 14	750-031088	YH8446	MPC Type 2 3D Q
CPU	REV 06	711-030884	YH9612	MPC PMB 2G
MIC 0				
MIC 1	REV 10	750-036132	ZP7062	2x0C12/8x0C3 CC-CE
PIC 2		BUILTIN	BUILTIN	2x0C12/8x0C3 CC-CE
Xcvr 0		NON-JNPR	23393-00492	UNKNOWN
Xcvr 1		NON-JNPR	23393-00500	UNKNOWN
Xcvr 2		NON-JNPR	23393-00912	UNKNOWN
Xcvr 3	REV 01	740-015638	22216-00575	Load SFP
Xcvr 4	REV 01	740-015638	24145-00110	Load SFP
Xcvr 5	REV 01	740-015638	24145-00016	Load SFP
Xcvr 6	REV 01	740-015638	24145-00175	Load SFP
Xcvr 7		NON-JNPR	23393-00627	UNKNOWN
QXM 0	REV 05	711-028408	YF4681	MPC QXM
QXM 1	REV 05	711-028408	YF4817	MPC QXM
Fan Tray 0	REV 01	710-021113	XL3645	MX240 Fan Tray

### show chassis hardware (MX240, MX480, MX960 routers with Application Services Modular Line Card)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN11D969BAFA	MX960
Midplane	REV 03	710-013698	ACAA2362	MX960 Backplane
FPM Board	REV 03	710-014974	ZR0639	Front Panel Display
PDM	Rev 03	740-013110	QCS152250SX	Power Distribution Module
PEM 0	Rev 10	740-013683	QCS1512718W	DC Power Entry Module
PEM 1	Rev 10	740-013683	QCS1512702Y	DC Power Entry Module
Routing Engine 0	REV 15	740-013063	9012024667	RE-S-2000
Routing Engine 1	REV 15	740-013063	9012024649	RE-S-2000
CB 0	REV 14	750-031391	ZJ7749	Enhanced MX SCB
CB 1	REV 14	750-031391	ZJ7750	Enhanced MX SCB
CB 2	REV 14	750-031391	ZY9233	Enhanced MX SCB
FPC 0	REV 17	750-031089	YR7434	MPC Type 2 3D
CPU				

FPC 1	REV 11	750-037207	ZW9727	AS-MCC
CPU	REV 04	711-038173	ZW4817	AS-MCC-PMB
MIC 0	REV 01	750-037214	ZH3764	AS-MSC
PIC 0		BUILTIN	BUILTIN	AS-MSC
MIC 1	REV 01	711-028408	JZ9200	AS-MXC
PIC 2		BUILTIN	BUILTIN	AS-MXC
FPC 4	REV 30	750-028467	ABBN0232	MPC 3D 16x 10GE
CPU				
FPC 5	REV 04	750-037207	ZK9074	AS-MCC
CPU				
Fan Tray 0	REV 05	740-014971	VT5683	Fan Tray
Fan Tray 1	REV 05	740-014971	VT5684	Fan Tray

show chassis hardware extensive (MX240, MX480, MX960 routers with Application Services Modular Line Card)

user@host> show chassis hardware extensive

```
ID: AS-MCC                                FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 37 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU                                REV 04    711-038173    ZW4817    AS-MCC-PMB
Jedec Code: 0x7fb0                EEPROM Version: 0x02
P/N: 711-038173                  S/N: ZW4817
Assembly ID: 0x0b38              Assembly Version: 01.04
Date: 12-30-2011                 Assembly Flags: 0x00
Version: REV 04
ID: AS-MCC-PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 37 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0                                REV 01    750-037214    ZH3764    AS-MSC
Jedec Code: 0x7fb0                EEPROM Version: 0x02
P/N: 750-037214                  S/N: ZH3764
Assembly ID: 0x0a44              Assembly Version: 01.01
Date: 07-04-2011                 Assembly Flags: 0x00
Version: REV 01
ID: AS-MSC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 01 52 45 56 20 30 31 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 48 33 37 36 34 00 00 00 04 07 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 00 00 00 00 00 00 00 00 00 00 00
```

```

Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff f6 c0 03 e1 bc 00 00 00 00 00 00 00 00
PIC 0          BUILTIN      BUILTIN      AS-MS
FPC 4          REV 30      750-028467  ABBN0232      MPC 3D 16x 10GE
Jedec Code:    0x7fb0      EEPROM Version: 0x01

```

#### show chassis hardware (MX480 Router with MPC4E)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN10FF57BAFB  MX480
Midplane      REV 05    750-047849   Good           MX480 Midplane
FPM Board     REV 02    710-017254   KG2066         Front Panel Display
PEM 0         Rev 03    740-017330   QCS081590BJ    PS 1.2-1.7kW; 100-240V
AC in
PEM 1         Rev 03    740-017330   QCS0815908Z    PS 1.2-1.7kW; 100-240V
AC in
PEM 2         Rev 03    740-029970   QCS1001U001    PS 1.4-2.52kW; 90-264V
AC in
Routing Engine 0 REV 05    740-031116   9009089502     RE-S-1800x4
Routing Engine 1 REV 05    740-031116   9009089624     RE-S-1800x4
CB 0          REV 02    750-031391   YE8506         Enhanced MX SCB
CB 1          REV 14    750-031391   ZK8265         Enhanced MX SCB
FPC 2         REV 05    750-037358   ZT0638         MPC4E 3D 32XGE
CPU           REV 07    711-035209   ZK3187         HMPD PMB 2G
PIC 0         BUILTIN   BUILTIN       8X10GE SFPP
PIC 1         BUILTIN   BUILTIN       8X10GE SFPP
PIC 2         BUILTIN   BUILTIN       8X10GE SFPP
PIC 3         BUILTIN   BUILTIN       8X10GE SFPP
FPC 3         REV 06    750-037355   CAAB1144       MPC4E 3D 2CGE+8XGE
CPU           REV 08    711-035209   CAAB1278       HMPD PMB 2G
PIC 0         BUILTIN   BUILTIN       4x10GE SFPP
Xcvr 0        REV 01    740-031980   B11E01439     SFP+-10G-SR
Xcvr 1        REV 01    740-031980   B11D05809     SFP+-10G-SR
PIC 1         BUILTIN   BUILTIN       1X100GE CFP
Xcvr 0        NON-JNPR   D5418         UNKNOWN
PIC 2         BUILTIN   BUILTIN       4x10GE SFPP
PIC 3         BUILTIN   BUILTIN       1X100GE CFP
Xcvr 0        NON-JNPR   X12J00362     CFP-100G-SR10
FPC 4         REV 12.3.10 750-033205   YR9445         MPCE Type 3 3D
CPU
Fan Tray                               Enhanced Left Fan Tray

```

#### show chassis hardware (MX2020 Router with MPC4E)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11E188CAFJ  MX2020
Midplane      REV 04    711-032387   ABAC7474       Lower Backplane
Midplane 1    REV 04    711-032386   ABAC7408       Upper Backplane
PMP 1         REV 03    711-032428   ACAJ1137       Upper Power Midplane
PMP 0         REV 03    711-032426   ACAJ1016       Lower Power Midplane
FPM Board     REV 06    760-040242   ABBT8832       Front Panel Display
PSM 3         REV 0C    740-033727   VK00255        DC 52V Power Supply
Module
PSM 4         REV 0C    740-033727   VJ00148        DC 52V Power Supply
Module
PSM 5         REV 0C    740-033727   VK00207        DC 52V Power Supply

```

Module				
PSM 6	REV 0C	740-033727	VK00319	DC 52V Power Supply
Module				
PSM 7	REV 0C	740-033727	VK00264	DC 52V Power Supply
Module				
PSM 8	REV 0B	740-033727	VG00025	DC 52V Power Supply
Module				
PSM 13	REV 0C	740-033727	VK00274	DC 52V Power Supply
Module				
PSM 14	REV 0C	740-033727	VJ00167	DC 52V Power Supply
Module				
PSM 15	REV 0C	740-033727	VK00299	DC 52V Power Supply
Module				
PSM 16	REV 0C	740-033727	VK00213	DC 52V Power Supply
Module				
PSM 17	REV 0C	740-033727	VK00253	DC 52V Power Supply
Module				
PDM 0	REV 0B	740-038109	VJ00040	DC Power Dist Module
PDM 2	REV 0B	740-038109	VJ00025	DC Power Dist Module
Routing Engine 0	REV 02	740-041821	9009089735	RE-S-1800x4
Routing Engine 1	REV 02	740-041821	9009089731	RE-S-1800x4
CB 0	REV 04	750-040257	ZT2846	Control Board
CB 1	REV 04	750-040257	ZT2877	Control Board
SPMB 0	REV 01	711-041855	ZS2282	PMB Board
SPMB 1	REV 01	711-041855	ZS2261	PMB Board
SFB 0	REV 07	711-032385	ZZ2582	Switch Fabric Board
SFB 1	REV 04	711-032385	ZV4229	Switch Fabric Board
SFB 2	REV 07	711-032385	CAAB4902	Switch Fabric Board
SFB 3	REV 07	711-032385	CAAB4891	Switch Fabric Board
SFB 4	REV 07	711-032385	CAAB4883	Switch Fabric Board
SFB 5	REV 07	711-032385	CAAB4889	Switch Fabric Board
SFB 6	REV 06	711-032385	ZV1818	Switch Fabric Board
SFB 7	REV 07	711-032385	CAAB4897	Switch Fabric Board
FPC 0	REV 34	750-031090	ZT9799	MPC Type 2 3D EQ
CPU	REV 06	711-030884	ZS1122	MPC PMB 2G
MIC 0	REV 11	750-033535	CAAD7674	MIC-3D-10C192-XFP
PIC 0		BUILTIN	BUILTIN	MIC-3D-10C192-XFP
Xcvr 0	REV 01	740-014279	753019A00404	XFP-0C192-SR
MIC 1	REV 14	750-031967	ZM6103	MIC-3D-80C30C12-40C48
PIC 2		BUILTIN	BUILTIN	MIC-3D-80C30C12-40C48
Xcvr 0	REV 01	740-011615	PEF1AZP	SFP-IR
Xcvr 1	REV 01	740-011615	PEF1AZN	SFP-IR
Xcvr 2	REV 01	740-021308	ANA0N8S	SFP+-10G-SR
QXM 0	REV 06	711-028408	ZT9339	MPC QXM
QXM 1	REV 06	711-028408	ZT9237	MPC QXM
FPC 9	REV 34	750-031090	ZT9770	MPC Type 2 3D EQ
CPU	REV 06	711-030884	ZS1302	MPC PMB 2G
MIC 0	REV 24	750-028387	YJ3950	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0		NON-JNPR	T09M52516	XFP-10G-SR
Xcvr 1		NON-JNPR	CA49BK095	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 02	740-014289	C834XU01T	XFP-10G-SR
Xcvr 1		NON-JNPR	T09M52515	XFP-10G-SR
MIC 1	REV 11	750-033535	CAAD7681	MIC-3D-10C192-XFP
PIC 2		BUILTIN	BUILTIN	MIC-3D-10C192-XFP
Xcvr 0	REV 01	740-014279	KBQ02BE	XFP-0C192-SR
QXM 0	REV 06	711-028408	ZT9151	MPC QXM
QXM 1	REV 06	711-028408	ZT9116	MPC QXM
FPC 10	REV 27	750-033205	ZL6215	MPCE Type 3 3D
CPU	REV 07	711-035209	ZK9038	HMPC PMB 2G

MIC 0	REV 18	750-028380	YG6885	3D 2x 10GE XFP
PIC 0		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 01	740-014289	C706XU0AG	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	1x 10GE XFP
Xcvr 0	REV 02	740-014289	T08L84366	XFP-10G-SR
FPC 14	REV 09	750-037355	CAAF1534	MPC4E 3D 2CGE+8XGE
CPU	REV 08	711-035209	CAAB9879	HMPC PMB 2G
PIC 0		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	21T511100436	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AHPOGPM	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	123363A00032	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	19T511100477	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12J00260	CFP-100G-SR10
PIC 2		BUILTIN	BUILTIN	4x10GE SFPP
Xcvr 0	REV 01	740-021308	21T511104086	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	21T511104627	SFP+-10G-SR
Xcvr 3	REV 01	740-021308	21T511104644	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	1X100GE CFP
FPC 19	REV 32	750-028467	ZR2008	MPC 3D 16x 10GE
CPU	REV 10	711-029089	ZT6933	AMPC PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	19T511100291	SFP+-10G-SR
Xcvr 1	REV 01	740-021308	AMH02VE	SFP+-10G-SR
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	23T511102128	SFP+-10G-SR
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-021308	AMS15PP	SFP+-10G-SR
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	123363A00716	SFP+-10G-SR
ADC 0	REV 05	750-043596	CAAC2072	Adapter Card
ADC 9	REV 01	750-043596	ZV4111	Adapter Card
ADC 10	REV 05	750-043596	CAAC2058	Adapter Card
ADC 14	REV 02	750-043596	ZW1561	Adapter Card
ADC 19	REV 01	750-043596	ZV4127	Adapter Card
Fan Tray 0	REV 03	760-046960	ACAY0124	172mm FanTray - 6 Fans
Fan Tray 1	REV 2A	760-046960	ACAY0022	172mm FanTray - 6 Fans
Fan Tray 2	REV 2A	760-046960	ACAY0023	172mm FanTray - 6 Fans
Fan Tray 3	REV 2A	760-046960	ACAY0025	172mm FanTray - 6 Fans

show chassis hardware (MX5, MX10, MX40, MX80, MX240, MX480, and MX960 routers with Enhanced 20-port Gigabit Ethernet MIC)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			F3434	MX80-P
Midplane	REV 01	711-044315	ZK2681	MX80-P
PEM 0	Rev 04	740-028288	VE05267	AC Power Entry Module
PEM 1	Rev 04	740-028288	VE05270	AC Power Entry Module
Routing Engine		BUILTIN	BUILTIN	Routing Engine
TFEB 0		BUILTIN	BUILTIN	Forwarding Engine
Processor				
QXM 0	REV 05	711-028408	ZK0952	MPC QXM
FPC 0		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0		BUILTIN	BUILTIN	4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	4x 10GE XFP
FPC 1		BUILTIN	BUILTIN	MPC BUILTIN
MIC 0	REV 02	750-049846	CAAV2153	3D 20x 1GE(LAN)-E,SFP
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) -E SFP
Xcvr 0	REV 01	740-011613	AM0816S9B81	SFP-SX

Xcvr 1	REV 02	740-011613	AM0925SBLK7	SFP-SX
Xcvr 2	REV 01	740-011613	UAQ0005	SFP-SX
Xcvr 3	REV 01	740-011613	UAQ000C	SFP-SX
Xcvr 4	REV 01	740-011613	P9F195E	SFP-SX
Xcvr 5	REV 01	740-011613	UAQ0003	SFP-SX
Xcvr 6	REV 01	740-031851	AM1041SU1LD	SFP-SX
Xcvr 8	REV 02	740-013111	B101501	SFP-T
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) -E SFP
Xcvr 0	REV 01	740-011613	PFM1ML7	SFP-SX
Xcvr 4	REV 01	740-011613	PE729P6	SFP-SX
Xcvr 6	REV 02	740-011613	AM1014SGC84	SFP-SX
Xcvr 9	REV 01	740-011613	AM0812S8UK3	SFP-SX
MIC 1	REV 26	750-028392	ZY0187	3D 20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	P9F1AN9	SFP-SX
Xcvr 5	REV 02	740-011613	AM1003SFUF4	SFP-SX
Xcvr 9	REV 01	740-031851	AM1041SU1LM	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 4	REV 01	740-011613	PAJ4MYT	SFP-SX
Xcvr 7	+	NON-JNPR	XG32A024	SFP-SX
Xcvr 8		NON-JNPR	PFROV6J	SFP-SX
Xcvr 9	REV 01	740-031851	AM1041SU02U	SFP-SX
Fan Tray				

#### show chassis hardware models (MX5, MX10, MX40, MX80, MX240, MX480, and MX960 routers with Enhanced 20-port Gigabit Ethernet MIC)

```
user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
PEM 0         Rev 04    740-028288  VE05267       PWR-MX80-AC-S
PEM 1         Rev 04    740-028288  VE05270       PWR-MX80-AC-S
Routing Engine
TFEB 0        BUILTIN   BUILTIN
FPC 0         BUILTIN   BUILTIN
FPC 1         BUILTIN   BUILTIN
MIC 0         REV 02    750-049846  CAAV2153      MIC-3D-20GE-SFP-E
MIC 1         REV 26    750-028392  ZY0187        MIC-3D-20GE-SFP
Fan Tray      FANTRAY-MX80-S
```

#### show chassis hardware (T320 Router)

```
user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis       19093    T320
Midplane      REV 04    710-004339  BC1436        T320 Backplane
FPM GBUS      REV 03    710-004461  BC1407        T320 FPM Board
FPM Display   REV 04    710-002897  BE0763        FPM Display
CIP           REV 05    710-002895  BB2311        T Series CIP
PEM 0         Rev 01    740-004359  NB12546       Power Entry Module
SCG 0         REV 06    710-004455  AY4522        T320 Sonet
Clock Gen.
Routing Engine 0
CB 0          REV 13    710-002728  BC1577        unknown
T Series
Control Board
CB 1          REV 13    710-002728  BC1595        T Series
Control Board
FPC 1         REV 09    710-007531  HS1572        FPC Type 2
CPU           REV 15    710-001726  HR8763        FPC CPU
PIC 0         REV 01    750-010618  CB5579        4x G/E SFP,
```

```

1000 BASE
  SFP 0      REV 01  740-007326  P5809Z1      SFP-SX
  SFP 1      REV 01  740-007326  P4Q10XU      SFP-SX
  SFP 2              NON-JNPR    RA45020031   SFP-SX
  SFP 3              NON-JNPR    RA45020032   SFP-SX
  PIC 1      REV 01  750-010618  CD9587       4x G/E SFP,
1000 BASE
  SFP 0              NON-JNPR    P5A08QZ      SFP-T
  SFP 1      REV 01  740-007326  P4Q133K      SFP-SX
  SFP 2      REV 01  740-007326  P5809YY      SFP-SX
  SFP 3      REV 01  740-007327  4C81704      SFP-LX
  MMB 1      REV 03  710-005555  HR9401       MMB-288mbit
  PPB 0      REV 04  710-003758  HR2886       PPB Type 2
  FPC 2      REV 07  710-005860  HP2392       FPC Type 1
  CPU        REV 14  710-001726  HP7797       FPC CPU
  PIC 0      REV 02  750-007643  HM0853       1x G/E QPP,
1000 BASE
  SFP 0      REV 01  740-007326  P11E9JJ      SFP-SX
  MMB 1      REV 02  710-005555  HN2379       MMB-288mbit
  PPB 0      REV 04  710-003758  HP8092       PPB Type 2
  FPC 3      REV 07  710-005860  HP2393       FPC Type 1
  CPU        REV 14  710-001726  HP0968       FPC CPU
  PIC 0      REV 01  750-010240  CB5363       1x G/E SFP,
1000 BASE
  SFP 0      REV 01  740-007326  P4R0PNH      SFP-SX
  PIC 1      REV 03  750-003034  HD2832       4x OC-3 SONET,
SMIR
  MMB 1      REV 02  710-005555  HN6307       MMB-288mbit
  PPB 0      REV 04  710-003758  HP5051       PPB Type 2
  FPC 4      REV 01  710-010845  JD3872       FPC Type 4
  CPU        REV 02  710-011481  JB6042       FPC CPU
  5          REV 01  710-005802  BC1566       FPC Type 2
  CPU        REV 09  710-001726  AY4922       FPC CPU
  PIC 0      REV 02  750-008155  BE2114       2x G/E QPP,
1000 BASE
  SFP 0      REV 01  740-007326  P4R0PMQ      SFP-SX
  SFP 1      REV 01  740-007326  P4R0PN9      SFP-SX
  PIC 1      REV 01  750-008155  BE2116       2x G/E QPP,
1000 BASE
  SFP 0      REV 01  740-007326  P4R0PNZ      SFP-SX
  SFP 1              NON-JNPR    2908         SFP-T
  MMB 1      REV 01  710-005555  AZ2246       MMB-288mbit
  PPB 0      REV 03  710-003758  AY4839       PPB Type 2
  FPC 7      REV 01  710-005803  AZ2123       FPC Type 3
...

```

### show chassis hardware (T640 Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               19182         T640
Midplane      REV 04   710-002726  AX5608        T640 Backplane
FPM GBUS      REV 02   710-002901  HE3064        T640 FPM Board
FPM Display   REV 02   710-002897  HE7864        FPM Display
CIP           REV 05   710-002895  HA5024        T Series CIP
PEM 0         Rev 02   740-029522  VH26235       AC PEM 10kW US
PEM 1         Rev 02   740-029522  VH26230       AC PEM 10kW US
SCG 0         REV 03   710-003423  HA4508        T640 Sonet Clock Gen.
Routing Engine 0 REV 02   740-005022  210865700483 RE-3.0 (RE-600)
CB 0          REV 01   710-002728  HD3044        T Series Control Board

```



FPC 2	REV 04	710-001721	HD5572	FPC Type 3
CPU	REV 06	710-001726	HA4712	FPC CPU
PIC 1	REV 03	750-009567	HV2331	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202R103	XENPAK-SR
PIC 2	REV 03	750-009567	HV2332	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-011268	USC202R112	XENPAK-ZR
PIC 3	REV 03	750-009567	HX4416	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012056	434TC004	XENPAK-CX4
PIC 4	REV 03	750-009567	HX4420	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-012058	434TC124	XENPAK-LX4
FPC 5	REV 01	710-013553	JE4839	E2-FPC Type 1
CPU	REV 01	710-013569	JW9163	FPC CPU
PIC 0	REV 01	750-009567	HX4419	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009898	USC202RT05	XENPAK-LR
PIC 1	REV 03	750-009567	HN7426	1x 10GE(LAN),XENPAK
SFP 0	REV 01	740-009550	03L90051	XENPAK-ER
PIC 2	REV 03	750-009467	HT7423	1x 10GE(LAN),XENPAK
SFP 0		NON-JNPR		UNKNOWN
PIC 3	REV 04	750-005100	AY4850	1x 10GE(LAN),DWDM
FPC 4	REV 01	710-010845	JD3872	FPC Type 4
CPU	REV 02	710-011481	JB6042	FPC CPU
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

#### show chassis hardware models (T640 Router)

```

user@host> show chassis hardware models
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-002726		CHAS-BP-T640-S
FPM Display	REV 02	710-002897		CRAFT-T640-S
CIP	REV 05	710-002895		CIP-L-T640-S
PEM 0	Rev 01	740-002595		PWR-T-DC-S
SCG 0	REV 04	710-003423		SCG-T-S
SCG 1	REV 04	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-005022		RE-600-2048-S
Routing Engine 1	REV 07	740-005022		RE-600-2048-S
CB 0	REV 06	710-002726		CHAS-BP-T640-S
CB 1	REV 06	710-002728		CB-L-T-S
FPC 5	REV 05	710-007527		T640-FPC2
PIC 0	REV 05	750-002510		PB-2GE-SX
PIC 1	REV 05	750-001901		PB-40C12-SON-SMIR
FPC 6	REV 03	710-001721		T640-FPC3
PIC 1	REV 01	750-009553		PC-40C48-SON-SFP
SIB 4	REV 02	750-005486		SIB-I-T640-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FAN-REAR-TX-T640-S

#### show chassis hardware extensive (T640 Router)

```

user@host> show chassis hardware extensive
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis				T640
Jedec Code:	0x7fb0	EEPROM Version:	0x01	
P/N:	.....	S/N:	.....	
Assembly ID:	0x0507	Assembly Version:	00.00	
Date:	00-00-0000	Assembly Flags:	0x00	
Version:	.....			

```

ID: Gibson LCC Chassis
Board Information Record:
  Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 05 07 00 00 00 00 00 00 00 00 00 00
  Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x20: ff ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00 00 00
  Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Midplane      REV 04   710-002726   AX5633
Jedec Code:   0x7fb0           EEPROM Version:   0x01
P/N:          710-002726.      S/N:           AX5633.
Assembly ID:  0x0127          Assembly Version: 01.04
Date:         06-27-2001      Assembly Flags: 0x00
Version:      REV 04.....
ID: Gibson Backplane
Board Information Record:
  Address 0x00: ad 01 08 00 00 90 69 0e f8 00 ff ff ff ff ff ff
I2C Hex Data:
  Address 0x00: 7f b0 01 ff 01 27 01 04 52 45 56 20 30 34 00 00
  Address 0x10: 00 00 00 00 37 31 30 2d 30 30 32 37 32 36 00 00
  Address 0x20: 53 2f 4e 20 41 58 35 36 33 33 00 00 00 1b 06 07
  Address 0x30: d1 ff ff ff ad 01 08 00 00 90 69 0e f8 00 ff ff
  Address 0x40: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
FPM GBUS      REV 02   710-002901   HE3245
...
FPM Display   REV 02   710-002897   HA4873
...
CIP           REV 05   710-002895   HA4729
...
PEM 1         RevX02   740-002595   MD21815           Power Entry Module
...
SCG 0         REV 04   710-003423   HF6023
...
SCG 1         REV 04   710-003423   HF6061
...
Routing Engine 0 REV 01   740-005022   210865700292     RE-3.0
...
CB 0          REV 06   710-002728   HE3614
...
FPC 1         REV 01   710-002385   HE3009           FPC Type 1
...
              REV 06   710-001726   HC0010

```

### show chassis hardware (T4000 Router)

```

user@host> show chassis hardware
Hardware inventory:

```

Item	Version	Part number	Serial number	Description
Chassis			JN1172F25AHA	T4000
Midplane	REV 01	710-027486	RC8355	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAE0927	T640 FPM Board
FPM Display	REV 01	710-021387	EF6764	T1600 FPM Display
CIP	REV 06	710-002895	BBAD9210	T-series CIP
PEM 0	REV 01	740-036442	VA00016	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAD7248	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAE3874	T640 Sonet Clock Gen.
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-1800
CB 0	REV 09	710-022597	ED0295	LCC Control Board
CB 1	REV 09	710-022597	EA6050	LCC Control Board
FPC 0	REV 26	750-032819	EK1173	FPC Type 5-3D

CPU	REV 12	711-030686	EJ8584	SNG PMB
PIC 0	REV 07	750-034624	EF6837	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	123363A01145	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	123363A01147	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01P3	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M03256	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01M2	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	123363A01137	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01PN	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01NW	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	123363A01139	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01KE	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01336	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10M01325	SFP+-10G-SR
PIC 1	REV 07	750-034624	EF6800	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJJ01SA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01QZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJH0217	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ01TE	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01KV	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJJ01MU	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01R0	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01TC	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ0364	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJD0GV3	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B10M03343	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01QJ	SFP+-10G-SR
LMB 0	REV 05	711-034381	EJ8490	Type-0 LMB
LMB 1	REV 04	711-035774	EJ8517	Type-1 LMB
LMB 2	REV 05	711-034381	EJ8489	Type-0 LMB
FPC 3	REV 07	750-032819	EG3637	FPC Type 5-3D
CPU	REV 09	711-030686	EG0150	SNG PMB
PIC 0	REV 08	750-035293	EF3657	1x100GE
Xcvr 0	REV 01	740-032210	C22CQNJ	CFP-100G-LR4
PIC 1	REV 10	750-034624	BBAN4098	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR
LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP

Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

### show chassis hardware (T4000 Router with 16 GB line card chassis (LCC) Routing Engine)

```
user@host> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11BDF2CAHA	T1600
Midplane	REV 01	710-027486	ACAJ0774	T640 Backplane
FPM GBUS	REV 13	710-002901	BBAL6812	T640 FPM Board
FPM Display	REV 04	710-021387	BBAP2679	T1600 FPM Display
CIP	REV 06	710-002895	BBAP4758	T-series CIP
PEM 0	Rev 03	740-026384	XF86421	Power Entry Module 3x80
PEM 1	Rev 03	740-026384	XF86429	Power Entry Module 3x80
SCG 0	REV 18	710-003423	BBAP1896	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAN8659	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-042243	737F-002238	RE-DUO-1800-16G
Routing Engine 1	REV 01	740-042243	737F-002403	RE-DUO-1800-16G
CB 1	REV 11	710-022597	EK4526	LCC Control Board
CB 1	REV 11	710-022597	EK4527	LCC Control Board
FPC 0	REV 05	710-033871	EK5644	FPC Type 4-ES
CPU	REV 11	710-016744	EK3428	ST-PMB2
PIC 0	REV 20	750-017405	EJ3041	4x 10GE (LAN/WAN) XFP

PIC 1	REV 17	750-026962	EH7536	10x10GE(LAN/WAN) SFPP
MMB 0	REV 07	710-025563	EK6039	ST-MMB2
MMB 1	REV 07	710-025563	EK6086	ST-MMB2
FPC 1	REV 05	710-033871	EK6583	FPC Type 4-ES
CPU	REV 11	710-016744	EK3401	ST-PMB2
PIC 0	REV 17	750-026962	EJ8948	10x10GE(LAN/WAN) SFPP
MMB 0	REV 07	710-025563	EK6202	ST-MMB2
MMB 1	REV 07	710-025563	EK6112	ST-MMB2
SPMB 1	REV 05	710-023321	EK4900	LCC Switch CPU
SIB 0	REV 11	710-013074	EK5958	SIB-I8-SF
SIB 1	REV 11	710-013074	EK4606	SIB-I8-SF
SIB 2	REV 11	710-013074	EK5971	SIB-I8-SF
SIB 3	REV 11	710-013074	EK4609	SIB-I8-SF
SIB 4	REV 11	710-013074	EK4602	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 2

#### show chassis hardware (T4000 Router with LSR FPC)

```
user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1173A24AHA  T4000
FPC 3         REV     750-048373  AN7797         FPC Type 5-LSR
CPU           REV 10  711-030686  AN6649         SNG PMB
PIC 0         REV 07  750-034624  EF6830         12x10GE (LAN/WAN) SFPP
```

#### show chassis hardware clei-models (T4000 Router)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code  FRU model number
Midplane      REV 01  710-027486  IPMJ700DRD CHAS-BP-T1600-S
FPM Display   REV 01  710-021387                CRAFT-T1600-S
CIP           REV 06  710-002895                CIP-L-T640-S
PEM 0         REV 01  740-036442  IPUPAG6KAA PWR-T-6-60-DC
SCG 0         REV 18  710-003423                SCG-T-S
SCG 1         REV 18  710-003423                SCG-T-S
Routing Engine 0 REV 05  740-026941                RE-DUO-C1800-8G-S
Routing Engine 1 REV 06  740-026941                RE-DUO-C1800-8G-S
CB 0          REV 09  710-022597                CB-LCC-S
CB 1          REV 09  710-022597                CB-LCC-S
FPC 3
PIC 0         REV 08  750-035293  XXXXXXXXBB PF-1CGE-CFP
PIC 1         REV 10  750-034624  XXXXXXXXCC PF-12XGE-SFPP
FPC 5         REV 03  710-033871  IPUCAMBCTD T1600-FPC4-ES
PIC 1         REV 03  750-034781  IPUIBKLMAA PD-1CE-CFP-FPC4
FPC 6
PIC 0         REV 10  750-034624  XXXXXXXXCC PF-12XGE-SFPP
Fan Tray 0    FANTRAY-T-S
Fan Tray 1    FANTRAY-T4000-S
Fan Tray 2    FANTRAY-TXP-R-S
```

#### show chassis hardware detail (T4000 Router)

```
user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1172F25AHA  T4000
Midplane      REV 01  710-027486  RC8355         T-series Backplane
FPM GBUS      REV 13  710-002901  BBAE0927       T640 FPM Board
```

FPM Display	REV 01	710-021387	EF6764	T1600 FPM Display
CIP	REV 06	710-002895	BBAD9210	T-series CIP
PEM 0	REV 01	740-036442	VA00016	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAD7248	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAE3874	T640 Sonet Clock Gen.
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-1800
ad0 3823 MB	SMART CF		2009121602A661576157	Compact Flash
ad1 59690 MB	STEC MACH-8 SSD		STM000103FDB	Disk 1
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-1800
ad0 3823 MB	SMART CF		201011150153F52CF52C	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		2010110900150A880A88	Disk 1
CB 0	REV 09	710-022597	ED0295	LCC Control Board
CB 1	REV 09	710-022597	EA6050	LCC Control Board
FPC 0	REV 26	750-032819	EK1173	FPC Type 5-3D
CPU	REV 12	711-030686	EJ8584	SNG PMB
PIC 0	REV 07	750-034624	EF6837	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	123363A01145	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	123363A01147	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01P3	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B10M03256	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01M2	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	123363A01137	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01PN	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01NW	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	123363A01139	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01KE	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01336	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B10M01325	SFP+-10G-SR
PIC 1	REV 07	750-034624	EF6800	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJJ01SA	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01QZ	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJH0217	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ01TE	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	AJJ01KV	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJJ01MU	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01R0	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	AJJ01TC	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ0364	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJD0GV3	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B10M03343	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01QJ	SFP+-10G-SR
LMB 0	REV 05	711-034381	EJ8490	Type-0 LMB
LMB 1	REV 04	711-035774	EJ8517	Type-1 LMB
LMB 2	REV 05	711-034381	EJ8489	Type-0 LMB
FPC 3	REV 07	750-032819	EG3637	FPC Type 5-3D
CPU	REV 09	711-030686	EG0150	SNG PMB
PIC 0	REV 08	750-035293	EF3657	1x100GE
Xcvr 0	REV 01	740-032210	C22CQNJ	CFP-100G-LR4
PIC 1	REV 10	750-034624	BBAN4098	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04902	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04891	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01MX	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04183	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04894	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04184	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04897	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04899	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AJJ01TV	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04057	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ01M4	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04905	SFP+-10G-SR

LMB 0	REV 04	711-034381	EG1524	Type-0 LMB
LMB 1	REV 03	711-035774	EG0345	Type-1 LMB
LMB 2	REV 04	711-034381	EG1522	Type-0 LMB
FPC 5	REV 03	710-033871	BBAJ0768	FPC Type 4-ES
CPU	REV 11	710-016744	BBAH9342	ST-PMB2
PIC 0	REV 09	750-029262	EE6789	100GE
PIC 1	REV 03	750-034781	EE6655	100GE CFP
Xcvr 0	REV 01	740-032210	J11A22334	CFP-100G-LR4
BRIDGE 0	REV 03	711-029995	EE6572	100GE Bridge Board
MMB 0	REV 07	710-025563	BBAJ4657	ST-MMB2
MMB 1	REV 07	710-025563	BBAJ3073	ST-MMB2
FPC 6	REV 05	750-010153	EF4936	FPC Type 5-3D
CPU	REV 06	711-030686	EF4189	SNG PMB
PIC 0	REV 10	750-034624	BBAN4109	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	B11J04895	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11J04898	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	B11J04021	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	B11J04903	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B11J04311	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J04059	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11J04016	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11J04017	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B11J04887	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	B11J04297	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11J04893	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	B11J04022	SFP+-10G-SR
PIC 1	REV 02	750-034624	EE3711	12x10GE (LAN/WAN) SFPP
Xcvr 0	REV 01	740-031980	AJH033X	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	AJJ01N0	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AJJ01SV	SFP+-10G-SR
Xcvr 3	REV 01	740-031980	AJJ032L	SFP+-10G-SR
Xcvr 4	REV 01	740-031980	B10M01593	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	AJD0FF1	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	AJJ01NU	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	123363A01305	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	B10M00361	SFP+-10G-SR
Xcvr 9	REV 01	740-031980	AJJ01M7	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	AJJ032X	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AJJ01PG	SFP+-10G-SR
LMB 0	REV 04	711-034381	EF3838	Type-0 LMB
LMB 1	REV 03	711-035774	EF3821	Type-1 LMB
LMB 2	REV 04	711-034381	EF3834	Type-0 LMB
SPMB 0	REV 05	710-023321	ED1990	LCC Switch CPU
SPMB 1	REV 05	710-023321	EA2768	LCC Switch CPU
SIB 0	REV 02	711-036340	EF8802	SIB-HC-3D
SIB 1	REV 07	711-036340	EG2286	SIB-HC-3D
SIB 2	REV 07	711-036340	EG2252	SIB-HC-3D
SIB 3	REV 02	711-036340	EF1358	SIB-HC-3D
SIB 4	REV 02	711-036340	EF8806	SIB-HC-3D
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
-- Rev 2				
Fan Tray 2				Rear Fan Tray -- Rev 3

### show chassis hardware models (T4000 Router)

```
user@host> show chassis hardware models
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 01	710-027486	RC8355	CHAS-BP-T1600-S

FPM Display	REV 01	710-021387	EF6764	CRAFT-T1600-S
CIP	REV 06	710-002895	BBAD9210	CIP-L-T640-S
PEM 0	REV 01	740-036442	VA00016	PWR-T-6-60-DC
SCG 0	REV 18	710-003423	BBAD7248	SCG-T-S
SCG 1	REV 18	710-003423	BBAE3874	SCG-T-S
Routing Engine 0	REV 05	740-026941	P737F-002248	RE-DUO-C1800-8G-S
Routing Engine 1	REV 06	740-026941	P737F-002653	RE-DUO-C1800-8G-S
CB 0	REV 09	710-022597	ED0295	CB-LCC-S
CB 1	REV 09	710-022597	EA6050	CB-LCC-S
FPC 3				
PIC 0	REV 08	750-035293	EF3657	PF-1CGE-CFP
PIC 1	REV 10	750-034624	BBAN4098	PF-12XGE-SFPP
FPC 5	REV 03	710-033871	BBAJ0768	T1600-FPC4-ES
PIC 1	REV 03	750-034781	EE6655	PD-1CE-CFP-FPC4
FPC 6				
PIC 0	REV 10	750-034624	BBAN4109	PF-12XGE-SFPP
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T4000-S
Fan Tray 2				FAN-REAR-TXP-LCC

### show chassis hardware lcc (TX Matrix Router)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
```

#### Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			65751	T640
Midplane	REV 03	710-005608	RA1408	T640 Backplane
FPM GBUS	REV 09	710-002901	RA2784	T640 FPM Board
FPM Display	REV 05	710-002897	RA2825	FPM Display
CIP	REV 06	710-002895	HT0684	T Series CIP
PEM 0	Rev 11	740-002595	PM18483	Power Entry Module
PEM 1	Rev 11	740-002595	qb13984	Power Entry Module
SCG 0	REV 11	710-003423	HT0022	T640 Sonet Clock Gen.
Routing Engine 0	REV 13	740-005022	210865700363	RE-3.0 (RE-600)
CB 0	REV 03	710-007655	HW1195	Control Board (CB-T)
FPC 1	REV 05	710-007527	HM3245	FPC Type 2
CPU	REV 14	710-001726	HM1084	FPC CPU
PIC 0	REV 02	750-007218	AZ1112	2x OC-12 ATM2 IQ, SMIR
PIC 1	REV 02	750-007745	HG3462	4x OC-3 SONET, SMIR
PIC 2	REV 14	750-001901	BA5390	4x OC-12 SONET, SMIR
PIC 3	REV 09	750-008155	HS3012	2x G/E IQ, 1000 BASE
SFP 0		NON-JNPR	P1186TY	SFP-S
SFP 1	REV 01	740-007326	P11WLTF	SFP-SX
MMB 1	REV 02	710-005555	HL7514	MMB-288mbit
PPB 0	REV 04	710-003758	HM4405	PPB Type 2
PPB 1	REV 04	710-003758	AV1960	PPB Type 2
FPC 2	REV 08	710-010154	HZ3578	E-FPC Type 3
CPU	REV 05	710-010169	HZ3219	FPC CPU-Enhanced
PIC 0	REV 02	750-009567	HX2882	1x 10GE(LAN), XENPAK
SFP 0	REV 01	740-009898	USC202U709	XENPAK-LR
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 2	REV 01	750-004535	HC0235	1x OC-192 SM SR1
PIC 3	REV 07	750-007141	HX1699	10x 1GE(LAN), 1000 BASE
SFP 0	REV 01	740-007326	2441042	SFP-SX
SFP 1	REV 01	740-007326	2441027	SFP-SX
MMB 0	REV 03	710-010171	HV2365	MMB-5M3-288mbit
MMB 1	REV 03	710-010171	HZ3888	MMB-5M3-288mbit
SPMB 0	REV 09	710-003229	HW5245	T Series Switch CPU



SIB 3	REV 07	710-005781	HR5927	SIB-L8-F16
B Board	REV 06	710-005782	HR5971	SIB-L8-F16 (B)
SIB 4	REV 07	710-005781	HR5903	SIB-L8-F16
B Board	REV 06	710-005782	HZ5275	SIB-L8-F16 (B)

### show chassis hardware scc (TX Matrix Router)

```
user@host> show chassis hardware scc
scc-re0:
```

```
-----
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Midplane         REV 04    710-004396   RB0014         SCC Midplane
FPM GBUS         REV 04    710-004617   HW9141         SCC FPM Board
FPM Display      REV 04    710-004619   HS5950         SCC FPM
CIP 0            REV 01    710-010218   HV9151         SCC CIP
CIP 1            REV 01    710-010218   HV9152         SCC CIP
PEM 1            Rev 11    740-002595   QB13977        Power Entry Module
Routing Engine 0 REV 05    740-008883   P11123900153   RE-4.0 (RE-1600)
CB 0             REV 01    710-011709   HR5964         Control Board (CB-TX)
SPMB 0           REV 09    710-003229   HW5293         T Series Switch CPU
SIB 3
SIB 4            REV 01    710-005839   HW1177         SIB-S8-F16
B Board          REV 01    710-005840   HW1202         SIB-S8-F16 (B)
```

### show chassis hardware (T1600 Router)

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item             Version  Part number  Serial number  Description
Chassis
Midplane         REV 03    710-005608   RC4137         T640 Backplane
FPM GBUS         REV 10    710-002901   DT7062         T640 FPM Board
FPM Display      REV 05    710-002897   DS3067         FPM Display
CIP              REV 06    710-002895   DT3386         T-series CIP
PEM 0            Rev 07    740-017906   UA26344        Power Entry Module 3x80
PEM 1            Rev 18    740-002595   UF38441        Power Entry Module
SCG 0            REV 15    710-003423   DV0941         T640 Sonet Clock Gen.
Routing Engine 0 REV 08    740-014082   9009014502     RE-A-2000
Routing Engine 1 REV 07    740-014082   9009009591     RE-A-2000
CB 0             REV 05    710-007655   JA9360         Control Board (CB-T)
CB 1             REV 03    710-017707   DT3251         Control Board (CB-T)
FPC 0            REV 07    710-013558   DR4253         E2-FPC Type 2
CPU              REV 05    710-013563   DS3902         FPC CPU-Enhanced
PIC 0            REV 01    750-010618   CB5446         4x G/E SFP, 1000 BASE
Xcvr 0           REV 01    740-011613   P9F11CW        SFP-SX
Xcvr 1           REV 01    740-011613   P9F15C2        SFP-SX
Xcvr 2           REV 01    740-011782   PB94K0L        SFP-SX
PIC 1            REV 06    750-001900   HB6399         1x OC-48 SONET, SMSR
PIC 2            REV 14    750-001901   AP1092         4x OC-12 SONET, SMIR
PIC 3            REV 07    750-001900   AR8275         1x OC-48 SONET, SMSR
MMB 1            REV 07    710-010171   DS1524         MMB-5M3-288mbit
FPC 1            REV 06    710-013553   DL9067         E2-FPC Type 1
CPU              REV 04    710-013563   DM1685         FPC CPU-Enhanced
PIC 0            REV 08    750-001072   AB1688         1x G/E, 1000 BASE-SX
PIC 1            REV 10    750-012266   JX5519         4x 1GE(LAN), IQ2
Xcvr 0           REV 01    740-011613   AM0812S8UK6    SFP-SX
Xcvr 2           REV 01    740-011613   AM0812S8UK1    SFP-SX
Xcvr 3           REV 01    740-011782   P8N1YHG        SFP-SX
PIC 2            REV 22    750-005634   DP0083         1x CHOC12 IQ SONET, SMIR
```

MMB 1	REV 07	710-008923	DN1862	MMB 3M 288-bit
FPC 2	REV 01	710-005548	HJ9899	FPC Type 3
CPU	REV 06	710-001726	HC0586	FPC CPU
PIC 0	REV 16	750-007141	NC9660	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8XAR	SFP-SX
Xcvr 1	REV 01	740-011782	P920E7B	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XAU	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8XAK	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8XAA	SFP-SX
Xcvr 6	REV 01	740-011613	PAJ4NKY	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UJW	SFP-SX
Xcvr 8	REV 01	740-011782	PB81X89	SFP-SX
Xcvr 9	REV 01	740-011613	AM0812S8UJX	SFP-SX
PIC 1	REV 06	750-015217	DK3280	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8P0A3T	SFP-SX
Xcvr 1	REV 01	740-013111	5090002	SFP-T
Xcvr 2	REV 01	740-011613	AM0814S93BQ	SFP-SX
Xcvr 4		NON-JNPR	PDE0FAN	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q20XY	SFP-SX
Xcvr 6	REV 01	740-011613	AM0812S8UJV	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8UP7	SFP-SX
PIC 2	REV 05	750-004695	HT4383	1x Tunnel
PIC 3	REV 17	750-009553	RL0204	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T23	SFP-SR
Xcvr 1	REV 01	740-011785	P6Q0F3E	SFP-SR
MMB 0	REV 03	710-004047	HD5843	MMB-288mbit
MMB 1	REV 03	710-004047	HE3208	MMB-288mbit
PPB 0	REV 02	710-002845	HA4524	PPB Type 3
PPB 1	REV 02	710-002845	HA4766	PPB Type 3
FPC 3	REV 01	710-010154	HR0863	E-FPC Type 3
CPU	REV 01	710-010169	HN3422	FPC CPU-Enhanced
PIC 0	REV 07	750-012793	WF5096	1x 10GE(LAN/WAN) IQ2
Xcvr 0		NON-JNPR	M64294TP	XFP-10G-LR
PIC 1	REV 25	750-007141	DV2127	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011613	PFA6LTJ	SFP-SX
Xcvr 1	REV 01	740-011782	P9P0XV4	SFP-SX
Xcvr 2	REV 01	740-011782	P9M0TNX	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0TTP	SFP-SX
Xcvr 5		NON-JNPR	PBS4LED	SFP-SX
PIC 2	REV 17	750-009553	RL0212	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	PDS3T8G	SFP-SR
PIC 3	REV 32	750-003700	DL1279	1x OC-192 12xMM VSR
MMB 0	REV 01	710-010171	HR0821	MMB-288mbit
MMB 1	REV 01	710-010171	HR0818	MMB-288mbit
FPC 4	REV 16	710-013037	EB4919	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA4382	ST-PMB2
PIC 0	REV 03	711-029996	EB1569	100GE
PIC 1	REV 05	711-029999	EB9983	100GE CFP
Xcvr 0	REV 0	740-032210	J10G80746	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2235	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA7112	ST-MMB2
MMB 1	REV 04	710-025563	BBAA7149	ST-MMB2
FPC 5	REV 02	710-013037	DE3407	FPC Type 4-ES
CPU	REV 04	710-016744	DA2124	ST-PMB2
PIC 0	REV 16	750-012518	DF2554	4x OC-192 SONET XFP
Xcvr 0	REV 01	740-014279	AA0745N1FX8	XFP-OC192-SR
Xcvr 1	REV 01	740-014279	AA0748N1HN5	XFP-OC192-SR
Xcvr 2	REV 01	740-014279	AA0748N1HT6	XFP-OC192-SR

Xcvr 3	REV 01	740-014279	AA0744N1EC9	XFP-OC192-SR
PIC 1	REV 01	750-010850	JA0329	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DE9577	ST-MMB2
MMB 1	REV 04	710-016036	DK4060	ST-MMB2
FPC 6	REV 14	710-013037	DV1431	FPC Type 4-ES
CPU	REV 09	710-016744	DT9020	ST-PMB2
PIC 0	REV 11	750-017405	DM6261	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014289	C701XU05Q	XFP-10G-SR
Xcvr 1	REV 01	740-014279	AA0748N1HPT	XFP-10G-LR
Xcvr 2	REV 01	740-014289	T08E19189	XFP-10G-SR
Xcvr 3	REV 01	740-014289	C715XU058	XFP-10G-SR
PIC 1	REV 13	750-017405	DP8772	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 02	740-011571	C850XJ037	XFP-10G-SR
Xcvr 1	REV 02	740-014289	C839XU0L9	XFP-10G-SR
Xcvr 2	REV 02	740-014289	C834XU05A	XFP-10G-SR
Xcvr 3	REV 02	740-014289	C810XU0CE	XFP-10G-SR
MMB 0	REV 01	710-025563	DT8454	ST-MMB2
MMB 1	REV 01	710-025563	DT8366	ST-MMB2
FPC 7	REV 09	710-007529	HZ7624	FPC Type 3
CPU	REV 15	710-001726	HZ1413	FPC CPU
PIC 0	REV 10	750-012793	DM5627	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 02	740-011571	C831XJ062	XFP-10G-SR
PIC 1	REV 01	750-015217	JT6762	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q25JU	SFP-SX
Xcvr 1	REV 01	740-011782	P9B0U0K	SFP-SX
PIC 2	REV 01	750-015217	JS4268	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8XBZ	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAP	SFP-SX
Xcvr 2	REV 01	740-011613	AM0812S8XBY	SFP-SX
Xcvr 3	REV 01	740-011613	AM0812S8XBX	SFP-SX
Xcvr 4	REV 01	740-011613	P9F1652	SFP-SX
Xcvr 5	REV 01	740-011782	P8Q21YC	SFP-SX
Xcvr 6	REV 01	740-011782	P8Q27HQ	SFP-SX
Xcvr 7	REV 01	740-011613	P8E2SSU	SFP-SX
PIC 3	REV 15	750-009450	NB6790	1x OC-192 SM SR2
MMB 0	REV 03	710-005555	HZ3450	MMB-288mbit
MMB 1	REV 03	710-005555	HZ3415	MMB-288mbit
PPB 0	REV 04	710-002845	HP0887	PPB Type 3
PPB 1	REV 04	710-002845	HW5255	PPB Type 3
SPMB 0	REV 10	710-003229	HX3699	T-series Switch CPU
SPMB 1	REV 12	710-003229	DT3091	T-series Switch CPU
SIB 0	REV 07	710-013074	DS4747	SIB-I8-SF
SIB 1	REV 07	710-013074	DS4942	SIB-I8-SF
SIB 2	REV 07	710-013074	DS4965	SIB-I8-SF
SIB 3	REV 07	710-013074	DS4990	SIB-I8-SF
SIB 4	REV 07	710-013074	DS4944	SIB-I8-SF
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 2

### show chassis hardware (TX Matrix Plus Router)

```
user@host> show chassis hardware
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN113186EAHB	TXP
Midplane	REV 05	710-022574	TS3822	SFC Midplane
FPM Display	REV 03	710-024027	DW4701	TXP FPM Display
CIP 0	REV 05	710-023792	DW7998	TXP CIP

CIP 1	REV 05	710-023792	DW7999	TXP CIP
PEM 0	Rev 04	740-027463	UM26367	Power Entry Module
PEM 1	Rev 04	740-027463	UM26346	Power Entry Module
Routing Engine 0	REV 06	740-026942	737A-1081	RE-DUO-2600
Routing Engine 1	REV 06	740-026942	737A-1043	RE-DUO-2600
CB 0	REV 05	710-022606	DW4435	SFC Control Board
CB 1	REV 09	710-022606	DW6100	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	750-024564	DW5764	F13 SIB
B Board	REV 03	710-023431	DW9053	F13 SIB Mezz
SIB F13 3	REV 04	750-024564	DW5785	F13 SIB
B Board	REV 03	710-023431	DW9030	F13 SIB Mezz
SIB F13 6				
SIB F13 8	REV 04	750-024564	DW5752	F13 SIB
B Board	REV 03	710-023431	DW9051	F13 SIB Mezz
SIB F13 11	REV 04	750-024564	DW5782	F13 SIB
B Board	REV 03	710-023431	DW9058	F13 SIB Mezz
SIB F13 12	REV 03	750-024564	DT9466	F13 SIB
B Board	REV 02	710-023431	DT6556	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7898	F2S SIB
B Board	REV 05	710-023787	DW7625	F2S SIB Mezz
SIB F2S 0/2	REV 05	710-022603	DW7811	F2S SIB
B Board	REV 05	710-023787	DW7550	F2S SIB Mezz
SIB F2S 0/4	REV 04	710-022603	DW4873	F2S SIB
B Board	REV 05	710-023787	DW8509	F2S SIB Mezz
SIB F2S 0/6	REV 04	710-022603	DW4867	F2S SIB
B Board	REV 05	710-023787	DW8472	F2S SIB Mezz
SIB F2S 1/0	REV 04	710-022603	DW4871	F2S SIB
B Board	REV 05	710-023787	DW8497	F2S SIB Mezz
SIB F2S 1/2	REV 05	710-022603	DW7868	F2S SIB
B Board	REV 05	710-023787	DW7551	F2S SIB Mezz
SIB F2S 1/4	REV 04	710-022603	DW4854	F2S SIB
B Board	REV 05	710-023787	DW8496	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7889	F2S SIB
B Board	REV 05	710-023787	DW7496	F2S SIB Mezz
SIB F2S 2/0	REV 04	710-022603	DW4852	F2S SIB
B Board	REV 05	710-023787	DW8498	F2S SIB Mezz
SIB F2S 2/2	REV 04	710-022603	DW4845	F2S SIB
B Board	REV 05	710-023787	DW8457	F2S SIB Mezz
SIB F2S 2/4	REV 05	710-022603	DW7802	F2S SIB
B Board	REV 05	710-023787	DW7562	F2S SIB Mezz
SIB F2S 2/6	REV 04	710-022603	DW4822	F2S SIB
B Board	REV 05	710-023787	DW8467	F2S SIB Mezz
SIB F2S 3/0	REV 05	710-022603	DW7815	F2S SIB
B Board	REV 05	710-023787	DW7518	F2S SIB Mezz
SIB F2S 3/2	REV 03	710-022603	DV0068	F2S SIB
B Board	REV 03	710-023787	DT9974	F2S SIB Mezz
SIB F2S 3/4	REV 05	710-022603	DW7874	F2S SIB
B Board	REV 05	710-023787	DW7601	F2S SIB Mezz
SIB F2S 3/6	REV 03	710-022603	DV0033	F2S SIB
B Board	REV 03	710-023787	DT9969	F2S SIB Mezz
SIB F2S 4/0	REV 03	710-022603	DV0043	F2S SIB
B Board	REV 03	710-023787	DT9948	F2S SIB Mezz
SIB F2S 4/2	REV 05	710-022603	DW5446	F2S SIB
B Board	REV 05	710-023787	DW7611	F2S SIB Mezz
SIB F2S 4/4	REV 04	710-022603	DW4826	F2S SIB
B Board	REV 05	710-023787	DW8458	F2S SIB Mezz
SIB F2S 4/6	REV 03	710-022603	DV0026	F2S SIB
B Board	REV 03	710-023787	DT9963	F2S SIB Mezz
Fan Tray 0	REV 02	760-024497	DR8290	Front Fan Tray

Fan Tray 1	REV 02	760-024497	DR8293	Front Fan Tray
Fan Tray 2	REV 05	760-024502	DR8280	Rear Fan Tray
Fan Tray 3				
Fan Tray 4	REV 05	760-024502	DR8276	Rear Fan Tray
Fan Tray 5	REV 02	760-024502	DP5643	Rear Fan Tray

lcc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11036F8AHA	T1600
Midplane	REV 03	710-017247	RC3799	T-series Backplane
FPM GBUS	REV 10	710-002901	DP7009	T640 FPM Board
FPM Display	REV 01	710-021387	DN7026	T1600 FPM Display
CIP	REV 06	710-002895	DP6024	T-series CIP
PEM 1	Rev 02	740-023211	WA50019	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DR6757	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DS2225	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1040	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1016	RE-DUO-1800
CB 0	REV 06	710-022597	DX4011	LCC Control Board
CB 1	REV 06	710-022597	DX4017	LCC Control Board
FPC 1	REV 07	710-013035	DN5847	FPC Type 3-ES
CPU	REV 08	710-016744	DP2570	ST-PMB2
PIC 0	REV 05	750-015217	DB0418	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P8Q27ZG	SFP-SX
Xcvr 1		NON-JNPR	PDA1U0D	SFP-SX
Xcvr 2	REV 01	740-011613	P9F1ALW	SFP-SX
Xcvr 3	REV 01	740-011782	PBA403V	SFP-SX
Xcvr 4		NON-JNPR	PDE09DP	SFP-SX
Xcvr 5	REV 01	740-011782	PCH2P4K	SFP-SX
Xcvr 6	REV 01	740-011782	PB94K0F	SFP-SX
Xcvr 7	REV 01	740-011782	PBA2R2A	SFP-SX
PIC 1	REV 03	750-004424	HJ4020	1x 10GE(LAN),DWDM
PIC 2	REV 01	750-003336	HG6073	4x OC-48 SONET, SMSR
MMB 0	REV 04	710-016036	DP3401	ST-MMB2
FPC 3	REV 12	710-013037	DR1169	FPC Type 4-ES
CPU	REV 08	710-016744	DP9429	ST-PMB2
PIC 0	REV 02	750-010850	JA0332	1x OC-768 SONET SR
MMB 0	REV 04	710-016036	DR0628	ST-MMB2
MMB 1	REV 04	710-016036	DR0592	ST-MMB2
FPC 4	REV 05	710-021534	DR7350	FPC Type 1-ES
CPU	REV 08	710-016744	DP8096	ST-PMB2
PIC 0	REV 04	750-014627	DP9171	4x OC-3 1x OC-12 SFP
Xcvr 0	REV 02	740-011615	PDE2RVR	SFP-SR
PIC 1	REV 22	750-005634	DS5815	1x CHOC12 IQ SONET, SMIR
PIC 2	REV 09	750-002911	CF4539	4x F/E, 100 BASE-TX
PIC 3	REV 08	750-021652	DR2827	1x CHOC12 IQE SONET
Xcvr 0		NON-JNPR	8	UNKNOWN
MMB 0	REV 04	710-016036	DR0809	ST-MMB2
FPC 5	REV 07	710-007529	HS5608	FPC Type 3
CPU	REV 15	710-001726	HX4351	FPC CPU
PIC 0	REV 14	750-009567	WJ8961	1x 10GE(LAN),XENPAK
Xcvr 0	REV 01	740-013170	J05K05961	XENPAK-LR
PIC 1	REV 16	750-007141	JJ8146	10x 1GE(LAN), 1000 BASE
Xcvr 1	REV 01	740-011613	P9F117T	SFP-SX
Xcvr 2	REV 01	740-011782	PBA2VCL	SFP-SX
Xcvr 3	REV 01	740-011782	PB83DRB	SFP-SX
Xcvr 4	REV 01	740-011613	AM0812S8UP8	SFP-SX

PIC 2	REV 12	750-009567	WF3566	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T07C94489	XENPAK-LR
MMB 0	REV 03	710-005555	HZ1907	MMB-288mbit
MMB 1	REV 03	710-005555	HW5283	MMB-288mbit
PPB 0	REV 04	710-002845	HZ7717	PPB Type 3
PPB 1	REV 04	710-002845	HS0110	PPB Type 3
FPC 6	REV 07	710-013035	DP7486	FPC Type 3-ES
CPU	REV 08	710-016744	DP2545	ST-PMB2
PIC 0	REV 09	750-009567	NE6323	1x 10GE(LAN), XENPAK
Xcvr 0	REV 02	740-013170	T09C71959	XENPAK-LR
PIC 1	REV 06	750-015217	DN4775	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011782	P7E0T6M	SFP-SX
Xcvr 1	REV 01	740-011613	AM0812S8XAY	SFP-SX
Xcvr 2	REV 01	740-011782	P7E0T6J	SFP-SX
Xcvr 3	REV 01	740-011782	PCH2P7D	SFP-SX
Xcvr 4	REV 01	740-011782	P9B0QYT	SFP-SX
Xcvr 5	REV 01	740-011613	AM0812S8WQJ	SFP-SX
Xcvr 6	REV 02	740-013111	9301220	SFP-T
Xcvr 7	REV 01	740-011782	P9B0TZ5	SFP-SX
PIC 2	REV 06	750-015217	DM6747	8x 1GE(TYPE3), IQ2
Xcvr 0	REV 01	740-011613	PAP0ZB2	SFP-SX
Xcvr 1	REV 01	740-013111	70191002	SFP-T
Xcvr 6	REV 01	740-011782	PBA29H8	SFP-SX
Xcvr 7	REV 01	740-011613	AM0812S8WQG	SFP-SX
MMB 0	REV 04	710-016036	DP3238	ST-MMB2
FPC 7	REV 03	710-021540	DV3154	FPC Type 2-ES
CPU	REV 09	710-016744	DT9053	ST-PMB2
PIC 0	REV 13	750-001901	HB4225	4x OC-12 SONET, SMIR
PIC 1	REV 05	750-001900	AD3644	1x OC-48 SONET, SMSR
PIC 2	REV 10	750-008155	HV0335	2x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011782	PCH2UKF	SFP-SX
Xcvr 1	REV 01	740-011782	PCH2V19	SFP-SX
PIC 3	REV 03	750-014638	JS9493	1x OC-48-12-3 SFP
Xcvr 0	REV 01	740-011785	P6Q0ENK	SFP-SR
MMB 0	REV 05	710-016036	DP3323	ST-MMB2
SPMB 0	REV 04	710-023321	DX3004	LCC Switch CPU
SPMB 1	REV 04	710-023321	DX3009	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4195	LCC SIB
B Board	REV 07	710-023185	DW3930	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4179	LCC SIB
B Board	REV 07	710-023185	DW3919	LCC SIB Mezz
SIB 2				
SIB 3	REV 06	710-022594	DT8251	LCC SIB
B Board	REV 06	710-023185	DT5792	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8014	LCC SIB
B Board	REV 07	710-023185	DW3917	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

lcc1-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1102270AHA	T1600
Midplane	REV 04	710-017247	RC5358	T-series Backplane
FPM GBUS	REV 10	710-002901	DS3443	T640 FPM Board
FPM Display	REV 01	710-021387	DS6411	T1600 FPM Display
CIP	REV 06	710-002895	DS4235	T-series CIP
PEM 0	Rev 02	740-023211	VM82438	Power Entry Module 4x60A
SCG 0	REV 15	710-003423	DS6649	T640 Sonet Clock Gen.

SCG 1	REV 15	710-003423	DR6775	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026941	737F-1083	RE-DUO-1800
Routing Engine 1	REV 01	740-026941	737F-1104	RE-DUO-1800
CB 0	REV 06	710-022597	DW8542	LCC Control Board
CB 1	REV 06	710-022597	DW8530	LCC Control Board
FPC 0	REV 02	710-010845	JE2392	FPC Type 4
CPU	REV 02	710-011481	JF6820	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP7259	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 01	740-014279	AA0741N1C8T	XFP-10G-LR
Xcvr 1	REV 01	740-014279	AA0746N1GAM	XFP-10G-LR
Xcvr 2	REV 01	740-014279	AA0747N1H0B	XFP-10G-LR
Xcvr 3	REV 01	740-014279	AA0748N1HZ5	XFP-10G-LR
MMB 0	REV 03	710-010842	HY7601	ST-MMB
FPC 1	REV 16	710-013037	BBAA7398	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA2329	ST-PMB2
PIC 0	REV 03	711-029996	EB1575	100GE
PIC 1	REV 06	750-034781	EB9980	100GE CFP
MMB 0	REV 04	710-025563	BBAA5325	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5444	ST-MMB2
FPC 2	REV 16	710-013037	BBAA7185	FPC Type 4-ES
CPU	REV 09	710-016744	BBAA3522	ST-PMB2
PIC 0	REV 03	711-029996	EB1557	100GE
PIC 1	REV 05	750-034781	EB4660	100GE CFP
Xcvr 0	REV 0	740-032210	J10F73666	CFP-100G-LR4
BRIDGE 0	REV 02	711-029995	EB2237	100GE Bridge Board
MMB 0	REV 04	710-025563	BBAA5347	ST-MMB2
MMB 1	REV 04	710-025563	BBAA5401	ST-MMB2
FPC 3	REV 10	710-021534	DZ0941	FPC Type 1-ES
CPU	REV 09	710-016744	DY6364	ST-PMB2
PIC 0	REV 13	750-012266	DK9192	4x 1GE(LAN), IQ2
Xcvr 0	REV 01	740-011613	AM0812S8WVD	SFP-SX
Xcvr 1		NON-JNPR	PDD63Q4	SFP-SX
Xcvr 2		NON-JNPR	PDE4G54	SFP-SX
Xcvr 3		NON-JNPR	PD4OMAG	SFP-SX
PIC 1	REV 01	750-007641	HJ2003	1x G/E IQ, 1000 BASE
Xcvr 0	REV 01	740-011613	AM0812S8WVG	SFP-SX
PIC 3	REV 17	750-007444	JB6873	1x CHSTM1 IQ SDH, SMIR
MMB 0	REV 04	710-025563	DZ0281	ST-MMB2
FPC 4	REV 06	710-013035	DK0614	FPC Type 3-ES
CPU	REV 07	710-016744	DK1616	ST-PMB2
PIC 0	REV 22	750-007141	DM1870	10x 1GE(LAN), 1000 BASE
Xcvr 0	REV 01	740-011782	PCL3UKW	SFP-SX
Xcvr 1	REV 01	740-011782	P7E0T73	SFP-SX
Xcvr 2	REV 01	740-007326	P4TOWLR	SFP-SX
Xcvr 3	REV 01	740-011782	PAR1LLRL	SFP-SX
Xcvr 4	REV 01	740-011782	P9M0U3Z	SFP-SX
Xcvr 5	REV 01	740-011782	P9M0U0C	SFP-SX
Xcvr 6	REV 01	740-011782	P9M0TLG	SFP-SX
Xcvr 7	REV 01	740-011782	P9M0U0F	SFP-SX
Xcvr 8	REV 01	740-011613	PFA6LAP	SFP-SX
Xcvr 9	REV 01	740-011782	PCH2P0U	SFP-SX
PIC 1	REV 16	750-009450	CV2565	1x OC-192 SM SR2
PIC 2	REV 05	750-004424	HH3057	1x 10GE(LAN), 10GBASE-LR
PIC 3	REV 12	750-013423	DP0403	MultiServices 500
MMB 0	REV 04	710-016036	DK1988	ST-MMB2
FPC 5	REV 07	710-013560	DR0004	E2-FPC Type 3
CPU	REV 05	710-013563	DR0089	FPC CPU-Enhanced
PIC 0	REV 11	750-012793	DR6107	1x 10GE(LAN/WAN) IQ2
Xcvr 0	REV 01	740-014289	C743XU074	XFP-10G-SR

PIC 1	REV 01	750-004695	HD5980	1x Tunnel
PIC 2	REV 32	750-003700	DL3770	1x OC-192 12xMM VSR
PIC 3	REV 12	750-009553	WB8901	4x OC-48 SONET
Xcvr 0	REV 01	740-011785	P9D1GTQ	SFP-SR
Xcvr 1	REV 01	740-011785	PDSOMMB	SFP-SR
Xcvr 3	REV 01	740-011785	PDE1KXP	SFP-SR
MMB 0	REV 07	710-010171	DP7374	MMB-5M3-288mbit
MMB 1	REV 07	710-010171	DP7404	MMB-5M3-288mbit
FPC 6	REV 07	710-013035	DM0994	FPC Type 3-ES
CPU	REV 07	710-016744	DM3651	ST-PMB2
PIC 0	REV 07	750-015217	DN4743	8x 1GE(TYPE3), IQ2
Xcvr 3	REV 01	740-011613	AM0812S8XB0	SFP-SX
Xcvr 4	REV 01	740-011782	PB829RB	SFP-SX
Xcvr 5	REV 01	740-011782	P8J1SYX	SFP-SX
PIC 1	REV 03	750-003336	HJ9954	4x OC-48 SONET, SMSR
PIC 3	REV 02	750-012793	JM7665	1x 10GE(LAN/WAN) IQ2
MMB 0	REV 04	710-016036	DN6913	ST-MMB2
FPC 7	REV 08	710-010845	JM3958	FPC Type 4
CPU	REV 04	710-011481	JK3669	FPC CPU-Enhanced
PIC 0	REV 11	750-017405	DP8837	4x 10GE (LAN/WAN) XFP
Xcvr 1	REV 01	740-014279	753019A00277	XFP-10G-LR
Xcvr 2	REV 02	740-011571	C850XJ00P	XFP-10G-SR
Xcvr 3	REV 01	740-014279	AA0813N1RTG	XFP-10G-LR
MMB 0	REV 04	710-010842	JN1971	ST-MMB
SPMB 0	REV 04	710-023321	DW3629	LCC Switch CPU
SPMB 1	REV 04	710-023321	DW3621	LCC Switch CPU
SIB 0	REV 07	710-022594	DW4200	LCC SIB
B Board	REV 07	710-023185	DW3932	LCC SIB Mezz
SIB 1	REV 07	710-022594	DW4193	LCC SIB
B Board	REV 07	710-023185	DW3904	LCC SIB Mezz
SIB 2				
SIB 3	REV 07	710-022594	DW4210	LCC SIB
B Board	REV 06	710-023185	DT5780	LCC SIB Mezz
SIB 4	REV 08	710-022594	DW8019	LCC SIB
B Board	REV 06	710-023185	DT5795	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 3

### show chassis hardware sfc (TX Matrix Plus Router)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP
Midplane	REV 05	710-022574	TS4027	SFC Midplane
FPM Display	REV 03	710-024027	DX0282	TXP FPM Display
CIP 0	REV 04	710-023792	DW4889	TXP CIP
CIP 1	REV 04	710-023792	DW4887	TXP CIP
PEM 0	Rev 07	740-027463	UM26368	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1064	SFC RE
Routing Engine 1	REV 01	740-026942	737A-1082	SFC RE
CB 0	REV 09	710-022606	DW6099	SFC Control Board
CB 1	REV 09	710-022606	DW6096	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 04	710-022600	DX0841	F13 SIB
B Board	REV 03	710-023431	DX0966	F13 SIB Mezz
SIB F13 1	REV 04	750-024564	DW5776	F13 SIB



B Board	REV 03	710-023431	DW9028	F13 SIB
SIB F13 3	REV 04	750-024564	DW5762	F13 SIB
B Board	REV 03	710-023431	DW9059	F13 SIB
SIB F13 4	REV 04	750-024564	DW5797	F13 SIB
B Board	REV 03	710-023431	DW9041	F13 SIB
SIB F13 6	REV 04	750-024564	DW5770	F13 SIB
B Board	REV 03	710-023431	DW9079	F13 SIB Mezz
SIB F13 7	REV 04	750-024564	DW5758	F13 SIB
B Board	REV 03	710-023431	DW9047	F13 SIB
SIB F13 8	REV 04	750-024564	DW5761	F13 SIB
B Board	REV 03	710-023431	DW9043	F13 SIB Mezz
SIB F13 9	REV 04	750-024564	DW5754	F13 SIB
B Board	REV 03	710-023431	DW9078	F13 SIB Mezz
SIB F13 11	REV 04	710-022600	DX0826	F13 SIB
B Board	REV 03	710-023431	DX0967	F13 SIB Mezz
SIB F13 12	REV 04	750-024564	DW5794	F13 SIB
B Board	REV 03	710-023431	DW9044	F13 SIB Mezz
SIB F2S 0/0	REV 05	710-022603	DW7897	F2S SIB
B Board	REV 05	710-023787	DW7657	NEO PMB
SIB F2S 0/2	REV 05	710-022603	DW7833	F2S SIB
B Board	REV 05	710-023787	DW7526	NEO PMB
SIB F2S 0/4	REV 05	710-022603	DW7875	F2S SIB
B Board	REV 05	710-023787	DW7588	NEO PMB
SIB F2S 0/6	REV 05	710-022603	DW7860	F2S SIB
B Board	REV 05	710-023787	DW7589	NEO PMB
SIB F2S 1/0	REV 04	710-022603	DW4820	F2S SIB
B Board	REV 05	710-023787	DW8510	NEO PMB
SIB F2S 1/2	REV 05	710-022603	DW7849	F2S SIB
B Board	REV 05	710-023787	DW7525	NEO PMB
SIB F2S 1/4	REV 05	710-022603	DW7927	F2S SIB
B Board	REV 05	710-023787	DW7556	F2S SIB Mezz
SIB F2S 1/6	REV 05	710-022603	DW7866	F2S SIB
B Board	REV 05	710-023787	DW7651	NEO PMB
SIB F2S 2/0	REV 05	710-022603	DW7880	F2S SIB
B Board	REV 05	710-023787	DW7523	NEO PMB
SIB F2S 2/2	REV 05	710-022603	DW7895	F2S SIB
B Board	REV 05	710-023787	DW7591	NEO PMB
SIB F2S 2/4	REV 05	710-022603	DW7907	F2S SIB
B Board	REV 05	710-023787	DW7590	NEO PMB
SIB F2S 2/6	REV 05	710-022603	DW7785	F2S SIB
B Board	REV 05	710-023787	DW7524	NEO PMB
SIB F2S 3/0	REV 05	710-022603	DW7782	F2S SIB
B Board	REV 05	710-023787	DW7634	NEO PMB
SIB F2S 3/2	REV 05	710-022603	DW7793	F2S SIB
B Board	REV 05	710-023787	DW7548	NEO PMB
SIB F2S 3/4	REV 05	710-022603	DW7779	F2S SIB
B Board	REV 05	710-023787	DW7587	NEO PMB
SIB F2S 3/6	REV 05	710-022603	DW7930	F2S SIB
B Board	REV 05	710-023787	DW7505	NEO PMB
SIB F2S 4/0	REV 05	710-022603	DW7867	F2S SIB
B Board	REV 05	710-023787	DW7656	NEO PMB
SIB F2S 4/2	REV 05	710-022603	DW7917	F2S SIB
B Board	REV 05	710-023787	DW7640	NEO PMB
SIB F2S 4/4	REV 05	710-022603	DW7929	F2S SIB
B Board	REV 05	710-023787	DW7643	NEO PMB
SIB F2S 4/6	REV 05	710-022603	DW7870	F2S SIB
B Board	REV 05	710-023787	DW7635	NEO PMB
Fan Tray 0	REV 06	760-024497	DV7831	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9614	Front Fan Tray
Fan Tray 2	REV 06	760-024502	DV9618	Rear Fan Tray
Fan Tray 3	REV 06	760-024502	DV9616	Rear Fan Tray

Fan Tray 4	REV 06	760-024502	DV7807	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7828	Rear Fan Tray

### show chassis hardware extensive (TX Matrix Plus Router)

```
user@host> show chassis hardware extensive
sfc0-re0:
```

#### ----- Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN112F007AHB	TXP
Jedec Code:	0x7fb0		EEPROM Version:	0x02
			S/N:	JN112F007AHB
Assembly ID:	0x052c		Assembly Version:	00.00
Date:	00-00-0000		Assembly Flags:	0x00

ID: TXP

#### Board Information Record:

Address 0x00: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

#### I2C Hex Data:

Address 0x00: 7f b0 02 ff 05 2c 00 00 00 00 00 00 00 00 00 00  
 Address 0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x20: 4a 4e 31 31 32 46 30 30 37 41 48 42 00 00 00 00  
 Address 0x30: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00  
 Address 0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Midplane	REV 05	710-022574	TS4027	SFC Midplane
----------	--------	------------	--------	--------------

Jedec Code:	0x7fb0		EEPROM Version:	0x01
P/N:	710-022574		S/N:	TS4027
Assembly ID:	0x0962		Assembly Version:	01.05
Date:	03-23-2009		Assembly Flags:	0x00
Version:	REV 05			

ID: SFC Midplane

#### Board Information Record:

Address 0x00: ad 01 ff ff 00 1d b5 14 00 00 ff ff ff ff ff ff

#### I2C Hex Data:

Address 0x00: 7f b0 01 ff 09 62 01 05 52 45 56 20 30 35 00 00  
 Address 0x10: 00 00 00 00 37 31 30 2d 30 32 32 35 37 34 00 00  
 Address 0x20: 53 2f 4e 20 54 53 34 30 32 37 00 00 00 17 03 07  
 Address 0x30: d9 ff ff ff ad 01 ff ff 00 1d b5 14 00 00 ff ff  
 Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff  
 Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
 Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
 Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

FPM Display	REV 03	710-024027	DX0282	TXP FPM Display
-------------	--------	------------	--------	-----------------

Jedec Code:	0x7fb0		EEPROM Version:	0x01
P/N:	710-024027		S/N:	DX0282
Assembly ID:	0x096c		Assembly Version:	01.03
Date:	02-10-2009		Assembly Flags:	0x00
Version:	REV 03			

ID: TXP FPM Display      FRU Model Number: CRAFT-TXP

#### Board Information Record:

Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

#### I2C Hex Data:

Address 0x00: 7f b0 01 ff 09 6c 01 03 52 45 56 20 30 33 00 00  
 Address 0x10: 00 00 00 00 37 31 30 2d 30 32 34 30 32 37 00 00  
 Address 0x20: 53 2f 4e 20 44 58 30 32 38 32 00 00 00 0a 02 07  
 Address 0x30: d9 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff  
 Address 0x40: ff ff ff ff 01 00 00 00 00 00 00 00 00 00 00 43  
 Address 0x50: 52 41 46 54 2d 54 58 50 00 00 00 00 00 00 00 00

```

Address 0x60: 00 00 00 00 00 00 ff ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
CIP 0          REV 04   710-023792   DW4889          TXP CIP
Jedec Code:    0x7fb0          EEPROM Version:    0x01
P/N:           710-023792      S/N:              DW4889
Assembly ID:   0x0969          Assembly Version:  01.04
Date:          01-26-2009      Assembly Flags:    0x00
Version:       REV 04
ID: TXP CIP          FRU Model Number: CIP-TXP
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

### show chassis hardware clei-models (TX Matrix Plus Router)

```

user@host> show chassis hardware clei-models
sfc0-re0:

```

```

-----
Hardware inventory:

```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 05	710-022574		CHAS-BP-TXP-S
FPM Display	REV 03	710-024027		CRAFT-TXP-S
CIP 0	REV 05	710-023792		CIP-TXP-S
CIP 1	REV 05	710-023792		CIP-TXP-S
PEM 0	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
PEM 1	Rev 04	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC
Routing Engine 0	REV 06	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 05	710-022606		CB-TXP-S
CB 1	REV 09	710-022606		CB-TXP-S
SIB F13 0	REV 04	750-024564		SIB-TXP-F13
SIB F13 3	REV 04	750-024564		SIB-TXP-F13
SIB F13 8	REV 04	750-024564		SIB-TXP-F13
SIB F13 11	REV 04	750-024564		SIB-TXP-F13
SIB F13 12	REV 03	750-024564		SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 0/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 0/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 1/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 1/6	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/0	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/2	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 2/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 2/6	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 3/0	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 3/2	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 3/4	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 3/6	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 4/0	REV 03	710-022603		SIB-TXP-F2S-S
SIB F2S 4/2	REV 05	710-022603		SIB-TXP-F2S-S
SIB F2S 4/4	REV 04	710-022603		SIB-TXP-F2S-S
SIB F2S 4/6	REV 03	710-022603		SIB-TXP-F2S-S
Fan Tray 0	REV 02	760-024497		FANTRAY-TXP-H-S
Fan Tray 1	REV 02	760-024497		FANTRAY-TXP-H-S
Fan Tray 2	REV 05	760-024502		FANTRAY-TXP-V-S
Fan Tray 3				
Fan Tray 4	REV 05	760-024502		FANTRAY-TXP-V-S
Fan Tray 5	REV 02	760-024502		FANTRAY-TXP-V-S

## lcc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 03	710-017247		CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 1	Rev 02	740-023211	IPUPAC8KTA	PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 1	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 05	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-004424		PC-1XGE-LR
PIC 2	REV 01	750-003336		PC-40C48-SON-SMSR
FPC 3	REV 12	710-013037		T1600-FPC4-ES
PIC 0	REV 02	750-010850		PD-10C768-SON-SR
FPC 4	REV 05	710-021534		T640-FPC1-ES
PIC 0	REV 04	750-014627		PB-40C3-10C12-SON-SFP
PIC 1	REV 22	750-005634		PB-1CHOC12SMIR-QPP
PIC 2	REV 09	750-002911		PB-4FE-TX
PIC 3	REV 08	750-021652		PB-1CHOC12-STM4-IQE-SFP
FPC 5	REV 07	710-007529		T640-FPC3
PIC 0	REV 14	750-009567		PC-1XGE-XENPAK
PIC 1	REV 16	750-007141		PC-10GE-SFP
PIC 2	REV 12	750-009567		PC-1XGE-XENPAK
FPC 6	REV 07	710-013035		T640-FPC3-ES
PIC 0	REV 09	750-009567		PC-1XGE-XENPAK
PIC 1	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
PIC 2	REV 06	750-015217		PC-8GE-TYPE3-SFP-IQ2
FPC 7	REV 03	710-021540		T640-FPC2-ES
PIC 0	REV 13	750-001901		PB-40C12-SON-SMIR
PIC 1	REV 05	750-001900		PB-10C48-SON-SMSR
PIC 2	REV 10	750-008155		PB-2GE-SFP-QPP
PIC 3	REV 03	750-014638		PB-10C48-SON-B-SFP
SIB 0	REV 07	710-022594		SIB-TXP-T1600-S
SIB 1	REV 07	710-022594		SIB-TXP-T1600-S
SIB 3	REV 06	710-022594		SIB-TXP-T1600-S
SIB 4	REV 08	710-022594		SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

## lcc1-re0:

-----  
Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 04	710-017247		CHAS-BP-T1600-S
FPM Display	REV 01	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	Rev 02	740-023211	IPUPAC8KTA	PWR-T1600-4-60-DC-S
SCG 0	REV 15	710-003423		SCG-T-S
SCG 1	REV 15	710-003423		SCG-T-S
Routing Engine 0	REV 01	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 01	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 06	710-022597		CB-LCC-S
CB 1	REV 06	710-022597		CB-LCC-S
FPC 0	REV 02	710-010845		T640-FPC4-ES

PIC 0	REV 11	750-017405	PD-4XGE-XFP
FPC 1	REV 16	710-013037	T1600-FPC4-ES
PIC 1	REV 06	750-034781	PD-1CE-CFP
FPC 2	REV 16	710-013037	T1600-FPC4-ES
PIC 1	REV 05	750-034781	PD-1CE-CFP
FPC 3	REV 10	710-021534	T640-FPC1-ES
PIC 0	REV 13	750-012266	PB-4GE-TYPE1-SFP-IQ2
PIC 1	REV 01	750-007641	PE-1GE-SFP-QPP
PIC 3	REV 17	750-007444	PB-1CHSTM1-SMIR-QPP
FPC 4	REV 06	710-013035	T640-FPC3-ES
PIC 0	REV 22	750-007141	PC-10GE-SFP
PIC 1	REV 16	750-009450	PC-10C192-SON-SR2
PIC 2	REV 05	750-004424	PC-1XGE-LR
PIC 3	REV 12	750-013423	PC-MS-500-3
FPC 5	REV 07	710-013560	T640-FPC3-E2
PIC 0	REV 11	750-012793	PC-1XGE-TYPE3-XFP-IQ2
PIC 1	REV 01	750-004695	PC-TUNNEL
PIC 2	REV 32	750-003700	PC-10C192-SON-VSR
PIC 3	REV 12	750-009553	PC-40C48-SON-SFP
FPC 6	REV 07	710-013035	T640-FPC3-ES
PIC 0	REV 07	750-015217	PC-8GE-TYPE3-SFP-IQ2
PIC 1	REV 03	750-003336	PC-40C48-SON-SMSR
PIC 3	REV 02	750-012793	PC-1XGE-TYPE3-XFP-IQ2
FPC 7	REV 08	710-010845	T640-FPC4-ES
PIC 0	REV 11	750-017405	PD-4XGE-XFP
SIB 0	REV 07	710-022594	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	SIB-TXP-T1600-S
Fan Tray 0			FANTRAY-T-S
Fan Tray 1			FANTRAY-T-S
Fan Tray 2			FANTRAY-TXP-R-S

### show chassis hardware detail (TX Matrix Plus Router)

```
user@host> show chassis hardware detail
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN111B023AHB	TXP
Midplane	REV 01	710-022574	TR7990	SFC Midplane
FPM Display	REV 03	710-024027	DW4699	TXP FPM Display
CIP 0	REV 01	710-023792	DR1437	TXP CIP
CIP 1	REV 02	710-023792	DS4564	TXP CIP
PEM 0	Rev 07	740-027463	UM26360	Power Entry Module
Routing Engine 0	REV 01	740-026942	737A-1024	SFC RE
ad0	3887 MB	SMART CF	200811050193CEB1CEB1	Compact Flash
ad1	30533 MB	SAMSUNG	MCBQE32G8MPP-0V SY814A0762	Disk 1
Routing Engine 1	REV 01	740-026942	737A-1024	SFC RE
ad0	3887 MB	SMART CF	20081105004C19A019A0	Compact Flash
ad1	30533 MB	SAMSUNG	MCBQE32G8MPP-0V SY814A0794	Disk 1
CB 0	REV 03	710-022606	DR7134	SFC Control Board
CB 1	REV 01	710-022606	DP8890	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 03	750-024564	DT9478	F13 SIB
B Board	REV 02	710-023431	DT6554	F13 SIB
SIB F13 1	REV 03	750-024564	DT9454	F13 SIB
B Board	REV 02	710-023431	DT6551	F13 SIB
SIB F2S 0/0	REV 02	710-022603	DT2838	F2S SIB

B Board	REV 02	710-023787	DT1725	NEO PMB
SIB F2S 0/2	REV 02	710-022603	DT2824	F2S SIB
B Board	REV 02	710-023787	DT1706	NEO PMB
SIB F2S 0/4	REV 02	710-022603	DT2822	F2S SIB
B Board	REV 02	710-023787	DT1696	NEO PMB
SIB F2S 0/6	REV 02	710-022603	DT2823	F2S SIB
B Board	REV 02	710-023787	DT1717	NEO PMB
SIB F2S 1/0	REV 03	710-022603	DV0059	F2S SIB
B Board	REV 03	710-023787	DT9942	NEO PMB
SIB F2S 1/2	REV 02	710-022603	DT2826	F2S SIB
B Board	REV 02	710-023787	DT1713	NEO PMB
SIB F2S 1/4	REV 03	710-022603	DV0092	F2S SIB
B Board	REV 03	710-023787	DV0000	NEO PMB
SIB F2S 1/6	REV 03	710-022603	DV0079	F2S SIB
B Board	REV 03	710-023787	DT9972	NEO PMB
SIB F2S 2/0	REV 03	710-022603	DV0100	F2S SIB
B Board	REV 03	710-023787	DT9925	NEO PMB
SIB F2S 2/2	REV 03	710-022603	DV0050	F2S SIB
B Board	REV 03	710-023787	DV0005	NEO PMB
SIB F2S 2/4	REV 03	710-022603	DV0097	F2S SIB
B Board	REV 03	710-023787	DT9936	NEO PMB
Fan Tray 0	REV 02	760-024497	DR8286	Front Fan Tray
Fan Tray 1	REV 06	760-024497	DV9624	Front Fan Tray
Fan Tray 2	REV 02	760-024502	DR8259	Rear Fan Tray
Fan Tray 3	REV 02	760-024502	DR8270	Rear Fan Tray
Fan Tray 4	REV 02	760-024502	DR8284	Rear Fan Tray
Fan Tray 5	REV 06	760-024502	DV7813	Rear Fan Tray

lcc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1101F27AHA	T1600
Midplane	REV 04	710-017247	RC5317	T Series Backplane
FPM GBUS	REV 10	710-002901	DS8197	T640 FPM Board
FPM Display	REV 01	710-021387	DS6433	T1600 FPM Display
CIP	REV 06	710-002895	DS1493	T Series CIP
PEM 0	Rev 08	740-017906	UD26601	Power Entry Module 3x80
SCG 0	REV 15	710-003423	DP5847	T640 Sonet Clock Gen.
SCG 1	REV 15	710-003423	DR0924	T640 Sonet Clock Gen.
Routing Engine 0	REV 01	740-026942	737F-1024	LCC RE
ad0	3887 MB	SMART CF	2008110502B63E513E51	Compact Flash
ad1	30533 MB	SAMSUNG	MCBQE32G8MPP-0V SY814A1208	Disk 1
Routing Engine 1	REV 01	740-026942	737F-1024	LCC RE
ad0	3887 MB	SMART CF	2008110500F9A8A8A8A8	Compact Flash
ad1	30533 MB	SAMSUNG	MCBQE32G8MPP-0V SY814A1076	Disk 1
CB 0	REV 05	710-022597	DV4264	LCC Control Board
CB 1	REV 03	710-022597	DP8558	LCC Control Board
FPC 0	REV 14	710-013037	DS9967	FPC Type 4-ES
CPU	REV 08	710-016744	DS3989	ST-PMB2
PIC 0	REV 12	750-013198	DL7506	1x Tunnel
PIC 1	REV 12	750-013198	DL7505	1x Tunnel
MMB 0	REV 01	710-025563	DS8524	ST-MMB2
MMB 1	REV 01	710-025563	DS8373	ST-MMB2
FPC 1	REV 14	710-013037	DT0027	FPC Type 4-ES
CPU	REV 09	710-016744	DS7684	ST-PMB2
PIC 0	REV 12	750-013198	DL7512	1x Tunnel
PIC 1	REV 12	750-013198	DL7498	1x Tunnel
MMB 0	REV 01	710-025563	DS8494	ST-MMB2
MMB 1	REV 01	710-025563	DS8436	ST-MMB2
SPMB 0	REV 04	710-023321	DV3867	LCC Switch CPU

SPMB 1	REV 02	710-023321	DP0238	LCC Switch CPU
SIB 0	REV 06	710-022594	DT8268	LCC SIB
B Board	REV 06	710-023185	DT5791	LCC SIB Mezz
SIB 1	REV 06	710-022594	DT8261	LCC SIB
B Board	REV 06	710-023185	DT5769	LCC SIB Mezz
SIB 2	REV 04	710-022594	DS2315	LCC SIB
B Board	REV 06	710-023185	DT5788	LCC SIB Mezz
SIB 3	REV 06	710-022594	DT8253	LCC SIB
B Board	REV 06	710-023185	DT5811	LCC SIB Mezz
SIB 4	REV 06	710-022594	DT8248	LCC SIB
B Board	REV 06	710-023185	DT5812	LCC SIB Mezz
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray

### show chassis hardware models (TX Matrix Plus Router)

```
user@host> show chassis hardware models
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	FRU model number
FPM Display	REV 03	710-024027	DX0282	CRAFT-TXP
CIP 0	REV 04	710-023792	DW4889	CIP-TXP
CIP 1	REV 04	710-023792	DW4887	CIP-TXP
PEM 0	Rev 07	740-027463	UM26368	yyyyyyyyyyyyyyyyyyyyyyyyyyyyyy
Routing Engine 0	REV 01	740-026942	737A-1064	RE-TXP-SFC-DUO-2600-16G
Routing Engine 1	REV 01	740-026942	737A-1082	RE-TXP-SFC-DUO-2600-16G
CB 0	REV 09	710-022606	DW6099	CB-TXP
CB 1	REV 09	710-022606	DW6096	CB-TXP
SIB F13 1	REV 04	750-024564	DW5776	SIB-TXP-F13
SIB F13 3	REV 04	750-024564	DW5762	SIB-TXP-F13
SIB F13 4	REV 04	750-024564	DW5797	SIB-TXP-F13
SIB F13 6	REV 04	750-024564	DW5770	SIB-TXP-F13
SIB F13 7	REV 04	750-024564	DW5758	SIB-TXP-F13
SIB F13 8	REV 04	750-024564	DW5761	SIB-TXP-F13
SIB F13 9	REV 04	750-024564	DW5754	SIB-TXP-F13
SIB F13 12	REV 04	750-024564	DW5794	SIB-TXP-F13
SIB F2S 0/0	REV 05	710-022603	DW7897	
SIB F2S 0/2	REV 05	710-022603	DW7833	
SIB F2S 0/4	REV 05	710-022603	DW7875	
SIB F2S 0/6	REV 05	710-022603	DW7860	
SIB F2S 1/0	REV 04	710-022603	DW4820	
SIB F2S 1/2	REV 05	710-022603	DW7849	
SIB F2S 1/4	REV 05	710-022603	DW7927	SIB-TXP-F2S
SIB F2S 1/6	REV 05	710-022603	DW7866	
SIB F2S 2/0	REV 05	710-022603	DW7880	
SIB F2S 2/2	REV 05	710-022603	DW7895	
SIB F2S 2/4	REV 05	710-022603	DW7907	
SIB F2S 2/6	REV 05	710-022603	DW7785	
SIB F2S 3/0	REV 05	710-022603	DW7782	
SIB F2S 3/2	REV 05	710-022603	DW7793	
SIB F2S 3/4	REV 05	710-022603	DW7779	
SIB F2S 3/6	REV 05	710-022603	DW7930	
SIB F2S 4/0	REV 05	710-022603	DW7867	
SIB F2S 4/2	REV 05	710-022603	DW7917	
SIB F2S 4/4	REV 05	710-022603	DW7929	
SIB F2S 4/6	REV 05	710-022603	DW7870	
Fan Tray 0	REV 06	760-024497	DV7831	FANTRAY-TXP-F
Fan Tray 1	REV 06	760-024497	DV9614	FANTRAY-TXP-F
Fan Tray 2	REV 06	760-024502	DV9618	FANTRAY-TXP-R

Fan Tray 3	REV 06	760-024502	DV9616	FANTRAY-TXP-R
Fan Tray 4	REV 06	760-024502	DV7807	FANTRAY-TXP-R
Fan Tray 5	REV 06	760-024502	DV7828	FANTRAY-TXP-R

lcc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3765	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN5441	CRAFT-T1600-S
CIP	REV 06	710-002895	DP6021	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UA26384	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UA26296	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DR0875	SCG-T-S
CB 0	REV 06	710-022597	DW8534	CB-LCC
CB 1	REV 06	710-022597	DW8527	CB-LCC
FPC 4	REV 12	710-013037	DJ8717	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8795	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8794	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS5335	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7634	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7637	PD-4XGE-XFP
FPC 7	REV 07	710-013035	DM0990	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8067	PC-10GE-SFP
PIC 1	REV 08	750-015749	WE9598	PC-10C192-SON-XFP
PIC 2	REV 10	750-009450	HX6466	PC-10C192-SON-SR2
SIB 0	REV 08	710-022594	DW8033	SIB-TXP-T1600-S
SIB 1	REV 08	710-022594	DW8044	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8020	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8063	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8064	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc1-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5361	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6430	CRAFT-T1600-S
CIP	REV 06	710-002895	DS4239	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26649	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5820	SCG-T-S
CB 0	REV 06	710-022597	DW8523	CB-LCC
CB 1	REV 06	710-022597	DW8528	CB-LCC
FPC 4	REV 12	710-013037	DP8509	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8808	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP7263	PD-4XGE-XFP
FPC 6	REV 14	710-013037	DS9961	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS5532	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7639	PD-4XGE-XFP
FPC 7	REV 03	710-013035	DF5564	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8063	PC-10GE-SFP
SIB 0	REV 08	710-022594	DW8035	SIB-TXP-T1600-S
SIB 1	REV 10	710-022594	DX7672	SIB-TXP-T1600-S
SIB 2	REV 08	710-022594	DW8060	SIB-TXP-T1600-S
SIB 3	REV 08	710-022594	DW8072	SIB-TXP-T1600-S
SIB 4	REV 08	710-022594	DW8043	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S



## Fan Tray 2

FANTRAY-TXP-R-S

lcc2-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 03	710-017247	RC3956	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DN7030	CRAFT-T1600-S
CIP	REV 06	710-002895	DM3962	CIP-L-T640-S
PEM 0	Rev 08	740-017906	UD26519	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26601	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP0277	SCG-T-S
CB 0	REV 06	710-022597	DW8524	CB-LCC
CB 1	REV 06	710-022597	DW8536	CB-LCC
FPC 4	REV 12	710-013037	DR1194	T1600-FPC4-ES
PIC 0	REV 11	750-017405	DP8811	PD-4XGE-XFP
PIC 1	REV 11	750-017405	DP8823	PD-4XGE-XFP
FPC 5	REV 12	710-013037	DR1184	T1600-FPC4-ES
PIC 1	REV 11	750-017405	DP4744	PD-4XGE-XFP
FPC 6	REV 12	710-013037	DN8622	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9924	PD-40C192-SON-XFP
PIC 1	REV 11	750-017405	DP8776	PD-4XGE-XFP
FPC 7	REV 04	710-013560	JR3968	T640-FPC3-E2
PIC 0	REV 16	750-007141	NC9330	PC-10GE-SFP
SIB 0	REV 07	710-022594	DW4217	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4213	SIB-TXP-T1600-S
SIB 2	REV 07	710-022594	DW4189	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4173	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4201	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

lcc3-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	FRU model number
Midplane	REV 04	710-017247	RC5319	CHAS-BP-T1600-S
FPM Display	REV 01	710-021387	DS6402	CRAFT-T1600-S
CIP	REV 06	710-002895	DR9973	CIP-L-T640-S
PEM 0	Rev 07	740-017906	UC26496	PWR-T1600-3-80-DC-S
PEM 1	Rev 07	740-017906	UC26599	PWR-T1600-3-80-DC-S
SCG 0	REV 15	710-003423	DP5831	SCG-T-S
CB 0	REV 06	710-022597	DW8533	CB-LCC
CB 1	REV 06	710-022597	DW8538	CB-LCC
FPC 0	REV 14	710-013037	DS5345	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7641	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS5479	PD-4XGE-XFP
FPC 1	REV 14	710-013037	DS7338	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7631	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7632	PD-4XGE-XFP
FPC 2	REV 14	710-013037	DS9962	T1600-FPC4-ES
PIC 0	REV 13	750-017405	DS7581	PD-4XGE-XFP
PIC 1	REV 13	750-017405	DS7627	PD-4XGE-XFP
FPC 4	REV 10	710-010845	JZ6573	T640-FPC4-ES
PIC 0	REV 14	750-012518	JT5124	PD-40C192-SON-XFP
FPC 5	REV 14	710-013037	DT0016	T1600-FPC4-ES
PIC 0	REV 14	750-012518	JY9918	PD-40C192-SON-XFP
FPC 7	REV 07	710-013035	DM0967	T1600-FPC3-ES
PIC 0	REV 16	750-007141	JJ8059	PC-10GE-SFP
PIC 1	REV 13	750-004695	DM5712	PC-TUNNEL

SIB 0	REV 07	710-022594	DW4174	SIB-TXP-T1600-S
SIB 1	REV 07	710-022594	DW4207	SIB-TXP-T1600-S
SIB 2	REV 06	710-022594	DT8231	SIB-TXP-T1600-S
SIB 3	REV 07	710-022594	DW4175	SIB-TXP-T1600-S
SIB 4	REV 07	710-022594	DW4209	SIB-TXP-T1600-S
Fan Tray 0				FANTRAY-T-S
Fan Tray 1				FANTRAY-T-S
Fan Tray 2				FANTRAY-TXP-R-S

### show chassis hardware (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module
Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module

Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

lcc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP

Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

lcc2-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B3975AHA	T1600
Midplane	REV 01	710-027486	RC9826	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5124	T640 FPM Board
FPM Display	REV 03	710-021387	BBAJ1112	T1600 FPM Display
CIP	REV 06	710-002895	BBAL3744	T-series CIP
PEM 0	REV 05	740-036442	1G022060081	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060188	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAH8775	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7272	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002992	RE-DUO-1800
Routing Engine 1	REV 07	740-026941	P737F-002938	RE-DUO-1800
CB 0	REV 11	710-022597	EH4805	LCC Control Board
CB 1	REV 11	710-022597	EH4786	LCC Control Board
FPC 1	REV 01	710-033873	BBAH0320	FPC Type 3-ES
CPU	REV 11	710-016744	BBAF3281	ST-PMB2
MMB 0	REV 06	710-025563	BBAF5061	ST-MMB2
FPC 5	REV 04	710-033871	BBAM5070	FPC Type 4-ES
CPU	REV 11	710-016744	BBAM6653	ST-PMB2
PIC 1	REV 20	750-017405	BBAM1296	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10B42981	XFP-10G-SR
MMB 0	REV 07	710-025563	BBAN2631	ST-MMB2
MMB 1	REV 07	710-025563	BBAN2538	ST-MMB2
SPMB 0	REV 05	710-023321	EH3903	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3902	LCC Switch CPU
SIB 0	REV 01	750-041657	EH8019	LCC SIB 3D
B Board	REV 01	711-042424	EH7680	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB04F	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB04S	CXP Module

Xcvr 4	REV 01	740-047547	XB48FB04B	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB043	CXP Module
SIB 1	REV 01	750-041657	EH8012	LCC SIB 3D
B Board	REV 01	711-042424	EH7658	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05E	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01Z	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB018	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB054	CXP Module
SIB 2	REV 01	750-041657	EH7993	LCC SIB 3D
B Board	REV 01	711-042424	EH7678	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05C	CXP Module
Xcvr 2	REV 01	740-047547	XB47FB00N	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB05U	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05L	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

### show chassis hardware clei-models (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware clei-models
sfc0-re0:
```

```
-----
Hardware inventory:
```

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 05	710-022574		CHAS-BP-TXP-S
FPM Display	REV 09	710-024027		CRAFT-TXP-S
CIP 0	REV 12	710-023792		CIP-TXP-S
CIP 1	REV 12	710-023792		CIP-TXP-S
PEM 0	Rev 06	740-027463	IPUPAFGKTA	PWR-TXP-7-60-DC-S
Routing Engine 0	REV 07	740-026942		RE-DUO-C2600-16G-S
Routing Engine 1	REV 07	740-026942		RE-DUO-C2600-16G-S
CB 0	REV 13	710-022606		CB-TXP-S
CB 1	REV 14	710-022606		CB-TXP-S
SIB F13 0	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 1	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
Xcvr 0	REV 01	740-048813		

Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-048813		
Xcvr 10	REV 01	740-048813		
Xcvr 12	REV 01	740-048813		
Xcvr 14	REV 01	740-048813		
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 6	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 7	REV 10	750-035002	PROTOXCLEI	SIB-TXP-3D-F13-S
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D

SIB F13 9	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 11	REV 10	750-035002	PROTOXCLEI	750-035002
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-048813		
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F13 12	REV 16	750-035002	PROTOXCLEI	SIB-TXP-3D-F13
Xcvr 0	REV 01	740-047547		CXP-TXP-3D
Xcvr 1	REV 01	740-047547		CXP-TXP-3D
Xcvr 2	REV 01	740-047547		CXP-TXP-3D
Xcvr 3	REV 01	740-047547		CXP-TXP-3D
Xcvr 4	REV 01	740-047547		CXP-TXP-3D
Xcvr 5	REV 01	740-047547		CXP-TXP-3D
Xcvr 6	REV 01	740-047547		CXP-TXP-3D
Xcvr 7	REV 01	740-047547		CXP-TXP-3D
Xcvr 8	REV 01	740-047547		CXP-TXP-3D
Xcvr 10	REV 01	740-047547		CXP-TXP-3D
Xcvr 12	REV 01	740-047547		CXP-TXP-3D
Xcvr 14	REV 01	740-047547		CXP-TXP-3D
SIB F2S 0/0	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/2	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 0/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/2	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/4	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 1/6	REV 08	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/0	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/4	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 2/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 3/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/0	REV 07	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/2	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/4	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
SIB F2S 4/6	REV 06	750-034978	PROTOXCLEI	SIB-TXP-3D-F2S
Fan Tray 0	REV 10	760-024497		FANTRAY-TXP-H-S
Fan Tray 1	REV 10	760-024497		FANTRAY-TXP-H-S
Fan Tray 2	REV 10	760-024502		FANTRAY-TXP-V-S

Fan Tray 3	REV 10	760-024502	FANTRAY-TXP-V-S
Fan Tray 4	REV 10	760-024502	FANTRAY-TXP-V-S
Fan Tray 5	REV 10	760-024502	FANTRAY-TXP-V-S

lcc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	CLEI code	FRU model number
Midplane	REV 01	710-027486	IPMJ700DRD	CHAS-BP-T1600-S
FPM Display	REV 04	710-021387		CRAFT-T1600-S
CIP	REV 06	710-002895		CIP-L-T640-S
PEM 0	REV 05	740-036442	IPUPAG6KAA	PWR-T-6-60-DC-S
PEM 1	REV 05	740-036442	IPUPAG6KAA	PWR-T-6-60-DC-S
SCG 0	REV 18	710-003423		SCG-T-S
SCG 1	REV 18	710-003423		SCG-T-S
Routing Engine 0	REV 10	740-026941		RE-DUO-C1800-8G-S
Routing Engine 1	REV 07	740-026941		RE-DUO-C1800-8G-S
CB 0	REV 11	710-022597		CB-LCC-S
CB 1	REV 11	710-022597		CB-LCC-S
FPC 0	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 3	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 13	750-033423	XXXXXXXXXD	PF-12-24XGE-SFPP
FPC 4	REV 02	750-045173	IP9IAL4DAC	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 5	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
FPC 6	REV 01	750-045173	IP9IAL4DAB	T4000-FPC5-3D
PIC 0	REV 17	750-034624	IP9IAL2DAA	PF-12XGE-SFPP
PIC 1	REV 10	750-035293	IP9IAL3DAA	PF-1CGE-CFP
SIB 0	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 1	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 2	REV 06	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC
Xcvr 0	REV 01	740-048813		
Xcvr 1	REV 01	740-048813		
Xcvr 2	REV 01	740-048813		
Xcvr 3	REV 01	740-048813		
Xcvr 4	REV 01	740-048813		
Xcvr 5	REV 01	740-048813		
Xcvr 6	REV 01	740-048813		
Xcvr 7	REV 01	740-048813		
SIB 3	REV 07	750-041657	PROTOXCLEI	SIB-TXP-3D-LCC



```

Xcvr 0      REV 01  740-048813
Xcvr 1      REV 01  740-048813
Xcvr 2      REV 01  740-048813
Xcvr 3      REV 01  740-048813
Xcvr 4      REV 01  740-048813
Xcvr 5      REV 01  740-048813
Xcvr 6      REV 01  740-048813
Xcvr 7      REV 01  740-048813
SIB 4       REV 06  750-041657  PROTOXCLEI  SIB-TXP-3D-LCC
Xcvr 0      REV 01  740-048813
Xcvr 1      REV 01  740-048813
Xcvr 2      REV 01  740-048813
Xcvr 3      REV 01  740-048813
Xcvr 4      REV 01  740-048813
Xcvr 5      REV 01  740-048813
Xcvr 6      REV 01  740-048813
Xcvr 7      REV 01  740-048813
Fan Tray 0
Fan Tray 1
Fan Tray 2
[Output Truncated]
FANTRAY-T-S
FANTRAY-T-S
FANTRAY-TXP3D-LCC-R-S

```

#### show chassis hardware detail (TX Matrix Plus router with 3D SIBs)

```

user@host> show chassis hardware detail
sfc0-re0:
-----
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis
Midplane      REV 05  710-022574  ABAC4696      SFC Midplane
FPM Display   REV 09  710-024027  EH3138        TXP FPM Display
CIP 0         REV 12  710-023792  EF6349        TXP CIP
CIP 1         REV 12  710-023792  EG5294        TXP CIP
PEM 0         Rev 06  740-027463  XH04595       Power Entry Module
PEM 1         Rev 06  740-027463  XH04592       Power Entry Module
Routing Engine 0 REV 07  740-026942  P737A-002541  RE-DUO-2600
  ad0 3823 MB SMART CF 2011030400062C132C13 Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 201105100009A452A452 Disk 1
Routing Engine 1 REV 07  740-026942  P737A-002602  RE-DUO-2600
  ad0 3823 MB SMART CF 20110508085EE471E471 Compact Flash
  ad1 62720 MB SMART Lite SATA Drive 201110210089DF39DF39 Disk 1
CB 0          REV 15  710-022606  EH4376        SFC Control Board
CB 1          REV 15  710-022606  EH4379        SFC Control Board
SPMB 0        BUILTIN
SPMB 1        BUILTIN
SIB F13 0     REV 10  750-035002  EM9305        F13 SIB 3D
  B Board     REV 06  711-035082  EM9667        F13 SIB 3D Mezz
  P Board     REV 05  711-043544  EM9708        F13 SIB 3D Power
Xcvr 0        REV 01  740-047547  XB34FB00S     CXP Module
Xcvr 2        REV 01  740-047547  XB48FB01H     CXP Module
Xcvr 4        REV 01  740-047547  XB34FB02W     CXP Module
Xcvr 6        REV 01  740-047547  XB34FB01T     CXP Module
Xcvr 8        REV 01  740-047547  XB48FB00W     CXP Module
Xcvr 10       REV 01  740-047547  XB34FB01S     CXP Module
Xcvr 12       REV 01  740-047547  XB34FB03H     CXP Module
Xcvr 14       REV 01  740-047547  XB34FB023     CXP Module
SIB F13 3     REV 01  710-035001  EJ2612        F13 SIB 3D
  B Board     REV 01  711-035082  EJ3815        F13 SIB 3D Mezz
  P Board     REV 01  711-043544  EJ2678        F13 SIB 3D Power
Xcvr 0        REV 01  740-047547  XB48FB04C     CXP Module

```

Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray
Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

1cc0-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
ad0	3823 MB	SMART CF	201103030490604E604E	Compact Flash
ad1	62720 MB	SMART Lite SATA Drive	20110729028B11D411D4	Disk 1
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800
ad0	3823 MB	SMART CF	2011010504EB99649964	Compact Flash
ad1	62720 MB	SMART Lite SATA Drive	201102140058934A934A	Disk 1
CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

lcc2-re0:

-----  
Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN1B3975AHA	T1600
Midplane	REV 01	710-027486	RC9826	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5124	T640 FPM Board
FPM Display	REV 03	710-021387	BBAJ1112	T1600 FPM Display
CIP	REV 06	710-002895	BBAL3744	T-series CIP
PEM 0	REV 05	740-036442	1G022060081	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060188	Power Entry Module 6x60

SCG 0	REV 18	710-003423	BBAH8775	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7272	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002992	RE-DUO-1800
ad0	3823 MB	SMART CF	201103030356329E329E	Compact Flash
ad1	62720 MB	SMART Lite SATA Drive	2011051000488D8B8D8B	Disk 1
Routing Engine 1	REV 07	740-026941	P737F-002938	RE-DUO-1800
ad0	3823 MB	SMART CF	20110304000F02680268	Compact Flash
ad1	62720 MB	SMART Lite SATA Drive	201105300A70F325F325	Disk 1
CB 0	REV 11	710-022597	EH4805	LCC Control Board
CB 1	REV 11	710-022597	EH4786	LCC Control Board
FPC 1	REV 01	710-033873	BBAH0320	FPC Type 3-ES
CPU	REV 11	710-016744	BBAF3281	ST-PMB2
MMB 0	REV 06	710-025563	BBAF5061	ST-MMB2
FPC 5	REV 04	710-033871	BBAM5070	FPC Type 4-ES
CPU	REV 11	710-016744	BBAM6653	ST-PMB2
PIC 1	REV 20	750-017405	BBAM1296	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10B42981	XFP-10G-SR
MMB 0	REV 07	710-025563	BBAN2631	ST-MMB2
MMB 1	REV 07	710-025563	BBAN2538	ST-MMB2
SPMB 0	REV 05	710-023321	EH3903	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3902	LCC Switch CPU
SIB 0	REV 01	750-041657	EH8019	LCC SIB 3D
B Board	REV 01	711-042424	EH7680	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB04F	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB04S	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04B	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB043	CXP Module
SIB 1	REV 01	750-041657	EH8012	LCC SIB 3D
B Board	REV 01	711-042424	EH7658	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05E	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01Z	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB018	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB054	CXP Module
SIB 2	REV 01	750-041657	EH7993	LCC SIB 3D
B Board	REV 01	711-042424	EH7678	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05C	CXP Module
Xcvr 2	REV 01	740-047547	XB47FB00N	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB05U	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05L	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

### show chassis hardware lcc (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware lcc 0
lcc0-re0:
```

#### Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11B23FEAHA	T1600
Midplane	REV 01	710-027486	RC9787	T-series Backplane
FPM GBUS	REV 13	710-002901	BBAG5132	T640 FPM Board
FPM Display	REV 04	710-021387	BBAL9612	T1600 FPM Display
CIP	REV 06	710-002895	BBAN0605	T-series CIP
PEM 0	REV 05	740-036442	1G022060143	Power Entry Module 6x60
PEM 1	REV 05	740-036442	1G022060011	Power Entry Module 6x60
SCG 0	REV 18	710-003423	BBAL7318	T640 Sonet Clock Gen.
SCG 1	REV 18	710-003423	BBAL7255	T640 Sonet Clock Gen.
Routing Engine 0	REV 07	740-026941	P737F-002933	RE-DUO-1800
Routing Engine 1	REV 06	740-026941	P737F-002749	RE-DUO-1800

CB 0	REV 11	710-022597	EH3611	LCC Control Board
CB 1	REV 11	710-022597	EH4798	LCC Control Board
FPC 5	REV 17	710-013037	BBAC5333	FPC Type 4-ES
CPU	REV 10	710-016744	BBAB7619	ST-PMB2
PIC 0	REV 18	750-017405	BBAE3420	4x 10GE (LAN/WAN) XFP
Xcvr 0	REV 03	740-014289	T10C90659	XFP-10G-SR
MMB 0	REV 05	710-025563	BBAB9538	ST-MMB2
MMB 1	REV 05	710-025563	BBAB9502	ST-MMB2
FPC 7	REV 01	750-045173	BBAV0032	FPC Type 5-3D
CPU				
SPMB 0	REV 05	710-023321	EG9434	LCC Switch CPU
SPMB 1	REV 05	710-023321	EH3878	LCC Switch CPU
SIB 0	REV 01	750-041657	EH7997	LCC SIB 3D
B Board	REV 01	711-042424	EH7674	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB014	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB05A	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB052	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB01B	CXP Module
SIB 1	REV 01	750-041657	EH8023	LCC SIB 3D
B Board	REV 01	711-042424	EH7659	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB48FB05J	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01E	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB01J	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB02S	CXP Module
SIB 2	REV 03	750-041657	EJ6554	LCC SIB 3D
B Board	REV 02	711-042424	EJ5756	LCC SIB 3D Mezz
Xcvr 0	REV 01	740-047547	XB34FB01Z	CXP Module
Xcvr 2	REV 01	740-047547	XB34FB013	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB04Z	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB05N	CXP Module
Fan Tray 0				Front Top Fan Tray
Fan Tray 1				Front Bottom Fan Tray
Fan Tray 2				Rear Fan Tray -- Rev 4

### show chassis hardware sfc (TX Matrix Plus router with 3D SIBs)

```
user@host> show chassis hardware sfc 0
sfc0-re0:
```

```
-----
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN11CAAA4AHB	TXP
Midplane	REV 05	710-022574	ABAC4696	SFC Midplane
FPM Display	REV 09	710-024027	EH3138	TXP FPM Display
CIP 0	REV 12	710-023792	EF6349	TXP CIP
CIP 1	REV 12	710-023792	EG5294	TXP CIP
PEM 0	Rev 06	740-027463	XH04595	Power Entry Module
PEM 1	Rev 06	740-027463	XH04592	Power Entry Module
Routing Engine 0	REV 07	740-026942	P737A-002541	RE-DUO-2600
Routing Engine 1	REV 07	740-026942	P737A-002602	RE-DUO-2600
CB 0	REV 15	710-022606	EH4376	SFC Control Board
CB 1	REV 15	710-022606	EH4379	SFC Control Board
SPMB 0		BUILTIN		SFC Switch CPU
SPMB 1		BUILTIN		SFC Switch CPU
SIB F13 0	REV 10	750-035002	EM9305	F13 SIB 3D
B Board	REV 06	711-035082	EM9667	F13 SIB 3D Mezz
P Board	REV 05	711-043544	EM9708	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB34FB00S	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB01H	CXP Module
Xcvr 4	REV 01	740-047547	XB34FB02W	CXP Module
Xcvr 6	REV 01	740-047547	XB34FB01T	CXP Module

Xcvr 8	REV 01	740-047547	XB48FB00W	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01S	CXP Module
Xcvr 12	REV 01	740-047547	XB34FB03H	CXP Module
Xcvr 14	REV 01	740-047547	XB34FB023	CXP Module
SIB F13 3	REV 01	710-035001	EJ2612	F13 SIB 3D
B Board	REV 01	711-035082	EJ3815	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2678	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB04C	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB00Z	CXP Module
Xcvr 4	REV 01	740-047547	XB47FB036	CXP Module
Xcvr 6	REV 01	740-047547	XB47FB029	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02N	CXP Module
Xcvr 10	REV 01	740-047547	XB42FB0CS	CXP Module
Xcvr 12	REV 01	740-047547	XB47FB01X	CXP Module
Xcvr 14	REV 01	740-047547	XB48FB02F	CXP Module
SIB F13 6	REV 05	750-035002	EK2675	F13 SIB 3D
B Board	REV 03	711-035082	EK2612	F13 SIB 3D Mezz
P Board	REV 04	711-043544	EK1179	F13 SIB 3D Power
Xcvr 0	REV 01	740-047547	XB48FB01T	CXP Module
Xcvr 2	REV 01	740-047547	XB48FB02M	CXP Module
Xcvr 4	REV 01	740-047547	XB48FB031	CXP Module
Xcvr 6	REV 01	740-047547	XB48FB04P	CXP Module
Xcvr 8	REV 01	740-047547	XB48FB02T	CXP Module
Xcvr 10	REV 01	740-047547	XB34FB01V	CXP Module
Xcvr 12	REV 01	740-047547	XB48FB02C	CXP Module
Xcvr 14		NON-JNPR		No Module
SIB F13 12	REV 01	710-035001	EJ2631	F13 SIB 3D
B Board	REV 01	711-035082	EJ3808	F13 SIB 3D Mezz
P Board	REV 01	711-043544	EJ2676	F13 SIB 3D Power
SIB F2S 0/0	REV 01	711-034977	EH9829	F2S SIB 3D
B Board	REV 01	711-034979	EH9927	F2S SIB 3D Mezz
SIB F2S 0/2	REV 01	711-034977	EH9791	F2S SIB 3D
B Board	REV 01	711-034979	EH9852	F2S SIB 3D Mezz
SIB F2S 0/4	REV 01	711-034977	EH9803	F2S SIB 3D
B Board	REV 01	711-034979	EH9915	F2S SIB 3D Mezz
SIB F2S 0/6	REV 01	711-034977	EH9763	F2S SIB 3D
B Board	REV 01	711-034979	EH9880	F2S SIB 3D Mezz
SIB F2S 1/0	REV 01	711-034977	EH9757	F2S SIB 3D
B Board	REV 01	711-034979	EH9889	F2S SIB 3D Mezz
SIB F2S 1/2	REV 01	711-034977	EH9815	F2S SIB 3D
B Board	REV 01	711-034979	EH9890	F2S SIB 3D Mezz
SIB F2S 1/4	REV 08	750-034978	EN1954	F2S SIB 3D
B Board	REV 02	711-034979	EN1436	F2S SIB 3D Mezz
SIB F2S 1/6	REV 01	711-034977	EJ7054	F2S SIB 3D
B Board	REV 01	711-034979	EJ8238	F2S SIB 3D Mezz
SIB F2S 2/0	REV 01	711-034977	EH9830	F2S SIB 3D
B Board	REV 01	711-034979	EH9844	F2S SIB 3D Mezz
SIB F2S 2/2	REV 01	711-034977	EH9818	F2S SIB 3D
B Board	REV 01	711-034979	EH9888	F2S SIB 3D Mezz
SIB F2S 2/4	REV 01	711-034977	EH9795	F2S SIB 3D
B Board	REV 01	711-034979	EH9869	F2S SIB 3D Mezz
SIB F2S 2/6	REV 01	711-034977	EJ7026	F2S SIB 3D
B Board	REV 01	711-034979	EJ8273	F2S SIB 3D Mezz
SIB F2S 3/0	REV 01	711-034977	EH9811	F2S SIB 3D
B Board	REV 01	711-034979	EH9892	F2S SIB 3D Mezz
SIB F2S 3/2	REV 01	711-034977	EH9812	F2S SIB 3D
B Board	REV 01	711-034979	EH9877	F2S SIB 3D Mezz
SIB F2S 3/4	REV 08	750-034978	EN1947	F2S SIB 3D
B Board	REV 02	711-034979	EN1471	F2S SIB 3D Mezz
Fan Tray 0	REV 10	760-024497	EH3313	Front Fan Tray
Fan Tray 1	REV 10	760-024497	EH3290	Front Fan Tray

Fan Tray 2	REV 10	760-024502	EH3292	Rear Fan Tray
Fan Tray 3	REV 10	760-024502	EH3287	Rear Fan Tray
Fan Tray 4	REV 10	760-024502	EH3286	Rear Fan Tray
Fan Tray 5	REV 10	760-024502	EH3285	Rear Fan Tray

### show chassis hardware (16-Port 10-Gigabit Ethernet MPC with SFP+ Optics [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN112D865AFA	MX960
Midplane	REV 03	710-013698	TS3339	MX960 Backplane
FPM Board	REV 03	710-014974	WW6267	Front Panel Display
PDM	Rev 03	740-013110	QCS12485026	Power Distribution
Module				
PEM 0	Rev 04	740-013682	QCS12434086	PS 1.7kW; 200-240VAC
in				
PEM 1	Rev 04	740-013682	QCS1243408Z	PS 1.7kW; 200-240VAC
in				
PEM 2	Rev 04	740-013682	QCS1243407X	PS 1.7kW; 200-240VAC
in				
Routing Engine 0	REV 07	740-015113	9009009677	RE-S-1300
Routing Engine 1	REV 07	740-015113	9009011510	RE-S-1300
CB 0	REV 03	710-021523	XF0394	MX SCB
CB 1	REV 03	710-021523	XF0550	MX SCB
CB 2	REV 03	710-021523	XD7455	MX SCB
FPC 4	REV 02	750-028467	JR6127	MPC M 16x 10GE
CPU	REV 02	711-029089	JX0129	AS PMB
PIC 0		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 1		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 2		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
PIC 3		BUILTIN	BUILTIN	4x 10GE(LAN) SFP+
Fan Tray 0	REV 05	740-014971	TP9990	Fan Tray
Fan Tray 1	REV 05	740-014971	VS1709	Fan Tray

### show chassis hardware (MPC3E [MX Series Routers])

```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN1101AFEAFB	MX480
Midplane	REV 05	710-017414	TR4444	MX480 Midplane
FPM Board	REV 02	710-017254	KG6056	Front Panel Display
PEM 0	Rev 03	740-017330	QCS082090FC	PS 1.2-1.7kW; 100-240V
PEM 1	Rev 03	740-017330	QCS082090FD	PS 1.2-1.7kW; 100-240V
Routing Engine 0	REV 07	740-013063	9009004124	RE-S-2000
Routing Engine 1	REV 07	740-013063	9009005569	RE-S-2000
CB 0	REV 07	710-021523	XZ3587	MX SCB
CB 1	REV 03	710-021523	KH8306	MX SCB
FPC 1	REV 04.1.07	750-033205	P1240	MPC Type 3
CPU	REV 01	711-035209	YL0504	HMPC PMB 2G
MIC 1	REV 10	750-033199	YX4495	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	C22CQNE	CFP-100G-LR4
FPC 2	REV 26	750-016670	KH0045	DPCE 40x 1GE R EQ
CPU	REV 07	710-013713	KF5448	DPC PMB
PIC 0		BUILTIN	BUILTIN	10x 1GE(LAN) EQ

Xcvr 0	REV 01	740-011613	PF21JHU	SFP-SX
PIC 1		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 9	REV 01	740-011613	AM0813S8ZL6	SFP-SX
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 0	REV 02	740-011613	PGL2KYF	SFP-SX
Xcvr 2	REV 01	740-011613	AM0806S8N4P	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) EQ
Xcvr 5	REV 01	740-011613	AM0815S967N	SFP-SX
Xcvr 7	REV 01	740-011613	AM0806S8N1X	SFP-SX
Xcvr 8	REV 01	740-011613	AM0815S967J	SFP-SX
Xcvr 9	REV 01	740-011613	AM0815S967M	SFP-SX
FPC 3	REV 12.2.09	750-033205	YR9443	MPC Type 3
CPU	REV 03	711-035209	YL6931	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3269	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULHOKG3	CFP-100G-LR4
MIC 1	REV 02	750-033199	YG3245	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-032210	ULHOKGF	CFP-100G-LR4
FPC 4	REV 12.3.09	750-033205	YR9437	MPC Type 3
CPU	REV 03	711-035209	YT5857	HMPC PMB 2G
MIC 0	REV 05	750-033199	YR3295	1X100GE CFP
PIC 0		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0		NON-JNPR	X12000187	CFP-100G-SR10
MIC 1	REV 10	750-033199	YX4518	1X100GE CFP
PIC 2		BUILTIN	BUILTIN	1X100GE CFP
Xcvr 0	REV 01	740-035329	X12J00008	CFP-100G-SR10
FPC 5	REV 06	750-024884	JW9769	MPC Type 2 3D EQ
CPU	REV 02	711-028401	JR6158	MPC PMB 2G Proto
MIC 0	REV 05	750-028387	JR6197	3D 4x 10GE XFP
PIC 0		BUILTIN	BUILTIN	2x 10GE XFP
Xcvr 0	REV 01	740-014289	T07M71112	XFP-10G-SR
Xcvr 1	REV 02	740-014289	T08L85610	XFP-10G-SR
PIC 1		BUILTIN	BUILTIN	2x 10GE XFP
MIC 1	REV 22	750-028392	YM0053	3D 20x 1GE(LAN) SFP
PIC 2		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 0	REV 01	740-011613	AM0703S005B	SFP-SX
Xcvr 1	REV 01	740-011613	E07L01352	SFP-SX
PIC 3		BUILTIN	BUILTIN	10x 1GE(LAN) SFP
Xcvr 5	REV 01	740-013111	6500217	SFP-T
Xcvr 9	REV 02	740-013111	8499527	SFP-T
Fan Tray				Left Fan Tray

The PIC number for MIC 1 always starts from 2 (even if the first MIC is a 1X100GE CFP or a legacy MIC).

#### show chassis hardware (QFX3500 Switches)

```
user@switch> show chassis hardware
```

Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis				QFX3500
Routing Engine 0				QFX Routing Engine
FPC 0	REV 04	750-044071	BBAR3902	QFX3500-48S4Q-AFI
CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0		BUILTIN	BUILTIN	48x 10G-SFP+
PIC 1		BUILTIN	BUILTIN	15x 10G-SFP+
MGMT BRD	REV 02	750-044063	BBAR0398	QFX3500-MGMT-SFP-AF0
Xcvr 0	REV 01	740-011614	AC0946S0BD1	SFP-LX10
Xcvr 1	REV 02	740-013111	A281922	SFP-T



Power Supply 0	Rev 04	740-032091	UI00677	JPSU-650W-AC-AFI
Power Supply 1	REV 00	740-041741	VJ00162	JPSU-650W-AC-AFO
Fan Tray 0				QFX Fan Tray, Back to
Front Airflow				
Fan Tray 1				QFX Fan Tray, Back to
Front Airflow				
Fan Tray 2				QFX Fan Tray, Back to
Front Airflow				

### show chassis hardware detail (QFX3500 Switches)

```
user@switch> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			JN000TEST5	QFX3500
Routing Engine 0		BUILTIN	BUILTIN	QFX Routing Engine
FPC 0	REV 05	750-036931	EE0823	QFX3500-48S4Q-AFI

CPU		BUILTIN	BUILTIN	FPC CPU
PIC 0		BUILTIN	BUILTIN	48x 10G-SFP+
Xcvr 0	REV 01	740-030589	S99E270079	SFP+-10G-LPBK
Xcvr 1	REV 01	740-030589	S9AK450099	SFP+-10G-LPBK
Xcvr 2	REV 01	740-030589	S99E270078	SFP+-10G-LPBK
Xcvr 3	REV 01	740-030589	S9AK450098	SFP+-10G-LPBK
Xcvr 4	REV 01	740-030589	S99E270075	SFP+-10G-LPBK
Xcvr 5	REV 01	740-030589	S9AK450093	SFP+-10G-LPBK
Xcvr 6	REV 01	740-030589	S9AK450097	SFP+-10G-LPBK
Xcvr 7	REV 01	740-030589	S9AK450095	SFP+-10G-LPBK
Xcvr 8	REV 01	740-030589	S99E270072	SFP+-10G-LPBK
Xcvr 9	REV 01	740-030589	S99E270073	SFP+-10G-LPBK
Xcvr 10	REV 01	740-030589	S99E270080	SFP+-10G-LPBK
Xcvr 11	REV 01	740-030589	S9AK450169	SFP+-10G-LPBK
Xcvr 12	REV 01	740-030589	S99E270076	SFP+-10G-LPBK
Xcvr 13	REV 01	740-030589	S9AK450167	SFP+-10G-LPBK
Xcvr 14	REV 01	740-030589	S9AK450170	SFP+-10G-LPBK
Xcvr 15	REV 01	740-030589	S9AK450166	SFP+-10G-LPBK
Xcvr 16	REV 01	740-030589	S9AK450092	SFP+-10G-LPBK
Xcvr 17	REV 01	740-030589	S9AK450163	SFP+-10G-LPBK
Xcvr 18	REV 01	740-030589	S9AK450094	SFP+-10G-LPBK
Xcvr 19	REV 01	740-030589	S9AK450100	SFP+-10G-LPBK
Xcvr 20	REV 01	740-030589	S9AK450168	SFP+-10G-LPBK
Xcvr 21	REV 01	740-030589	S9AK450165	SFP+-10G-LPBK
Xcvr 22	REV 01	740-030589	S9AK450073	SFP+-10G-LPBK
Xcvr 23	REV 01	740-030589	S9AK450164	SFP+-10G-LPBK
Xcvr 24	REV 01	740-030589	S9AK450074	SFP+-10G-LPBK
Xcvr 25	REV 01	740-030589	SA62270195	SFP+-10G-LPBK
Xcvr 26	REV 01	740-030589	S9AK450078	SFP+-10G-LPBK
Xcvr 27	REV 01	740-030589	S9AK450024	SFP+-10G-LPBK
Xcvr 28	REV 01	740-030589	S9AK450027	SFP+-10G-LPBK
Xcvr 29	REV 01	740-030589	S9AK450080	SFP+-10G-LPBK
Xcvr 30	REV 01	740-030589	S9AK450030	SFP+-10G-LPBK
Xcvr 31	REV 01	740-030589	S9AK450025	SFP+-10G-LPBK
Xcvr 32	REV 01	740-030589	S9AK450023	SFP+-10G-LPBK
Xcvr 33	REV 01	740-030589	S9AK450075	SFP+-10G-LPBK
Xcvr 34	REV 01	740-030589	S9AK450161	SFP+-10G-LPBK
Xcvr 35	REV 01	740-030589	S9AK450071	SFP+-10G-LPBK
Xcvr 36	REV 01	740-030589	S9AK450072	SFP+-10G-LPBK
Xcvr 37	REV 01	740-030589	S9AK450022	SFP+-10G-LPBK
Xcvr 38	REV 01	740-030589	S9AK450021	SFP+-10G-LPBK
Xcvr 39	REV 01	740-030589	S9AK450175	SFP+-10G-LPBK

Xcvr 40	REV 01	740-030589	S9AK450162	SFP+-10G-LPBK
Xcvr 41	REV 01	740-030589	S99E270074	SFP+-10G-LPBK
Xcvr 42	REV 01	740-030589	S9AK450174	SFP+-10G-LPBK
Xcvr 43	REV 01	740-030589	S9AK450077	SFP+-10G-LPBK
Xcvr 44	REV 01	740-030589	S9AK450076	SFP+-10G-LPBK
Xcvr 45	REV 01	740-030589	S9AK450026	SFP+-10G-LPBK
Xcvr 46	REV 01	740-030589	S9AK450079	SFP+-10G-LPBK
Xcvr 47	REV 01	740-030589	S9AK450029	SFP+-10G-LPBK
PIC 1		BUILTIN	BUILTIN	15x 10G-SFP+
Xcvr 1	REV 01	740-032986	QA170087	QSFP+-40G-SR4
Xcvr 4	REV 01	740-032986	QA360442	QSFP+-40G-SR4
Xcvr 8	REV 01	740-032986	QA170091	QSFP+-40G-SR4
Xcvr 12	REV 01	740-032986	QA170042	QSFP+-40G-SR4
MGMT BRD	REV 08	750-036946	EE0731	QFX3500-MB
Power Supply 0	Rev 04	740-032091	UI00690	QFX PS 650W AC
Power Supply 1	Rev 04	740-032091	UI00679	QFX PS 650W AC
Fan Tray 0				QFX Fan Tray
Fan Tray 1				QFX Fan Tray

**show chassis hardware models (QFX3500 Switches)**

```

user@switch> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Routing Engine 0          BUILTIN    BUILTIN
FPC 0          REV 02    711-032234  EC4074
Power Supply 0  PSMI 2C  11-d65800  --

```

**show chassis hardware clei-models (QFX3500 Switches)**

```

user@switch> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Routing Engine 0          BUILTIN
FPC 0          REV 02    711-032234
Power Supply 0  PSMI 2C  11-d65800

```

**show chassis hardware clei-models (QFX5100 Switches)**

```

user@switch> show chassis hardware clei-models
Hardware inventory:
Item          Version  Part number  CLEI code      FRU model number
Routing Engine 0          BUILTIN      CMMNV10BRA
FPC 0          REV 01    611-053010  CMMNV10BRA
PIC 0          BUILTIN      CMMNV10BRA
Power Supply 0  REV 03    740-053352  MUPABHBAA      JPSU-850W-AC-AFO
Power Supply 1  REV 03    740-053352  MUPABHBAA      JPSU-850W-AC-AFO
Fan Tray 0          QFX5100-96S-FANAFO
Fan Tray 1          QFX5100-96S-FANAFO
Fan Tray 2          QFX5100-96S-FANAFO

```

**show chassis hardware interconnect-device (QFabric Systems)**

```

user@switch> show chassis hardware interconnect-device interconnect1
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis          REV 07
Midplane          REV 07    750-021261  BH0208188289  QFX Midplane
CB 0             REV 07    750-021261  BH0208188289  QFXIC08-CB4S

```

## show chassis hardware node-device (QFabric Systems)

```

user@switch> show chassis hardware node-device node1
Routing Engine 0    BUILTIN    BUILTIN    QFX Routing Engine
node1              REV 05    711-032234 ED3694      QFX3500-48S4Q-AFI

CPU
PIC 0              BUILTIN    BUILTIN
Xcvr 8            REV 01    740-030658 AD0946A028B FPC CPU
                                     48x 10G-SFP+
                                     SFP+-10G-USR
...

```

## show chassis hardware (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN11D1FD7AJA  PTX5000
Midplane      REV 03    711-031896  ABAC5589      Midplane-8S
FPM           REV 08    760-030647  EG1679        Front Panel Display
PDU 0         Rev 05    740-032019  ZE00006       DC Power Dist Unit
  PSM 0        Rev 05    740-032022  ZJ00018       DC 12V Power Supply
  PSM 1        Rev 04    740-032022  ZC00052       DC 12V Power Supply
  PSM 2        Rev 04    740-032022  ZD00051       DC 12V Power Supply
  PSM 3        Rev 05    740-032022  ZJ00060       DC 12V Power Supply
CCG 0         REV 04    750-030653  EG3703        Clock Generator
CCG 1         REV 04    750-030653  EG3698        Clock Generator
Routing Engine 0 REV 05    740-026942  P737A-002231  RE-DUO-2600
Routing Engine 1 REV 06    740-026942  P737A-002438  RE-DUO-2600
CB 0          REV 08    750-030625  EG5519        Control Board
CB 1          REV 08    750-030625  EG5516        Control Board
FPC 0         REV 18    750-036844  EJ3080        FPC
  CPU         REV 12    711-030686  EJ3260        SNG PMB
FPC 2         REV 13    750-036844  EG5065        FPC
  CPU         REV 09    711-030686  EG4082        SNG PMB
  PIC 0       REV 14    750-031913  EG5127        24x 10GE(LAN) SFP+
    Xcvr 0    REV 01    740-031980  143363A00240 SFP+-10G-SR
    Xcvr 1    REV 01    740-031981  UK90PZ1       SFP+-10G-LR
    Xcvr 2    REV 01    740-031980  AD1141A04XH   SFP+-10G-SR
    Xcvr 3    REV 01    740-031981  UK90Q46       SFP+-10G-LR
    Xcvr 4    REV 01    740-031980  AD1141A04X4   SFP+-10G-SR
    Xcvr 6    REV 01    740-031980  B11H02560     SFP+-10G-SR
    Xcvr 7    REV 01    740-031980  B11C01589     SFP+-10G-SR
    Xcvr 8    REV 01    740-031980  AD1141A04XF   SFP+-10G-SR
    Xcvr 10   REV 01    740-031980  123363A01094  SFP+-10G-SR
    Xcvr 11   REV 01    740-031980  AK80LKF       SFP+-10G-SR
    Xcvr 12   REV 01    740-031980  183363A01528  SFP+-10G-SR
    Xcvr 14   REV 01    740-031980  193363A01079  SFP+-10G-SR
    Xcvr 15   REV 01    740-031980  AK80MC8       SFP+-10G-SR
    Xcvr 16   REV 01    740-031980  AJC0BHC       SFP+-10G-SR
    Xcvr 19   REV 01    740-021309  J08D26856     SFP+-10G-LR
    Xcvr 21   REV 01    740-031980  AK80KCT       SFP+-10G-SR
    Xcvr 22   REV 01    740-031981  UK90PZL       SFP+-10G-LR
    Xcvr 23   REV 01    740-031980  AK80N1V       SFP+-10G-SR
FPC 3         REV 13    750-036844  EG5074        FPC
  CPU         REV 09    711-030686  EG4064        SNG PMB
  PIC 1       REV 10    750-031903  EG0325        SNG Load
FPC 5         REV 06    750-036844  EH3198        FPC
  CPU
  PIC 0       REV 14    750-031913  EG5134        24x 10GE(LAN) SFP+
    Xcvr 0    REV 01    740-031980  AK80LBH       SFP+-10G-SR

```

Xcvr 1	REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC
CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

#### show chassis hardware (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```
user@host> show chassis hardware
```

Hardware inventory:				
Item	Version	Part number	Serial number	Description
Chassis			JN1204FC0AJA	PTX5000
Midplane	REV 11	750-035893	ACAB8038	Midplane-8S
FPM	REV 12	760-030647	BBBD5619	Front Panel
Display				
PDU 0	Rev 04	740-048336	1GB93470043	High Capacity DC PDU
PSM 0	Rev 04	740-046988	1GB63500184	High Capacity DC PSM
PSM 2	Rev 04	740-046988	1GB63500169	High Capacity DC PSM
PSM 4	Rev 04	740-046988	1GB63500306	High Capacity DC PSM
PSM 6	Rev 04	740-046988	1GB63500074	High Capacity DC PSM
PDU 1	Rev 04	740-048336	1GB93470045	High Capacity DC PDU
PSM 1	Rev 04	740-046988	1GB63500193	High Capacity DC PSM
PSM 3	Rev 04	740-046988	1GB63500143	High Capacity DC PSM
PSM 5	Rev 04	740-046988	1GB63500146	High Capacity DC PSM
PSM 7	Rev 04	740-046988	1GB63500192	High Capacity DC PSM
CCG 0	REV 09	750-030653	BBBC1909	Clock Generator
CCG 1	REV 09	750-030653	BBBD2970	Clock Generator
...				

**show chassis hardware clei-models (PTX5000 Packet Transport Router)**

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code  FRU model number
FPM                 REV 08    760-030647  PROTOXCLEI CRAFT-PTX5000-S
PDU 0              Rev 05    740-032019  IPUPAHLKAA  PWR-SAN-PDU-DC
  PSM 0             Rev 05    740-032022  IPUPAHNKAA  PSM-PTX-DC-120-S
  PSM 1             Rev 04    740-032022  032022XXXX  PWR-SAN-12-DC
  PSM 2             Rev 04    740-032022  032022XXXX  PWR-SAN-12-DC
  PSM 3             Rev 05    740-032022  IPUPAHNKAA  PSM-PTX-DC-120-S
CCG 0              REV 04    750-030653  PROTOXCLEI CCG-PTX-S
CCG 1              REV 04    750-030653  PROTOXCLEI CCG-PTX-S
Routing Engine 0   REV 05    740-026942  RE-DUO-C2600-16G-S
Routing Engine 1   REV 06    740-026942  RE-DUO-C2600-16G-S
CB 0               REV 08    750-030625  PROTOXCLEI CB-PTX-S
CB 1               REV 08    750-030625  PROTOXCLEI CB-PTX-S
FPC 0              REV 18    750-036844  PROTOXCLEI FPC-PTX-P1-A
FPC 2              REV 13    750-036844  PROTOXCLEI FPC-PTX-P1-A
  PIC 0             REV 14    750-031913  PROTOXCLEI P1-PTX-24-10GE-SFPP
FPC 3              REV 13    750-036844  PROTOXCLEI FPC-PTX-P1-A
FPC 5
  PIC 0             REV 14    750-031913  PROTOXCLEI P1-PTX-24-10GE-SFPP
FPC 6              REV 18    750-036844  PROTOXCLEI FPC-PTX-P1-A
FPC 7              REV 18    750-036844  PROTOXCLEI FPC-PTX-P1-A
SIB 0              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 1              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 2              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 3              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 4              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 5              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 6              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 7              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
SIB 8              REV 07    750-030631  PROTOXCLEI SIB-I-PTX5008
Fan Tray 1         REV 04    760-030642  PROTOXCLEI FAN-PTX-H-S

```

**show chassis hardware clei-models (PTX5000 Packet Transport Router with FPC2-PTX-P1A)**

```

user@host> show chassis hardware clei-models
Hardware inventory:
Item                Version  Part number  CLEI code  FRU model number
Midplane            REV 11    750-035893  IPMUN00ARA  CHAS-MP-PTX5000-S
FPM                 REV 12    760-030647  IPUCA7SCAA  CRAFT-PTX5000-S
PDU 0              Rev 04    740-048336  IPUPAL7KAA  PDU2-PTX-DC-S
  PSM 0             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
  PSM 2             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
  PSM 4             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
  PSM 6             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
PDU 1              Rev 04    740-048336  IPUPAL7KAA  PDU2-PTX-DC-S
  PSM 1             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
  PSM 3             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
  PSM 5             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
  PSM 7             Rev 04    740-046988  IPUPAL8KAA  PSM2-PTX-DC-S
CCG 0              REV 09    750-030653  IPUCA7DCAA  CCG-PTX-S
CCG 1              REV 09    750-030653  IPUCA7DCAA  CCG-PTX-S
...

```

**show chassis hardware detail (PTX5000 Packet Transport Router)**

```

user@host> show chassis hardware detail

```

## Hardware inventory:

Item	Version	Part number	Serial number	Description
Chassis			JN11D1FD7AJA	PTX5000
Midplane	REV 03	711-031896	ABAC5589	Midplane-8S
FPM	REV 08	760-030647	EG1679	Front Panel Display
PDU 0	Rev 05	740-032019	ZE00006	DC Power Dist Unit
PSM 0	Rev 05	740-032022	ZJ00018	DC 12V Power Supply
PSM 1	Rev 04	740-032022	ZC00052	DC 12V Power Supply
PSM 2	Rev 04	740-032022	ZD00051	DC 12V Power Supply
PSM 3	Rev 05	740-032022	ZJ00060	DC 12V Power Supply
CCG 0	REV 04	750-030653	EG3703	Clock Generator
CCG 1	REV 04	750-030653	EG3698	Clock Generator
Routing Engine 0	REV 05	740-026942	P737A-002231	RE-DUO-2600
ad0 3823 MB	SMART CF		201006190039C02DC02D	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		2011042300CF4C6B4C6B	Disk 1
Routing Engine 1	REV 06	740-026942	P737A-002438	RE-DUO-2600
ad0 3823 MB	SMART CF		20100619053455F055F0	Compact Flash
ad1 62720 MB	SMART Lite SATA Drive		20110423000AE8E7E8E7	Disk 1
CB 0	REV 08	750-030625	EG5519	Control Board
CB 1	REV 08	750-030625	EG5516	Control Board
FPC 0	REV 18	750-036844	EJ3080	FPC
CPU	REV 12	711-030686	EJ3260	SNG PMB
FPC 2	REV 13	750-036844	EG5065	FPC
CPU	REV 09	711-030686	EG4082	SNG PMB
PIC 0	REV 14	750-031913	EG5127	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	143363A00240	SFP+-10G-SR
Xcvr 1	REV 01	740-031981	UK90PZ1	SFP+-10G-LR
Xcvr 2	REV 01	740-031980	AD1141A04XH	SFP+-10G-SR
Xcvr 3	REV 01	740-031981	UK90Q46	SFP+-10G-LR
Xcvr 4	REV 01	740-031980	AD1141A04X4	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	B11H02560	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11C01589	SFP+-10G-SR
Xcvr 8	REV 01	740-031980	AD1141A04XF	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	123363A01094	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LKF	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	183363A01528	SFP+-10G-SR
Xcvr 14	REV 01	740-031980	193363A01079	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	AK80MC8	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	AJC0BHC	SFP+-10G-SR
Xcvr 19	REV 01	740-021309	J08D26856	SFP+-10G-LR
Xcvr 21	REV 01	740-031980	AK80KCT	SFP+-10G-SR
Xcvr 22	REV 01	740-031981	UK90PZL	SFP+-10G-LR
Xcvr 23	REV 01	740-031980	AK80N1V	SFP+-10G-SR
FPC 3	REV 13	750-036844	EG5074	FPC
CPU	REV 09	711-030686	EG4064	SNG PMB
PIC 1	REV 10	750-031903	EG0325	SNG Load
FPC 5	REV 06	750-036844	EH3198	FPC
CPU				
PIC 0	REV 14	750-031913	EG5134	24x 10GE(LAN) SFP+
Xcvr 0	REV 01	740-031980	AK80LBH	SFP+-10G-SR
Xcvr 1	REV 01	740-031980	B11B03724	SFP+-10G-SR
Xcvr 2	REV 01	740-031980	AK80FMH	SFP+-10G-SR
Xcvr 5	REV 01	740-031980	B11J00818	SFP+-10G-SR
Xcvr 6	REV 01	740-031980	193363A00743	SFP+-10G-SR
Xcvr 7	REV 01	740-031980	B11B06125	SFP+-10G-SR
Xcvr 10	REV 01	740-031980	B11H02529	SFP+-10G-SR
Xcvr 11	REV 01	740-031980	AK80LFB	SFP+-10G-SR
Xcvr 12	REV 01	740-031980	193363A01061	SFP+-10G-SR
Xcvr 15	REV 01	740-031980	B11J00687	SFP+-10G-SR
Xcvr 16	REV 01	740-031980	193363A00738	SFP+-10G-SR
Xcvr 18	REV 01	740-031980	AK80MQX	SFP+-10G-SR

Xcvr 19	REV 01	740-021309	J08C17257	SFP+-10G-LR
Xcvr 22	REV 01	740-031980	B11J00730	SFP+-10G-SR
Xcvr 23	REV 01	740-031980	AK80KEE	SFP+-10G-SR
PIC 1	REV 08	750-036710	EG3105	2x 40GE CFP
Xcvr 0	REV 01	740-034554	B260HLT	CFP-40G-LR4
Xcvr 1	REV 01	740-034554	B11C02847	CFP-40G-LR4
FPC 6	REV 18	750-036844	EJ4391	FPC
CPU	REV 12	711-030686	EJ3257	SNG PMB
FPC 7	REV 18	750-036844	EJ4382	FPC
CPU	REV 12	711-030686	EJ3238	SNG PMB
SPMB 0	REV 10	711-030686	EG5418	SNG PMB
SPMB 1	REV 09	711-030686	EG5373	SNG PMB
SIB 0	REV 07	750-030631	EG4858	SIB-I-8S
SIB 1	REV 07	750-030631	EG4872	SIB-I-8S
SIB 2	REV 07	750-030631	EG4866	SIB-I-8S
SIB 3	REV 07	750-030631	EG6011	SIB-I-8S
SIB 4	REV 07	750-030631	EG4907	SIB-I-8S
SIB 5	REV 07	750-030631	EG4879	SIB-I-8S
SIB 6	REV 07	750-030631	EG4864	SIB-I-8S
SIB 7	REV 07	750-030631	EG4899	SIB-I-8S
SIB 8	REV 07	750-030631	EG4880	SIB-I-8S
Fan Tray 0	REV 04	760-032784	EG1496	Vertical Fan Tray
Fan Tray 1	REV 04	760-030642	EG1335	Horizontal Fan Tray
Fan Tray 2	REV 02	760-030642	ED4952	Horizontal Fan Tray

#### show chassis hardware detail (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis hardware detail
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1204FC0AJA   PTX5000
Midplane      REV 11   750-035893   ACAB8038      Midplane-8S
FPM           REV 12   760-030647   BBBD5619      Front Panel
Display
PDU 0         Rev 04   740-048336   1GB93470043   High Capacity DC PDU
PSM 0         Rev 04   740-046988   1GB63500184   High Capacity DC PSM
PSM 2         Rev 04   740-046988   1GB63500169   High Capacity DC PSM
PSM 4         Rev 04   740-046988   1GB63500306   High Capacity DC PSM
PSM 6         Rev 04   740-046988   1GB63500074   High Capacity DC PSM
PDU 1         Rev 04   740-048336   1GB93470045   High Capacity DC PDU
PSM 1         Rev 04   740-046988   1GB63500193   High Capacity DC PSM
PSM 3         Rev 04   740-046988   1GB63500143   High Capacity DC PSM
PSM 5         Rev 04   740-046988   1GB63500146   High Capacity DC PSM
PSM 7         Rev 04   740-046988   1GB63500192   High Capacity DC PSM
CCG 0         REV 09   750-030653   BBBC1909      Clock Generator
CCG 1         REV 09   750-030653   BBBD2970      Clock Generator
...

```

#### show chassis hardware models (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
FPM           REV 08   760-030647   EG1679        CRAFT-PTX5000-S
PDU 0         Rev 05   740-032019   ZE00006       PWR-SAN-PDU-DC
PSM 0         Rev 05   740-032022   ZJ00018       PSM-PTX-DC-120-S
PSM 1         Rev 04   740-032022   ZC00052       PWR-SAN-12-DC
PSM 2         Rev 04   740-032022   ZD00051       PWR-SAN-12-DC
PSM 3         Rev 05   740-032022   ZJ00060       PSM-PTX-DC-120-S
CCG 0         REV 04   750-030653   EG3703        CCG-PTX-S
CCG 1         REV 04   750-030653   EG3698        CCG-PTX-S

```

Routing Engine 0	REV 05	740-026942	P737A-002231	RE-DUO-C2600-16G-S
Routing Engine 1	REV 06	740-026942	P737A-002438	RE-DUO-C2600-16G-S
CB 0	REV 08	750-030625	EG5519	CB-PTX-S
CB 1	REV 08	750-030625	EG5516	CB-PTX-S
FPC 0	REV 18	750-036844	EJ3080	FPC-PTX-P1-A
FPC 2	REV 13	750-036844	EG5065	FPC-PTX-P1-A
PIC 0	REV 14	750-031913	EG5127	P1-PTX-24-10GE-SFPP
FPC 3	REV 13	750-036844	EG5074	FPC-PTX-P1-A
FPC 5				
PIC 0	REV 14	750-031913	EG5134	P1-PTX-24-10GE-SFPP
FPC 6	REV 18	750-036844	EJ4391	FPC-PTX-P1-A
FPC 7	REV 18	750-036844	EJ4382	FPC-PTX-P1-A
SIB 0	REV 07	750-030631	EG4858	SIB-I-PTX5008
SIB 1	REV 07	750-030631	EG4872	SIB-I-PTX5008
SIB 2	REV 07	750-030631	EG4866	SIB-I-PTX5008
SIB 3	REV 07	750-030631	EG6011	SIB-I-PTX5008
SIB 4	REV 07	750-030631	EG4907	SIB-I-PTX5008
SIB 5	REV 07	750-030631	EG4879	SIB-I-PTX5008
SIB 6	REV 07	750-030631	EG4864	SIB-I-PTX5008
SIB 7	REV 07	750-030631	EG4899	SIB-I-PTX5008
SIB 8	REV 07	750-030631	EG4880	SIB-I-PTX5008
Fan Tray 1	REV 04	760-030642	EG1335	FAN-PTX-H-S

#### show chassis hardware models (PTX5000 Packet Transport Router with FPC2-PTX-P1A)

```

user@host> show chassis hardware models
Hardware inventory:
Item          Version  Part number  Serial number  FRU model number
Midplane      REV 11    750-035893   ACAB8038       CHAS-MP-PTX5000-S
FPM           REV 12    760-030647   BBBD5619       CRAFT-PTX5000-S
PDU 0         Rev 04    740-048336   1GB93470043    PDU2-PTX-DC-S
  PSM 0        Rev 04    740-046988   1GB63500184    PSM2-PTX-DC-S
  PSM 2        Rev 04    740-046988   1GB63500169    PSM2-PTX-DC-S
  PSM 4        Rev 04    740-046988   1GB63500306    PSM2-PTX-DC-S
  PSM 6        Rev 04    740-046988   1GB63500074    PSM2-PTX-DC-S
PDU 1         Rev 04    740-048336   1GB93470045    PDU2-PTX-DC-S
  PSM 1        Rev 04    740-046988   1GB63500193    PSM2-PTX-DC-S
  PSM 3        Rev 04    740-046988   1GB63500143    PSM2-PTX-DC-S
  PSM 5        Rev 04    740-046988   1GB63500146    PSM2-PTX-DC-S
  PSM 7        Rev 04    740-046988   1GB63500192    PSM2-PTX-DC-S
CCG 0         REV 09    750-030653   BBBC1909       CCG-PTX-S
CCG 1         REV 09    750-030653   BBBD2970       CCG-PTX-S
...

```

#### show chassis hardware extensive (PTX5000 Packet Transport Router)

```

user@host> show chassis hardware extensive
Hardware inventory:
Item          Version  Part number  Serial number  Description
.....
PDU 0         Rev 04    740-032019   UE0003         DC Power Dist Unit
Jedec Code:   0x7fb0          EEPROM Version: 0x02
P/N:          740-032019        S/N:          UE0003
Assembly ID:  0x043d          Assembly Version: 04.00
Date:         11-29-2010      Assembly Flags: 0x00
Version:      Rev 04          CLEI Code:    032022XXXX
ID: DC Power Dist Unit        FRU Model Number: PWR-SAN-PDU-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 3d 04 00 52 65 76 20 30 34 00 00

```



```

Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 31 39 00 00
Address 0x20: 53 2f 4e 20 55 45 30 30 30 33 00 00 00 1d 0b 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 50 44 55 2d 44 43 00 00 00
Address 0x60: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x70: 00 00 00 a3 ff ff ff ff ff ff ff ff ff ff ff
PSM 0          Rev 04    740-032022  YG00065          DC 12V Power Supply
Module
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           740-032022      S/N:              YG00065
Assembly ID:   0x0440          Assembly Version:  04.00
Date:          07-30-2010      Assembly Flags:    0x00
Version:       Rev 04          CLEI Code:         032022XXXX
ID: DC 12V Power Supply Module FRU Model Number: PWR-SAN-12-DC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 04 40 04 00 52 65 76 20 30 34 00 00
Address 0x10: 00 00 00 00 37 34 30 2d 30 33 32 30 32 32 00 00
Address 0x20: 53 2f 4e 20 59 47 30 30 30 36 35 00 00 1e 07 07
Address 0x30: da ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 30 33 32 30 32 32 58 58 58 58 50
Address 0x50: 57 52 2d 53 41 4e 2d 31 32 2d 44 43 20 20 20 20
Address 0x60: 20 20 20 20 20 20 01 00 ff ff ff ff ff ff ff ff
Address 0x70: ff ff ff 0c ff ff ff ff ff ff ff ff ff ff ff ff

```

#### show chassis hardware (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               JN1100FB1AFB  MX480
Midplane      REV 05   710-017414   TR3310         MX480 Midplane
FPM Board     REV 02   710-017254   KG1872         Front Panel Display
PEM 2         Rev 02   740-017343   QCS0812A00N    DC Power Entry Module
PEM 3         Rev 02   740-017343   QCS0812A00U    DC Power Entry Module
Routing Engine 0 REV 07   740-015113   1000740938     RE-S-1300
CB 0          REV 03   710-021523   KF4630         MX SCB
FPC 1         REV 11   750-037207   ZW9726         AS-MCC
CPU           REV 04   711-038173   ZW4819         AS-MCC PMB
MIC 0         REV 06   750-037214   ZW3574         AS-MSC
PIC 0         BUILTIN BUILTIN      AS-MSC
MIC 1         REV 00   750-037211   BUILTIN        AS-MXC
PIC 2         BUILTIN BUILTIN      AS-MXC

```

#### show chassis hardware extensive (MX Routers with Media Services Blade [MSB])

```

user@switch> show chassis hardware extensive
FPC 1          REV 11   750-037207   ZW9726         AS-MCC
Jedec Code:    0x7fb0          EEPROM Version:    0x02
P/N:           750-037207      S/N:              ZW9726
Assembly ID:   0x0b37          Assembly Version:  01.11
Date:          02-17-2012      Assembly Flags:    0x00
Version:       REV 11          CLEI Code:         PROTOXCLEI
ID: AS-MCC     FRU Model Number: 750-037207
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 37 01 0b 52 45 56 20 31 31 00 00

```

```

Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 30 37 00 00
Address 0x20: 53 2f 4e 20 5a 57 39 37 32 36 00 00 00 11 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 30 37 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 31 31 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 5e ff ff ff ff ff ff ff ff ff ff ff ff
CPU          REV 04    711-038173    ZW4819          AS-MCC-PMB
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         711-038173      S/N:         ZW4819
Assembly ID: 0x0b38          Assembly Version: 01.04
Date:        12-30-2011      Assembly Flags: 0x00
Version:     REV 04
ID: AS-MCC PMB
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0b 38 01 04 52 45 56 20 30 34 00 00
Address 0x10: 00 00 00 00 37 31 31 2d 30 33 38 31 37 33 00 00
Address 0x20: 53 2f 4e 20 5a 57 34 38 31 39 00 00 00 1e 0c 07
Address 0x30: db ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 31 31 2d 30 33 38 31 37 33 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 34 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 00 00 00 00 00 00 00 00 00 00 00 00
MIC 0          REV 06    750-037214    ZW3574          AS-MS
Jedec Code:  0x7fb0          EEPROM Version:  0x02
P/N:         750-037214      S/N:         ZW3574
Assembly ID: 0x0a44          Assembly Version: 01.06
Date:        02-19-2012      Assembly Flags: 0x00
Version:     REV 06          CLEI Code:      PROTOXCLEI
ID: AS-MS      FRU Model Number: 750-037214
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 02 ff 0a 44 01 06 52 45 56 20 30 36 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 34 00 00
Address 0x20: 53 2f 4e 20 5a 57 33 35 37 34 00 00 00 13 02 07
Address 0x30: dc ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 01 50 52 4f 54 4f 58 43 4c 45 49 37
Address 0x50: 35 30 2d 30 33 37 32 31 34 00 00 00 00 00 00 00
Address 0x60: 00 00 00 00 00 00 30 36 00 ff ff ff ff ff ff ff
Address 0x70: ff ff ff 60 c0 03 e5 f4 00 00 00 00 00 00 00 00
PIC 0          BUILTIN    BUILTIN          AS-MS
MIC 1          REV 00    750-037211          AS-MXC
Jedec Code:  0x7fb0          EEPROM Version:  0x01
P/N:         750-037211
Assembly ID: 0x0a43          Assembly Version: 01.00
Date:        255-255-65535    Assembly Flags: 0x00
Version:     REV 00
ID: AS-MXC
Board Information Record:
Address 0x00: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
I2C Hex Data:
Address 0x00: 7f b0 01 ff 0a 43 01 00 52 45 56 20 30 30 00 00
Address 0x10: 00 00 00 00 37 35 30 2d 30 33 37 32 31 31 00 00
Address 0x20: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Address 0x30: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x40: ff ff ff ff 00 ff ff ff ff ff ff ff ff ff ff ff
Address 0x50: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff
Address 0x60: ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff

```

```

Address 0x70: ff ff ff ff c0 02 e6 6c 7f b0 02 ff 0a 44 01 06
PIC 2                BUILTIN        BUILTIN        AS-MXC

```

### show chassis hardware (QFX3500 Switch running Enhanced Layer 2 Software)

```

user@switch> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis
Pseudo CB 0
Routing Engine 0    BUILTIN    BUILTIN      QFX3500
FPC 0              REV 16     750-036931   P3566-C        QFX3500-48S4Q
CPU                BUILTIN    BUILTIN      FPC CPU
PIC 0              BUILTIN    BUILTIN      48x 10G-SFP+
  Xcvr 12          REV 01     740-030658   AD1125A0438    SFP+-10G-USR
  Xcvr 13          REV 01     740-030658   AD1125A02GN    SFP+-10G-USR
PIC 1              BUILTIN    BUILTIN      4x 40G-QSFP+
PIC 2
MGMT BRD           REV 10     750-036946   BBAW0328       QFX3500-MGMT-RJ45-AFI
Power Supply 0     Rev 05     740-032091   WA13035        JPSU-650W-AC-AFI
Power Supply 1
Fan Tray 0
  to Back Airflow  QFX3500 Fan Tray, Front
Fan Tray 1
  to Back Airflow  QFX3500 Fan Tray, Front
Fan Tray 2
  to Back Airflow  QFX3500 Fan Tray, Front

```

### show chassis hardware (QFX5100 Switch running Enhanced Layer 2 Software)

```

user@switch> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis
Pseudo CB 0
Routing Engine 0    BUILTIN    BUILTIN      QFX5100-24Q-2P
FPC 0              REV 02     650-049942   TB3113280048   QFX5100-24Q-2P
CPU                BUILTIN    BUILTIN      FPC CPU
PIC 0              BUILTIN    BUILTIN      24x 40G-QSFP
  Xcvr 8           REV 01     740-032986   QA470143        QSFP+-40G-SR4
  Xcvr 14          REV 01     740-032986   QB500525        QSFP+-40G-SR4
PIC 1              REV 02     611-049555   RR3113310169   QFX-EM-4Q
  Xcvr 0           REV 01     740-032986   QC440904        QSFP+-40G-SR4
  Xcvr 1           REV 01     740-032986   QB240154        QSFP+-40G-SR4
  Xcvr 2           REV 01     740-035085   018110105       QSFP+-40G-LPBK
PIC 2              REV 02     611-049555   RR3113310209   QFX-EM-4Q
  Xcvr 0           REV 01     740-032986   QB190270        QSFP+-40G-SR4
  Xcvr 1           REV 01     740-035085   018110063       QSFP+-40G-LPBK
  Xcvr 2           REV 01     740-032986   QB210034        QSFP+-40G-SR4
Power Supply 0     REV 03     740-041741   1GA23110973    JPSU-650W-AC-AFO
Power Supply 1     REV 03     740-041741   1GA23090878    JPSU-650W-AC-AFO
Fan Tray 0
  to Back Airflow - AFO  QFX5100 Fan Tray 0, Front
Fan Tray 1
  to Back Airflow - AFO  QFX5100 Fan Tray 1, Front
Fan Tray 2
  to Back Airflow - AFO  QFX5100 Fan Tray 2, Front
Fan Tray 3
  to Back Airflow - AFO  QFX5100 Fan Tray 3, Front

```

Fan Tray 4  
to Back Airflow - AFO

QFX5100 Fan Tray 4, Front

## show chassis in-service-upgrade

**Syntax** `show chassis in-service-upgrade`

**Release Information** Command introduced in Junos OS Release 9.0.  
 Command introduced in Junos OS Release 12.3R2, 13.1R2, and 13.2R1 for TX Matrix Plus routers.  
 Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.  
 Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.  
 Command introduced in Junos OS Release 13.2 for PTX5000 routers.  
 Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

**Description** Display the status of Flexible PIC Concentrators (FPCs) and their corresponding PICs after the most recent unified in-service software upgrade (ISSU). This command must be issued on the master Routing Engine.



**NOTE:** Only Intelligent Queuing (IQ) PICs are displayed by this command output. Unified ISSU status for other PIC types is controlled internally by the FPC.

**Options** This command has no options.

**Required Privilege Level** view

**Related Documentation**

- *request system software abort*
- [request system software in-service-upgrade on page 436](#)
- *Unified ISSU Concepts*
- *Performing a Unified ISSU*

**List of Sample Output** [show chassis in-service-upgrade on page 852](#)  
[show chassis in-service-upgrade \(MX2010 Router\) on page 852](#)  
[show chassis in-service-upgrade \(MX2020 Router\) on page 852](#)  
[show chassis in-service-upgrade \(TX Matrix Plus Router\) on page 853](#)  
[show chassis in-service-upgrade \(QFX5100 Switch\) on page 854](#)

**Output Fields** [Table 36 on page 851](#) lists the output fields for the `show chassis in-service-upgrade` command. Output fields are listed in the approximate order in which they appear.

**Table 36: show chassis in-service-upgrade Output Fields**

Field Name	Field Description
Item	Flexible PIC Concentrator (FPC) slot number.

Table 36: show chassis in-service-upgrade Output Fields (*continued*)

Field Name	Field Description
<b>Status</b>	FPC and corresponding PIC state. State can be either of the following: <ul style="list-style-type: none"> <li>• <b>Online</b>—FPC is online and running.</li> <li>• <b>Offline</b>—FPC is powered down.</li> </ul>
<b>Reason</b>	Reason for the state (if offline).

## Sample Output

### show chassis in-service-upgrade

```

user@host> show chassis in-service-upgrade
Item           Status           Reason
FPC 0          Online
FPC 1          Online
FPC 2          Online
  PIC 0        Online
  PIC 1        Online
FPC 3          Offline          Offlined by CLI command
FPC 4          Online
  PIC 1        Online
FPC 5          Online
  PIC 0        Online
FPC 6          Online
  PIC 3        Online
FPC 7          Online

```

### show chassis in-service-upgrade (MX2010 Router)

```

user@host> show chassis in-service-upgrade
Item           Status           Reason
FPC 0          Online
FPC 1          Online
FPC 8          Online
FPC 9          Online

```

### show chassis in-service-upgrade (MX2020 Router)

```

user@host> show chassis in-service-upgrade
Item           Status           Reason
FPC 0          Online
FPC 1          Online
FPC 2          Online
FPC 3          Online
FPC 4          Online
FPC 5          Online
FPC 6          Online
FPC 7          Online
FPC 8          Online
FPC 9          Online
FPC 10         Online
FPC 11         Online
FPC 12         Online
FPC 13         Online

```

```

FPC 14      Online
FPC 15      Online
FPC 16      Online
FPC 17      Online
FPC 18      Online
FPC 19      Online

```

### show chassis in-service-upgrade (TX Matrix Plus Router)

```

user@host> show chassis in-service-upgrade
1cc0-re0:

```

Item	Status	Reason
FPC 1	Online	
PIC 0	Online	
FPC 2	Online	
FPC 3	Online	
PIC 1	Online	
FPC 4	Online	
FPC 6	Online	
FPC 7	Online	

```

1cc1-re0:

```

Item	Status	Reason
FPC 0	Online	
PIC 3	Online	
FPC 1	Online	
FPC 2	Online	
FPC 4	Online	
FPC 6	Online	
FPC 7	Online	

```

1cc2-re0:

```

Item	Status	Reason
FPC 0	Online	
FPC 2	Online	
FPC 3	Online	
PIC 0	Online	
FPC 4	Online	
FPC 6	Online	
FPC 7	Online	
PIC 1	Online	

```

1cc3-re0:

```

Item	Status	Reason
FPC 0	Online	
PIC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Online	
PIC 2	Online	
FPC 4	Online	
FPC 5	Online	
FPC 6	Online	
FPC 7	Online	
PIC 1	Online	

### show chassis in-service-upgrade (QFX5100 Switch)

```
user@switch> show chassis in-service-upgrade
```

Item	Status	Reason
FPC 0	Online (ISSU)	



## show chassis lcd

<b>List of Syntax</b>	<a href="#">show chassis lcd (EX Series) on page 855</a> <a href="#">show chassis lcd (QFX Series and QFabric Systems) on page 855</a>
<b>show chassis lcd (EX Series)</b>	<pre>show chassis lcd &lt;fpc-slot <i>fpc-slot-number</i>&gt; &lt;menu &lt;(all-members   local   member <i>member-id</i>)&gt;&gt;</pre>
<b>show chassis lcd (QFX Series and QFabric Systems)</b>	<pre>show chassis lcd &lt;fpc-slot <i>fpc-slot-number</i>&gt; &lt;interconnect-device <i>device-id</i>&gt; &lt;node-device <i>device-id</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>menu</b> option introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.1 for QFabric systems.</p>
<b>Description</b>	<p>Display the information that appears on the LCD panel of EX3200, EX3300, EX4200, EX4500, EX6200, and EX8200 switches, XRE200 External Routing Engines, QFX Series standalone switches, and Interconnect devices and Node devices within a QFabric system. Display the status of the currently selected port parameter of the Status LED for each network port on the device.</p>
<b>Options</b>	<p><b>none</b>—Display the information that appears on the LCD panel (for any EX Series member switch in a Virtual Chassis or for XRE200 External Routing Engines, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.</p> <p><b>fpc-slot &lt;<i>fpc-slot-number</i>&gt;</b>—(Optional) Display the information as follows:</p> <ul style="list-style-type: none"> <li>(EX3200, EX3300, EX4200, and EX4500 switches, or the QFX Series) Display the information that appears on the LCD panel for either an FPC slot with no <i>fpc-slot-number</i> value specified or for the FPC slot specified by <b>fpc-slot 0</b>. <b>fpc-slot</b> refers to the switch itself and <b>0</b> is the only valid value for <i>fpc-slot-number</i>. Output for these options is the same as for the <b>none</b> option.</li> </ul> <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> <li>(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) If no <i>fpc-slot-number</i> value is specified, display the information that appears on the LCD panel for all members of the Virtual Chassis. Output for this option is the same as for the <b>none</b> option. If the <i>fpc-slot-number</i> value is specified (it equals the <i>member-id</i> value), display the information for the specified member.</li> </ul> <p>Also display the status of the currently selected port parameter of the Status LED for each network port.</p> <ul style="list-style-type: none"> <li>(EX6200 or EX8200 switches)—Display the information that appears on the LCD panel for the line card in the line-card slot specified by the <i>fpc-slot-number</i> value.</li> </ul>

Also display the status of the currently selected port parameter of the Status LED for each network port.

**interconnect-device *device-id***—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Interconnect device.

**menu**—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

**menu all-members**—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

**menu local**—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

**menu member *member-id***—(EX Series Virtual Chassis member switches or XRE200 External Routing Engines) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

**node-device *device-id***—(QFabric systems only) (Optional) Display the front panel contents and LED status of all the ports on the Node device.

**Required Privilege Level**

view

**Related Documentation**

- *LCD Panel in EX3200 Switches*
- *LCD Panel in EX4200 Switches*
- *LCD Panel in EX4500 Switches*
- *LCD Panel in an EX8200 Switch*
- *LCD Panel in an XRE200 External Routing Engine*
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- *set chassis display message*

**List of Sample Output**

[show chassis lcd \(Two-Member EX4200 Virtual Chassis\) on page 857](#)  
[show chassis lcd fpc-slot 1 \(EX4200 Virtual Chassis\) on page 859](#)  
[show chassis lcd \(EX8200 Switch\) on page 859](#)  
[show chassis lcd fpc-slot 2 \(EX8200 Switch\) on page 861](#)  
[show chassis lcd menu \(EX4200 Switch\) on page 861](#)  
[show chassis lcd menu \(EX8200 Switch\) on page 861](#)  
[show chassis lcd \(QFX3500 Switches\) on page 862](#)  
[show chassis lcd \(XRE200 External Routing Engine in EX8200 Virtual Chassis\) on page 862](#)  
[show chassis lcd interconnect-device \(QFabric Systems\) on page 865](#)

[show chassis lcd node-device \(QFabric Systems\) on page 867](#)

**Output Fields** [Table 37 on page 857](#) lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

**Table 37: show chassis lcd Output Fields**

Field Name	Field Description
<b>membernumber</b> (XRE200 External Routing Engine)	Member ID of the device whose content is being displayed.
<b>Front panel contents for slot</b>	FPC slot number of the switch whose content is being displayed. The number is always <b>0</b> , except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
<b>Front panel contents</b> (EX6200, EX8200 switch, XRE200 External Routing Engine, and QFX Series)	<p>On EX6200 switches, EX8200 switches, and XRE200 External Routing Engines, no slot number is displayed.</p> <p>On XRE200 External Routing Engines, this field appears under the <b>member number</b> field for each member device in the EX8200 Virtual Chassis.</p>
<b>LCD screen</b>	<p>The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for EX8200 switches, displays <b>RE</b> and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are:</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>POE</b>—Power over Ethernet (EX3200 and EX4200 switches only)</li> </ul>
<b>LEDs status</b>	Current state of the Alarms, System, and Master LEDs (chassis status LEDs).
<b>Interface</b>	Names of the interfaces on the switch.
<b>LED (ADM/SPD/DPX/POE)</b>	<p>State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p><b>NOTE:</b> The XRE200 External Routing Engine always displays the <b>NA</b> parameter. The QFX Series products do not have any of the port parameters listed below.</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>NA</b>—Not applicable.</li> <li>• <b>POE</b>—Power over Ethernet</li> </ul>
<b>fpcx</b>	On standalone EX Series and QFX Series switches, always <b>0</b> . On EX Series Virtual Chassis member switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

## Sample Output

### show chassis lcd (Two-Member EX4200 Virtual Chassis)

```
user@switch> show chassis lcd
```

## Front panel contents for slot: 0

-----  
LCD screen:  
    00:BK switch1  
    LED:SPD ALARM 00  
LEDs status:  
    Alarms LED: Off  
    System LED: Green  
    Master LED: Off  
Interface        LED(ADM/SPD/DPX/POE)  
-----  
ge-0/0/0          Off  
ge-0/0/1          Off  
ge-0/0/2          Off  
ge-0/0/3          Off  
ge-0/0/4          Off  
ge-0/0/5          Off  
ge-0/0/6          Off  
ge-0/0/7          Off  
ge-0/0/8          Off  
ge-0/0/9          Off  
ge-0/0/10         Off  
ge-0/0/11         Off  
ge-0/0/12         Off  
ge-0/0/13         Off  
ge-0/0/14         Off  
ge-0/0/15         Off  
ge-0/0/16         Off  
ge-0/0/17         Off  
ge-0/0/18         Off  
ge-0/0/19         Off  
ge-0/0/20         Off  
ge-0/0/21         Off  
ge-0/0/22         Off  
ge-0/0/23         Off

## Front panel contents for slot: 1

-----  
LCD screen:  
    01:RE switch2  
    LED:SPD ALARM 01  
LEDs status:  
    Alarms LED: Yellow  
    System LED: Green  
    Master LED: Green  
Interface        LED(ADM/SPD/DPX/POE)  
-----  
ge-1/0/0          Off  
ge-1/0/1          Off  
ge-1/0/2          Off  
ge-1/0/3          Off  
ge-1/0/4          Off  
ge-1/0/5          Off  
ge-1/0/6          Off  
ge-1/0/7          Off  
ge-1/0/8          Off  
ge-1/0/9          Off  
ge-1/0/10         Off  
ge-1/0/11         Off  
ge-1/0/12         Off  
ge-1/0/13         Off  
ge-1/0/14         Off

```

ge-1/0/15      Off
ge-1/0/16      Off
ge-1/0/17      Off
ge-1/0/18      Off
ge-1/0/19      Off
ge-1/0/20      Off
ge-1/0/21      Off
ge-1/0/22      Off
ge-1/0/23      Off

```

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

#### show chassis lcd fpc-slot 1 (EX4200 Virtual Chassis)

```

user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
-----
LCD screen:
  01:RE switch2
  LED:SPD ALARM 01
LEDs status:
  Alarms LED: Yellow
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-1/0/0      Off
ge-1/0/1      Off
ge-1/0/2      Off
ge-1/0/3      Off
ge-1/0/4      Off
ge-1/0/5      Off
ge-1/0/6      Off
ge-1/0/7      Off
ge-1/0/8      Off
ge-1/0/9      Off
ge-1/0/10     Off
ge-1/0/11     Off
ge-1/0/12     Off
ge-1/0/13     Off
ge-1/0/14     Off
ge-1/0/15     Off
ge-1/0/16     Off
ge-1/0/17     Off
ge-1/0/18     Off
ge-1/0/19     Off
ge-1/0/20     Off
ge-1/0/21     Off
ge-1/0/22     Off
ge-1/0/23     Off

```

#### show chassis lcd (EX8200 Switch)

```

user@switch> show chassis lcd
Front panel contents:
-----
LCD screen:
  RE st-8200-r
  LED:ADM ALARM 01

```

## LEDs status:

Alarms LED: Yellow

System LED: Yellow

Master LED: Green

Interface	LED(ADM/SPD/DPX)
-----------	------------------

-----	
ge-0/0/0	Off
ge-0/0/1	Off
ge-0/0/2	Off
ge-0/0/3	Off
ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	Off
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	Off
ge-0/0/41	Off
ge-0/0/42	Off
ge-0/0/43	Off
ge-0/0/44	Off
ge-0/0/45	Off
ge-0/0/46	Off
ge-0/0/47	Off
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off

```

xe-2/0/7      Off
xe-3/0/0      Off
xe-3/0/1      Off
xe-3/0/2      Off
xe-3/0/3      Off
xe-3/0/4      Off
xe-3/0/5      Off
xe-3/0/6      Off
xe-3/0/7      Off
xe-5/0/0      Off
xe-5/0/1      Off
xe-5/0/2      Off
xe-5/0/3      Off
xe-5/0/4      Off
xe-5/0/5      Off
xe-5/0/6      On
xe-5/0/7      On
xe-7/0/5      Off

```

#### show chassis lcd fpc-slot 2 (EX8200 Switch)

```
show chassis lcd fpc-slot 2
```

Interface	LED (ADM/SPD/DPX)
xe-2/0/0	Off
xe-2/0/1	Off
xe-2/0/2	Off
xe-2/0/3	Off
xe-2/0/4	Off
xe-2/0/5	Off
xe-2/0/6	Off
xe-2/0/7	Off

#### show chassis lcd menu (EX4200 Switch)

```
user@switch> show chassis lcd menu
fpc0:
```

```

-----
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default

```

On an EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu** **all-members** command is the same as the output for the **show chassis lcd menu** command.

#### show chassis lcd menu (EX8200 Switch)

```

user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu

```

```
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

### show chassis lcd (QFX3500 Switches)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
-----
LCD screen:
00:RE switch
ALARM 01
LEDs status:
Status/Beacon LED: Yellow Blinking
Interface STATUS LED ACTIVITY LED
-----
fte-0/1/0 Off Off
```

### show chassis lcd (XRE200 External Routing Engine in EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis lcd
member0:
-----
Front panel contents:
-----
LCD screen:
  RE ex8200-member0
  LED:ADM ALARM 04
LEDs status:
  Alarms LED: Red
  System LED: Yellow
  Master LED: Green

member1:
-----

member8:
-----
Front panel contents:
-----
LCD screen:
  BACKUP

member9:
-----
Front panel contents:
-----
LCD screen:
  09:RE xre200-member9
  LED: NA ALARM 01
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      On
ge-0/0/1      On
ge-0/0/2      On
ge-0/0/3      On
```



ge-0/0/4	Off
ge-0/0/5	Off
ge-0/0/6	Off
ge-0/0/7	Off
ge-0/0/8	Off
ge-0/0/9	Off
ge-0/0/10	On
ge-0/0/11	Off
ge-0/0/12	Off
ge-0/0/13	Off
ge-0/0/14	Off
ge-0/0/15	Off
ge-0/0/16	Off
ge-0/0/17	Off
ge-0/0/18	Off
ge-0/0/19	Off
ge-0/0/20	Off
ge-0/0/21	Off
ge-0/0/22	Off
ge-0/0/23	Off
ge-0/0/24	Off
ge-0/0/25	Off
ge-0/0/26	Off
ge-0/0/27	Off
ge-0/0/28	Off
ge-0/0/29	Off
ge-0/0/30	Off
ge-0/0/31	Off
ge-0/0/32	Off
ge-0/0/33	Off
ge-0/0/34	Off
ge-0/0/35	Off
ge-0/0/36	Off
ge-0/0/37	Off
ge-0/0/38	Off
ge-0/0/39	Off
ge-0/0/40	On
ge-0/0/41	On
ge-0/0/42	On
ge-0/0/43	On
ge-0/0/44	On
ge-0/0/45	On
ge-0/0/46	On
ge-0/0/47	On
ge-16/0/0	On
ge-16/0/1	Off
ge-16/0/2	On
ge-16/0/3	Off
ge-16/0/4	On
ge-16/0/5	Off
ge-16/0/6	On
ge-16/0/7	Off
ge-16/0/8	Off
ge-16/0/9	Off
ge-16/0/10	Off
ge-16/0/11	Off
ge-16/0/12	Off
ge-16/0/13	On
ge-16/0/14	Off
ge-16/0/15	On
ge-16/0/16	Off

ge-16/0/17	On
ge-16/0/18	On
ge-16/0/19	On
ge-16/0/20	On
ge-16/0/21	Off
ge-16/0/22	On
ge-16/0/23	Off
ge-16/0/24	Off
ge-16/0/25	Off
ge-16/0/26	On
ge-16/0/27	Off
ge-16/0/28	Off
ge-16/0/29	Off
ge-16/0/30	On
ge-16/0/31	Off
ge-16/0/32	On
ge-16/0/33	On
ge-16/0/34	On
ge-16/0/35	Off
ge-16/0/36	On
ge-16/0/37	Off
ge-16/0/38	Off
ge-16/0/39	Off
ge-16/0/40	Off
ge-16/0/41	Off
ge-16/0/42	On
ge-16/0/43	Off
ge-16/0/44	Off
ge-16/0/45	Off
ge-16/0/46	Off
ge-16/0/47	Off
xe-19/0/0	Off
xe-19/0/1	On
xe-19/0/2	On
xe-19/0/3	On
xe-19/0/4	On
xe-19/0/5	On
ge-22/0/0	Off
ge-22/0/1	Off
ge-22/0/2	On
ge-22/0/3	Off
ge-22/0/4	On
ge-22/0/5	On
ge-22/0/6	On
ge-22/0/7	On
ge-22/0/8	Off
ge-22/0/9	Off
ge-22/0/10	Off
ge-22/0/11	Off
ge-22/0/12	Off
ge-22/0/13	Off
ge-22/0/14	Off
ge-22/0/15	Off
ge-22/0/16	On
ge-22/0/17	Off
ge-22/0/18	On
ge-22/0/19	Off
ge-22/0/20	On
ge-22/0/21	Off
ge-22/0/22	On
ge-22/0/23	Off

```

ge-22/0/24      On
ge-22/0/25      Off
ge-22/0/26      Off
ge-22/0/27      Off
ge-22/0/28      Off
ge-22/0/29      Off
ge-22/0/30      Off
ge-22/0/31      Off
ge-22/0/32      On
ge-22/0/33      Off
ge-22/0/34      On
ge-22/0/35      Off
ge-22/0/36      Off
ge-22/0/37      Off
ge-22/0/38      Off
ge-22/0/39      Off
ge-22/0/40      Off
ge-22/0/41      Off
ge-22/0/42      Off
ge-22/0/43      Off
ge-22/0/44      Off
ge-22/0/45      Off
ge-22/0/46      Off
ge-22/0/47      Off

```

#### show chassis lcd interconnect-device (QFabric Systems)

```

show chassis lcd interconnect-device IC-F1012
      Front Panel Module Information
      -----
      LCD screen:
      IC-F1012      3 Alarms active

LEDs status:
  Status LED: Green
  Power LED : Green
  Major Alarm LED: off
  Minor Alarm LED: Yellow
  Fan 0 LED : Green
  Fan 1 LED : Green
  Fan 2 LED : Green
  Fan 3 LED : Green
  Fan 4 LED : Green
  Fan 5 LED : Green
  Fan 6 LED : Green
  Fan 7 LED : Green
  Fan 8 LED : Green
  Fan 9 LED : Green
  PEM 0 LED : Green
  PEM 1 LED : Green
  PEM 2 LED : Green
  PEM 3 LED : off
  PEM 4 LED : off
  PEM 5 LED : off

      LED info for: CB - 0
      -----

LEDs status:
  Status LED: Green
  Mastership LED: Green

Interface      STATUS LED      LINK/ACTIVITY LED

```

```

-----
IC-F1012:pme0 :      Green      N/A
IC-F1012:pme1 :      Green      N/A
IC-F1012:pme2 :      off        N/A
IC-F1012:pme3 :      off        N/A

```

```

LED info for: CB - 1
-----

```

```

LEDs status:
  Status LED: Green
  Mastership LED: Amber

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:pme0 :	Green	N/A
IC-F1012:pme1 :	Green	N/A
IC-F1012:pme2 :	off	N/A
IC-F1012:pme3 :	off	N/A

```

LED info for: FC 0 FPC - 0
-----

```

```

LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-0/0/0	Green	N/A
IC-F1012:fte-0/0/1	Green	N/A
IC-F1012:fte-0/0/2	Green	N/A
IC-F1012:fte-0/0/3	Green	N/A
IC-F1012:fte-0/0/4	Green	N/A

```

LED info for: FC 1 FPC - 1
-----

```

```

LEDs status:
  Status LED: Green

```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F1012:fte-1/0/0	Green	N/A
IC-F1012:fte-1/0/1	Green	N/A
IC-F1012:fte-1/0/2	Green	N/A
IC-F1012:fte-1/0/3	Green	N/A
IC-F1012:fte-1/0/4	Green	N/A

```

LED info for: RC 0 FPC - 8
-----

```

```

LEDs status:
  Status LED: Green

```

```

LED info for: RC 1 FPC - 9
-----

```

```

LEDs status:
  Status LED: Green

```

```

LED info for: RC 2 FPC - 10
-----

```

```

LEDs status:
  Status LED: Green

```

```

LED info for: RC 3 FPC - 11

```

```

-----
LEDs status:
  Status LED: Green

  LED info for: RC 4 FPC - 12
-----
LEDs status:
  Status LED: Green

  LED info for: RC 5 FPC - 13
-----
LEDs status:
  Status LED: Green

  LED info for: RC 6 FPC - 14
-----
LEDs status:
  Status LED: Green

  LED info for: RC 7 FPC - 15
-----
LEDs status:
  Status LED: Green

```

#### show chassis lcd node-device (QFabric Systems)

```

show chassis lcd node-device P3774-C
  Front panel contents for: P3774-C
-----
  LCD screen:
  P3774-C

LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	STATUS LED	LINK/ACTIVITY LED
P3774-C:xe-0/0/6	Green	Green
P3774-C:xe-0/0/7	Green	Green
P3774-C:ge-0/0/10	Green	Green
P3774-C:ge-0/0/11	Green	Green Blinking
P3774-C:ge-0/0/12	Green	Off
P3774-C:ge-0/0/13	Green	Green Blinking
P3774-C:ge-0/0/20	Green	Green
P3774-C:ge-0/0/21	Green	Green
P3774-C:ge-0/0/22	Green	Green Blinking
P3774-C:ge-0/0/23	Green	Off
P3774-C:ge-0/0/30	Green	Green
P3774-C:ge-0/0/31	Green	Green
P3774-C:ge-0/0/32	Green	Green Blinking
P3774-C:ge-0/0/33	Green	Green Blinking
P3774-C:fte-0/1/0	Green	Green
P3774-C:fte-0/1/1	Green	Green Blinking
P3774-C:fte-0/1/2	Green	Green Blinking
P3774-C:fte-0/1/3	Green	Green

## show chassis led

---

<b>List of Syntax</b>	<a href="#">show chassis led (EX Series) on page 868</a> <a href="#">show chassis led (QFX Series) on page 868</a>
<b>show chassis led (EX Series)</b>	<code>show chassis led</code> <code>&lt;fpc-slot &lt;<i>fpc-slot-number</i>&gt;&gt;</code>
<b>show chassis led (QFX Series)</b>	<code>show chassis led</code> <code>&lt;fpc-slot &lt;<i>fpc-slot-number</i>&gt;&gt;</code> <code>interconnect-device <i>name</i></code> <code>node-device <i>name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.1 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the status and colors of the chassis LEDs on the front panel of the switch. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.
<b>Options</b>	<p><b>none</b>—Display the status of the chassis status LEDs (for EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members).</p> <p><b>fpc-slot &lt;<i>fpc-slot-number</i>&gt;</b>—(Optional) (Not on EX2200 switches) Display the information as follows:</p> <ul style="list-style-type: none"><li>(EX3200, standalone EX4200, standalone QFX3500, and EX4500 switches) Display the status of the chassis status LEDs for either an FPC slot with no <i>fpc-slot-number</i> value specified or for the FPC slot specified by <b>fpc-slot 0</b>. <i>fpc-slot</i> refers to the switch itself and <b>0</b> is the only valid value for <i>fpc-slot-number</i>. Output for these options is the same as for the <b>none</b> option.</li><li>(EX4200 switches in a Virtual Chassis with two or more members) If no <i>fpc-slot-number</i> value is specified, display the status of the chassis status LEDs for all members of the Virtual Chassis. Output for this option is the same as for the <b>none</b> option. If the <i>fpc-slot-number</i> value is specified (it equals the <i>member-id</i> value), display the status of the chassis status LEDs for the specified member.</li><li>(EX8200 switches)—Display the status of the chassis status LEDs for the line card in the line-card slot specified by the <i>fpc-slot-number</i> value.</li></ul> <p><b>interconnect-device <i>name</i></b>—</p> <p>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Interconnect device.</p> <p><b>node-device <i>name</i></b>— (QFabric systems only) (Optional) Display the status of the chassis and interface status LEDs for the Node device.</p>

**Required Privilege Level** view

- Related Documentation**
- *Chassis Status LEDs in EX2200 Switches*
  - *Chassis Status LEDs in EX3200 Switches*
  - *Chassis Status LEDs in EX4200 Switches*
  - *Chassis Status LEDs in EX4500 Switches*
  - *Chassis Status LEDs in an EX8200 Switch*
  - *Chassis Status LEDs on a QFX3500 Device*
  - *Chassis Status LEDs in the QFX3600 and QFX3600-I Device*
  - *Management Port LEDs on a QFX3500 Device*
  - *Management Port LEDs in the QFX3600 and QFX3600-I Device*
  - *Chassis Status LEDs on a QFX3008-I Interconnect Device*
  - *Control Board LEDs on a QFX3008-I Interconnect Device*

**List of Sample Output**

[show chassis led \(EX2200 Switch\) on page 872](#)  
[show chassis led on page 873](#)  
[show chassis led fpc-slot 0 on page 874](#)  
[show chassis led \(EX Series\) on page 874](#)  
[show chassis led node-device \(QFabric System Node Device\) on page 875](#)  
[show chassis led interconnect-device \(QFabric System - QFX3600-I Interconnect Device\) on page 875](#)  
[show chassis led interconnect-device \(QFabric System - QFX3008-I Interconnect Device\) on page 876](#)

**Output Fields** [Table 26 on page 509](#) lists the output fields for the **show chassis led** command. Output fields are listed in the approximate order in which they appear.

**Table 38: show chassis led Output Fields**

Field Name	Field Description
<b>Front panel contents for slot</b>	FPC slot number of the device whose content is being displayed. The number is always 0, except for EX4200 switches in a Virtual Chassis, where it is the member ID value.
<b>Front panel contents</b> (EX8200 Switches)	
<b>Front Panel Module Information</b> (QFabric system QFX3008-I Interconnect device)	On EX8200 switches, no slot number is displayed.
<b>Front panel contents for</b> (QFabric system Node devices and QFX3600-I Interconnect devices)	On QFabric system Node devices, the name of the Node device whose content is being displayed.

Table 38: show chassis led Output Fields (*continued*)

Field Name	Field Description
<b>Alarms LED</b>	<p>(EX Series switches only) Displays status of the ALM LED:</p> <ul style="list-style-type: none"> <li>• Off—No alarm has been configured.</li> <li>• Green—No alarm has been triggered.</li> <li>• Red—Major alarm.</li> <li>• Yellow—Minor alarm</li> </ul>
<b>System LED</b>	<p>(EX Series switches only) Displays status of the SYS LED:</p> <ul style="list-style-type: none"> <li>• Off—Switch is powered off.</li> <li>• Green—Switch is operating normally.</li> <li>• Yellow—Switch is booting.</li> </ul>
<b>Master LED:</b>	<p>Displays status of the MST LED (on EX3200, EX4200, and EX8200 switches):</p> <ul style="list-style-type: none"> <li>• Green—On an EX4200 Virtual Chassis switch, indicates the switch is the master in the Virtual Chassis configuration. On other switches, indicates that the Routing Engine is operational.</li> <li>• Off <ul style="list-style-type: none"> <li>• On an EX4200 Virtual Chassis switch, indicates that this switch is not the master in the Virtual Chassis configuration.</li> <li>• On EX3200, standalone EX4200, and EX8200 switches, indicates that the Routing Engine is not operational.</li> </ul> </li> </ul>
<b>Mode LED:</b>	<p>(EX Series switches only) On an EX2200 switch only, displays the currently selected port parameter of the Status LED:</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>POE</b>—Power over Ethernet</li> </ul>
<b>Status/Beacon LED</b>	<p>(QFX Series only) Displays the system status as indicated by the Status LED on the chassis. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Chassis Status LEDs on a QFX3500 Device</i></li> <li>• <i>Chassis Status LEDs in the QFX3600 and QFX3600-I Device</i></li> </ul>
<b>LINK/SPEED LED</b>	<p>(QFX Series only) Displays the link status and speed of a management port. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Management Port LEDs on a QFX3500 Device</i></li> <li>• <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i></li> </ul>
<b>ACTIVITY LED</b>	<p>(QFX Series only) Displays the activity status of a management port. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Management Port LEDs on a QFX3500 Device</i></li> <li>• <i>Management Port LEDs in the QFX3600 and QFX3600-I Device</i></li> </ul>



Table 38: show chassis led Output Fields (*continued*)

Field Name	Field Description
<b>STATUS LED</b>	<p>(QFX Series only) Displays the link status of an interface as indicated by the ST LED. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Control Board LEDs on a QFX3008-I Interconnect Device</i></li> <li>• <i>Access Port and Uplink Port LEDs on a QFX3500 Device</i></li> <li>• <i>Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device</i></li> </ul>
<b>LINK/ACTIVITY LED</b>	<p>(QFX Series only) Displays link activity or faults on an interface as indicated by the LA LED. For more information, see:</p> <ul style="list-style-type: none"> <li>• <i>Access Port and Uplink Port LEDs on a QFX3500 Device</i></li> <li>• <i>Access Port and Uplink Port LEDs on a QFX3600 or QFX3600-I Device</i></li> </ul>
<b>Status LED</b>	<p>(QFX3008-I Interconnect device only)</p> <ul style="list-style-type: none"> <li>• Displays the system status as indicated by the STATUS LED on the front panel of the chassis. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</li> <li>• Displays the status of a Control Board as indicated by the STATUS LED on the Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i>.</li> </ul>
<b>Power LED</b>	<p>(QFX3008-I Interconnect device only) Displays the status of system power on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Major Alarm LED</b>	<p>(QFX3008-I Interconnect device only) Displays whether a critical error condition that requires immediate action exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Minor Alarm LED</b>	<p>(QFX3008-I Interconnect device only) Displays whether a noncritical condition that requires monitoring or maintenance exists on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Fan 0 LED</b>	<p>(QFX3008-I Interconnect device only) Displays the status of fan trays on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i>.</p>
<b>Fan 1 LED</b>	
<b>Fan 2 LED</b>	
<b>Fan 3 LED</b>	
<b>Fan 4 LED</b>	
<b>Fan 5 LED</b>	
<b>Fan 6 LED</b>	
<b>Fan 7 LED</b>	
<b>Fan 8 LED</b>	

Table 38: show chassis led Output Fields (*continued*)

Field Name	Field Description
PEM 0 LED	(QFX3008-I Interconnect device only) Displays the status of power supplies on the device. For more information, see <i>Chassis Status LEDs on a QFX3008-I Interconnect Device</i> .
PEM 1 LED	
PEM 2 LED	
PEM 3 LED	
PEM 4 LED	
LED info for	(QFX3008-I Interconnect device only) Displays the LED information for a Control Board.
Mastership LED	(QFX3008-I Interconnect device only) Displays status of the MASTER LED on a Control Board. For more information, see <i>Control Board LEDs on a QFX3008-I Interconnect Device</i> .
Interface	Names of the interfaces on the device.
LED (ADM/SPD/DPX/POE)	<p>(EX Series switches only) State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are:</p> <p><b>NOTE:</b> EX4500 and EX8200 switches do not have the POE port parameter.</p> <ul style="list-style-type: none"> <li>• <b>ADM</b>—Administrative</li> <li>• <b>SPD</b>—Speed</li> <li>• <b>DPX</b>—Duplex</li> <li>• <b>POE</b>—Power over Ethernet</li> </ul>

## Sample Output

### show chassis led (EX2200 Switch)

```

user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  System LED: Green
  Mode LED : Duplex
Interface    LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0    Off
ge-0/0/1    Full Duplex
ge-0/0/2    Full Duplex
ge-0/0/3    Off
ge-0/0/4    Off
ge-0/0/5    Full Duplex
ge-0/0/6    Full Duplex
ge-0/0/7    Full Duplex
ge-0/0/8    Full Duplex
ge-0/0/9    Full Duplex
ge-0/0/10   Full Duplex
ge-0/0/11   Full Duplex

```

```

ge-0/0/12      Full Duplex
ge-0/0/13      Full Duplex
ge-0/0/14      Full Duplex
ge-0/0/15      Full Duplex
ge-0/0/16      Full Duplex
ge-0/0/17      Full Duplex
ge-0/0/18      Full Duplex
ge-0/0/19      Full Duplex
ge-0/0/20      Full Duplex
ge-0/0/21      Full Duplex
ge-0/0/22      Off
ge-0/0/23      Off
ge-0/0/24      Full Duplex
ge-0/0/25      Full Duplex
ge-0/0/26      Off
ge-0/0/27      Off
ge-0/0/28      Full Duplex
ge-0/0/29      Full Duplex

```

### show chassis led

```
user@switch> show chassis led
```

```
Front panel contents for slot: 0
```

```
-----
LEDs status:
```

```
  Alarms LED: Off
```

```
  System LED: Green
```

```
  Master LED: Green
```

```
Interface      LED (ADM/SPD/DPX/POE)
```

```
-----
ge-0/0/0      Off
ge-0/0/1      Full Duplex
ge-0/0/2      Full Duplex
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Full Duplex
ge-0/0/6      Full Duplex
ge-0/0/7      Full Duplex
ge-0/0/8      Full Duplex
ge-0/0/9      Full Duplex
ge-0/0/10     Full Duplex
ge-0/0/11     Full Duplex
ge-0/0/12     Full Duplex
ge-0/0/13     Full Duplex
ge-0/0/14     Full Duplex
ge-0/0/15     Full Duplex
ge-0/0/16     Full Duplex
ge-0/0/17     Full Duplex
ge-0/0/18     Full Duplex
ge-0/0/19     Full Duplex
ge-0/0/20     Full Duplex
ge-0/0/21     Full Duplex
ge-0/0/22     Off
ge-0/0/23     Off
ge-0/0/24     Full Duplex
ge-0/0/25     Full Duplex
ge-0/0/26     Off
ge-0/0/27     Off
ge-0/0/28     Full Duplex
ge-0/0/29     Full Duplex

```

### show chassis led fpc-slot 0

```
user@switch> show chassis led fpc-slot 0
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Red
  System LED: Green
  Master LED: Green
Interface      LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0      Off
ge-0/0/1      Off
ge-0/0/2      Off
ge-0/0/3      Off
ge-0/0/4      Off
ge-0/0/5      Off
ge-0/0/6      Off
ge-0/0/7      Off
ge-0/0/8      Off
ge-0/0/9      Off
ge-0/0/10     Off
ge-0/0/11     Off
ge-0/0/12     Off
ge-0/0/13     Off
ge-0/0/14     Off
ge-0/0/15     Off
ge-0/0/16     Off
ge-0/0/17     Off
ge-0/0/18     Off
ge-0/0/19     Off
ge-0/0/20     Off
ge-0/0/21     Off
ge-0/0/22     Off
ge-0/0/23     Off
```

### show chassis led (EX Series)

```
user@switch> show chassis led
Front panel contents for slot: 0
-----
LEDs status:
  Alarms LED: Amber
  Status LED: Green
  Mode LED : Duplex
Interface LED(ADM/SPD/DPX/POE)
-----
ge-0/0/0 Off
ge-0/0/1 Full Duplex
ge-0/0/2 Full Duplex
ge-0/0/3 Off
ge-0/0/4 Off
ge-0/0/5 Full Duplex
ge-0/0/6 Full Duplex
ge-0/0/7 Full Duplex
ge-0/0/8 Full Duplex
ge-0/0/9 Full Duplex
ge-0/0/10 Full Duplex
ge-0/0/11 Full Duplex
ge-0/0/12 Full Duplex
ge-0/0/13 Full Duplex
```

```

ge-0/0/14 Full Duplex
ge-0/0/15 Full Duplex
ge-0/0/16 Full Duplex
ge-0/0/17 Full Duplex
ge-0/0/18 Full Duplex
ge-0/0/19 Full Duplex
ge-0/0/20 Full Duplex
ge-0/0/21 Full Duplex
ge-0/0/22 Off
ge-0/0/23 Off
ge-0/0/24 Full Duplex
ge-0/0/25 Full Duplex
ge-0/0/26 Off
ge-0/0/27 Off
ge-0/0/28 Full Duplex
ge-0/0/29 Full Duplex

```

#### show chassis led node-device (QFabric System Node Device)

```

user@switch> show chassis led node-device node1
Front panel contents for: node1
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
node1:me5	Green	N/A
node1:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
node1:xe-0/0/8	Green	Green
node1:ge-0/0/10	Green	Green
node1:ge-0/0/12	Green	Green
node1:ge-0/0/24	Green	Green
node1:ge-0/0/25	Green	Green
node1:ge-0/0/26	Green	Green
node1:ge-0/0/27	Green	Green
node1:ge-0/0/28	Green	Green
node1:ge-0/0/29	Green	Green
node1:ge-0/0/30	Green	Green
node1:ge-0/0/31	Green	Green
node1:ge-0/0/32	Green	Green
node1:ge-0/0/33	Green	Green
node1:ge-0/0/34	Green	Green
node1:ge-0/0/35	Green	Green
node1:ge-0/0/36	Green	Green
node1:ge-0/0/37	Green	Green
node1:ge-0/0/38	Green	Green
node1:ge-0/0/39	Green	Green
node1:fte-0/1/0	Green	Green Blinking
node1:fte-0/1/2	Green	Green Blinking

#### show chassis led interconnect-device (QFabric System - QFX3600-I Interconnect Device)

```

user@switch> show chassis led interconnect-device IC-EG0712
Front panel contents for: FPC 0
-----
LEDs status:
  Status/Beacon LED: Yellow Blinking

```

Interface	LINK/SPEED LED	ACTIVITY LED
IC-EG0712:me5	Green	N/A
IC-EG0712:me6	Green	N/A

Interface	STATUS LED	LINK/ACTIVITY LED
IC-EG0712:fte-0/1/0	Green	Green
IC-EG0712:fte-0/1/1	Green	Green Blinking
IC-EG0712:fte-0/1/2	Green	Green
IC-EG0712:fte-0/1/3	Green	Green Blinking
IC-EG0712:fte-0/1/4	Green	Green
IC-EG0712:fte-0/1/5	Green	Green Blinking
IC-EG0712:fte-0/1/6	Green	Green
IC-EG0712:fte-0/1/7	Green	Green
IC-EG0712:fte-0/1/8	Green	Green Blinking
IC-EG0712:fte-0/1/9	Green	Green Blinking
IC-EG0712:fte-0/1/10	Green	Green Blinking

### show chassis led interconnect-device (QFabric System - QFX3008-I Interconnect Device)

```
user@switch> show chassis led interconnect-device IC-EG0712
Front Panel Module Information
```

#### LEDs status:

```
Status LED: Green
Power LED : Yellow Blinking
Major Alarm LED: Red
Minor Alarm LED: Yellow
Fan 0 LED : Green
Fan 1 LED : Green
Fan 2 LED : Green
Fan 3 LED : Green
Fan 4 LED : Green
Fan 5 LED : Green
Fan 6 LED : Green
Fan 7 LED : Green
Fan 8 LED : Green
Fan 9 LED : Green
PEM 0 LED : Green
PEM 1 LED : Green
PEM 2 LED : Green
PEM 3 LED : off
PEM 4 LED : Yellow Blinking
PEM 5 LED : off
```

```
LED info for: CB - 0
```

#### LEDs status:

```
Status LED: Green
Mastership LED: Green
```

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:pme0 :	Green	N/A
IC-F4899:pme1 :	off	N/A
IC-F4899:pme2 :	off	N/A
IC-F4899:pme3 :	off	N/A

```
LED info for: CB - 1
```

## LEDs status:

Status LED: Green

Mastership LED: Amber

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:pme0 :	Green	N/A
IC-F4899:pme1 :	off	N/A
IC-F4899:pme2 :	off	N/A
IC-F4899:pme3 :	off	N/A

LED info for: FC 0 FPC - 0

## LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-0/0/0	Green	N/A
IC-F4899:fte-0/0/1	Green	N/A
IC-F4899:fte-0/0/2	Green	N/A
IC-F4899:fte-0/0/3	Green	N/A
IC-F4899:fte-0/0/4	Green	N/A
IC-F4899:fte-0/0/5	Green	N/A
IC-F4899:fte-0/0/6	Green	N/A
IC-F4899:fte-0/0/7	Green	N/A
IC-F4899:fte-0/0/8	Green	N/A
IC-F4899:fte-0/0/9	Green	N/A
IC-F4899:fte-0/0/10	Green	N/A
IC-F4899:fte-0/0/11	Green	N/A
IC-F4899:fte-0/0/12	Green	N/A
IC-F4899:fte-0/0/13	Green	N/A
IC-F4899:fte-0/0/14	Green	N/A
IC-F4899:fte-0/0/15	Green	N/A

LED info for: FC 1 FPC - 1

## LEDs status:

Status LED: Green

Interface	STATUS LED	LINK/ACTIVITY LED
IC-F4899:fte-1/0/0	Green	N/A
IC-F4899:fte-1/0/1	Green	N/A

LED info for: RC 2 FPC - 10

## LEDs status:

Status LED: Green

LED info for: RC 3 FPC - 11

## LEDs status:

Status LED: Green

## show chassis location

---

<b>List of Syntax</b>	<a href="#">Syntax on page 878</a> <a href="#">Syntax (TX Matrix Router) on page 878</a> <a href="#">Syntax (TX Matrix Plus Router) on page 878</a> <a href="#">Syntax (MX Series Router) on page 878</a> <a href="#">Syntax (QFX Series) on page 878</a>
<b>Syntax</b>	show chassis location
<b>Syntax (TX Matrix Router)</b>	show chassis location <fpc   interface (by-name <i>name</i>   by-slot fpc number lcc number)   lcc number   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis location <fpc   interface (by-name <i>name</i>   by-slot fpc number lcc number)   lcc number   sfc number>
<b>Syntax (MX Series Router)</b>	show chassis location <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show chassis location <interconnect-device <i>name</i> > <node-device <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the physical location of the chassis. This command can only be used on the master Routing Engine.
<b>Options</b>	<b>none</b> —Display all information about the physical location of the chassis. On a TX Matrix router, display all information about the physical location of the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display all information about the physical location of the TX Matrix Plus router and its attached routers.  <b>all-members</b> —(MX Series routers only) (Optional) Display the physical location of the chassis for all the member routers in the Virtual Chassis configuration.  <b>fpc</b> —(TX Matrix router and TX Matrix Plus router only) (Optional) Display the physical location of all Flexible PIC Concentrators (FPCs).  <b>interconnect-device <i>name</i></b> —(QFabric systems only) (Optional) Display the physical location of the Interconnect device.  <b>interface by-name <i>name</i></b> —(TX Matrix and TX Matrix Plus routers only) (Optional) Display the physical location of a specified interface name. On a TX Matrix router, this option displays the FPC number and T640 router (line-card chassis) number associated



with the specified interface. On a TX Matrix Plus router, this option displays the FPC number and router (line-card chassis) number associated with the specified interface.

**interface by-slot fpc *number* lcc *number***—(TX Matrix and TX Matrix Plus router only)

(Optional) On a TX Matrix router, display the global FPC number of an interface by specifying its local FPC number and T640 router (line-card chassis) number. On a TX Matrix Plus router, display the global FPC number of an interface by specifying its local FPC number and router (line-card chassis) number.

- The global FPC number is the FPC slot number when all the FPC slots in the routing matrix are considered: **0** through **31**. On TX Matrix Plus router with 3D SIBs, the value is **0** through **63**. The local FPC number is the FPC slot number on a particular T640 router.
- For **fpc**, replace *number* with a value from **0** through **7**.
- For **lcc**, replace *number* with a value from **0** through **7**.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the physical location of a specified T640 router (line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display the physical location of a specified router (line-card chassis) that is connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display the physical location of the chassis for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the physical location of the chassis for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display the physical location of the Node device.

**scc**—(TX Matrix routers only) (Optional) Display the physical location of the TX Matrix router (switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display the physical location of the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level** view

**Related Documentation** • *Displaying Chassis Physical Locations for a Routing Matrix with a TX Matrix Plus Router*

**List of Sample Output** [show chassis location on page 880](#)  
[show chassis location fpc \(TX Matrix Router\) on page 881](#)  
[show chassis location interface by-slot \(TX Matrix Router\) on page 881](#)  
[show chassis location fpc \(TX Matrix Plus Router\) on page 881](#)  
[show chassis location interface by-slot \(TX Matrix Plus Router\) on page 881](#)  
[show chassis location \(QFX3500 Switches\) on page 881](#)  
[show chassis location \(QFabric Systems\) on page 881](#)

**Output Fields** [Table 39 on page 880](#) lists the output fields for the **show chassis location** command. Output fields are listed in the approximate order in which they appear.

**Table 39: show chassis location Output Fields**

Field Name	Field Description
country-code	Country code information.
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the routing matrix are considered. The range of values is 0 through 31. On TX Matrix Plus router with 3D SIBs the value is 0 through 63.
LATA	Local access transport area information.
LCC	Line-card chassis number. On a TX Matrix router, the number of a particular T640 router connected to the TX Matrix router. On a TX Matrix Plus router, the number of a particular router connected to the TX Matrix Plus router.
Local FPC	Local FPC number. On a TX Matrix router, the FPC slot number on a particular T640 router. On a TX Matrix Plus router, the FPC slot number on a particular router.

## Sample Output

**show chassis location**

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

**show chassis location fpc (TX Matrix Router)**

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    17         2        1
    21         2        5
```

**show chassis location interface by-slot (TX Matrix Router)**

```
user@host> show chassis location interface by-slot fpc 1 lcc 1
Global FPC: 9
```

**show chassis location fpc (TX Matrix Plus Router)**

```
user@host> show chassis location fpc
Global FPC    LCC    Local FPC
    0         0        0
    1         0        1
```

**show chassis location interface by-slot (TX Matrix Plus Router)**

```
user@host> show chassis location interface by-slot fpc 2 lcc 1
Global FPC: 10
```

**show chassis location (QFX3500 Switches)**

```
user@switch> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

**show chassis location (QFabric Systems)**

```
user@switch> show chassis location interconnect-device interconnect1
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

## show chassis mac-addresses

---

<b>List of Syntax</b>	<a href="#">Syntax on page 882</a> <a href="#">Syntax (TX Matrix Router) on page 882</a> <a href="#">Syntax (TX Matrix Plus Router) on page 882</a> <a href="#">Syntax (MX Series Router) on page 882</a> <a href="#">Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers) on page 882</a> <a href="#">Syntax (QFX Series) on page 882</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 882</a>
<b>Syntax</b>	show chassis mac-addresses
<b>Syntax (TX Matrix Router)</b>	show chassis mac-addresses <lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show chassis mac-addresses <lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show chassis mac-addresses <all-members> <local> <member <i>member-id</i> >
<b>Syntax (MX104, MX2010, and MX2020 3D Universal Edge Routers)</b>	show chassis mac-addresses
<b>Syntax (QFX Series)</b>	show chassis mac-addresses <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Syntax (ACX Series Universal Access Routers)</b>	show chassis mac-addresses
<b>Release Information</b>	Command introduced before JUNOS Release 7.4. Command introduced in JUNOS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers. Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers. Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.
<b>Description</b>	Display the media access control (MAC) addresses for the router, switch chassis, or switch.
<b>Options</b>	<b>none</b> —(TX Matrix, TX Matrix Plus routers, and the QFX Series) Display the MAC addresses for the router chassis or switch. On a TX Matrix router, display MAC addresses on the

TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display MAC addresses on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display the MAC addresses for all the member routers of the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display the MAC addresses for the Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display MAC addresses for a specified T640 router (line-card chassis) that is connected to the TX Matrix Plus router. On a TX Matrix Plus router, display MAC addresses for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display the MAC addresses for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display the MAC addresses for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display the MAC addresses for the specified Node group.

**scc**—(TX Matrix routers only) (Optional) Display MAC addresses for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display MAC addresses for the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level**

view

**Related Documentation**

- *ACX2000 and ACX2100 Routers Hardware and CLI Terminology Mapping*

**List of Sample Output**

[show chassis mac-addresses on page 884](#)  
[show chassis mac-addresses \(MX104 Router\) on page 884](#)  
[show chassis mac-addresses \(MX2010 Router\) on page 884](#)

[show chassis mac-addresses \(MX2020 Router\) on page 885](#)  
[show chassis mac-addresses \(TX Matrix Router\) on page 885](#)  
[show chassis mac-addresses \(TX Matrix Plus Router\) on page 885](#)  
[show chassis mac-addresses \(QFX3500 Switches\) on page 886](#)  
[show chassis mac-addresses interconnect-device \(QFabric Systems\) on page 886](#)  
[show chassis mac-addresses node-group \(QFabric Systems\) on page 886](#)  
[show chassis mac-addresses \(ACX2000 Universal Access Router\) on page 886](#)

**Output Fields** Table 40 on page 884 lists the output fields for the **show chassis mac-addresses** command. Output fields are listed in the approximate order in which they appear.

**Table 40: show chassis mac-addresses Output Fields**

Field Name	Field Description
MAC address information	
Public base address	Base address of the MAC addresses allocated to this router or switch.
Public count	Number of allocated public addresses.
Private base address	Base address of the private MAC addresses allocated to this router or switch.
Private count	Number of allocated private addresses.

## Sample Output

### show chassis mac-addresses

```

user@host> show chassis mac-addresses
MAC address information
  Public base address  0:90:69:0:4:0
  Public count         1008
  Private base address 0:90:69:0:7:f0
  Private count        16

```

### show chassis mac-addresses (MX104 Router)

```

user@host > show chassis mac-addresses
MAC address information:
  Public base address  b0:a8:6e:a1:e8:58
  Public count         2032
  Private base address b0:a8:6e:a1:f0:48
  Private count        16

```

### show chassis mac-addresses (MX2010 Router)

```

user@host> show chassis mac-addresses
MAC address information:
  Public base address  64:87:88:04:50:00
  Public count         1984
  Private base address 64:87:88:04:57:c0
  Private count        64

```

**show chassis mac-addresses (MX2020 Router)**

```

user@host> show chassis mac-addresses
MAC address information:
  Public base address    2c:21:72:70:20:00
  Public count           4032
  Private base address   2c:21:72:70:2f:c0
  Private count          64

```

**show chassis mac-addresses (TX Matrix Router)**

```

user@host> show chassis mac-addresses
scc-re0:
-----
MAC address information:
  Public base address    00:05:85:9e:cc:00
  Public count           8064
  Private base address   00:05:85:9e:eb:80
  Private count          128
lcc0-re0:
-----
MAC address information:
  Public base address    00:05:85:68:98:00
  Public count           2032
  Private base address   00:05:85:68:9f:f0
  Private count          16
lcc2-re0:
-----
MAC address information:
  Public base address    00:05:85:68:78:00
  Public count           2032
  Private base address   00:05:85:68:7f:f0
  Private count          16

```

**show chassis mac-addresses (TX Matrix Plus Router)**

```

user@host> show chassis mac-addresses
sfc0-re0:
-----
MAC address information:
  Public base address    00:1d:b5:14:00:00
  Public count           65023
  Private base address   00:1d:b5:14:fd:ff
  Private count          512
lcc0-re0:
-----
MAC address information:
  Public base address    00:1f:12:7a:84:00
  Public count           2032
  Private base address   00:1f:12:7a:8b:f0
  Private count          16
lcc1-re0:
-----
MAC address information:
  Public base address    00:22:83:42:48:00
  Public count           2032
  Private base address   00:22:83:42:4f:f0
  Private count          16

```

lcc2-re0:

-----  
MAC address information:

Public base address	00:1f:12:c3:58:00
Public count	2032
Private base address	00:1f:12:c3:5f:f0
Private count	16

lcc3-re0:

-----  
MAC address information:

Public base address	00:21:59:ef:b8:00
Public count	2032
Private base address	00:21:59:ef:bf:f0
Private count	16

#### show chassis mac-addresses (QFX3500 Switches)

```
user@switch> show chassis mac-addresses
```

MAC address information:

Public base address	02:00:08:00:00:00
Public count	512
Private base address	02:00:00:00:00:00
Private count	64

#### show chassis mac-addresses interconnect-device (QFabric Systems)

```
user@switch> show chassis mac-addresses interconnect-device interconnect1
```

MAC address information:

Public base address	00:1f:12:30:9c:c0
Public count	58
Private base address	00:1f:12:30:9c:fa
Private count	6

#### show chassis mac-addresses node-group (QFabric Systems)

```
user@switch> show chassis mac-addresses node-group NW-NG-0
```

MAC address information:

-----  
RE:

FC MAC base	00:11:00:00:00:00
FC MAC count	2
VLAN MAC	00:11:00:00:00:09

EC6007

Base address	00:00:01:76:00:00
Count	64

EC6008

Base address	00:22:83:22:52:ae
Count	260

#### show chassis mac-addresses (ACX2000 Universal Access Router)

```
user@switch> show chassis mac-addresses
```

MAC address information:

Public base address	84:18:88:c0:2b:00
Public count	112
Private base address	84:18:88:c0:2b:70
Private count	16



## show chassis nonstop-upgrade

<b>Syntax</b>	<b>show chassis nonstop-upgrade</b>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
<b>Description</b>	(EX6200 switches, EX8200 switches, EX8200 Virtual Chassis, QFX3500 and QFX3600 Virtual Chassis, and Virtual Chassis Fabric only) Display the status of the line cards or Virtual Chassis members in the linecard role after the most recent nonstop software upgrade (NSSU). This command must be issued on the master Routing Engine.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request system software nonstop-upgrade on page 449</a></li> <li>• <i>Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade (CLI Procedure)</i></li> <li>• <i>Upgrading Software on QFX3500, QFX3600, and QFX5100 Virtual Chassis Using Nonstop Software Upgrade</i></li> <li>• <a href="#">Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade on page 139</a></li> <li>• <i>Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade (CLI Procedure)</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show chassis nonstop-upgrade (EX8200 Switch) on page 888</a> <a href="#">show chassis nonstop-upgrade (EX8200 Virtual Chassis) on page 888</a> <a href="#">show chassis nonstop-upgrade (Virtual Chassis Fabric) on page 888</a>
<b>Output Fields</b>	Table 41 on page 887 lists the output fields for the <b>show chassis nonstop-upgrade</b> command. Output fields are listed in the approximate order in which they appear.

**Table 41: show chassis nonstop-upgrade Output Fields**

Field Name	Field Description
<b>Item</b>	Line card slot number.
<b>Status</b>	State of line card: <ul style="list-style-type: none"> <li>• <b>Error</b>—Line card is in an error state.</li> <li>• <b>Offline</b>—Line card is powered down.</li> <li>• <b>Online</b>—Line card is online and running.</li> </ul>
<b>Reason</b>	Reason for the state (if the line card is offline).

## Sample Output

### show chassis nonstop-upgrade (EX8200 Switch)

```
user@switch> show chassis nonstop-upgrade
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Offline	Offlined by CLI command
FPC 4	Online	
FPC 5	Online	
FPC 6	Online	
FPC 7	Online	

### show chassis nonstop-upgrade (EX8200 Virtual Chassis)

```
user@external-routing-engine> show chassis nonstop-upgrade
member0:
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 5	Online	

```
member1:
```

Item	Status	Reason
FPC 0	Online	
FPC 1	Offline	Offlined due to config
FPC 2	Online	
FPC 3	Online	
FPC 4	Online	
FPC 5	Online	
FPC 7	Online	

### show chassis nonstop-upgrade (Virtual Chassis Fabric)

Item	Status	Reason
FPC 0	Online	
FPC 1	Online	
FPC 2	Online	
FPC 3	Online	
FPC 4	Online	
FPC 5	Online	

## show chassis pic

<b>List of Syntax</b>	<a href="#">Syntax on page 889</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 889</a> <a href="#">Syntax (MX Series Routers) on page 889</a> <a href="#">Syntax (MX104, MX2010 and MX2020 3D Universal Edge Routers) on page 889</a> <a href="#">Syntax (PTX Series Packet Transport Router) on page 889</a> <a href="#">Syntax (QFX Series) on page 889</a> <a href="#">Syntax (ACX Series Universal Access Routers) on page 889</a>
<b>Syntax</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> &lt;lcc <i>number</i>&gt;</code>
<b>Syntax (MX Series Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i> &lt;all-members&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt;</code>
<b>Syntax (MX104, MX2010 and MX2020 3D Universal Edge Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (PTX Series Packet Transport Router)</b>	<code>show chassis pic transport fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Syntax (QFX Series)</b>	<code>show chassis pic &lt;interconnect-device <i>name</i> (fpc-slot <i>slot-number</i>   pic-slot <i>slot-number</i>)&gt; &lt;node-device <i>name</i> pic-slot <i>slot-number</i>&gt;</code>
<b>Syntax (ACX Series Universal Access Routers)</b>	<code>show chassis pic fpc-slot <i>slot-number</i> pic-slot <i>slot-number</i></code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX104 3D Universal Edge Routers.</p>
<b>Description</b>	Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.
<b>Options</b>	<b>fpc-slot <i>slot-number</i></b> —Display information about the PIC in this particular FPC slot:

- On a TX Matrix router, if you specify the number of the T640 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31.

Likewise, on a TX Matrix Plus router, if you specify the number of the T1600 router by using the **lcc number** option (the recommended method), replace **slot-number** with a value from 0 through 7. Otherwise, replace **slot-number** with a value from 0 through 31. For example, the following commands have the same result:

```
user@host> show chassis pic fpc-slot 1 lcc 1 pic-slot 1
user@host> show chassis pic fpc-slot 9 pic-slot 1
```

- M120 routers only—Replace **slot-number** with a value from 0 through 5.
- MX80 routers only—Replace **slot-number** with a value from 0 through 1.
- MX104 routers only—Replace **slot-number** with a value from 0 through 2.
- MX240 routers only—Replace **slot-number** with a value from 0 through 2.
- MX480 routers only—Replace **slot-number** with a value from 0 through 5.
- MX960 routers only—Replace **slot-number** with a value from 0 through 11.
- MX2010 routers only—Replace **slot-number** with a value from 0 through 9.
- MX2020 routers only—Replace **slot-number** with a value from 0 through 19.
- Other routers—Replace **slot-number** with a value from 0 through 7.
- EX Series switches:
  - EX3200 switches and EX4200 standalone switches—Replace **slot-number** with 0.
  - EX4200 switches in a Virtual Chassis configuration—Replace **slot-number** with a value from 0 through 9 (switch's member ID).
  - EX8208 switches—Replace **slot-number** with a value from 0 through 7 (line card).
  - EX8216 switches—Replace **slot-number** with a value from 0 through 15 (line card).
- QFX Series:
  - QFX3500 and QFX5100 standalone switches—Replace **slot-number** with 0. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.
  - QFabric systems—Replace **slot-number** with any number between 0 and 15. In the command output, FPC refers to a line card. The FPC number equals the slot number for the line card.

**all-members**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for all member routers in the Virtual Chassis configuration.

**interconnect-device *name***—(QFabric systems only) (Optional) Display PIC information for a specified Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display PIC information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display PIC information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers and EX Series switches only) (Optional) Display PIC information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**node-device *name***—(QFabric systems only) (Optional) Display PIC information for a specified Node device.

**pic-slot *slot-number***—Display information about the PIC in this particular PIC slot. For routers, replace *slot-number* with a value from 0 through 3. For EX3200 and EX4200 switches, replace *slot-number* with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For EX8208 and EX8216 switches, replace *slot-number* with 0. For the QFX3500 standalone switch and the QFabric system, replace *slot-number* with 0 or 1.

**transport**—Display PIC information for optical transport network.

**Required Privilege Level**

view

**Related Documentation**

- [request chassis pic on page 388](#)
- [show chassis hardware on page 676](#)
- *Configuring the PIC Type*
- *100-Gigabit Ethernet Type 4 PIC with CFP Overview*

**List of Sample Output**    [show chassis pic fpc-slot pic-slot on page 894](#)

[show chassis pic fpc-slot pic-slot \(PIC Offline\) on page 895](#)  
[show chassis pic fpc-slot pic-slot \(FPC Offline\) on page 895](#)  
[show chassis pic fpc-slot pic-slot \(FPC Not Present\) on page 895](#)  
[show chassis pic fpc-slot pic-slot \(PIC Not Present\) on page 895](#)  
[show chassis pic fpc-slot pic-slot \(M120 Router\) on page 895](#)  
[show chassis pic fpc-slot pic-slot \(MX104 Router\) on page 895](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router Bidirectional Optics\) on page 896](#)  
[show chassis pic fpc-slot pic-slot \(MX480 Router with 100-Gigabit Ethernet MIC\) on page 896](#)  
[show chassis pic fpc-slot pic-slot \(MX240, MX480, MX960 Routers with Application Services Modular Line Card\) on page 896](#)  
[show chassis pic fpc-slot pic-slot \(MX960 Router with MPC5EQ\) on page 897](#)  
[show chassis pic fpc-slot pic-slot \(MX480 Routers with MPC4E\) on page 897](#)  
[show chassis pic fpc-slot pic-slot \(MX480 routers with OTN Interfaces\) on page 897](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Routers with OTN Interfaces\) on page 897](#)  
[show chassis pic fpc-slot pic-slot \(MX2010 Routers\) on page 898](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers\) on page 898](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers with MPC5EQ and MPC6E\) on page 898](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers with MPC6E and OTN MIC\) on page 899](#)  
[show chassis pic fpc-slot pic-slot \(MX2020 Routers with MPC4E\) on page 899](#)  
[show chassis pic fpc-slot pic-slot \(T1600 Router with 100-Gigabit Ethernet PIC\) on page 899](#)  
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Router\) on page 900](#)  
[show chassis pic fpc-slot pic-slot lcc \(TX Matrix Plus Router\) on page 900](#)  
[show chassis pic fpc-slot pic-slot \(Next-Generation SONET/SDH SFP\) on page 900](#)  
[show chassis pic fpc-slot pic-slot \(12-Port T1/E1\) on page 900](#)  
[show chassis pic fpc-slot pic-slot \(4x CHOC3 SONET CE SFP\) on page 901](#)  
[show chassis pic fpc-slot pic-slot \(SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 901](#)  
[show chassis pic fpc-slot pic-slot \(8-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 901](#)  
[show chassis pic fpc-slot pic-slot \(4-port Channelized SONET/SDH OC3/STM1 \[Multi-Rate\] MIC with SFP\) on page 902](#)  
[show chassis pic fpc-slot pic-slot \(1-port OC192/STM64 MIC with XFP\) on page 902](#)  
[show chassis pic fpc-slot 1 pic-slot 2 \(8-port DS3/E3 MIC\) on page 902](#)  
[show chassis pic fpc-slot pic-slot \(OTN\) on page 902](#)  
[show chassis pic fpc-slot pic-slot \(QFX3500 Switch\) on page 902](#)  
[show chassis pic fpc-slot pic-slot \(QFX5100 Standalone Switch\) on page 903](#)  
[show chassis pic interconnect-device fpc-slot pic-slot \(QFabric Systems\) on page 903](#)  
[show chassis pic node-device fpc-slot pic-slot \(QFabric System\) on page 903](#)  
[show chassis pic fpc-slot pic-slot \(ACX2000 Universal Access Router\) on page 904](#)  
[show chassis pic fpc-slot pic-slot \(MX Routers with Media Services Blade \[MSB\]\) on page 904](#)  
[show chassis pic FPC slot PIC slot \(MX Routers with Media Services Blade \[MSB\]\) on page 904](#)  
[show chassis pic transport fpc-slot pic-slot \(PTX Series Packet Transport Routers\) on page 904](#)

**Output Fields** Table 42 on page 893 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

**Table 42: show chassis pic Output Fields**

Field Name	Field Description
Type	<p>PIC type.</p> <p><b>NOTE:</b> On the 1-port OC192/STM64 MICs with the SDH framing mode, the type is displayed as <b>MIC-3D-1STM64-XFP</b> and with the SONET framing mode, the type is displayed as <b>MIC-3D-1OC192-XFP</b>. By default, the 1-port OC192/STM64 MICs displays the type as <b>MIC-3D-1OC192-XFP</b>.</p>
Account Layer2 Overhead	(MX Series routers) Indicates whether functionality to count the Layer 2 overhead bytes in the interface statistics at the PIC level is enabled or disabled.
ASIC type	Type of ASIC on the PIC.
State	<p>Status of the PIC. State is displayed only when a PIC is in the slot.</p> <ul style="list-style-type: none"> <li>• <b>Online</b>— PIC is online and running.</li> <li>• <b>Offline</b>—PIC is powered down.</li> </ul>
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.
Package	(Multiservices PICs only) Services package supported: <b>Layer-2</b> or <b>Layer-3</b> .
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: <b>LH</b> , <b>LX</b> , or <b>SX</b> .
PIC Port Information (MX480 Router 100-Gigabit Ethernet CFP)	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>• Port—Port number</li> <li>• Cable type—Type of optical transceiver installed.</li> <li>• Fiber type—Type of fiber. SM is single-mode.</li> <li>• Xcvr vendor—Transceiver vendor name.</li> <li>• Xcvr vendor part number—Transceiver vendor part number.</li> <li>• Wavelength—Wavelength of the transmitted signal. Uplinks and downlinks are always 1550 nm. There is a separate fiber for each direction</li> </ul>

Table 42: show chassis pic Output Fields (*continued*)

Field Name	Field Description
<b>PIC Port Information (MX960 Router Bidirectional Optics )</b>	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>Port—Port number</li> <li>Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.</li> <li>Fiber type—Type of fiber. SM is single-mode.</li> <li>Xcvr vendor—Transceiver vendor name.</li> <li>Xcvr vendor part number—Transceiver vendor part number. <ul style="list-style-type: none"> <li>BX10-10-km bidirectional optics.</li> <li>BX40-40-km bidirectional optics.</li> <li>SFP-LX-40-km SFP optics.</li> </ul> </li> <li>Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.</li> </ul>
<b>PIC Port Information (Next-Generation SONET/SDH SFP)</b>	<p>Port-level information for the next-generation SONET/SDH SFP PIC.</p> <ul style="list-style-type: none"> <li>Port—Port number.</li> <li>Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed.</li> <li>Fiber type—Type of fiber: <b>SM</b> (single-mode) or <b>MM</b> (multimode).</li> <li>Xcvr vendor—Transceiver vendor name.</li> <li>Xcvr vendor part number—Transceiver vendor part number.</li> <li>Wavelength—Wavelength of the transmitted signal. Next-generation SONET/SDH SFPs use 1310 nm.</li> </ul>
<b>Pic port information (MX104 router)</b>	<p>Port-level information for the PIC.</p> <ul style="list-style-type: none"> <li>Port—Port number</li> <li>Cable type—Type of optical transceiver installed.</li> <li>Fiber type—Type of fiber. SM is single-mode.</li> <li>Xcvr vendor—Transceiver vendor name.</li> <li>Xcvr vendor part number—Transceiver vendor part number.</li> <li>Wavelength—Wavelength of the transmitted signal.</li> <li>Xcvr Firmware—Firmware version of the transceiver.</li> </ul>
<b>Multirate Mode</b>	Rate-selectability status for the MIC: <b>Enabled</b> or <b>Disabled</b> .
<b>Channelization</b>	Indicates whether channelization is enabled or disabled on the DS3/E3 MIC.

## Sample Output

### show chassis pic fpc-slot pic-slot

```

user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
Type                               10x 1GE(LAN), 1000 BASE

```



```

ASIC type           H chip
State               Online
PIC version         1.1
Uptime              1 day, 50 minutes, 58 seconds
PIC Port Information:
  Port      Cable      Xcvr      Xcvr Vendor
  Number    Type       Vendor Name  Part Number
  0         GIGE 1000EX  FINISAR CORP.  FTRJ8519P1BNL-J3
  1         GIGE 1000EX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

#### show chassis pic fpc-slot pic-slot (PIC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
PIC fpc slot 1 pic slot 0 information:
  State               Offline

```

#### show chassis pic fpc-slot pic-slot (FPC Offline)

```

user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online

```

#### show chassis pic fpc-slot pic-slot (FPC Not Present)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty

```

#### show chassis pic fpc-slot pic-slot (PIC Not Present)

```

user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty

```

#### show chassis pic fpc-slot pic-slot (M120 Router)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
PC slot 3, PIC slot 0 information:
  Type               2x G/E IQ, 1000 BASE
  ASIC type          IQ GE 2 VLAN-TAG FPGA
  State              Online
  PIC version         1.16
  Uptime              3 hours, 3 minutes

PIC Port Information:
  Port      Cable      Xcvr      Xcvr Vendor
  Number    Type       Vendor Name  Part Number
  0         GIGE 1000SX  FINISAR CORP.  FTRJ8519P1BNL-J3
  1         GIGE 1000SX  FINISAR CORP.  FTRJ-8519-7D-JUN

```

#### show chassis pic fpc-slot pic-slot (MX104 Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1
FPC slot 1, PIC slot 1 information:
  Type               10x 1GE(LAN) -E SFP
  State              Online
  PIC version         1.1
  Uptime              1 hour, 30 minutes, 59 seconds

PIC port information:
  Fiber      Xcvr vendor      Wave-      Xcvr
  Port Cable type      type Xcvr vendor      part number      length
Firmware
  3  GIGE 1000T      n/a  Methode Elec.      SP7041-M1-JN      n/a      0.0

```

6	GIGE 1000LX10	SM	FINISAR CORP.	FTLF1318P2BTL-J1	1310 nm	0.0
8	GIGE 1000T	n/a	Methode Elec.	SP7041-M1-JN	n/a	0.0
9	GIGE 1000T	n/a	Methode Elec.	SP7041-M1-JN	n/a	0.0

### show chassis pic fpc-slot pic-slot (MX960 Router Bidirectional Optics)

```

user@host> show chassis pic fpc-slot 4 pic-slot 1
FPC slot 4, PIC slot 1 information:
  Type                               10x 1GE(LAN)
  Account Layer2 Overhead            Enabled
  State                               Online
  PIC version                         0.0
  Uptime                             18 days, 5 hours, 41 minutes, 54 seconds

PIC port information:

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
1	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
2	SFP-1000BASE-BX10-D	SM	SumitomoElectric	SBP6H44-J3-BW-49	1490 nm
3	SFP-1000BASE-BX10-D	SM	OCF	TRXBG1LXDBVM2-JW	1490 nm
4	SFP-1000BASE-BX10-D	SM	OCF	TRXBG1LXDBVM2-JW	1490 nm
5	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
6	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm
7	SFP-1000BASE-BX10-U	SM	OCF	TRXBG1LXDBBMH-J1	1310 nm
8	SFP-1000BASE-BX10-U	SM	OCF	TRXBG1LXDBBMH-J1	1310 nm
9	SFP-1000BASE-BX10-U	SM	SumitomoElectric	SBP6H44-J3-BW-31	1310 nm

### show chassis pic fpc-slot pic-slot (MX480 Router with 100-Gigabit Ethernet MIC)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                               1X100GE CFP
  State                               Online
  PIC version                         2.10
  Uptime                             4 minutes, 48 seconds

PIC port information:
  Fiber

```

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	100GBASE LR4	SM	FINISAR CORP.	FTLC1181RDN3-J3	1310 nm

```

  Xcvr vendor
  firmware version
  1.8

```

### show chassis pic fpc-slot pic-slot (MX240, MX480, MX960 Routers with Application Services Modular Line Card)

```

user@host> show chassis pic fpc-slot 1 pic-slot 2
FPC slot 1, PIC slot 2 information:
  Type                               AS-MXC
  State                               Online
  PIC version                         1.0
  Uptime                             11 hours, 18 minutes, 3 seconds

```

**show chassis pic fpc-slot pic-slot (MX960 Router with MPC5EQ)**

```

user@host> show chassis pic fpc-slot 0 pic-slot 3
FPC slot 0, PIC slot 3 information:
  Type                1X100GE CFP2 OTN
  State                Online
  PIC version          0.0
  Uptime               1 hour, 22 minutes, 42 seconds

PIC port information:

```

	Fiber	Xcvr vendor	Wave-	Xcvr
Port Cable type	type	Xcvr vendor	part number	length
Firmware				
0	10GBASE LR4	n/a	Oclaro Inc.	TRB5E20FNF-LF150 1309 nm 1.0

**show chassis pic fpc-slot pic-slot (MX480 Routers with MPC4E)**

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 3, PIC slot 0 information:
  Type                4x10GE SFPP
  State                Online
  PIC version          0.0
  Uptime               41 seconds

PIC port information:

```

	Fiber	Xcvr vendor	Wave-	Xcvr
Port Cable type	type	Xcvr vendor	part number	length
Firmware				
0	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014 850 nm 0.0
1	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014 850 nm 0.0

**show chassis pic fpc-slot pic-slot (MX480 routers with OTN Interfaces)**

```

user@host> show chassis pci fpc-slot 4 pic-slot 0
FPC slot 4, PIC slot 0 information:
  Type                12X10GE SFPP OTN
  State                Online
  PIC version          0.0
  Uptime               5 hours, 28 minutes, 23 seconds

PIC port information:

```

	Fiber	Xcvr vendor	Wave-	Xcvr
Port Cable type	type	Xcvr vendor	part number	length
Firmware				
0	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BNL-J1 850 nm 0.0
1	10GBASE SR	MM	FINISAR CORP.	FTLX8571D3BCL-J1 850 nm 0.0
2	10GBASE SR	MM	OPNEXT, INC.	TRS2001EM-0014 850 nm 0.0

**show chassis pic fpc-slot pic-slot (MX2010 Routers with OTN Interfaces)**

```

user@host> show chassis pic fpc-slot 9 pic-slot 0

```

FPC slot 9, PIC slot 0 information:

```
Type                2X100GE CFP2 OTN
State                Online
PIC version          1.9
Uptime               3 hours, 56 minutes, 16 seconds
```

PIC port information:

		Fiber	Xcvr vendor		Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	
Firmware						
0	100GBASE LR4-D	SM	FUJITSU	FIM37300/222	1310 nm	1.3
1	100GBASE SR10	MM	AVAGO	AFBR-8420Z	n/a	1.0

#### show chassis pic fpc-slot pic-slot (MX2010 Routers)

```
user@host> show chassis pic fpc-slot 9 pic-slot 3
```

FPC slot 9, PIC slot 3 information:

```
Type                1X100GE CFP
Account Layer2 Overhead Enabled
State                Online
PIC version          0.0
Uptime               14 hours, 51 seconds
```

#### show chassis pic fpc-slot pic-slot (MX2020 Routers)

```
user@host> show chassis pic fpc-slot 19 pic-slot 3
```

FPC slot 19, PIC slot 3 information:

```
Type                4x 10GE(LAN) SFP+
Account Layer2 Overhead Enabled
State                Online
PIC version          0.0
Uptime               1 day, 11 hours, 26 minutes, 36 seconds
```

PIC port information:

		Fiber	Xcvr vendor		Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	
Firmware						
0	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
1	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
2	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0
3	10GBASE SR	MM	SumitomoElectric	SPP5200SR-J6-M	850 nm	0.0

#### show chassis pic fpc-slot pic-slot (MX2020 Routers with MPC5EQ and MPC6E)

```
user@host> show chassis pic fpc-slot 18 pic-slot 2
```

FPC slot 18, PIC slot 2 information:

```
Type                3X40GE QSFP
State                Online
PIC version          0.0
Uptime               6 minutes, 31 seconds
```

PIC port information:

		Fiber	Xcvr vendor		Wave-	Xcvr
Port	Cable type	type	Xcvr vendor	part number	length	

```

Firmware
 0  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0
 1  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0
 2  40GBASE SR4      MM  AVAGO          AFBR-79E4Z-D-JU2  850 nm  0.0

```

### show chassis pic fpc-slot pic-slot (MX2020 Routers with MPC6E and OTN MIC)

```

user@host> show chassis pic fpc-slot 3 pic-slot 0
FPC slot 0, PIC slot 1 information:
  Type                24X10GE SFPP OTN
  State                Online
  PIC version          1.1
  Uptime                1 hour, 33 minutes, 59 seconds

PIC port information:

  Port Cable type      Fiber                Xcvr vendor      Wave-    Xcvr
  type                type Xcvr vendor      part number      length
Firmware
 7  10GBASE SR         MM  SumitomoElectric SPP5200SR-J6-M   850 nm  0.0
 9  10GBASE SR         MM  FINISAR CORP.    FTLX8571D3BNL-J1 850 nm  0.0
12  10GBASE LR         SM  FINISAR CORP.    FTLX1472M3BNL-J3 1310 nm 0.0
20  10GBASE ZR         SM  FINISAR CORP.    FTLX1871M3BNL-J3 1550 nm 0.0
21  10GBASE ER         SM  FINISAR CORP.    FTLX1671D3BTL-J4 1550 nm 0.0
22  10GBASE LR         SM  SOURCEPHOTONICS SPP10SLREDFCJNP  1310 nm 0.0
23  10GBASE LR         SM  FINISAR CORP.    FTLX1471D3BNL-J1 1310 nm 0.0

```

### show chassis pic fpc-slot pic-slot (MX2020 Routers with MPC4E)

```

user@host> show chassis pic fpc-slot 14 pic-slot 0
FPC slot 14, PIC slot 2 information:
  Type                4x10GE SFPP
  State                Online
  PIC version          0.0
  Uptime                1 day, 14 hours, 49 minutes, 9 seconds

PIC port information:

  Port Cable type      Fiber                Xcvr vendor      Wave-    Xcvr
  type                type Xcvr vendor      part number      length
Firmware
 0  10GBASE SR         MM  SumitomoElectric SPP5100SR-J3     850 nm  0.0
 1  10GBASE SR         MM  SumitomoElectric SPP5100SR-J3     850 nm  0.0
 3  10GBASE SR         MM  SumitomoElectric SPP5100SR-J3     850 nm  0.0

```

### show chassis pic fpc-slot pic-slot (T1600 Router with 100-Gigabit Ethernet PIC)

```

user@host> run show chassis pic fpc-slot 3 pic-slot 1
FPC slot 3, PIC slot 1 information:
  Type                100GE SLOT1

```

```

ASIC type           Brooklyn 100GE FPGA
State               Online
PIC version         1.3
Uptime              10 minutes, 44 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	100GBASE LR4	SM	Opnext Inc.	TRC5E20ENFSF000F	1310 nm

### show chassis pic fpc-slot pic-slot lcc (TX Matrix Router)

```

user@host> show chassis pic fpc-slot 1 pic-slot 1 lcc 0
lcc0-re0:

```

-----

PIC fpc slot 1 pic slot 1 information:

```

Type               4x OC-3 SONET, SMIR
ASIC type          D chip
State              Online
PIC version         1.2
Uptime              5 days, 2 hours, 12 minutes, 8 seconds

```

### show chassis pic fpc-slot pic-slot lcc (TX Matrix Plus Router)

```

user@host> show chassis pic pic-slot 0 fpc-slot 8
lcc0-re0:

```

-----

FPC slot 8, PIC slot 0 information:

```

Type               1x 10GE(LAN/WAN)
State              Online
Uptime              2 hours, 46 minutes, 23 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	part number	Wavelength
0	10GBASE ZR	SM	Opnext Inc.	TRF7061BN-LF150	1550 nm
0	10GBASE ZR	SM	FINISAR CORP.	FTRX-1811-3-J2	1550 nm

### show chassis pic fpc-slot pic-slot (Next-Generation SONET/SDH SFP)

```

user@host> show chassis pic fpc-slot 4 pic-slot 0

```

FPC slot 4, PIC slot 0 information:

```

Type               4x OC-3 1x OC-12 SFP
ASIC type          D FPGA
State              Online
PIC version         1.3
Uptime              1 day, 50 minutes, 4 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC48 short reach	SM	FINISAR CORP.	FTRJ1321P18TL-J2	1310 nm
1	OC3 short reach	MM	QCP	TRPA03MM3BAS-JE	1310 nm
2	OC3 short reach	MM	QCP	TRXA03MM3BAS-JW	1310 nm
3	OC12 inter reach	SM	FINISAR CORP.	FTLF1322P18TR	1310 nm

### show chassis pic fpc-slot pic-slot (12-Port T1/E1)

```

user@host> show chassis pic fpc-slot 0 pic-slot 3

```

FPC slot 0, PIC slot 3 information:

```

Type                12x T1/E1 CE
State               Online
PIC version         1.1
CPU load average    1 percent
Interrupt load average 0 percent
Total DRAM size     128 MB
Memory buffer utilization 100 percent
Memory heap utilization 4 percent
Uptime              1 day, 22 hours, 28 minutes, 12 seconds
Internal Clock Synchronization Normal

```

#### show chassis pic fpc-slot pic-slot (4x CHOC3 SONET CE SFP)

user@host> show chassis pic fpc-slot 0 pic-slot 1

FPC slot 0, PIC slot 1 information:

```

Type                4x CHOC3 SONET CE SFP
State               Online
PIC version         1.3
CPU load average    1 percent
Interrupt load average 0 percent
Total DRAM size     128 MB
Memory buffer utilization 99 percent
Memory heap utilization 4 percent
Uptime              1 day, 22 hours, 55 minutes, 37 seconds
Internal Clock Synchronization Normal

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
1	OC3 short reach	MM	AVAGO	HFBR-57E0P-JU2	n/a
3	OC3 long reach	SM	OPNEX INC	TRF5456AVLB314	1310 nm

#### show chassis pic fpc-slot pic-slot (SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

user@host> show chassis pic fpc-slot 0 pic-slot 0

FPC slot 0, PIC slot 0 information:

```

Type                MIC-3D-80C30C12-40C48
State               Online
PIC version         1.8
Uptime              3 days, 22 hours, 3 minutes, 50 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm
7	OC12 inter reach	SM	FINISAR CORP	FTRJ1322P1BTR-J3	1310 nm

Multirate Mode Enabled

#### show chassis pic fpc-slot pic-slot (8-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

user@host> show chassis pic fpc-slot 3 pic-slot 0

FPC slot 3, PIC slot 0 information:

```

Type                MIC-3D-8CHOC3-4CHOC12
State               Online
PIC version         1.9
Uptime              1 hour, 21 minutes, 24 seconds

```

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
------	------------	------------	-------------	-------------------------	------------

0	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
1	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J2	1310 nm
4	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
5	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
6	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
7	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

#### show chassis pic fpc-slot pic-slot (4-port Channelized SONET/SDH OC3/STM1 [Multi-Rate] MIC with SFP)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
```

FPC slot 5, PIC slot 0 information:

Type	MIC-3D-4CHOC3-2CHOC12
State	Online
PIC version	1.9
Uptime	1 hour, 21 minutes

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
1	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
2	OC12 inter reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm
3	OC12 short reach	SM	FINISAR CORP.	FTRJ1322P1BTR-J3	1310 nm

#### show chassis pic fpc-slot pic-slot (1-port OC192/STM64 MIC with XFP)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
```

FPC slot 1, PIC slot 0 information:

Type	MIC-3D-10C192-XFP
State	Online
PIC version	1.2
Uptime	1 day, 11 hours, 4 minutes, 6 seconds

PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	OC192 short reach	n/a	FINISAR CORP.	FTLX1412M3BCL-J3	1310 nm

#### show chassis pic fpc-slot 1 pic-slot 2 (8-port DS3/E3 MIC)

```
user@host> show chassis pic fpc-slot 1 pic-slot 2
```

FPC slot 1, PIC slot 2 information:

Type	MIC-3D-8DS3-E3
State	Online
PIC version	1.10
Uptime	4 days, 1 hour, 29 minutes, 19 seconds
Channelization Mode	Disabled

#### show chassis pic fpc-slot pic-slot (OTN)

```
user@host> show chassis pic fpc-slot 5 pic-slot 0
```

PIC fpc slot 5 pic slot 0 information:

Type	1x10GE(LAN),OTN
ASIC type	H chip
State	Online
PIC version	1.0
Uptime	5 minutes, 50 seconds

#### show chassis pic fpc-slot pic-slot (QFX3500 Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
```



```
FPC slot 0, PIC slot 0 information:
Type 48x 10G-SFP+ Builtin
State Online
Uptime 3 days, 3 hours, 5 minutes, 20 seconds
```

#### show chassis pic fpc-slot pic-slot (QFX5100 Standalone Switch)

```
user@switch> show chassis pic fpc-slot 0 pic-slot 0
FPC slot 0, PIC slot 0 information:
Type                               Unknown Builtin
State                              Online
Uptime                             1 day, 17 hours, 5 minutes, 9 seconds
```

#### show chassis pic interconnect-device fpc-slot pic-slot (QFabric Systems)

```
user@switch> show chassis pic interconnect-device interconnect1 fpc-slot 9 pic-slot 0
FPC slot 9, PIC slot 0 information:
Type                               16x 40G-GE Builtin
State                              Online
Uptime                             2 hours, 47 minutes, 40 seconds
```

#### show chassis pic node-device fpc-slot pic-slot (QFabric System)

```
user@switch> show chassis pic node-device node1 pic-slot 0
FPC slot node1, PIC slot 0 information:
Type                               48x 10G-SFP+ Builtin
State                              Online
Uptime                             2 hours, 52 minutes, 37 seconds
```

#### PIC port information:

Port	Cable type	Fiber type	Xcvr vendor	Xcvr vendor part number	Wavelength
0	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
1	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
2	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
3	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
4	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
5	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
6	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
7	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
8	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
9	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
10	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
11	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
12	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
13	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
14	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
15	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
16	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
17	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
18	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
19	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
20	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
21	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
22	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
23	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
24	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
25	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
26	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
27	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
28	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
29	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

30	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
31	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
32	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
33	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
34	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
35	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
36	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
37	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
38	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
39	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
40	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
41	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
42	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
43	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
44	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
45	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
46	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm
47	10GBASE SR	MM	SumitomoElectric	SPP5101SR-J3	850 nm

#### show chassis pic fpc-slot pic-slot (ACX2000 Universal Access Router)

```
user@host> show chassis pic fpc-slot 0 pic-slot 1
FPC slot 0, PIC slot 1 information:
  Type                8x 1GE(LAN) RJ45 Built-in
  State                Online
  Uptime               6 days, 2 hours, 51 minutes, 11 seconds
```

#### show chassis pic fpc-slot pic-slot (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 0
FPC slot 1, PIC slot 0 information:
  Type                AS-MSB
  State                Online
  PIC version          1.6
  Uptime               11 hours, 17 minutes, 56 seconds
```

#### show chassis pic FPC slot PIC slot (MX Routers with Media Services Blade [MSB])

```
user@switch> show chassis pic fpc-slot 1 pic-slot 2
Type                AS-MXC
State                Online
PIC version          1.0
Uptime               11 hours, 18 minutes, 3 seconds
```

#### show chassis pic transport fpc-slot pic-slot (PTX Series Packet Transport Routers)

```
user@host> show chassis pic transport fpc-slot 2 pic-slot 0
Administrative State: In Service
Operational State:   Normal
```

## show chassis routing-engine

**List of Syntax**    [Syntax on page 905](#)  
                           [Syntax \(EX Series Switches\) on page 905](#)  
                           [Syntax \(T Series routers\) on page 905](#)  
                           [Syntax \(TX Matrix Routers\) on page 905](#)  
                           [Syntax \(TX Matrix Plus Routers\) on page 905](#)  
                           [Syntax \(QFX Series\) on page 905](#)  
                           [Syntax \(MX Series Routers\) on page 905](#)  
                           [Syntax \(MX2010 3D Universal Edge Routers\) on page 905](#)  
                           [Syntax \(MX2020 3D Universal Edge Routers\) on page 905](#)  
                           [Syntax \(MX104 3D Universal Edge Routers\) on page 905](#)  
                           [Syntax \(ACX Series Universal Access Routers\) on page 906](#)

**Syntax**    show chassis routing-engine  
                   <bios | *slot*>

**Syntax (EX Series Switches)**    show chassis routing-engine  
                                           <*slot*>

**Syntax (T Series routers)**    show chassis routing-engine  
                                           <bios | *slot*>

**Syntax (TX Matrix Routers)**    show chassis routing-engine  
                                           <bios | *slot*>  
                                           <lcc *number* | scc>

**Syntax (TX Matrix Plus Routers)**    show chassis routing-engine  
                                           <bios | *slot*>  
                                           <lcc *number* | sfc *number*>

**Syntax (QFX Series)**    show chassis routing-engine  
                                   <interconnect-device *name*>  
                                   <node-device *name*>

**Syntax (MX Series Routers)**    show chassis routing-engine  
                                           <bios | *slot*>  
                                           <all-members>  
                                           <local>  
                                           <member *member-id*>

**Syntax (MX2010 3D Universal Edge Routers)**    show chassis routing-engine  
                                           <bios | *slot*>

**Syntax (MX2020 3D Universal Edge Routers)**    show chassis routing-engine  
                                           <bios | *slot*>

**Syntax (MX104 3D Universal Edge Routers)**    show chassis routing-engine

**Syntax (ACX Series Universal Access Routers)** `show chassis routing-engine`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
**sfc** option introduced for the TX Matrix Plus router in Junos OS Release in 9.6.  
Command introduced in Junos OS Release 11.1 for QFX Series.  
Command introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.  
Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.  
Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.  
Command introduced in Junos OS Release 13.2 for MX104 3D Univesral Edge Routers.

**Description** Display the status of the Routing Engine.

**Options** **none**—Display information about one or more Routing Engines. On a TX Matrix router, display information about all Routing Engines on the TX Matrix router and its attached T640 routers. On a TX Matrix Plus router, display information about all Routing Engines on the TX Matrix Plus router and its attached routers.

**all-members**—(MX Series routers only) (Optional) Display Routing Engine information for all members of the Virtual Chassis configuration.

**bios**—(Optional) Display the (BIOS) firmware version.

**interconnect-device *number***—(QFabric systems only) (Optional) Display Routing Engine information for a specified Interconnect device.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Routing Engine information for a specified T640 router (line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, display Routing Engine information for a specified router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display Routing Engine information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display Routing Engine information for the specified member of the Virtual Chassis configuration. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-device *number***—(QFabric systems only) (Optional) Display Routing Engine information for a specified Node device.

**scc**—(TX Matrix routers only) (Optional) Display Routing Engine information for the TX Matrix router (switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display Routing Engine information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**slot**—(Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace *slot* with 0 or 1. For QFX3500 switches, there is only one Routing Engine, so you do not need to specify the slot number.

**Required Privilege Level** view

**Related Documentation**

- [request chassis routing-engine master on page 392](#)
- *Configuring Routing Engine Redundancy*
- *Switching the Global Master and Backup Roles in a Virtual Chassis Configuration*

**List of Sample Output**

- [show chassis routing-engine \(M5 Router\) on page 909](#)
- [show chassis routing-engine \(M10 Router\) on page 910](#)
- [show chassis routing-engine \(M20 Router\) on page 910](#)
- [show chassis routing-engine \(M40 Router\) on page 911](#)
- [show chassis routing-engine \(M120 Router\) on page 911](#)
- [show chassis routing-engine \(M160 Router\) on page 912](#)
- [show chassis routing-engine \(MX104 Router\) on page 912](#)
- [show chassis routing-engine \(MX240 Router\) on page 913](#)
- [show chassis routing-engine \(MX480 Router\) on page 914](#)
- [show chassis routing-engine \(MX960 Router\) on page 914](#)
- [show chassis routing-engine \(MX2010 Router\) on page 914](#)
- [show chassis routing-engine \(MX2020 Router\) on page 915](#)
- [show chassis routing-engine \(T320 router\) on page 916](#)
- [show chassis routing-engine \(T640 router\) on page 917](#)
- [show chassis routing-engine \(T1600 router\) on page 917](#)
- [show chassis routing-engine \(T4000 router\) on page 918](#)
- [show chassis routing-engine \(TX Matrix Router\) on page 919](#)
- [show chassis routing-engine lcc \(TX Matrix Router\) on page 920](#)
- [show chassis routing-engine bios \(TX Matrix Router\) on page 920](#)
- [show chassis routing-engine \(TX Matrix Plus Router\) on page 921](#)
- [show chassis routing-engine lcc \(TX Matrix Plus Router\) on page 922](#)
- [show chassis routing-engine bios \(TX Matrix Plus Router\) on page 923](#)
- [show chassis routing-engine \(QFX Series\) on page 923](#)
- [show chassis routing-engine interconnect-device \(QFabric systems\) on page 923](#)
- [show chassis routing-engine \(PTX Series Packet Transport Switch\) on page 924](#)

[show chassis routing-engine \(EX9200 Switch\) on page 925](#)

[show chassis routing-engine \(ACX2000 Universal Access Router\) on page 925](#)

[show chassis routing-engine \(ACX1000 Universal Access Router\) on page 926](#)

**Output Fields** [Table 43 on page 908](#) lists the output fields for the **show chassis routing-engine** command. Output fields are listed in the approximate order in which they appear.

**Table 43: show chassis routing-engine Output Fields**

Field Name	Field Description
<b>Slot</b>	(Systems with single and multiple Routing Engines) Slot number.
<b>Current state</b>	(Systems with multiple Routing Engines) Current state of the Routing Engine: <b>Master</b> , <b>Backup</b> , or <b>Disabled</b> .
<b>Election priority</b>	(Systems with multiple Routing Engines) Election priority for the Routing Engine: <b>Master</b> or <b>Backup</b> .
<b>Temperature</b>	Temperature of the air flowing past the Routing Engine.
<b>CPU Temperature</b>	Temperature of the CPU.
<b>DRAM</b>	Total DRAM available to the Routing Engine's processor.  Starting with Junos OS Release 12.3R1, the DRAM field displays both available memory and installed memory.
<b>Memory utilization</b>	Percentage of Routing Engine memory being used.
<b>CPU utilization</b>	Information about the Routing Engine's CPU utilization: <ul style="list-style-type: none"> <li>• <b>User</b>—Percentage of CPU time being used by user processes.</li> <li>• <b>Background</b>—Percentage of CPU time being used by background processes.</li> <li>• <b>Kernel</b>—Percentage of CPU time being used by kernel processes.</li> <li>• <b>Interrupt</b>—Percentage of CPU time being used by interrupts.</li> <li>• <b>Idle</b>—Percentage of CPU time that is idle.</li> </ul>
<b>Model</b>	Routing Engine model number.
<b>Serial ID</b>	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
<b>Start time</b>	Time at which the Routing Engine started running.
<b>Uptime</b>	How long the Routing Engine has been running.
<b>Routing Engine BIOS Version</b>	BIOS version being run by the Routing Engine.

Table 43: show chassis routing-engine Output Fields (*continued*)

Field Name	Field Description
Last reboot reason	<p>Reason for last reboot, including:</p> <ul style="list-style-type: none"> <li><b>power cycle/failure</b>—Halt of the Routing Engine using the <b>halt</b> command, powering down using the power button on the chassis or any other method (such as removal of the control board or Routing Engine), and then powering back the Routing Engine. A halt of the operating system also occurs if you enter the <b>request system halt</b> command. You can enter this command to halt the system operations on the chassis or specific Routing Engines. To restart the software, press any key on the keyboard.</li> <li><b>watchdog</b>—Reboot due to a hardware watchdog. A watchdog is a hardware monitoring process that examines the health and performance of the router to enable the device to recover from failures. A watchdog checks for problems at certain intervals, and reboots the routing engine if a problem is encountered.</li> <li><b>reset-button reset</b>—(Not available on the J Series router or EX Series switch) Reboot due to pressing of the reset button on the Routing Engine.</li> <li><b>power-button hard power off</b>—Reboot due to pressing of the power button on the chassis. A powering down of the software also occurs if you enter the <b>request system power-off</b> command. You can enter this command to power down the chassis or specific Routing Engines; you can then restart the software.</li> <li><b>misc hardware reason</b>—Reboot due to miscellaneous hardware reasons.</li> <li><b>thermal shutdown</b>—Reboot due to the router or switch reaching a critical temperature at which point it is unsafe to continue operations.</li> <li><b>hard disk failure</b>—Reboot due to a hard disk or solid-state drive (SSD) failure.</li> <li><b>reset from debugger</b>—Reboot due to reset from the debugger.</li> <li><b>chassis control reset</b>—Restart the chassis process that manages PICs, FPCs, and other hardware components. The chassis control module that runs the Routing Engine performs management and monitoring functions, and it provides a single access point for operational and maintenance functions. A reset of the chassis management process occurs when you enter the <b>restart chassis-control</b> command.</li> <li><b>bios auto recovery reset</b>—Reboot due to a BIOS auto-recovery reset.</li> <li><b>could not be determined</b>—Reboot due to an undetermined reason.</li> <li><b>Router rebooted after a normal shutdown</b>—Reboot due to a normal shutdown. This reason is displayed if the Routing Engine is powered down by pushing and holding the online/offline button on the Routing Engine faceplate for 30 seconds, and then powered back. A reboot of the software also occurs if you enter the <b>request system reboot</b> command. You can enter this command to reboot the chassis or specific Routing Engines.</li> </ul>
Load averages	Routing Engine load averages for the last 1, 5, and 15 minutes.

## Sample Output

### show chassis routing-engine (M5 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature                25 degrees C / 77 degrees F
  DRAM                       768 MB
  Memory utilization         21 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                  0 percent

```

Idle	100 percent
Model	RE-2.0
Serial ID	31000007349bf701
Start time	2003-12-04 09:42:17 PST
Uptime	26 days, 1 hour, 12 minutes, 27 seconds
Last reboot reason	Router rebooted after a normal shutdown
Load averages:	1 minute 5 minute 15 minute
	0.00 0.01 0.00

#### show chassis routing-engine (M10 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
  Temperature      25 degrees C / 77 degrees F
  DRAM             768 MB
  Memory utilization 21 percent
  CPU utilization:
    User           0 percent
    Background     0 percent
    Kernel         0 percent
    Interrupt      0 percent
    Idle           100 percent
  Model            RE-2.0
  Serial ID        31000007349bf701
  Start time       2003-12-04 09:42:17 PST
  Uptime           26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:  1 minute 5 minute 15 minute
                  0.00      0.01      0.00
```

#### show chassis routing-engine (M20 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority  Master (default)
  Temperature        29 degrees C / 84 degrees F
  DRAM               768 MB
  Memory utilization 20 percent
  CPU utilization:
    User             1 percent
    Background       0 percent
    Kernel           2 percent
    Interrupt        0 percent
    Idle             97 percent
  Model              RE-2.0
  Serial ID          58000007348d9a01
  Start time         2003-12-30 07:05:47 PST
  Uptime             3 hours, 41 minutes, 14 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:    1 minute 5 minute 15 minute
                    0.00      0.02      0.00

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority  Backup (default)
  Temperature        29 degrees C / 84 degrees F
  DRAM               768 MB
  Memory utilization 0 percent
  CPU utilization:
```



```

User                0 percent
Background          0 percent
Kernel              1 percent
Interrupt           0 percent
Idle                99 percent
Model               RE-2.0
Serial ID            d800000734745701
Start time          2003-06-17 16:37:33 PDT
Uptime              195 days, 18 hours, 47 minutes, 9 seconds
Last reboot reason   Router rebooted after a normal shutdown

```

### show chassis routing-engine (M40 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature        25 degrees C / 77 degrees F
  DRAM                768 MB
  Memory utilization  21 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            0 percent
    Interrupt         0 percent
    Idle              100 percent
  Model              RE-2.0
  Serial ID          31000007349bf701
  Start time         2003-12-04 09:42:17 PST
  Uptime             26 days, 1 hour, 12 minutes, 27 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:     1 minute   5 minute  15 minute
                      0.00       0.01    0.00

```

### show chassis routing-engine (M120 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        46 degrees C / 114 degrees F
  CPU temperature     44 degrees C / 111 degrees F
  DRAM                2048 MB
  Memory utilization  18 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel            5 percent
    Interrupt         0 percent
    Idle              95 percent
  Model              RE-A-1000
  Serial ID          1000621154
  Start time         2006-10-31 17:10:05 PST
  Uptime             14 minutes, 31 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:     1 minute   5 minute  15 minute
                      0.02       0.07    0.07

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        45 degrees C / 113 degrees F

```

```
CPU temperature      42 degrees C / 107 degrees F
DRAM                2048 MB
Memory utilization   15 percent
CPU utilization:
  User              0 percent
  Background        0 percent
  Kernel            0 percent
  Interrupt          0 percent
  Idle              100 percent
Model               RE-A-1000
Serial ID           1000621151
Start time          2006-10-31 17:10:04 PST
Uptime              14 minutes, 30 seconds
Last reboot reason   Router rebooted after a normal shutdown
```

### show chassis routing-engine (M160 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority   Master (default)
  Temperature        43 degrees C / 109 degrees F
  DRAM               2048 MB
  Memory utilization  11 percent
  CPU utilization:
    User             1 percent
    Background       0 percent
    Kernel           2 percent
    Interrupt        0 percent
    Idle             97 percent
  Model              RE-3.0
  Serial ID          210865700403
  Start time         2003-12-23 12:25:55 PST
  Uptime             6 days, 22 hours, 33 minutes, 24 seconds
  Last reboot reason Router rebooted after a normal shutdown
  Load averages:    1 minute   5 minute   15 minute
                   0.24      0.13      0.04

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority   Backup (default)
  Temperature        40 degrees C / 104 degrees F
  DRAM               2048 MB
  Memory utilization  9 percent
  CPU utilization:
    User             0 percent
    Background       0 percent
    Kernel           0 percent
    Interrupt        0 percent
    Idle             100 percent
  Model              RE-3.0
  Serial ID          210865700332
  Start time         2003-12-23 12:25:55 PST
  Uptime             6 days, 22 hours, 33 minutes, 21 seconds
  Last reboot reason Router rebooted after a normal shutdown
```

### show chassis routing-engine (MX104 Router)

```
user@host> show chassis routing-engine
```

```

Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             32 degrees C / 89 degrees F
  CPU temperature         42 degrees C / 107 degrees F
  DRAM                   3840 MB (3840 MB installed)
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                3 percent
    Interrupt             2 percent
    Idle                  94 percent
  Model                   RE-MX-104
  Serial ID               CAAR5925
  Start time              2013-06-05 13:17:08 IST
  Uptime                  1 hour, 15 minutes, 8 seconds
  Last reboot reason      0x200:normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.87       0.90       0.41

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)
  Temperature             32 degrees C / 89 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                   3840 MB (3840 MB installed)
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                1 percent
    Interrupt             2 percent
    Idle                  97 percent
  Model                   RE-MX-104
  Serial ID               CAAM6369
  Start time              2013-06-05 13:07:37 IST
  Uptime                  1 hour, 24 minutes, 34 seconds
  Last reboot reason      0x200:normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.19       0.15       0.06

```

### show chassis routing-engine (MX240 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
  Temperature             40 degrees C / 104 degrees F
  CPU temperature         47 degrees C / 116 degrees F
  DRAM                   3584 MB
  Memory utilization      7 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                0 percent
    Interrupt             0 percent
    Idle                  100 percent
  Model                   RE-S-2000

```

Serial ID	1000703522
Start time	2007-12-19 10:35:40 PST
Uptime	16 days, 3 hours, 15 minutes, 23 seconds
Last reboot reason	Router rebooted after a normal shutdown

#### show chassis routing-engine (MX480 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             41 degrees C / 105 degrees F
  CPU temperature         38 degrees C / 100 degrees F
  DRAM                   2048 MB
  Memory utilization      13 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                2 percent
    Interrupt             0 percent
    Idle                  98 percent
  Model                  RE-S-1300
  Serial ID              1000697044
  Start time             2008-01-04 06:46:08 PST
  Uptime                 8 hours, 17 minutes, 16 seconds
  Last reboot reason      Router rebooted after a normal shutdown
```

#### show chassis routing-engine (MX960 Router)

```
user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             37 degrees C / 98 degrees F
  CPU temperature         37 degrees C / 98 degrees F
  DRAM                   2048 MB
  Memory utilization      18 percent
  CPU utilization:
    User                  0 percent
    Background            0 percent
    Kernel                4 percent
    Interrupt             0 percent
    Idle                  96 percent
  Model                  RE-S-1300
  Serial ID              1000617944
  Start time             2006-10-26 12:37:13 PDT
  Uptime                 6 days, 4 hours, 59 minutes, 40 seconds
  Last reboot reason      Router rebooted after a normal shutdown
  Load averages:         1 minute   5 minute   15 minute
                        0.16       0.08       0.02
```

#### show chassis routing-engine (MX2010 Router)

```
user@host> show chassis routing-engine

Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             3 degrees C / 37 degrees F
```

```

CPU temperature          3 degrees C / 37 degrees F
DRAM                    17152 MB
Memory utilization       13 percent
CPU utilization:
  User                   0 percent
  Background             0 percent
  Kernel                 4 percent
  Interrupt              2 percent
  Idle                   95 percent
Model                   RE-S-1800x4
Serial ID                9009099704
Start time              2012-10-02 14:33:32 PDT
Uptime                  14 hours, 39 minutes, 39 seconds
Last reboot reason      Router rebooted after a normal shutdown.
Load averages:          1 minute   5 minute   15 minute
                        0.06       0.05       0.01

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            1 degrees C / 33 degrees F
  CPU temperature        2 degrees C / 35 degrees F
  DRAM                   17152 MB
  Memory utilization     11 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-S-1800x4
  Serial ID              9009099706
  Start time             2012-10-02 10:36:06 PDT
  Uptime                  18 hours, 36 minutes, 57 seconds
  Last reboot reason      Router rebooted after a normal shutdown.
  Load averages:         1 minute   5 minute   15 minute
                        0.01       0.00       0.00

```

### show chassis routing-engine (MX2020 Router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            6 degrees C / 42 degrees F
  CPU temperature        6 degrees C / 42 degrees F
  DRAM                   17152 MB
  Memory utilization     14 percent
  CPU utilization:
    User                 1 percent
    Background           0 percent
    Kernel               7 percent
    Interrupt            2 percent
    Idle                 91 percent
  Model                  RE-S-1800x4
  Serial ID              9009089704
  Start time             2012-10-02 11:05:24 PDT
  Uptime                  2 days, 15 hours, 49 minutes, 13 seconds
  Last reboot reason      Router rebooted after a normal shutdown.
  Load averages:         1 minute   5 minute   15 minute

```

```

                                0.10      0.05      0.01
Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  7 degrees C / 44 degrees F
  CPU temperature              5 degrees C / 41 degrees F
  DRAM                        17152 MB
  Memory utilization           12 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                  0 percent
    Idle                      99 percent
  Model                       RE-S-1800x4
  Serial ID                   9009094138
  Start time                  2012-10-02 11:09:57 PDT
  Uptime                      2 days, 15 hours, 44 minutes, 27 seconds
  Last reboot reason          Router rebooted after a normal shutdown.
  Load averages:             1 minute  5 minute 15 minute
                                0.00      0.00      0.00

```

#### show chassis routing-engine (T320 router)

```

user@host> show chassis routing-engine
Slot 0:
  Current state                Master
  Election priority            Master (default)
  Temperature                  51 degrees C / 123 degrees F
  CPU temperature              55 degrees C / 131 degrees F
  DRAM                        3584 MB
  Memory utilization           11 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    2 percent
    Interrupt                  0 percent
    Idle                      97 percent
  Model                       RE-A-2000
  Serial ID                   9009010618
  Start time                  2012-10-10 01:24:05 PDT
  Uptime                      5 days, 10 hours, 49 minutes, 23 seconds
  Last reboot reason          0x1:power cycle/failure
  Load averages:             1 minute  5 minute 15 minute
                                0.00      0.05      0.04

Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  45 degrees C / 113 degrees F
  CPU temperature              48 degrees C / 118 degrees F
  DRAM                        3584 MB
  Memory utilization           9 percent
  CPU utilization:
    User                      0 percent
    Background                0 percent
    Kernel                    0 percent
    Interrupt                  0 percent
    Idle                      100 percent
  Model                       RE-A-2000

```

```

Serial ID          9009003642
Start time        2012-10-10 01:24:04 PDT
Uptime           5 days, 10 hours, 49 minutes, 28 seconds
Last reboot reason 0x1:power cycle/failure

```

### show chassis routing-engine (T640 router)

```
user@host> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

```

Current state      Master
Election priority  Master (default)
Temperature        50 degrees C / 122 degrees F
CPU temperature    58 degrees C / 136 degrees F
DRAM              3584 MB
Memory utilization 14 percent
CPU utilization:
  User            1 percent
  Background      0 percent
  Kernel          4 percent
  Interrupt       1 percent
  Idle           95 percent
Model            RE-A-2000
Serial ID        1000686556
Start time      2012-10-10 01:24:02 PDT
Uptime         5 days, 10 hours, 50 minutes, 27 seconds
Last reboot reason 0x1:power cycle/failure
Load averages:  1 minute   5 minute  15 minute
                  1.24      0.33     0.12

```

```
Routing Engine status:
```

```
Slot 1:
```

```

Current state      Backup
Election priority  Backup (default)
Temperature        44 degrees C / 111 degrees F
CPU temperature    49 degrees C / 120 degrees F
DRAM              3584 MB
Memory utilization 12 percent
CPU utilization:
  User            0 percent
  Background      0 percent
  Kernel          0 percent
  Interrupt       1 percent
  Idle           99 percent
Model            RE-A-2000
Serial ID        1000702739
Start time      2012-10-10 01:24:02 PDT
Uptime         5 days, 10 hours, 50 minutes, 26 seconds
Last reboot reason 0x1:power cycle/failure

```

### show chassis routing-engine (T1600 router)

```
user@host> show chassis routing-engine
```

```
Routing Engine status:
```

```
Slot 0:
```

```

Current state      Master
Election priority  Master (default)
Temperature        48 degrees C / 118 degrees F
CPU temperature    58 degrees C / 136 degrees F
DRAM              3584 MB
Memory utilization 13 percent
CPU utilization:

```

```

User                0 percent
Background          0 percent
Kernel              3 percent
Interrupt            1 percent
Idle                96 percent
Model               RE-A-2000
Serial ID            1000704521
Start time           2012-10-10 01:23:41 PDT
Uptime              5 days, 10 hours, 46 minutes, 56 seconds
Last reboot reason   0x1:power cycle/failure
Load averages:      1 minute   5 minute   15 minute
                    0.05       0.03       0.01

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority  Backup (default)
  Temperature        44 degrees C / 111 degrees F
  CPU temperature    48 degrees C / 118 degrees F
  DRAM               3584 MB
  Memory utilization 12 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             0 percent
    Interrupt          0 percent
    Idle              100 percent
  Model              RE-A-2000
  Serial ID           9009006579
  Start time          2012-10-10 01:23:42 PDT
  Uptime              5 days, 10 hours, 46 minutes, 54 seconds
  Last reboot reason  0x1:power cycle/failure

```

#### show chassis routing-engine (T4000 router)

```

user@host> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state      Master
  Election priority  Master (default)
  Temperature        33 degrees C / 91 degrees F
  CPU temperature    50 degrees C / 122 degrees F
  DRAM               8960 MB
  Memory utilization 18 percent
  CPU utilization:
    User              0 percent
    Background        0 percent
    Kernel             4 percent
    Interrupt          1 percent
    Idle              95 percent
  Model              RE-DUO-1800
  Serial ID           P737F-002248
  Start time          2012-02-09 22:49:53 PST
  Uptime              2 hours, 21 minutes, 35 seconds
  Last reboot reason  Router rebooted after a normal shutdown.
  Load averages:      1 minute   5 minute   15 minute
                    0.00       0.04       0.00

Routing Engine status:
Slot 1:
  Current state      Backup
  Election priority  Backup (default)
  Temperature        32 degrees C / 89 degrees F

```



```

CPU temperature      46 degrees C / 114 degrees F
DRAM                8960 MB
Memory utilization   24 percent
CPU utilization:
  User              0 percent
  Background        0 percent
  Kernel            0 percent
  Interrupt         0 percent
  Idle              99 percent
Model               RE-DUO-1800
Serial ID            P737F-002653
Start time           2012-02-08 20:12:51 PST
Uptime               1 day, 4 hours, 58 minutes, 28 seconds
Last reboot reason   Router rebooted after a normal shutdown.

```

### show chassis routing-engine (TX Matrix Router)

```

user@host> show chassis routing-engine
scc-re0:

```

#### Routing Engine status:

##### Slot 0:

```

Current state      Master
Election priority   Master (default)
Temperature         34 degrees C / 93 degrees F
CPU temperature     33 degrees C / 91 degrees F
DRAM               2048 MB
Memory utilization  12 percent
CPU utilization:
  User              0 percent
  Background        0 percent
  Kernel            2 percent
  Interrupt         0 percent
  Idle              98 percent
Model              RE-4.0
Serial ID           P11123900153
Start time          2004-08-05 18:42:05 PDT
Uptime              9 days, 22 hours, 49 minutes, 50 seconds
Last reboot reason   Router rebooted after a normal shutdown
Load averages:      1 minute   5 minute   15 minute
                    0.00       0.08       0.07

```

#### lcc0-re0:

#### Routing Engine status:

##### Slot 0:

```

Current state      Master
Election priority   Master (default)
Temperature         33 degrees C / 91 degrees F
CPU temperature     30 degrees C / 86 degrees F
DRAM               2048 MB
Memory utilization  12 percent
CPU utilization:
  User              0 percent
  Background        0 percent
  Kernel            1 percent
  Interrupt         0 percent
  Idle              98 percent
Model              RE-3.0
Serial ID           210865700363
Start time          2004-08-05 18:42:05 PDT

```

```

Uptime                9 days, 22 hours, 48 minutes, 20 seconds
Last reboot reason    Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                        0.00       0.02       0.00

```

```
lcc2-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

Current state          Master
Election priority      Master (default)
Temperature            34 degrees C / 93 degrees F
CPU temperature        35 degrees C / 95 degrees F
DRAM                  2048 MB
Memory utilization     12 percent
CPU utilization:
  User                 0 percent
  Background           0 percent
  Kernel               2 percent
  Interrupt            0 percent
  Idle                 98 percent
Model                 RE-4.0
Serial ID              P11123900126
Start time             2004-08-05 18:42:05 PDT
Uptime                9 days, 22 hours, 49 minutes, 4 seconds
Last reboot reason    Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                        0.01       0.01       0.0

```

### show chassis routing-engine lcc (TX Matrix Router)

```
user@host> show chassis routing-engine 0 lcc 0
```

```
lcc0-re0:
```

```
-----
Routing Engine status:
```

```
Slot 0:
```

```

Current state          Master
Election priority      Master (default)
Temperature            33 degrees C / 91 degrees F
CPU temperature        30 degrees C / 86 degrees F
DRAM                  2048 MB
Memory utilization     12 percent
CPU utilization:
  User                 0 percent
  Background           0 percent
  Kernel               1 percent
  Interrupt            0 percent
  Idle                 98 percent
Model                 RE-3.0
Serial ID              210865700363
Start time             2004-08-05 18:42:05 PDT
Uptime                7 days, 22 hours, 49 minutes, 6 seconds
Last reboot reason    Router rebooted after a normal shutdown
Load averages:        1 minute   5 minute   15 minute
                        0.00       0.00       0.00

```

### show chassis routing-engine bios (TX Matrix Router)

```
user@host> show chassis routing-engine bios
```

```
scc-re0:
```

```
Routing Engine BIOS Version: V1.0.0
1cc0-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.17
1cc2-re0:
```

```
-----
Routing Engine BIOS Version: V1.0.0
```

### show chassis routing-engine (TX Matrix Plus Router)

```
user@host> show chassis routing-engine
sfc0-re0:
```

```
-----
Routing Engine status:
```

Slot 0:

Current state	Master
Election priority	Master (default)
Temperature	27 degrees C / 80 degrees F
CPU temperature	42 degrees C / 107 degrees F
DRAM	3327 MB
Memory utilization	12 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	2 percent
Interrupt	0 percent
Idle	98 percent
Model	RE-TXP-SFC
Serial ID	737A-1024
Start time	2009-05-11 17:39:49 PDT
Uptime	3 hours, 45 minutes, 25 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute    5 minute    15 minute
	0.00        0.00        0.00

Routing Engine status:

Slot 1:

Current state	Backup
Election priority	Backup (default)
Temperature	29 degrees C / 84 degrees F
CPU temperature	43 degrees C / 109 degrees F
DRAM	3327 MB
Memory utilization	11 percent
CPU utilization:	
User	0 percent
Background	0 percent
Kernel	0 percent
Interrupt	0 percent
Idle	100 percent
Model	RE-TXP-SFC
Serial ID	737A-1024
Start time	2009-05-11 17:08:54 PDT
Uptime	4 hours, 16 minutes, 52 seconds
Last reboot reason	0x1:power cycle/failure

```
1cc0-re0:
```

```
-----
Routing Engine status:
```

Slot 0:

Current state	Master
Election priority	Master (default)
Temperature	30 degrees C / 86 degrees F

```

CPU temperature          43 degrees C / 109 degrees F
DRAM                    3327 MB
Memory utilization       9 percent
CPU utilization:
  User                   0 percent
  Background             0 percent
  Kernel                 2 percent
  Interrupt              0 percent
  Idle                   98 percent
Model                   RE-TXP-LCC
Serial ID                737F-1024
Start time              2009-05-11 17:40:32 PDT
Uptime                  3 hours, 44 minutes, 51 seconds
Last reboot reason      Router rebooted after a normal shutdown.
Load averages:          1 minute   5 minute   15 minute
                        0.00       0.00       0.00

Routing Engine status:
Slot 1:
  Current state          Backup
  Election priority      Backup (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               0 percent
    Interrupt            0 percent
    Idle                 100 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-06 17:31:32 PDT
  Uptime                 5 days, 3 hours, 54 minutes, 19 seconds
  Last reboot reason     Router rebooted after a normal shutdown.

```

### show chassis routing-engine lcc (TX Matrix Plus Router)

```

user@host> show chassis routing-engine 0 lcc 0
1cc0-re0:
-----
Routing Engine status:
Slot 0:
  Current state          Master
  Election priority      Master (default)
  Temperature            30 degrees C / 86 degrees F
  CPU temperature        43 degrees C / 109 degrees F
  DRAM                   3327 MB
  Memory utilization     9 percent
  CPU utilization:
    User                 0 percent
    Background           0 percent
    Kernel               2 percent
    Interrupt            0 percent
    Idle                 98 percent
  Model                  RE-TXP-LCC
  Serial ID              737F-1024
  Start time             2009-05-11 17:40:32 PDT
  Uptime                 3 hours, 45 minutes, 26 seconds
  Last reboot reason     Router rebooted after a normal shutdown.
  Load averages:        1 minute   5 minute   15 minute

```

```

                                0.00      0.00      0.00
Routing Engine status:
Slot 1:
  Current state                Backup
  Election priority            Backup (default)
  Temperature                  30 degrees C / 86 degrees F
  CPU temperature              43 degrees C / 109 degrees F
  DRAM                        3327 MB
  Memory utilization           9 percent
  CPU utilization:
    User                       0 percent
    Background                 0 percent
    Kernel                     0 percent
    Interrupt                   0 percent
    Idle                       100 percent
  Model                        RE-TXP-LCC
  Serial ID                    737F-1024
  Start time                   2009-05-06 17:31:32 PDT
  Uptime                       5 days, 3 hours, 54 minutes, 59 seconds
  Last reboot reason           Router rebooted after a normal shutdown.

```

#### show chassis routing-engine bios (TX Matrix Plus Router)

```

user@host> show chassis routing-engine bios
sfc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.Z

```

```

lcc0-re0:

```

```

-----
Routing Engine BIOS Version: V0.0.N

```

#### show chassis routing-engine (QFX Series)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state Master
  Election priority Master (default)
  DRAM 2820 MB
  Memory utilization 49 percent
  CPU utilization:
    User 1 percent
    Background 0 percent
    Kernel 1 percent
    Interrupt 0 percent
    Idle 97 percent
  Model QFX3500-48S4Q
  Serial ID S/N ED3709
  Uptime 3 days, 4 hours, 29 minutes, 42 seconds
  Last reboot reason 0x200:chassis control reset
  Load averages: 1 minute 5 minute 15 minute
0.37 0.26 0.19

```

#### show chassis routing engine interconnect-device (QFabric systems)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state                Master
  Election priority            Master (default)
  Temperature                  48 degrees C / 118 degrees F

```

```

DRAM                                3312 MB
Memory utilization                    63 percent
CPU utilization:
  User                               14 percent
  Background                         0 percent
  Kernel                             5 percent
  Interrupt                           0 percent
  Idle                               81 percent
Model                                RE-QFXC08-CB4S
Serial ID                            BUILTIN
Start time                           2011-07-06 13:26:15 UTC
Uptime                               11 hours, 24 minutes, 57 seconds
Last reboot reason                    0x4:reset-button reset
Load averages:                      1 minute   5 minute   15 minute
                                      2.62       2.31       2.28

Routing Engine status:
Slot 1:
  Current state                       Backup
  Election priority                   Backup (default)
  Temperature                         39 degrees C / 102 degrees F
  DRAM                                3312 MB
  Memory utilization                  59 percent
  CPU utilization:
    User                             9 percent
    Background                       0 percent
    Kernel                           1 percent
    Interrupt                         0 percent
    Idle                             91 percent
  Model                              RE-QFXC08-CB4S
  Serial ID                          BUILTIN
  Start time                         2011-07-06 13:24:58 UTC
  Uptime                             11 hours, 26 minutes, 18 seconds
  Last reboot reason                  0x4:reset-button reset

```

### show chassis routing-engine (PTX Series Packet Transport Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state                       Master
  Election priority                   Master (default)
  Temperature                         60 degrees C / 140 degrees F
  CPU temperature                     76 degrees C / 168 degrees F
  DRAM                                17152 MB
  Memory utilization                  11 percent
  CPU utilization:
    User                             0 percent
    Background                       0 percent
    Kernel                           4 percent
    Interrupt                         0 percent
    Idle                             95 percent
  Model                              RE-DUO-2600
  Serial ID                          P737A-002231
  Start time                         2011-12-21 16:54:37 PST
  Uptime                             25 minutes, 44 seconds
  Last reboot reason                  Router rebooted after a normal shutdown.
  Load averages:                    1 minute   5 minute   15 minute
                                      0.01       0.02       0.06

Routing Engine status:
Slot 1:

```

```

Current state           Backup
Election priority       Backup (default)
Temperature             50 degrees C / 122 degrees F
CPU temperature         64 degrees C / 147 degrees F
DRAM                   17152 MB
Memory utilization      10 percent
CPU utilization:
  User                  0 percent
  Background            0 percent
  Kernel                0 percent
  Interrupt             0 percent
  Idle                  99 percent
Model                  RE-DU0-2600
Serial ID               P737A-002438
Start time              2011-12-21 16:52:26 PST
Uptime                  27 minutes, 49 seconds
Last reboot reason      Router rebooted after a normal shutdown.

```

### show chassis routing-engine (EX9200 Switch)

```

user@switch> show chassis routing-engine
Routing Engine status:
Slot 0:
  Current state           Master
  Election priority       Master (default)
  Temperature             35 degrees C / 95 degrees F
  CPU temperature         33 degrees C / 91 degrees F
  DRAM                   8157 MB
  Installed Memory        8192 MB
  Memory utilization      18 percent
CPU utilization:
  User                    1 percent
  Background              0 percent
  Kernel                  4 percent
  Interrupt               1 percent
  Idle                    94 percent
Model                    RE-S-EX9200-1800X4
Serial ID                 9009119555
Start time                2014-03-12 14:58:05 UTC
Uptime                    1 hour, 41 minutes, 51 seconds
Last reboot reason        Router rebooted after a normal shutdown.
Load averages:            1 minute  5 minute  15 minute
                        0.02      0.02      0.00

Routing Engine status:
Slot 1:
  Current state           Backup
  Election priority       Backup (default)

```

[...Output truncated...]

### show chassis routing-engine (ACX2000 Universal Access Router)

```

user@host> show chassis routing-engine
Routing Engine status:
  Temperature             53 degrees C / 127 degrees F
  DRAM                   1536 MB
  Memory utilization      25 percent
CPU utilization:
  User                    0 percent
  Background              0 percent
  Kernel                  0 percent

```

Interrupt	1 percent
Idle	99 percent
Model	RE-ACX-2000
Start time	2012-05-09 00:57:07 PDT
Uptime	5 days, 3 hours, 16 minutes, 15 seconds
Last reboot reason	Router rebooted after a normal shutdown.
Load averages:	1 minute 5 minute 15 minute
	0.00 0.03 0.05

#### show chassis routing-engine (ACX1000 Universal Access Router)

```
user@host> show chassis routing-engine
Routing Engine status:
  Temperature          36 degrees C / 96 degrees F
  DRAM                 768 MB
  Memory utilization    50 percent
  CPU utilization:
    User               3 percent
    Background         0 percent
    Kernel             6 percent
    Interrupt          0 percent
    Idle               91 percent
  Model                RE-ACX-1000
  Start time           2012-05-10 07:12:23 PDT
  Uptime               4 days, 10 hours, 46 minutes, 53 seconds
  Last reboot reason    Router rebooted after a normal shutdown.
  Load averages:      1 minute 5 minute 15 minute
                      0.00 0.00 0.00
```



## show chassis zones

<b>List of Syntax</b>	<a href="#">Syntax on page 927</a> <a href="#">Syntax (QFX Series) on page 927</a>
<b>Syntax</b>	show chassis zones <detail>
<b>Syntax (QFX Series)</b>	show chassis zones <detail> <interconnect-device <i>name</i> >
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.3 for MX2020 3D Universal Edge Routers.</p> <p>Command introduced in Junos OS Release 12.3 for MX2010 3D Universal Edge Routers.</p>
<b>Description</b>	<p>(QFabric systems only) Display the status of the two cooling system zones on the Interconnect device. Zone 1 consists of eight (0 – 7) front cards, which are cooled by two fan trays. Zone 2 consists of two control boards and eight rear cards, which are cooled by eight (0 – 7) fan trays. On MX2010 and MX2020 routers, display the status of the cooling system zones of the chassis. Zone 0 consists of the Control Board, ten (0–9) FPCs, and their respective PICs, Switch Fabric Boards, and Adapter Cards. Zone 1 consists of the Routing Engine, Control Board, and Switch Processor Mezzanine Boards.</p>
<b>Options</b>	<p><b>detail</b>—(MX2010 and MX2020 routers only) (Optional) Display detailed status of the cooling system zones.</p> <p><b>detail <i>device-name</i></b>— (QFabric systems only) (Optional) Display detailed status of the two cooling systems on the Interconnect device.</p> <p><b>interconnect-device <i>name</i></b>— (QFabric systems only) (Optional) Display the status of the cooling zones on the Interconnect device.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request chassis beacon on page 382</a></li> <li>• <a href="#">show chassis fan on page 615</a></li> <li>• <a href="#">show chassis temperature-thresholds</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show chassis zones interconnect-device (QFabric System) on page 928</a> <a href="#">show chassis zones (MX2010 Router) on page 928</a> <a href="#">show chassis zones detail (MX2010 Router) on page 929</a> <a href="#">show chassis zones (MX2020 Router) on page 930</a> <a href="#">show chassis zones detail (MX2020 Router) on page 930</a> <a href="#">show chassis beacon interconnect-device (QFabric System) on page 931</a> <a href="#">show chassis beacon interconnect-device fpc (QFabric System) on page 932</a> <a href="#">show chassis beacon node-device (QFabric System) on page 932</a> <a href="#">show chassis beacon node-device fpc (QFabric System) on page 932</a>

**Output Fields** Table 26 on page 509 lists the output fields for the **show chassis zones** command. Output fields are listed in the approximate order in which they appear.

**Table 44: show chassis zones Output Fields**

Field Name	Field Description
Slot	FPC slot number of the device whose content is being displayed. On QFX3500 standalone switches, the number is always 0.
Beacon State	Status of the beacon state: <ul style="list-style-type: none"> <li>Off—The beacon is <b>OFF</b>.</li> <li>On—The beacon is <b>ON</b>.</li> </ul>
show chassis zones command output fields for MX2020 and MX2010 routers:	
Driving FRU	Field replaceable unit (FRU).
Temperature	Temperature of the specified FRU in degrees Celsius and degrees Fahrenheit.
Condition	Condition of the specified FRU. Condition can be <b>HIGH TEMP</b> , <b>WARM TEMP</b> , <b>OK</b> , and <b>Offline</b> .
Num Fans Missing	Number of fans or fan trays missing.
Num Fans Failed	Number of fans or fan trays that have failed.
Fan Duty Cycle	Fan duty cycle value.
show chassis zones detail command output fields for MX2020 and MX2010 routers:	
Item	Chassis component: <ul style="list-style-type: none"> <li>Information about the chassis, Routing Engines, Control Boards (CBs), Switch Fabric Boards (SFBs), PICs, Flexible PIC Concentrators (FPCs), and Adapter Cards (ADCs).</li> </ul>
Measurement	Fan tray speed utilization in percentage.
Status	Status of the specified item. Status can be <b>OK</b> , <b>Absent</b> , or <b>Offline</b> .

## Sample Output

### show chassis zones interconnect-device (QFabric System)

```

user@switch> show chassis zones interconnect-device interconnect1
Slot          Beacon State
FPC           0          OFF

```

### show chassis zones (MX2010 Router)

```

user@host> show chassis zones

```

```

ZONE 0 Status
  Driving FRU          FPC 6
  Temperature          81 degrees C / 177 degrees F
  Condition            HIGH TEMP
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

ZONE 1 Status
  Driving FRU          SFB 0 Exhaust-Zone1
  Temperature          71 degrees C / 159 degrees F
  Condition            WARM TEMP
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

```

#### show chassis zones detail (MX2010 Router)

```

user@host > show chassis zones
ZONE 0 Status
Item              Status              Measurement
CB 0              WARM TEMP
CB 1              WARM TEMP
FPC 0             HIGH TEMP
FPC 1             HIGH TEMP
FPC 2             WARM TEMP
FPC 3             HIGH TEMP
FPC 4             HIGH TEMP
FPC 5             HIGH TEMP
FPC 6             HIGH TEMP
FPC 7             HIGH TEMP
FPC 8             HIGH TEMP
FPC 9             HIGH TEMP
ADC 0             WARM TEMP
ADC 1             WARM TEMP
ADC 2             WARM TEMP
ADC 3             WARM TEMP
ADC 4             WARM TEMP
ADC 5             WARM TEMP
ADC 6             WARM TEMP
ADC 7             WARM TEMP
ADC 8             WARM TEMP
ADC 9             WARM TEMP
SFB 0             WARM TEMP
SFB 1             WARM TEMP
SFB 2             WARM TEMP
SFB 3             Offline
SFB 4             HIGH TEMP
SFB 5             WARM TEMP
SFB 6             HIGH TEMP
SFB 7             WARM TEMP
Fan Tray 0        OK                  Spinning at 98% fan tray speed
Fan Tray 1        OK                  Spinning at 98% fan tray speed

ZONE 1 Status
Item              Status              Measurement
CB 0              WARM TEMP
CB 1              WARM TEMP
Routing Engine 0  OK
Routing Engine 1  OK
SFB 0             WARM TEMP

```

SFB 1	WARM TEMP	
SFB 2	WARM TEMP	
SFB 3	Offline	
SFB 4	HIGH TEMP	
SFB 5	WARM TEMP	
SFB 6	HIGH TEMP	
SFB 7	WARM TEMP	
SPMB 0	OK	
SPMB 1	OK	
Fan Tray 2	OK	Spinning at 64% fan tray speed
Fan Tray 3	OK	Spinning at 64% fan tray speed

#### show chassis zones (MX2020 Router)

```
user@host> show chassis zones
ZONE 0 Status
  Driving FRU          FPC 0
  Temperature          31 degrees C / 87 degrees F
  Condition            OK
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30

ZONE 1 Status
  Driving FRU          FPC 19
  Temperature          32 degrees C / 89 degrees F
  Condition            OK
  Num Fans Missing     0
  Num Fans Failed      0
  Fan Duty Cycle       30
```

#### show chassis zones detail (MX2020 Router)

```
user@host> show chassis zones detail
ZONE 0 Status
Item              Status          Measurement
CB 0              OK
CB 1              OK
FPC 0             OK
FPC 1             OK
FPC 2             OK
FPC 3             OK
FPC 4             OK
FPC 5             OK
FPC 6             OK
FPC 7             OK
FPC 8             OK
FPC 9             OK
ADC 0             OK
ADC 1             OK
ADC 2             OK
ADC 3             OK
ADC 4             OK
ADC 5             OK
ADC 6             OK
ADC 7             OK
ADC 8             OK
ADC 9             OK
SFB 0             OK
SFB 1             OK
SFB 2             OK
```

SFB 3	OK	
SFB 4	OK	
SFB 5	OK	
SFB 6	OK	
SFB 7	OK	
Fan Tray 0	OK	Spinning at 38% fan tray speed
Fan Tray 1	OK	Spinning at 37% fan tray speed

## ZONE 1 Status

Item	Status	Measurement
CB 0	OK	
CB 1	OK	
Routing Engine 0	OK	
Routing Engine 1	OK	
FPC 10	OK	
FPC 11	OK	
FPC 12	OK	
FPC 13	OK	
FPC 14	OK	
FPC 15	OK	
FPC 16	OK	
FPC 17	OK	
FPC 18	OK	
FPC 19	OK	
ADC 10	OK	
ADC 11	OK	
ADC 12	OK	
ADC 13	OK	
ADC 14	OK	
ADC 15	OK	
ADC 16	OK	
ADC 17	OK	
ADC 18	OK	
ADC 19	OK	
SFB 0	OK	
SFB 1	OK	
SFB 2	OK	
SFB 3	OK	
SFB 4	OK	
SFB 5	OK	
SFB 6	OK	
SFB 7	OK	
SPMB 0	OK	
SPMB 1	OK	
Fan Tray 2	OK	Spinning at 38% fan tray speed
Fan Tray 3	OK	Spinning at 38% fan tray speed

## show chassis beacon interconnect-device (QFabric System)

```

user@switch> show chassis beacon interconnect-device interconnect1
Chassis          OFF
CB 0             OFF
CB 1             OFF
FC 0 FPC 0       OFF
FC 1 FPC 1       OFF
RC 0 FPC 8       OFF
RC 1 FPC 9       OFF

```

#### show chassis beacon interconnect-device fpc (QFabric System)

```
user@switch> show chassis beacon interconnect-device interconnect1 fpc 0
FPC 0                                ON
```

#### show chassis beacon node-device (QFabric System)

```
user@switch> show chassis beacon node-device node1
node1                                ON
```

#### show chassis beacon node-device fpc (QFabric System)

```
user@switch> show chassis beacon node-device node1 fpc 0
FPC 0                                ON
```

## show cli

<b>List of Syntax</b>	<a href="#">Syntax on page 933</a> <a href="#">Syntax (QFX Series) on page 933</a>
<b>Syntax</b>	show cli
<b>Syntax (QFX Series)</b>	show cli <authorization> <directory> <history <i>count</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display configured CLI settings.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show cli on page 934</a>
<b>Output Fields</b>	<a href="#">Table 45 on page 933</a> lists the output fields for the <b>show cli</b> command. Output fields are listed in the approximate order in which they appear.

**Table 45: show cli Output Fields**

Field Name	Field Description
CLI complete-on-space	Capability to complete a partial command entry when you type a space or a tab: <b>on</b> or <b>off</b> .
CLI idle-timeout	Maximum time that an individual session can be idle before the user is logged out from the router or switch. When this feature is enabled, the number of minutes is displayed. Otherwise, the state is <b>disabled</b> .
CLI restart-on-upgrade	CLI is set to prompt you to restart the router or switch after upgrading the software: <b>on</b> or <b>off</b> .
CLI screen-length	Number of lines of text that the terminal screen displays.
CLI screen-width	Number of characters in a line on the terminal screen.
CLI terminal	Terminal type.
CLI is operating in	Mode: <b>enhanced</b> .
CLI timestamp	Date and time format for the timestamp. If the timestamp is not set, the state is <b>disabled</b> .
CLI working directory	Pathname of the working directory.

## Sample Output

show cli

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 47
CLI screen-width set to 132
CLI terminal is 'vt100'
CLI is operating in enhanced mode
CLI timestamp disabled
CLI working directory is '/var/home/regress'
```



## show cli authorization

<b>Syntax</b>	show cli authorization
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the permissions for the current user.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show cli authorization on page 937</a>
<b>Output Fields</b>	<a href="#">Table 46 on page 935</a> lists the output fields for the <b>show cli authorization</b> command. In the table, all possible permissions are displayed and output fields are listed in alphabetical order.

**Table 46: show cli authorization Output Fields**

Field Name	Field Description
access	Can view access configuration information.
access-control	Can modify access configuration.
admin	Can view user account information.
admin-control	Can modify user account information.
clear	Can clear learned network information.
configure	Can enter configuration mode.
control	Can modify any configuration.
edit	Can edit configuration files.
field	Reserved for field (debugging) support.
firewall	Can view firewall configuration information.
firewall-control	Can modify firewall configuration information.
floppy	Can read from and write to removable media.
flow-tap	Can view flow-tap configuration information.

Table 46: show cli authorization Output Fields (*continued*)

Field Name	Field Description
<b>flow-tap-control</b>	Can configure flow-tap configuration information.
<b>idp-profiler-operation</b>	Can configure Profiler data.
<b>interface</b>	Can view interface configuration information.
<b>interface-control</b>	Can modify interface configuration information.
<b>maintenance</b>	Can perform system maintenance.
<b>network</b>	Can access the network by entering the <b>ping</b> , <b>ssh</b> , <b>telnet</b> , and <b>traceroute</b> commands.
<b>pgcp-session-mirroring</b>	Can view Packet Gateway Control Protocol session mirroring configuration.
<b>pgcp-session-mirroring-control</b>	Can modify Packet Gateway Control Protocol session mirroring configuration all-control.
<b>reset</b>	Can reset or restart interfaces and system processes.
<b>rollback</b>	Can roll back to previous configurations.
<b>routing</b>	Can view routing configuration information.
<b>routing-control</b>	Can modify routing configuration information.
<b>secret</b>	Can view passwords and authentication keys in the configuration.
<b>secret-control</b>	Can modify passwords and authentication keys in the configuration.
<b>security</b>	Can view security configuration information.
<b>security-control</b>	Can modify security configuration information.
<b>shell</b>	Can start a local shell.
<b>snmp</b>	Can view SNMP configuration information.
<b>snmp-control</b>	Can modify SNMP configuration information.
<b>system</b>	Can view system configuration information.
<b>system-control</b>	Can modify system configuration information.
<b>trace</b>	Can view trace file settings information.

Table 46: show cli authorization Output Fields (*continued*)

Field Name	Field Description
<b>trace-control</b>	Can modify trace file settings information.
<b>view</b>	Can view current values and statistics.
<b>view-configuration</b>	Can view all configuration information (not including secrets).

## Sample Output

### show cli authorization

```

user@host> show cli authorization
Current user: 'remote' login: 'user' class ''
Permissions:
  admin      -- Can view user accounts
  admin-control-- Can modify user accounts
  clear      -- Can clear learned network information
  configure  -- Can enter configuration mode
  control    -- Can modify any configuration
  edit       -- Can edit full files
  field      -- Special for field (debug) support
  floppy     -- Can read and write from the floppy
  interface  -- Can view interface configuration
  interface-control-- Can modify interface configuration
  network    -- Can access the network
  reset      -- Can reset/restart interfaces and daemons
  routing    -- Can view routing configuration
  routing-control-- Can modify routing configuration
  shell      -- Can start a local shell
  snmp       -- Can view SNMP configuration
  snmp-control-- Can modify SNMP configuration
  system     -- Can view system configuration
  system-control-- Can modify system configuration
  trace      -- Can view trace file settings
  trace-control-- Can modify trace file settings
  view       -- Can view current values and statistics
  maintenance -- Can become the super-user
  firewall   -- Can view firewall configuration
  firewall-control-- Can modify firewall configuration
  secret     -- Can view secret configuration
  secret-control-- Can modify secret configuration
  rollback   -- Can rollback to previous configurations
  security   -- Can view security configuration
  security-control-- Can modify security configuration
  access     -- Can view access configuration
  access-control-- Can modify access configuration
  view-configuration-- Can view all configuration (not including secrets)
  flow-tap   -- Can view flow-tap configuration
  flow-tap-control-- Can configure flow-tap service
Individual command authorization:
  Allow regular expression: none
  Deny regular expression: none
  Allow configuration regular expression: none
  Deny configuration regular expression: none

```



---

## show cli directory

---

<b>Syntax</b>	show cli directory
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the current working directory.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show cli directory on page 939</a>
<b>Output Fields</b>	<a href="#">Table 47 on page 939</a> lists the output fields for the <b>show cli directory</b> command. Output fields are listed in the approximate order in which they appear.

**Table 47: show cli directory Output Fields**

Field Name	Field Description
Current directory	Pathname of the current working directory.

## Sample Output

### show cli directory

```
user@host> show cli directory
Current directory: /var/home/regress
```

## show cli history

---

<b>Syntax</b>	<code>show cli history</code> <code>&lt;count&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a list of previous CLI commands.
<b>Options</b>	<b>none</b> —Display all previous CLI commands.  <b>count</b> —(Optional) Maximum number of commands to display.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show cli history on page 940</a>
<b>Output Fields</b>	<a href="#">Table 48 on page 940</a> lists the output fields for the <b>show cli history</b> command. Output fields are listed in the approximate order in which they appear.

**Table 48: show cli history Output Fields**

Field Name	Field Description
<i>timestamp</i>	Time at which the command was entered.
<i>command-syntax</i>	Command that was entered.

## Sample Output

### show cli history

```
user@host> show cli history
11:14:14 -- show arp
11:22:10 -- show cli authorization
11:27:12 -- show cli history
```

---

## show host

---

<b>Syntax</b>	<code>show host <i>hostname</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Domain Name System (DNS) hostname information.
<b>Options</b>	<i>hostname</i> —Hostname or address.
<b>Additional Information</b>	The <code>show host</code> command displays the raw data received from the DNS server.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show host on page 941</a>

### Sample Output

#### show host

```
user@host> show host snark
snark.boojum.net has address 192.168.1.254

user@host> show host 192.168.1.254
Name: snark.boojum.net
Address: 192.168.1.254
Aliases:
```

## show interfaces diagnostics optics

<b>Syntax</b>	<code>show interfaces diagnostics optics <i>interface-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Display diagnostics data and alarms for Gigabit Ethernet, 10-Gigabit Ethernet, and QSFP+ optical transceivers installed in a QFX Series product. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p>
<b>Options</b>	<i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 335</a></li> <li>• <a href="#">Installing a Transceiver in a QFX Series Device</a></li> <li>• <a href="#">Removing a Transceiver from a QFX Series Device</a></li> <li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver) on page 946</a> <a href="#">show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver) on page 947</a>
<b>Output Fields</b>	lists the output fields for the <code>show interfaces diagnostics optics</code> command. Output fields are listed in the approximate order in which they appear.

**Table 49: show interfaces diagnostics optics Output Fields**

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in volts.
(Not available for XFP transceivers)	



Table 49: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser rx power</b> (Not available for SFP and SFP+ transceivers)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
<b>Receiver signal average optical power</b> (Not available for XFP transceivers)	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
<b>Laser bias current high alarm</b>	Displays whether the laser bias power setting high alarm is <b>On</b> or <b>Off</b> .
<b>Laser bias current low alarm</b>	Displays whether the laser bias power setting low alarm is <b>On</b> or <b>Off</b> .
<b>Laser bias current high warning</b>	Displays whether the laser bias power setting high warning is <b>On</b> or <b>Off</b> .
<b>Laser bias current low warning</b>	Displays whether the laser bias power setting low warning is <b>On</b> or <b>Off</b> .
<b>Laser output power high alarm</b>	Displays whether the laser output power high alarm is <b>On</b> or <b>Off</b> .
<b>Laser output power low alarm</b>	Displays whether the laser output power low alarm is <b>On</b> or <b>Off</b> .
<b>Laser output power high warning</b>	Displays whether the laser output power high warning is <b>On</b> or <b>Off</b> .
<b>Laser output power low warning</b>	Displays whether the laser output power low warning is <b>On</b> or <b>Off</b> .
<b>Module temperature high alarm</b>	Displays whether the module temperature high alarm is <b>On</b> or <b>Off</b> .
<b>Module temperature low alarm</b>	Displays whether the module temperature low alarm is <b>On</b> or <b>Off</b> .
<b>Module temperature high warning</b>	Displays whether the module temperature high warning is <b>On</b> or <b>Off</b> .
<b>Module temperature low warning</b>	Displays whether the module temperature low warning is <b>On</b> or <b>Off</b> .
<b>Module voltage high alarm</b> (Not available for XFP transceivers)	Displays whether the module voltage high alarm is <b>On</b> or <b>Off</b> .
<b>Module voltage low alarm</b> (Not available for XFP transceivers)	Displays whether the module voltage low alarm is <b>On</b> or <b>Off</b> .
<b>Module voltage high warning</b> (Not available for XFP transceivers)	Displays whether the module voltage high warning is <b>On</b> or <b>Off</b> .
<b>Module voltage low warning</b> (Not available for XFP transceivers)	Displays whether the module voltage low warning is <b>On</b> or <b>Off</b> .
<b>Laser rx power high alarm</b>	Displays whether the receive laser power high alarm is <b>On</b> or <b>Off</b> .

Table 49: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power low alarm	Displays whether the receive laser power low alarm is <b>On</b> or <b>Off</b> .
Laser rx power high warning	Displays whether the receive laser power high warning is <b>On</b> or <b>Off</b> .
Laser rx power low warning	Displays whether the receive laser power low warning is <b>On</b> or <b>Off</b> .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Module not ready alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module not ready alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module has an operational fault.
Module power down alarm (Not available for SFP and SFP+ transceivers)	Displays whether the module power down alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module is in a limited power mode, low for normal operation.
Tx data not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is <b>On</b> or <b>Off</b> .
Tx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is <b>On</b> or <b>Off</b> .
Tx laser fault alarm (Not available for SFP and SFP+ transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is <b>On</b> or <b>Off</b> .
Tx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .
Rx not ready alarm (Not available for SFP and SFP+ transceivers)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is <b>On</b> or <b>Off</b> .
Rx loss of signal alarm (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When <b>on</b> , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is <b>On</b> or <b>Off</b> .
Rx CDR loss of lock alarm (Not available for SFP and SFP+ transceivers)	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.

Table 49: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold (Not available for XFP transceivers)	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser Rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser Rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser Rx power high warning.

Table 49: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser Rx power low warning.

## Sample Output

### show interfaces diagnostics optics xe-0/0/1 (SFP+ Transceiver)

```

user@host> show interfaces diagnostics optics xe-0/0/1
Physical interface: xe-0/0/1
  Laser bias current          : 4.968 mA
  Laser output power         : 0.4940 mW / -3.06 dBm
  Module temperature         : 27 degrees C / 81 degrees F
  Module voltage             : 3.2310 V
  Receiver signal average optical power : 0.0000
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm     : Off
  Module voltage low alarm      : Off
  Module voltage high warning   : Off
  Module voltage low warning    : Off
  Laser rx power high alarm     : Off
  Laser rx power low alarm      : On
  Laser rx power high warning   : Off
  Laser rx power low warning    : On
  Laser bias current high alarm threshold : 10.500 mA
  Laser bias current low alarm threshold  : 2.000 mA
  Laser bias current high warning threshold : 9.000 mA
  Laser bias current low warning threshold : 2.500 mA
  Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
  Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
  Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
  Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold     : 3.630 V
  Module voltage low alarm threshold      : 2.970 V
  Module voltage high warning threshold   : 3.465 V
  Module voltage low warning threshold    : 3.135 V
  Laser rx power high alarm threshold     : 1.5849 mW / 2.00 dBm
  Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
  Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
  Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

**show interfaces diagnostics optics node1:xe-0/0/1 (SFP+ Transceiver)**

```

user@host> show interfaces diagnostics optics node1:xe-0/0/1
Physical interface: node1:xe-0/0/1
  Laser bias current                : 4.968 mA
  Laser output power                : 0.4940 mW / -3.06 dBm
  Module temperature                : 27 degrees C / 81 degrees F
  Module voltage                    : 3.2310 V
  Receiver signal average optical power : 0.0000
  Laser bias current high alarm      : Off
  Laser bias current low alarm       : Off
  Laser bias current high warning    : Off
  Laser bias current low warning     : Off
  Laser output power high alarm      : Off
  Laser output power low alarm       : Off
  Laser output power high warning    : Off
  Laser output power low warning     : Off
  Module temperature high alarm      : Off
  Module temperature low alarm       : Off
  Module temperature high warning    : Off
  Module temperature low warning     : Off
  Module voltage high alarm          : Off
  Module voltage low alarm           : Off
  Module voltage high warning        : Off
  Module voltage low warning         : Off
  Laser rx power high alarm          : Off
  Laser rx power low alarm           : On
  Laser rx power high warning        : Off
  Laser rx power low warning         : On
  Laser bias current high alarm threshold : 10.500 mA
  Laser bias current low alarm threshold : 2.000 mA
  Laser bias current high warning threshold : 9.000 mA
  Laser bias current low warning threshold : 2.500 mA
  Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
  Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
  Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
  Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold : 3.630 V
  Module voltage low alarm threshold : 2.970 V
  Module voltage high warning threshold : 3.465 V
  Module voltage low warning threshold : 3.135 V
  Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
  Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
  Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
  Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

## show log

---

<b>List of Syntax</b>	<a href="#">Syntax on page 948</a> <a href="#">Syntax (QFabric System) on page 948</a> <a href="#">Syntax (TX Matrix Routers) on page 948</a>
<b>Syntax</b>	<code>show log</code> <code>&lt;filename   user &lt;username&gt;&gt;</code>
<b>Syntax (QFabric System)</b>	<code>show log filename</code> <code>&lt;device-type (device-id   device-alias)&gt;</code>
<b>Syntax (TX Matrix Routers)</b>	<code>show log</code> <code>&lt;all-lcc   lcc number   scc&gt;</code> <code>&lt;filename   user &lt;username&gt;&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <i>device-type (device-id   device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.
<b>Description</b>	List log files, display log file contents, or display information about users who have logged in to the router or switch.
<b>Options</b>	<b>none</b> —List all log files.  <b>&lt;all-lcc   lcc number   scc&gt;</b> —(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).  <b>device-type</b> —(QFabric system only) (Optional) Display log messages for only one of the following device types: <ul style="list-style-type: none"><li>• <b>director-device</b>—Display logs for Director devices.</li><li>• <b>infrastructure-device</b>—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).</li><li>• <b>interconnect-device</b>—Display logs for Interconnect devices.</li><li>• <b>node-device</b>—Display logs for Node devices.</li></ul>



**NOTE:** If you specify the *device-type* optional parameter, you must also specify either the *device-id* or *device-alias* optional parameter.

---

**(device-id | device-alias)**—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

**filename**—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



**NOTE:** The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

**user <username>**—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

**Required Privilege Level** trace

**List of Sample Output** [show log on page 949](#)  
[show log filename on page 949](#)  
[show log filename \(QFabric System\) on page 950](#)  
[show log user on page 950](#)

## Sample Output

### show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin        19656 Oct  1 19:37 wtmp
```

### show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

### show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

### show log user

```

user@host> show log user
darius  mg2546                Thu Oct  1 19:37   still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```



## show ntp associations

<b>Syntax</b>	<code>show ntp associations</code> <code>&lt;no-resolve&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Network Time Protocol (NTP) peers and their state.
<b>Options</b>	<b>none</b> —Display NTP peers and their state.  <b>no-resolve</b> —(Optional) Suppress symbolic addressing.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ntp status on page 953</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ntp associations on page 952</a>
<b>Output Fields</b>	<a href="#">Table 50 on page 951</a> describes the output fields for the <b>show ntp associations</b> command. Output fields are listed in the approximate order in which they appear.

**Table 50: show ntp associations Output Fields**

Field Name	Field Description
<b>remote</b>	Address or name of the remote NTP peer.
<b>refid</b>	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of <b>0.0.0.0</b> .
<b>st</b>	Stratum of the remote peer.
<b>t</b>	Type of peer: <b>b</b> (broadcast), <b>l</b> (local), <b>m</b> (multicast), or <b>u</b> (unicast).
<b>when</b>	When the last packet from the peer was received.
<b>poll</b>	Polling interval, in seconds.
<b>reach</b>	Reachability register, in octal.
<b>delay</b>	Current estimated delay of the peer, in milliseconds.
<b>offset</b>	Current estimated offset of the peer, in milliseconds.
<b>disp</b>	Current estimated dispersion of the peer, in milliseconds.

Table 50: show ntp associations Output Fields (*continued*)

Field Name	Field Description
<i>peer-name</i>	<p>Peer name and status of the peer in the clock selection process:</p> <ul style="list-style-type: none"> <li>• space—Discarded because of a high stratum value or failed sanity checks.</li> <li>• x—Designated "falseticker" by the intersection algorithm.</li> <li>• .—Culled from the end of the candidate list.</li> <li>• — —Discarded by the clustering algorithm.</li> <li>• +—Included in the final selection set.</li> <li>• #—Selected for synchronization, but the distance exceeds the maximum.</li> <li>• *—Selected for synchronization.</li> <li>• o—Selected for synchronization, but the packets-per-second (pps) signal is in use.</li> </ul>

## Sample Output

### show ntp associations

```

user@host> show ntp associations
      remote      refid      st t when poll reach  delay  offset  disp
=====
*wolfe-gw.junipe tick.ucla.edu  2 u  43  64  377   1.86   0.319   0.08

```

## show ntp status

<b>Syntax</b>	<code>show ntp status</code> <code>&lt;no-resolve&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the values of internal variables returned by Network Time Protocol (NTP) peers.
<b>Options</b>	<b>none</b> —Display the values of internal variables returned by NTP peers.  <b>no-resolve</b> —(Optional) Suppress symbolic addressing.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ntp associations on page 951</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ntp status on page 954</a>
<b>Output Fields</b>	<a href="#">Table 51 on page 953</a> describes the output fields for the <b>show ntp status</b> command. Output fields are listed in the approximate order in which they appear.

**Table 51: show ntp status Output Fields**

Field Name	Field Description
<b>status</b>	System status word, a code representing the status items listed.
<b>leap_none</b>	Indicates a normal synchronized state with no leap seconds imminent. Other options could be <b>leap_add_sec</b> , <b>leap_del_sec</b> , or <b>leap_alarm</b> , indicating a leap second will be added, deleted, or a leap second requirement is upcoming.
<b>sync_ntp</b>	Indicates the current synchronization source, in this case, an NTP server. Other options include <b>sync_alarm</b> and <b>sync_unspec</b> , both indicating that the router has not been synched.
<b>x events</b>	Indicates the number of events that have occurred since that last code change. An event is often the receipt of an NTP polling message.
<b>event_peer/strat_chg</b>	Describes the most recent event, in this case, the stratum of the peer server changed.
<b>version</b>	A detailed description of the version of NTP being used.
<b>processor</b>	Indicates the current hardware platform and version of the processor.
<b>system</b>	Detailed description of the name and version of the operating system in use.
<b>leap</b>	The number of leap seconds in use.

Table 51: show ntp status Output Fields (*continued*)

Field Name	Field Description
<b>stratum</b>	The stratum of the peer server. Anything greater than 1 is a secondary reference source, and the number roughly represents the number of hops away from the stratum 1 server.. Stratum 1 is a primary reference, such as an atomic clock.
<b>precision</b>	The precision of the peer clock, how precisely the frequency and time can be maintained with this particular timekeeping system.
<b>rootdelay</b>	The total roundtrip delay to the primary reference source, in seconds.
<b>rootdispersion</b>	The maximum error relative to the primary reference source, in seconds.
<b>peer</b>	An identification number of the peer in use.
<b>refid</b>	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0.
<b>reftime</b>	The local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
<b>poll</b>	The NTP broadcast message polling interval, in seconds.
<b>clock</b>	The current time on the local router clock.
<b>state</b>	The current mode of NTP operation, where 1 is symmetric active, 2 is symmetric passive, 3 is client, 4 is server, and 5 is broadcast.
<b>offset</b>	Current estimated offset of the peer, in milliseconds. Indicates the time difference between the reference clock and the local clock.
<b>frequency</b>	The frequency of the clock.
<b>jitter</b>	Indicates the magnitude of jitter, in milliseconds, between several time queries.
<b>stability</b>	A measure of how well this clock can maintain a constant frequency.

## Sample Output

### show ntp status

```

user@host> show ntp status
assID=0 status=0544 leap_none, sync_local_proto, 4 events, event_peer/strat_chg,
version="ntpd 4.2.2p1@1.1570-o Tue May 19 13:57:55 UTC 2009 (1)",
processor="x86_64", system="Linux/2.6.18-164.el5", leap=00, stratum=4,
precision=-10, rootdelay=0.000, rootdispersion=11.974, peer=59475,
refid=LOCAL(0),
reftime=d495c32c.0e71eaf2 Mon, Jan 7 2013 13:57:00.056, poll=10,
clock=d495c32c.cebd43bd Mon, Jan 7 2013 13:57:00.807, state=4,
offset=0.000, frequency=0.000, jitter=0.977, noise=0.977,
stability=0.000, tai=0

```



## show subscribers

---

**Syntax**    `show subscribers`  
              `<detail | extensive | terse>`  
              `<aci-interface-set-name aci-interface-set-name>`  
              `<address address>`  
              `<agent-circuit-identifier agent-circuit-identifier-substring>`  
              `<client-type client-type>`  
              `<count>`  
              `<id>`  
              `<interface interface>`  
              `<logical-system logical-system>`  
              `<mac-address mac-address>`  
              `<physical-interface physical-interface-name>`  
              `<profile-name profile-name>`  
              `<routing-instance routing-instance>`  
              `<stacked-vlan-id stacked-vlan-id>`  
              `<subscriber-state subscriber-state>`  
              `<user-name user-name>`  
              `<vci vci-identifier>`  
              `<vpi vpi-identifier>`  
              `<vlan-id vlan-id>`

**Release Information**    Command introduced in Junos OS Release 9.3.  
                              Command introduced in Junos OS Release 9.3 for EX Series switches.  
                              **client-type**, **mac-address**, **subscriber-state**, and **extensive** options introduced in Junos OS Release 10.2.  
                              **count** option usage with other options introduced in Junos OS Release 10.2.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              Options **aci-interface-set-name** and **agent-circuit-identifier** introduced in Junos OS Release 12.2.  
                              The **physical-interface** and **user-name** options introduced in Junos OS Release 12.3.  
                              Options **vci** and **vpi** introduced in Junos OS Release 12.3R3 and supported in later 12.3Rx releases.  
                              Options **vci** and **vpi** supported in Junos OS Release 13.2 and later releases. (Not supported in Junos OS Release 13.1.)

**Description**    Display information for active subscribers.

**Options**    **detail | extensive | terse**—(Optional) Display the specified level of output.

**aci-interface-set-name**—(Optional) Display all dynamic subscriber sessions that use the specified agent circuit identifier (ACI) interface set. Use the ACI interface set name generated by the router, such as `aci-1003-ge-1/0/0.4001`, and not the actual ACI value found in the DHCP or PPPoE control packets.

**address**—(Optional) Display subscribers whose IP address matches the specified address. You must specify the IPv4 or IPv6 address prefix without a netmask (for example, `192.168.17.1`). If you specify the IP address as a prefix with a netmask (for example, `192.168.17.1/32`), the router displays a message that the IP address is invalid, and rejects the command.

**agent-circuit-identifier-substring**—(Optional) Display all dynamic subscriber sessions whose ACI value matches the specified substring.

**client-type**—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, VLAN, or static).

**count**—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the **count** option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, or **vlan-id** options.

**id**—(Optional) Display a specific subscriber session whose session id matches the specified subscriber ID. You can display subscriber IDs by using the **show subscribers extensive** or the **show subscribers interface extensive** commands.

**interface**—(Optional) Display subscribers whose interface matches the specified interface.

**logical-system**—(Optional) Display subscribers whose logical system matches the specified logical system.

**mac-address**—(Optional) Display subscribers whose MAC address matches the specified MAC address.

**physical-interface-name**—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose physical interface matches the specified physical interface.

**profile-name**—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

**routing-instance**—(Optional) Display subscribers whose routing instance matches the specified routing instance.

**stacked-vlan-id**—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

**subscriber-state**—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

**user-name**—(M120, M320, and MX Series routers only) (Optional) Display subscribers whose username matches the specified subscriber name.

**vci-identifier**—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual circuit identifier (VCI) matches the specified VCI identifier. The range of values is **0** through **255**.

**vpi-identifier**—(MX Series routers with MPCs and ATM MICs with SFP only) (Optional) Display active ATM subscribers whose ATM virtual path identifier (VPI) matches the specified VPI identifier. The range of values is **0** through **65535**.

**vlan-id**—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID.



**NOTE:** Due to display limitations, logical system and routing instance output values are truncated when necessary.

<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>show subscribers summary</i></li> <li>• <i>Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show subscribers (IPv4) on page 962</a> <a href="#">show subscribers (IPv6) on page 962</a> <a href="#">show subscribers (IPv4 and IPv6 Dual Stack) on page 962</a> <a href="#">show subscribers (LNS on MX Series Routers) on page 963</a> <a href="#">show subscribers (L2TP Switched Tunnels) on page 963</a> <a href="#">show subscribers client-type dhcp detail on page 963</a> <a href="#">show subscribers count on page 963</a> <a href="#">show subscribers address detail (IPv6) on page 963</a> <a href="#">show subscribers detail (IPv4) on page 964</a> <a href="#">show subscribers detail (IPv6) on page 964</a> <a href="#">show subscribers detail (IPv6 Static Demux Interface) on page 965</a> <a href="#">show subscribers detail (L2TP LNS Subscribers on MX Series Routers) on page 965</a> <a href="#">show subscribers detail (L2TP Switched Tunnels) on page 965</a> <a href="#">show subscribers detail (Tunneled Subscriber) on page 966</a> <a href="#">show subscribers detail (IPv4 and IPv6 Dual Stack) on page 966</a> <a href="#">show subscribers detail (ACI Interface Set Session) on page 967</a> <a href="#">show subscribers detail (PPPoE Subscriber Session with ACI Interface Set) on page 967</a> <a href="#">show subscribers extensive on page 967</a> <a href="#">show subscribers extensive (RPF Check Fail Filter) on page 968</a> <a href="#">show subscribers extensive (L2TP LNS Subscribers on MX Series Routers) on page 968</a> <a href="#">show subscribers extensive (IPv4 and IPv6 Dual Stack) on page 968</a> <a href="#">show subscribers extensive (Effective Shaping-Rate) on page 969</a> <a href="#">show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set) on page 970</a> <a href="#">show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring) on page 970</a> <a href="#">show subscribers interface extensive on page 971</a> <a href="#">show subscribers logical-system terse on page 971</a> <a href="#">show subscribers physical-interface count on page 972</a> <a href="#">show subscribers routing-instance inst1 count on page 972</a> <a href="#">show subscribers stacked-vlan-id detail on page 972</a> <a href="#">show subscribers stacked-vlan-id vlan-id detail (Combined Output) on page 972</a> <a href="#">show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface) on page 972</a> <a href="#">show subscribers user-name detail on page 972</a> <a href="#">show subscribers vlan-id on page 973</a>



[show subscribers vlan-id detail on page 973](#)

[show subscribers vpi vci extensive \(PPPoE-over-ATM Subscriber Session\) on page 973](#)

**Output Fields** [Table 52 on page 959](#) lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

**Table 52: show subscribers Output Fields**

Field Name	Field Description
<b>Interface</b>	Interface associated with the subscriber. The router or switch displays subscribers whose interface matches or begins with the specified interface.  The * character indicates a continuation of addresses for the same session.
<b>IP Address/VLAN ID</b>	Subscriber IP address or VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i>  No IP address or VLAN ID is assigned to an L2TP tunnel-switched session. For these subscriber sessions the value is <b>Tunnel-switched</b> .
<b>User Name</b>	Name of subscriber.
<b>LS:RI</b>	Logical system and routing instance associated with the subscriber.
<b>Type</b>	Subscriber client type (DHCP, L2TP, PPP, PPPoE, STATIC-INTERFACE, VLAN).
<b>IP Address</b>	Subscriber IPv4 address.
<b>IP Netmask</b>	Subscriber IP netmask.
<b>Primary DNS Address</b>	IP address of primary DNS server.
<b>Secondary DNS Address</b>	IP address of secondary DNS server.
<b>Primary WINS Address</b>	IP address of primary WINS server.
<b>Secondary WINS Address</b>	IP address of secondary WINS server.
<b>IPv6 Address</b>	Subscriber IPv6 address, or multiple addresses.
<b>IPv6 Prefix</b>	Subscriber IPv6 prefix. If you are using DHCPv6 prefix delegation, this is the delegated prefix.
<b>IPv6 User Prefix</b>	IPv6 prefix obtained through ND/RA.
<b>IPv6 Address Pool</b>	Subscriber IPv6 address pool. The IPv6 address pool is used to allocate IPv6 prefixes to the DHCPv6 clients.
<b>IPv6 Network Prefix Length</b>	Length of the network portion of the IPv6 address.
<b>IPv6 Prefix Length</b>	Length of the subscriber IPv6 prefix.

Table 52: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>Logical System</b>	Logical system associated with the subscriber.
<b>Routing Instance</b>	Routing instance associated with the subscriber.
<b>Interface Type</b>	Whether the subscriber interface is <b>Static</b> or <b>Dynamic</b> .
<b>Interface Set</b>	Internally generated name of the dynamic ACI interface set used by the subscriber session.
<b>Interface Set Type</b>	Interface type of the ACI interface set: <b>Dynamic</b> . This is the only ACI interface set type currently supported.
<b>Interface Set Session ID</b>	Identifier of the dynamic ACI interface set entry in the session database.
<b>Underlying Interface</b>	Name of the underlying interface for the subscriber session.
<b>Dynamic Profile Name</b>	Dynamic profile used for the subscriber.
<b>Dynamic Profile Version</b>	Version number of the dynamic profile used for the subscriber.
<b>MAC Address</b>	MAC address associated with the subscriber.
<b>State</b>	Current state of the subscriber session ( <b>Init</b> , <b>Configured</b> , <b>Active</b> , <b>Terminating</b> , <b>Tunneled</b> ).
<b>L2TP State</b>	Current state of the L2TP session, <b>Tunneled</b> or <b>Tunnel-switched</b> . When the value is <b>Tunnel-switched</b> , two entries are displayed for the subscriber; the first entry is at the LNS interface on the LTS and the second entry is at the LAC interface on the LTS.
<b>Tunnel switch Profile Name</b>	Name of the L2TP tunnel switch profile that initiates tunnel switching.
<b>Local IP Address</b>	IP address of the local gateway (LAC).
<b>Remote IP Address</b>	IP address of the remote peer (LNS).
<b>VLAN Id</b>	VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
<b>Stacked VLAN Id</b>	Stacked VLAN ID associated with the subscriber in the form <i>tpid.vlan-id</i> .
<b>RADIUS Accounting ID</b>	RADIUS accounting ID associated with the subscriber.
<b>Agent Circuit ID</b>	Option 82 agent circuit ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
<b>Agent Remote ID</b>	Option 82 agent remote ID associated with the subscriber. The ID is displayed as an ASCII string unless the value has nonprintable characters, in which case it is displayed in hexadecimal format.
<b>DHCP Relay IP Address</b>	IP address used by the DHCP relay agent.

Table 52: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>ATM VPI</b>	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual path identifier (VPI) on the subscriber's physical interface.
<b>ATM VCI</b>	(MX Series routers with MPCs and ATM MICs with SFP only) ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface.
<b>Login Time</b>	Date and time at which the subscriber logged in.
<b>Effective shaping-rate</b>	Actual downstream traffic shaping rate for the subscriber, in kilobits per second.
<b>IPv4 rpf-check Fail Filter Name</b>	Name of the filter applied by the dynamic profile to IPv4 packets that fail the RPF check.
<b>IPv6 rpf-check Fail Filter Name</b>	Name of the filter applied by the dynamic profile to IPv6 packets that fail the RPF check.
<b>DHCP Options</b>	len = number of hex values in the message. The hex values specify the type, length, value (TLV) for DHCP options, as defined in RFC 2132.
<b>Session ID</b>	ID number for a subscriber service session.
<b>Underlying Session ID</b>	For DHCPv6 subscribers on a PPPoE network, displays the session ID of the underlying PPPoE interface.
<b>Service Sessions</b>	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
<b>Service Session Name</b>	Service session profile name.
<b>Session Timeout (seconds)</b>	Number of seconds of access provided to the subscriber before the session is automatically terminated.
<b>Idle Timeout (seconds)</b>	Number of seconds subscriber can be idle before the session is automatically terminated.
<b>IPv6 Delegated Address Pool</b>	Name of the pool used for DHCPv6 prefix delegation.
<b>IPv6 Delegated Network Prefix Length</b>	Length of the prefix configured for the IPv6 delegated address pool.
<b>IPv6 Interface Address</b>	Address assigned by the Framed-Ipv6-Prefix AAA attribute.
<b>IPv6 Framed Interface Id</b>	Interface ID assigned by the Framed-Interface-Id AAA attribute.
<b>ADF IPv4 Input Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.

Table 52: show subscribers Output Fields (*continued*)

Field Name	Field Description
<b>ADF IPv4 Output Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv4 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
<b>ADF IPv6 Input Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 input filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
<b>ADF IPv6 Output Filter Name</b>	Name assigned to the Ascend-Data-Filter (ADF) interface IPv6 output filter (client or service session). The filter name is followed by the rules (in hexadecimal format) associated with the ADF filter and the decoded rule in Junos OS filter style.
<b>IPv4 Input Filter Name</b>	Name assigned to the IPv4 input filter (client or service session).
<b>IPv4 Output Filter Name</b>	Name assigned to the IPv4 output filter (client or service session).
<b>IPv6 Input Filter Name</b>	Name assigned to the IPv6 input filter (client or service session).
<b>IPv6 Output Filter Name</b>	Name assigned to the IPv6 output filter (client or service session).
<b>IFL Input Filter Name</b>	Name assigned to the logical interface input filter (client or service session).
<b>IFL Output Filter Name</b>	Name assigned to the logical interface output filter (client or service session).

## Sample Output

### show subscribers (IPv4)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/3/0.1073741824 100                 WHOLESALE-CLIENT  default:default
demux0.1073741824   100.0.0.10         RETAILER1-CLIENT  test1:retailer1
demux0.1073741825   101.0.0.3          RETAILER2-CLIENT  test1:retailer2
demux0.1073741826   102.0.0.3

```

### show subscribers (IPv6)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name          LS:RI
ge-1/0/0.0         2001::c0:0:0:0/74  WHOLESALE-CLIENT  default:default
*                  2002::1/128        subscriber-25      default:default

```

### show subscribers (IPv4 and IPv6 Dual Stack)

```

user@host> show subscribers
Interface          IP Address/VLAN ID  User Name
LS:RI
demux0.1073741834  0x8100.1002 0x8100.1
default:default
demux0.1073741835  0x8100.1001 0x8100.1
default:default
pp0.1073741836     61.1.1.1        dualstackuser1@ISP1.com

```

```

default:ASP-1
*                2041:1:1::/48
*                2061:1:1:1::/64
pp0.1073741837   23.1.1.3                dualstackuser2@ISP1.com
default:ASP-1
*                2001:1:2:5::/64

```

#### show subscribers (LNS on MX Series Routers)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-4/0/0.1     192.168.4.1         xyz@example.com default:default

```

#### show subscribers (L2TP Switched Tunnels)

```

user@host> show subscribers
Interface      IP Address/VLAN ID  User Name      LS:RI
si-2/1/0.1073741842 Tunnel-switched     ap@lts.com     default:default

si-2/1/0.1073741843 Tunnel-switched     ap@lts.com     default:default

```

#### show subscribers client-type dhcp detail

```

user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT

```

#### show subscribers count

```

user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188

```

#### show subscribers address detail (IPv6)

```

user@host> show subscribers address 100.16.12.137 detail

```

```
Type: PPPoE
User Name: pppoeTerV6User1Svc
IP Address: 100.16.12.137
IP Netmask: 255.0.0.0
IPv6 User Prefix: 1016:0:0:c88::/64
Logical System: default
Routing Instance: default
Interface: pp0.1073745151
Interface type: Dynamic
Underlying Interface: demux0.8201
Dynamic Profile Name: pppoe-client-profile
MAC Address: 00:0d:02:01:00:01
Session Timeout (seconds): 31622400
Idle Timeout (seconds): 86400
State: Active
Radius Accounting ID: jnpr demux0.8201:6544
Session ID: 6544
Agent Circuit ID: if13720
Agent Remote ID: if13720
Login Time: 2012-05-21 13:37:27 PDT
Service Sessions: 1
```

#### show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Primary DNS Address: 192.168.17.1
Secondary DNS Address: 192.168.17.2
Primary WINS Address: 192.168.22.1
Secondary WINS Address: 192.168.22.2
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Idle Timeout (seconds): 600
Login Time: 2009-08-25 14:43:52 PDT
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 08 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 36 2f
33 2d 37 2d 30 37 05 01 06 0f 21 2c
Service Sessions: 2
```

#### show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03
State: Active
Radius Accounting ID: 1
```

```

Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00

```

#### show subscribers detail (IPv6 Static Demux Interface)

```

user@host> show subscribers detail
Type: STATIC-INTERFACE
User Name: demux0.1@jnpr.net
IPv6 Prefix: 1:2:3:4:5:6:7:aa/128
Logical System: default
Routing Instance: default
Interface: demux0.1
Interface type: Static
Dynamic Profile Name: junos-default-profile
State: Active
Radius Accounting ID: 185
Login Time: 2010-05-18 14:33:56 EDT

```

#### show subscribers detail (L2TP LNS Subscribers on MX Series Routers)

```

user@host> show subscribers detail
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST

```

#### show subscribers detail (L2TP Switched Tunnels)

```

user@host> show subscribers detail
Type: L2TP
User Name: ap@example.com
Logical System: default
Routing Instance: default
Interface: si-2/1/0.1073741842
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.50.1.1
Remote IP Address: 192.168.20.3
Radius Accounting ID: 21
Session ID: 21
Login Time: 2013-01-18 03:01:11 PST

Type: L2TP
User Name: ap@example.com
Logical System: default

```

```
Routing Instance: default
Interface: si-2/1/0.1073741843
Interface type: Dynamic
Dynamic Profile Name: dyn-lts-profile
State: Active
L2TP State: Tunnel-switched
Tunnel switch Profile Name: ce-lts-profile
Local IP Address: 10.30.1.1
Remote IP Address: 172.20.1.10
Session ID: 22
Login Time: 2013-01-18 03:01:14 PST
```

#### show subscribers detail (Tunneled Subscriber)

```
user@host> show subscribers detail
Type: PPPoE
User Name: user1@example.com
Logical System: default
Routing Instance: default
Interface: pp0.1
State: Active, Tunneled
Radius Accounting ID: 512
```

#### show subscribers detail (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
```



```

State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00 00
00 00

```

#### show subscribers detail (ACI Interface Set Session)

```

user@host> show subscribers detail
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ge-1/0/0
Interface Set: aci-1001-ge-1/0/0.2800
Interface Set Session ID: 0
Underlying Interface: ge-1/0/0.2800
Dynamic Profile Name: aci-vlan-set-profile-2
Dynamic Profile Version: 1
State: Active
Session ID: 1
Agent Circuit ID: aci-ppp-dhcp-20
Login Time: 2012-05-26 01:54:08 PDT

```

#### show subscribers detail (PPPoE Subscriber Session with ACI Interface Set)

```

user@host> show subscribers detail
Type: PPPoE
User Name: ppphint2
IP Address: 10.10.1.5
Logical System: default
Routing Instance: default
Interface: pp0.1073741825
Interface type: Dynamic
Interface Set: aci-1001-demux0.1073741824
Interface Set Type: Dynamic
Interface Set Session ID: 2
Underlying Interface: demux0.1073741824
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:64:39:01:02
State: Active
Radius Accounting ID: 3
Session ID: 3
Agent Circuit ID: aci-ppp-dhcp-dvlan-50
Login Time: 2012-03-07 13:46:53 PST

```

#### show subscribers extensive

```

user@host> show subscribers extensive
Type: DHCP
User Name: pd-user1
IPv6 Prefix: 2002:db2:ffff:1::/64
Logical System: default
Routing Instance: default
Interface: ge-3/1/3.2
Interface type: Static
MAC Address: 00:51:ff:ff:00:03

```

```
State: Active
Radius Accounting ID: 1
Session ID: 1
Login Time: 2011-08-25 12:12:26 PDT
DHCP Options: len 42
00 08 00 02 00 00 00 01 00 0a 00 03 00 01 00 51 ff ff 00 03
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Address Pool: pd_pool
IPv6 Network Prefix Length: 48
```

#### show subscribers extensive (RPF Check Fail Filter)

```
user@host> show subscribers extensive
...
Type: VLAN
Logical System: default
Routing Instance: default
Interface: ae0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof
State: Active
Session ID: 9
VLAN Id: 100
Login Time: 2011-08-26 08:17:00 PDT
IPv4 rpf-check Fail Filter Name: rpf-allow-dhcp
IPv6 rpf-check Fail Filter Name: rpf-allow-dhcpv6
...
```

#### show subscribers extensive (L2TP LNS Subscribers on MX Series Routers)

```
user@host> show subscribers extensive
Type: L2TP
User Name: user1@jnpr.net
IP Address: 10.1.32.58
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: si-5/2/0.1073749824
Interface type: Dynamic
Dynamic Profile Name: dyn-lns-profile2
Dynamic Profile Version: 1
State: Active
Radius Accounting ID: 8001
Session ID: 8001
Login Time: 2011-04-25 20:27:50 IST
IPv4 Input Filter Name: classify-si-5/2/0.1073749824-in
IPv4 Output Filter Name: classify-si-5/2/0.1073749824-out
```

#### show subscribers extensive (IPv4 and IPv6 Dual Stack)

```
user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlanProfile
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.1001
VLAN Id: 0x8100.1
```

```

Login Time: 2011-11-30 00:18:04 PST

Type: PPPoE
User Name: dualstackuser1@ISP1.com
IP Address: 61.1.1.1
IPv6 Prefix: 2041:1:1::/48
IPv6 User Prefix: 2061:1:1:1::/64
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Dynamic
Dynamic Profile Name: dualStack-Profile1
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
Login Time: 2011-11-30 00:18:05 PST
IPv6 Delegated Network Prefix Length: 48
IPv6 Interface Address: 2061:1:1:1::1/64
IPv6 Framed Interface Id: 1:1:2:2
IPv4 Input Filter Name: FILTER-IN-pp0.1073741825-in
IPv4 Output Filter Name: FILTER-OUT-pp0.1073741825-out
IPv6 Input Filter Name: FILTER-IN6-pp0.1073741825-in
IPv6 Output Filter Name: FILTER-OUT6-pp0.1073741825-out

Type: DHCP
IPv6 Prefix: 2041:1:1::/48
Logical System: default
Routing Instance: ASP-1
Interface: pp0.1073741825
Interface type: Static
MAC Address: 00:00:64:03:01:02
State: Active
Radius Accounting ID: jnpr :3
Session ID: 3
Underlying Session ID: 2
Login Time: 2011-11-30 00:18:35 PST
DHCP Options: len 42
00 08 00 02 0b b8 00 01 00 0a 00 03 00 01 00 00 64 03 01 02
00 06 00 02 00 19 00 19 00 0c 00 00 00 00 00 00 00 00 00
00 00
IPv6 Delegated Network Prefix Length: 48

```

### show subscribers extensive (Effective Shaping-Rate)

```

user@host> show subscribers extensive
Type: VLAN
Logical System: default
Routing Instance: default
Interface: demux0.1073741837
Interface type: Dynamic
Interface Set: ifset-1
Underlying Interface: ae1
Dynamic Profile Name: svlan-dhcp-test
State: Active
Session ID: 1
Stacked VLAN Id: 0x8100.201
VLAN Id: 0x8100.201
Login Time: 2011-11-30 00:18:04 PST

```

Effective shaping-rate: 31000000k

...

#### show subscribers aci-interface-set-name detail (Subscriber Sessions Using Specified ACI Interface Set)

```
user@host> show subscribers aci-interface-set-name aci-1003-ge-1/0/0.4001 detail
```

Type: VLAN  
Logical System: default  
Routing Instance: default  
Interface: ge-1/0/0.  
Underlying Interface: ge-1/0/0.4001  
Dynamic Profile Name: aci-vlan-set-profile  
Dynamic Profile Version: 1  
State: Active  
Session ID: 13  
Agent Circuit ID: aci-ppp-vlan-10  
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE  
User Name: ppphint2  
IP Address: 10.10.1.7  
Logical System: default  
Routing Instance: default  
Interface: pp0.1073741834  
Interface type: Dynamic  
**Interface Set: aci-1003-ge-1/0/0.4001**  
**Interface Set Type: Dynamic**  
**Interface Set Session ID: 13**  
Underlying Interface: ge-1/0/0.4001  
Dynamic Profile Name: aci-vlan-pppoe-profile  
Dynamic Profile Version: 1  
MAC Address: 00:00:65:26:01:02  
State: Active  
Radius Accounting ID: 14  
Session ID: 14  
Agent Circuit ID: aci-ppp-vlan-10  
Login Time: 2012-03-12 10:41:57 PDT

#### show subscribers agent-circuit-identifier detail (Subscriber Sessions Using Specified ACI Substring)

```
user@host> show subscribers agent-circuit-identifier aci-ppp-vlan detail
```

Type: VLAN  
Logical System: default  
Routing Instance: default  
Interface: ge-1/0/0.  
Underlying Interface: ge-1/0/0.4001  
Dynamic Profile Name: aci-vlan-set-profile  
Dynamic Profile Version: 1  
State: Active  
Session ID: 13  
**Agent Circuit ID: aci-ppp-vlan-10**  
Login Time: 2012-03-12 10:41:56 PDT

Type: PPPoE  
User Name: ppphint2  
IP Address: 10.10.1.7  
Logical System: default  
Routing Instance: default  
Interface: pp0.1073741834  
Interface type: Dynamic  
**Interface Set: aci-1003-ge-1/0/0.4001**

```

Interface Set Type: Dynamic
Interface Set Session ID: 13
Underlying Interface: ge-1/0/0.4001
Dynamic Profile Name: aci-vlan-pppoe-profile
Dynamic Profile Version: 1
MAC Address: 00:00:65:26:01:02
State: Active
Radius Accounting ID: 14
Session ID: 14
Agent Circuit ID: aci-ppp-vlan-10
Login Time: 2012-03-12 10:41:57 PDT

```

### show subscribers interface extensive

```

user@host> show subscribers interface demux0.1073741826 extensive
Type: VLAN
User Name: test1@test.com
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Dynamic
Dynamic Profile Name: profile-vdemux-relay-23qos
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 12
Session ID: 12
Stacked VLAN Id: 0x8100.1500
VLAN Id: 0x8100.2902
Login Time: 2011-10-20 16:21:59 EST

Type: DHCP
User Name: test1@test.com
IP Address: 172.16.200.6
IP Netmask: 255.255.255.0
Logical System: default
Routing Instance: testnet
Interface: demux0.1073741826
Interface type: Static
MAC Address: 00:00:6e:56:01:04
State: Active
Radius Accounting ID: 21
Session ID: 21
Login Time: 2011-10-20 16:24:33 EST
Service Sessions: 2

Service Session ID: 25
Service Session Name: SUB-QOS
State: Active

Service Session ID: 26
Service Session Name: service-cb-content
State: Active
IPv4 Input Filter Name: content-cb-in-demux0.1073741826-in
IPv4 Output Filter Name: content-cb-out-demux0.1073741826-out

```

### show subscribers logical-system terse

```

user@host> show subscribers logical-system test1 terse

```

Interface	IP Address/VLAN ID	User Name	LS:RI
demux0.1073741825	101.0.0.3	RETAILER1-CLIENT	test1:retailer1
demux0.1073741826	102.0.0.3	RETAILER2-CLIENT	test1:retailer2

#### show subscribers physical-interface count

```
user@host> show subscribers physical-interface ge-1/0/0 count
Total subscribers: 3998, Active Subscribers: 3998
```

#### show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

#### show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

#### show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

#### show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

#### show subscribers user-name detail

```
user@host> show subscribers user-name larry1 detail
Type: DHCP
User Name: larry1
IP Address: 100.0.0.37
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: ge-1/0/0.1
Interface type: Static
Dynamic Profile Name: foo
MAC Address: 00:10:94:00:00:01
State: Active
Radius Accounting ID: 1
Session ID: 1
```

```

Login Time: 2011-11-07 08:25:59 PST
DHCP Options: len 52
35 01 01 39 02 02 40 3d 07 01 00 10 94 00 00 01 33 04 00 00
00 3c 0c 15 63 6c 69 65 6e 74 5f 50 6f 72 74 20 2f 2f 32 2f
37 2d 30 2d 30 37 05 01 06 0f 21 2c

```

#### show subscribers vlan-id

```

user@host> show subscribers vlan-id 100
Interface          IP Address          User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825

```

#### show subscribers vlan-id detail

```

user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

```

#### show subscribers vpi vci extensive (PPPoE-over-ATM Subscriber Session)

```

user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 100.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102

```

## show system alarms

<b>Syntax</b>	show system alarms
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display active system alarms.
<b>Options</b>	This command has no options.
<b>Additional Information</b>	<p>System alarms are preset. They include a <i>configuration</i> alarm that appears when no rescue configuration alarm is set and a <i>license</i> alarm that appears when a software feature is configured and no valid license is configured for the feature. On EX6200 switches, an alarm can be triggered by an internal link error. For more information about system alarms, see the <i>Junos OS Administration Library for Routing Devices</i>.</p> <p>In Junos OS release 11.1 and later, alarms for fans also show the slot number of the malfunctioning fans in the CLI output.</p> <p>Starting with Junos OS Release 13.2, you can view degraded fabric alarms on a routing matrix based on TX Matrix Plus router with 3D SIBs. The alarm indicates that the source FPC is running with a degraded fabric condition. This alarm is an early warning of a possible fabric black-hole condition. When the degraded fabric alarm is raised on the source FPC, you can take remedial action to avoid a fabric black-hole condition. The degraded fabric alarm is raised on the source FPC if both the following conditions are met:</p> <ul style="list-style-type: none"> <li>• The active Packet Forwarding Engine destinations are reachable on one or no active switching planes.</li> <li>• At least one of the inactive switching planes has a fault that causes the destination Packet Forwarding Engine to become unreachable.</li> </ul>
<b>Required Privilege Level</b>	admin
<b>List of Sample Output</b>	<p><a href="#">show system alarms on page 975</a></p> <p><a href="#">show system alarms (Fan Tray) on page 975</a></p> <p><a href="#">show system alarms (QFX Series) on page 975</a></p> <p><a href="#">show system alarms (EX6200) on page 975</a></p> <p><a href="#">show system alarms (TX Matrix Plus router with 3D SIBs) on page 975</a></p>
<b>Output Fields</b>	Table 53 on page 974 lists the output fields for the <b>show system alarms</b> command. Output fields are listed in the approximate order in which they appear.

**Table 53: show system alarms Output Fields**

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.



Table 53: show system alarms Output Fields (*continued*)

Field Name	Field Description
<b>Class</b>	Severity class for this alarm: <b>Minor</b> or <b>Major</b> .
<b>Description</b>	Information about the alarm.

## Sample Output

### show system alarms

```

user@host> show system alarms
2 alarms currently active
Alarm time          Class    Description
2005-02-24 17:29:34 UTC  Minor    IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC  Minor    Rescue configuration is not sent

```

### show system alarms (Fan Tray)

```

user@host> show system alarms
4 alarms currently active
Alarm time          Class    Description
2010-11-11 20:27:38 UTC  Major    Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor    Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 0 Failure

```

### show system alarms (QFX Series)

```

user@switch> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent

```

### show system alarms (EX6200)

```

user@switch> show system alarms
2 alarms currently active
Alarm time          Class    Description
2013-04-05 16:51:41 PDT  Major    FPC 8 internal link errors detected
2013-04-04 18:05:35 PDT  Minor    Rescue configuration is not set

```

### show system alarms (TX Matrix Plus router with 3D SIBs)

```

user@router> show system alarms

sfc0-re0:
-----
2 alarms currently active
Alarm time          Class    Description
2013-05-08 18:13:58 UTC  Major    LCC 0 Major Errors
2013-05-08 17:48:46 UTC  Major    LCC 7 Major Errors

lcc0-re1:
-----
1 alarm currently active
Alarm time          Class    Description

```

2013-05-08 18:19:24 UTC Major FPC 1 degraded fabric condition detected

lcc7-re0:

-----  
1 alarm currently active

Alarm time	Class	Description
------------	-------	-------------

2013-05-08 18:19:24 UTC	Major	FPC 7 degraded fabric condition detected
-------------------------	-------	------------------------------------------

## show system audit

<b>List of Syntax</b>	<a href="#">Syntax on page 977</a> <a href="#">Syntax (EX Series Switch and MX Series Router) on page 977</a> <a href="#">Syntax (TX Matrix Router) on page 977</a> <a href="#">Syntax (TX Matrix Plus Router) on page 977</a> <a href="#">Syntax (QFX Series) on page 977</a>
<b>Syntax</b>	show system audit <root-only>
<b>Syntax (EX Series Switch and MX Series Router)</b>	show system audit <all-members> <local> <member <i>member-id</i> > <root-only>
<b>Syntax (TX Matrix Router)</b>	show system audit <all-lcc   lcc <i>number</i>   scc> <root-only>
<b>Syntax (TX Matrix Plus Router)</b>	show system audit <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <root-only>
<b>Syntax (QFX Series)</b>	show system audit <infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i>   root-only>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display the state and checksum values for file systems.
<b>Options</b>	<p><b>none</b>—Display the state and checksum values for all file systems.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for all of the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for all T1600 or T4000 routers connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on all members of the Virtual Chassis configuration.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display file system MD5 hash and permissions information for a specific T640 router</p>

that is connected to the TX Matrix router. On a TX Matrix Plus router, display file system MD5 hash and permissions information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**infrastructure *name***—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for a fabric control Routing Engine or a fabric control Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Interconnect device.

**local**—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the local Virtual Chassis member.

**member *member-id***—(EX4200 switch, QFX Series, and MX Series routers only) (Optional) Display file system MD5 hash and permissions information on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display file system MD5 hash and permissions information for the Node group

**root-only**—(Optional) Check only the root (/) file system. On a QFabric system, you can check the root (/) file system on the infrastructure (fabric manager Routing Engine and fabric control Routing Engine), Interconnect device, or Node group.

**scc**—(TX Matrix routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display file system MD5 hash and permissions information for the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.

**Additional Information** To redirect the output to a file, issue the following command:

***ssh device-name 'show system audit root-only' > output-file***

If you save the output of the **show system audit root-only** command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

By default, when you issue the **show system audit** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level	admin
List of Sample Output	<a href="#">show system audit root-only on page 979</a> <a href="#">show system audit lcc (TX Matrix Router) on page 980</a> <a href="#">show system audit lcc (TX Matrix Plus Router) on page 981</a> <a href="#">show system audit root-only (QFX3500 Switch) on page 983</a>

## Sample Output

### show system audit root-only

```

user@host> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
           md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
           mode=0744 size=1934552 time=944688902.0 \
           md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
           md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
           md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \
           md5digest=93d722493ed38477338a1405d7dcbb40
boot.help  uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
           md5digest=9b7126385734bcae753f4179ab59d8e5
compat     type=link mode=0777 size=11 time=915149058.0 \
           link=/usr/compat
kernel     mode=0444 size=1947607 time=950230892.0 \
           md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr uid=1112 size=1947642 time=950252597.0 \
           md5digest=82e1637682d58ec28964dfce7fccb62e
kernel.config \
           mode=0644 size=0 time=915149058.0 \
           md5digest=d41d8cd98f00b204e9800998ecf8427e
sys        type=link mode=0777 size=11 time=915149029.0 \
           link=/usr/src/sys

```

## show system audit lcc (TX Matrix Router)

```

user@host> show system audit lcc 2
lcc2-re0:
-----
#       user: root
#       machine: rodin-lcc2
#       tree: /
#       date: Mon Sep 13 11:55:33 2004

# .
/set type=file uid=0 gid=0 mode=0555 nlink=1 flags=none
.      type=dir nlink=20 size=512 time=1094982121.0
  COPYRIGHT mode=0644 size=4735 time=986012708.0 \
    md5digest=78396df1404ad742e6eb1be28f0cd63b
    kernel type=link mode=0700 size=17 time=1090266262.0 \
      link=/packages/jkernel

# ./altconfig
altconfig type=dir nlink=2 size=512 time=1089801320.0
# ./altconfig
..

# ./altroot
altroot type=dir nlink=2 size=512 time=1089801320.0
# ./altroot
..

# ./b
b type=dir mode=0755 nlink=2 size=512 time=1093961429.0
# ./b
..

# ./bin
/set type=file uid=0 gid=0 mode=0700 nlink=1 flags=none
bin type=dir mode=0755 nlink=2 size=512 time=1089843059.0
  [ type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/cat
  chmod type=link size=29 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/ed
  expr type=link size=28 time=1090266270.0 \
    link=/packages/mnt/jbase/bin/expr
  hostname type=link size=32 time=1090266270.0 \

```

```

link=/packages/mnt/jbase/bin/hostname
kill      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/kill
ln        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ln
ls        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1090266270.0 \
link=/packages/mnt/jbase/bin/mkdir
mv        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/mv
ps        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/ps
pwd       type=link size=27 time=1090266270.0 \
link=/packages/mnt/jbase/bin/pwd
rcp       type=link size=27 time=1090266270.0 \
link=/packages/mnt/jbase/bin/rcp
red       type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/red
rm        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1090266270.0 \
link=/packages/mnt/jbase/bin/rmdir
sh        type=link size=26 time=1090266270.0 \
link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1090266270.0 \
link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1090266270.0 \
link=/packages/mnt/jbase/bin/csh
test      type=link size=28 time=1090266270.0 \
link=/packages/mnt/jbase/bin/test
# ./bin
..

# ./boot
/set type=file uid=0 gid=0 mode=0444 nlink=1 flags=none
boot      type=dir mode=0555 nlink=3 size=512 time=1095069935.0
boot0     size=512 time=1094978286.0 \
md5digest=6f780822dd4ae482a20462b66e542cca
boot1     mode=0555 size=512 time=1094978294.0 \
md5digest=8d112b09df342cd0b60fdb9bdcde8e07
boot2     mode=0555 size=7680 time=1094978294.0 \
md5digest=28eb58c4068c6b85717e1484f9e028e4
cdboot    mode=0555 size=165888 time=1094978298.0 \
md5digest=1474c6b800dfc82ba552d7c36116d07d
kgzldr.o  size=5996 time=1094982121.0 \
md5digest=c53dc948eb07e2ea4eb0413e4c4634a3
loader    mode=0555 size=163840 time=1094978298.0 \
md5digest=82d9dc2d31033476bfb61bb7264c4fed
loader.4th size=9237 time=986013631.0 \
md5digest=43144391465ad50267d31e0a320be1de
...

```

#### show system audit lcc (TX Matrix Plus Router)

```
user@host> show system audit all-chassis
```

```

sfc0-re0:
-----
#       user: root
#       machine: finalfive
#       tree: /
#       date: Mon May 18 00:13:16 2009

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1 flags=none
.      type=dir nlink=23 size=512 time=1242347096.0
  COPYRIGHT mode=0644 size=6196 time=1168587741.0 \
    md5digest=bbad415e1c29bbdd9b383537100412c
    kernel type=link size=17 time=1242347011.0 link=/packages/jkernel
    staging type=link mode=0777 size=8 time=1242346935.0 link=/var/tmp

# ./snap
.snap type=dir mode=0775 nlink=2 size=512 time=1242346922.0
# ./snap
..

# ./altconfig
altconfig type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altconfig
..

# ./altroot
altroot type=dir mode=0500 nlink=2 size=512 time=1242319843.0
# ./altroot
..

# ./bin
bin type=dir nlink=2 size=512 time=1242346944.0
  \133 type=link size=28 time=1242346942.0 \
    link=/packages/mnt/jbase/bin/test
  cat type=link size=27 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/cat
  chflags type=link size=31 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/chflags
  chmod type=link size=29 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/chmod
  cp type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/cp
  csh type=link size=27 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/csh
  date type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/date
  dd type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/dd
  df type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/df
  echo type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/echo
  ed type=link size=26 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/ed
  expr type=link size=28 time=1242346941.0 \
    link=/packages/mnt/jbase/bin/expr
  hostname type=link size=32 time=1242346941.0 \

```



```

kill      link=/packages/mnt/jbase/bin/hostname
          type=link size=28 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/kill
ln        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ln
ls        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ls
mkdir     type=link size=29 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mkdir
mv        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/mv
pax       type=link size=27 time=1242346944.0 \
          link=/packages/mnt/jbase/bin/pax
ps        type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/ps
pwd       type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/pwd
rcp       type=link size=27 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rcp
red       type=link size=26 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/red
rm        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rm
rmdir     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/rmdir
sh        type=link size=26 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sh
sleep     type=link size=29 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sleep
stty      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/stty
sync      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/sync
tcsh      type=link size=27 time=1242346941.0 \
          link=/packages/mnt/jbase/bin/csh
test      type=link size=28 time=1242346942.0 \
          link=/packages/mnt/jbase/bin/test
# ./bin
...

```

#### show system audit root-only (QFX3500 Switch)

```

user@switch> show system audit root-only
#          user: root
#          machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.          type=dir nlink=23 size=1024 time=950252640.0
.cshrc     uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
          md5digest=f414e06fea6bd646244b98e13d6e6226
.kernel.jkernel.backup \
          mode=0744 size=1934552 time=944688902.0 \
          md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile   uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
          md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
COPYRIGHT  uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
          md5digest=7df8bc77dcee71382ea73eb0ec6a9243
boot.config mode=0644 size=3 time=945902618.0 \

```

```
boot.help      md5digest=93d722493ed38477338a1405d7dcbb40
                uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
                md5digest=9b7126385734bcae753f4179ab59d8e5
compat         type=link mode=0777 size=11 time=915149058.0 \
                link=/usr/compat
kernel         mode=0444 size=1947607 time=950230892.0 \
                md5digest=1a2a8aff2fec678a918ba0d6bf063980
kernel.avr     uid=1112 size=1947642 time=950252597.0 \
                md5digest=82e1637682d58ec28964dfee7fccb62e
kernel.config \
                mode=0644 size=0 time=915149058.0 \
                md5digest=d41d8cd98f00b204e9800998ecf8427e
sys            type=link mode=0777 size=11 time=915149029.0 \
                link=usr/src/sys
```

## show system boot-messages

<b>List of Syntax</b>	<a href="#">Syntax on page 985</a> <a href="#">Syntax (EX Series Switches) on page 985</a> <a href="#">Syntax (TX Matrix Router) on page 985</a> <a href="#">Syntax (TX Matrix Plus Router) on page 985</a> <a href="#">Syntax (MX Series Router) on page 985</a> <a href="#">Syntax (QFX Series) on page 985</a>
<b>Syntax</b>	show system boot-messages
<b>Syntax (EX Series Switches)</b>	show system boot-messages <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system boot-messages <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system boot-messages <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system boot-messages <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system boot-messages infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display initial messages generated by the system kernel upon startup. These messages are the contents of <code>/var/run/dmesg.boot</code> .
<b>Options</b>	<b>none</b> —Display all boot time messages.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display boot time messages for all of the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for all T640 routers connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for all connected T1600 or T4000 LCCs.  <b>all-members</b> —(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration.

**infrastructure *name***—(QFabric systems only) (Optional) Display boot time messages on the fabric control Routing Engine or fabric manager Routing engines.

**interconnect-device *name***—(QFabric systems only) (Optional) Display boot time messages on the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display boot time messages for a specific T640 router connected to a TX Matrix router. On a TX Matrix Plus router, display boot time messages for a specific router connected to a TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display boot time messages on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display boot time messages on the Node group.

**scc**—(TX Matrix routers only) (Optional) Display boot time messages for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display boot time messages for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system boot-messages** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

## Related Documentation

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

## List of Sample Output

- [show system boot-messages \(TX Matrix Router\) on page 987](#)
- [show system boot-messages lcc \(TX Matrix Router\) on page 988](#)
- [show system boot-messages \(TX Matrix Plus Router\) on page 989](#)
- [show system boot-messages \(QFX3500 Switch\) on page 989](#)

## Sample Output

### show system boot-messages (TX Matrix Router)

```

user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
    The Regents of the University of California. All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    tlim@single.juniper.net:/p/build/20000216-0905/4.1/release_kernel/sys/compil
e/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x66a Stepping=10
    Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on

pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on

pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300

```

```

ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SQFXB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)

```

### show system boot-messages lcc (TX Matrix Router)

```

user@host> show system boot-messages lcc 2
lcc2-re0:
-----
Copyright (c) 1996-2001, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2001 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 7.0-20040912.0 #0: 2004-09-12 09:16:32 UTC

builder@benten.juniper.net:/build/benten-b/7.0/20040912.0/obj-i386/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz
Timecounter "TSC" frequency 601368936 Hz
CPU: Pentium III/Pentium III Xeon/Celeron (601.37-MHz 686-class CPU)
    Origin = "GenuineIntel" Id = 0x68a Stepping = 10

Features=0x387f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,PAT,PSE36,PN,MMX,FXSR,SSE>
real memory = 2147467264 (2097136K bytes)
sio0: gdb debugging port
avail memory = 2084040704 (2035196K bytes)
Preloaded elf kernel "kernel" at 0xc06d9000.
DEVFS: ready for devices
Pentium Pro MTRR support enabled
md0: Malloc disk

```

```

DRAM Data Integrity Mode: ECC Mode with h/w scrubbing
npx0: <math processor> on motherboard
npx0: INT 16 interface
pcib0: <ServerWorks NB6635 3.0LE host to PCI bridge> on motherboard
pci0: <PCI bus> on pcib0
pcic-pci0: <TI PCI-1410 PCI-CardBus Bridge> irq 15 at device 1.0 on pci0
pcic-pci0: TI12XX PCI Config Reg: [pwr save][pci only]
fxp0: <Intel Embedded 10/100 Ethernet> port 0x1000-0x103f mem
0xfb800000-0xfb81ffff,0xfb820000-0xfb820fff irq 9 at device 3.0 on pci0
fxp1: <Intel Embedded 10/100 Ethernet> port 0x1040-0x107f mem
0xfb840000-0xfb85ffff,0xfb821000-0xfb821fff irq 11 at device 4.0 on pci0
...

```

### show system boot-messages (TX Matrix Plus Router)

```

user@host> show system boot-messages
sfc0-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6B3.3 #0: 2009-06-17 19:52:08 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/release/9.6B3.3/obj-i386/bsd/sys/compile/JUNIPER
MPTable: Timecounter "i8254" frequency 1193182 Hz quality 0 CPU: Intel(R) Xeon(R)
CPU          L5238 @ 2.66GHz (2660.01-MHz 686-class CPU)   Origin =
"GenuineIntel" Id = 0x1067a Stepping = 10   Features=0xbfebfbff
...
lcc1-re0:
-----
Copyright (c) 1996-2009, Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1992-2006 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
JUNOS 9.6-20090617.0 #0: 2009-06-17 04:15:14 UTC

builder@lanath.juniper.net:/volume/build/junos/9.6/production/20090617.0/obj-i386/bsd/sys/compile/JUNIPER
Timecounter "i8254" frequency 1193182 Hz quality 0
CPU: Intel(R) Xeon(R) CPU          @ 1.86GHz (1862.01-MHz 686-class CPU)

Origin = "GenuineIntel" Id = 0x1067a Stepping = 10
Features=0xbfebfbff
...

```

### show system boot-messages (QFX3500 Switch)

```

user@switch> show sytem boot-messages
getmemsize: msgbufp[size=32768] = 0x81d07fe4

System physical memory distribution:
-----
Total physical memory: 4160749568 (3968 MB)
Physical memory used: 3472883712 (3312 MB)
Physical memory allocated to kernel: 2130706432 (2032 MB)
Physical memory allocated to user BTLB: 1342177280 (1280 MB)
-----

Copyright (c) 1996-2010, Juniper Networks, Inc.

```

All rights reserved.

Copyright (c) 1992-2006 The FreeBSD Project.

Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994

The Regents of the University of California. All rights reserved.

JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
WARNING: debug.mpsafenet forced to 0 as ipsec requires Giant
JUNOS 11.1I #0: 2010-09-17 19:18:07 UTC
```

```
ssiano@svl-junos-pool125.juniper.net:/c/ssiano/DEV_QFX_SI_BRANCH/03/20100917.399988/
obj-xlr/bsd/sys/compile/JUNIPER-DCTOR
real memory = 3472883712 (3312MB)
avail memory = 1708171264 (1629MB)
cpuid: 0, bt1b_cpumap:0xffffffff8
FreeBSD/SMP: Multiprocessor System Detected: 12 CPUs
ETHERNET SOCKET BRIDGE initialising
Initializing QFX platform properties ..
cpu0 on motherboard
: RMI's XLR CPU Rev. 0.3 with no FPU implemented
  L1 Cache: I size 32kb(32 line), D size 32kb(32 line), eight way.
  L2 Cache: Size 1024kb, eight way
pic_lbus0: <XLR Local Bus>
pic_lbus0: <XLR Local Bus> on motherboard
Enter qfx control ethernet probe addr:0xc5eeec00
gmac4: <XLR GMAC GE Ethernet> on pic_lbus0
me0: Ethernet address 00:1d:b5:f7:68:40
Enter qfx control ethernet probe addr:0xc5eeeb40
gmac5: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:41
Enter qfx control ethernet probe addr:0xc5eeea80
gmac6: <XLR GMAC GE Ethernet> on pic_lbus0
me1: Ethernet address 00:1d:b5:f7:68:42
sio0 on pic_lbus0
Entering sioattach
sio0: type 16550A, console
xls_setup_intr: skip irq 3, xlr regs are set up somewhere else.
gblmem0 on pic_lbus0
ehci0: <RMI XLS USB 2.0 controller> on pic_lbus0
ehci_bus_attach: allocated resource. tag=1, base=bef24000
xls_ehci_init: endian hardware swapping NOT enabled.
usb0: EHCI version 1.0
usb0 on ehci0
usb0: USB revision 2.0
uhub0: vendor 0x0000 EHCI root hub, class 9/0, rev 2.00/1.00, addr 1
uhub0: 2 ports with 2 removable, self powered
umass0: USB USBFlashDrive, rev 2.00/11.00, addr 2
pcib0: PCIe link 0 up
pcib0: PCIe link 2 up
pcib0: PCIe link 3 up
pcib0: <XLS PCI Host Controller> on pic_lbus0
pci0: <PCI bus> on pcib0
pcib1: <PCI-PCI bridge> at device 0.0 on pci0
pci1: <PCI bus> on pcib1
pci1: <network, ethernet> at device 0.0 (no driver attached)
pcib2: <PCI-PCI bridge> at device 1.0 on pci0
pcib3: <PCI-PCI bridge> at device 2.0 on pci0
pci2: <PCI bus> on pcib3
pci2: <network, ethernet> at device 0.0 (no driver attached)
pcib4: <PCI-PCI bridge> at device 3.0 on pci0
```



```

pci3: <PCI bus> on pcib4
pci3: <network, ethernet> at device 0.0 (no driver attached)
cfi device address space at 0xbc000000
cfi0: <AMD/Fujitsu - 8MB> on pic_lbus0
cfi device address space at 0xbc000000
i2c0: <I2C bus controller> on pic_lbus0
i2c1: <I2C bus controller> on pic_lbus0
qfx_fmn0 on pic_lbus0
pool offset 1503776768
xlr_lbus0: <XLR Local Bus Controller> on motherboard
qfx_bcpld_probe[124]
qfx_bcpld_probe[138]: dev_type=0x0
qfx_bcpld_probe[124]
qfx_bcpld0: QFX BCPLD probe success
qfx_bcpld0qfx_bcpld_attach[174]
qfx_bcpld_attach[207] : bus_space_tag=0x0, bus_space_handle=0xbd900000
qfx_bcpld_probe[124]
qfx_bcpld1: QFX BCPLD probe success
qfx_bcpld1qfx_bcpld_attach[174]
tor_bcpld_slave_attach[1245] : bus_space_tag=0x0, bus_space_handle=0xbda00000
Initializing product: 96 ..
bmeb: bmeb_lib_init done 0xc60a5000, addr 0x809c99a0
bme0:Virtual BME driver initializing
Timecounter "mips" frequency 1200000000 Hz quality 0
Timecounter "xlr_pic_timer" frequency 66666666 Hz quality 1
Timecounters tick every 1.000 msec
Loading the NETPFE fc module
IPsec: Initialized Security Association Processing.
SMP: AP CPU #3 Launched!
SMP: AP CPU #1 Launched!
SMP: AP CPU #2 Launched!
SMP: AP CPU #4 Launched!
SMP: AP CPU #5 Launched!
SMP: AP CPU #7 Launched!
SMP: AP CPU #6 Launched!
SMP: AP CPU #11 Launched!
SMP: AP CPU #10 Launched!
SMP: AP CPU #9 Launched!
SMP: AP CPU #8 Launched!
da0 at umass-sim0 bus 0 target 0 lun 0
da0: <USB USBFlashDrive 1100> Removable Direct Access SCSI-0 device
da0: 40.000MB/s transfers
da0: 3920MB (8028160 512 byte sectors: 255H 63S/T 499C)
Trying to mount root from ufs:/dev/da0s1a

```

## show system buffers

---

<b>List of Syntax</b>	<a href="#">Syntax on page 992</a> <a href="#">Syntax (EX Series) on page 992</a> <a href="#">Syntax (TX Matrix Router) on page 992</a> <a href="#">Syntax (TX Matrix Plus Router) on page 992</a> <a href="#">Syntax (MX Series Router) on page 992</a> <a href="#">Syntax (QFX Series) on page 992</a>
<b>Syntax</b>	show system buffers
<b>Syntax (EX Series)</b>	show system buffers <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system buffers <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system buffers <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system buffers <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system buffers <infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i>   root-only (infrastructure <i>name</i>   interconnect-device <i>name</i>   node-group <i>name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router or switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).
<b>Options</b>	<b>none</b> —Show all buffer statistics.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for all routers connected to the TX Matrix Plus router.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show buffer statistics for all of the chassis.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for all members of the Virtual Chassis configuration.

**infrastructure *name***—(QFabric systems only) (Optional) Show buffer statistics for a fabric control Routing Engine or a fabric control Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Show buffer statistics for the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show buffer statistics for a specific T640 router (or line-card chassis) that is connected to the TX Matrix router. On a TX Matrix Plus router, show buffer statistics for a specific router (line-card chassis) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Show buffer statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Show buffer statistics for the Node group

**sfc**—(TX Matrix Plus routers only) (Optional) Show buffer statistics for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system buffers** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

A special type of memory buffer called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li></ul>
List of Sample Output	<a href="#">show system buffers on page 995</a> <a href="#">show system buffers scc (TX Matrix Router) on page 996</a> <a href="#">show system buffers sfc (TX Matrix Plus Router) on page 996</a> <a href="#">show system buffers all-chassis (TX Matrix Plus Router) on page 996</a> <a href="#">show system buffers node-group (QFabric System) on page 997</a>
Output Fields	<a href="#">Table 54 on page 995</a> describes the output fields for the <b>show system buffers</b> command. Output fields are listed in the approximate order in which they appear.

Table 54: show system buffers Output Fields

Field Name	Field Description
<b>mbufs in use</b>	Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. Each memory buffer has a type, and the output itemizes the amount allocated for each type. Types with no memory buffers allocated are not displayed.
<b>mbufs allocated to packet headers</b>	Number of memory buffers currently holding packet headers
<b>mbufs allocated to control blocks</b>	Number of memory buffers currently holding the state for sockets.
<b>mbufs allocated to send data</b>	Number of memory buffers currently holding socket send data.
<b>mbufs allocated to pfe refill data</b>	Number of memory buffers currently holding Packet Forwarding Engine refill data.
<b>mbufs allocated to fxp data</b>	Number of memory buffers currently holding fxp data.
<b>mbufs allocated to socket names and addresses</b>	Number of memory buffers currently holding addresses for sockets.
<b>mbuf clusters in use</b>	Allocation statistics for memory buffer clusters.
<b>allocated to network</b>	Total amount of memory in use by the networking and interprocess communication (IPC) code.
<b>requests for memory denied</b>	Number of times a memory allocation request within the IPC and networking code failed.
<b>requests for memory delayed</b>	Number of times a memory allocation request within the IPC and networking code was postponed.
<b>calls to protocol drain routines</b>	Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

## Sample Output

### show system buffers

```

user@host> show system buffers
397/893/1290 mbufs in use (current/cache/total)
395/331/726/30000 mbuf clusters in use (current/cache/total/max)
384/256 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
889K/885K/1774K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/5/1024 sfbufs in use (current/peak/max)

```

```
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

#### show system buffers scc (TX Matrix Router)

```
user@host> show system buffers scc
213 mbufs in use:
    11 mbufs allocated to packet headers
    26 mbufs allocated to socket names and addresses
    2 mbufs allocated to socket options
    17 mbufs allocated to socket send data
    2 mbufs allocated to pfe data
    155 mbufs allocated to fxp data (rx)
    511 mbufs allocated to <mbuf type 86>
    256 mbufs allocated to <mbuf type 92>
924/1162 mbuf clusters in use
2788 Kbytes allocated to network (75% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines
```

#### show system buffers sfc (TX Matrix Plus Router)

```
user@host> show system buffers sfc 0

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

#### show system buffers all-chassis (TX Matrix Plus Router)

```
user@host> show system buffers all-chassis

sfc0-re0:
-----
4363/2807/7170 mbufs in use (current/cache/total)
4358/1968/6326/30000 mbuf clusters in use (current/cache/total/max)
256/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
9806K/4637K/14444K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/10/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
```

```
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc0-re0:
```

```
-----
772/2558/3330 mbufs in use (current/cache/total)
772/598/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1737K/1835K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc1-re0:
```

```
-----
773/2437/3210 mbufs in use (current/cache/total)
773/453/1226/30000 mbuf clusters in use (current/cache/total/max)
768/384 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1739K/1515K/3254K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/7/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

```
lcc2-re0:
```

```
-----
816/2514/3330 mbufs in use (current/cache/total)
816/554/1370/30000 mbuf clusters in use (current/cache/total/max)
768/512 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1836K/1736K/3572K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
```

### show system buffers node-group (QFabric System)

```
user@switch> show system buffers node-group node1
node-group node1:
```

```
-----
2/2698/2700 mbufs in use (current/cache/total)
2/1520/1522/30000 mbuf clusters in use (current/cache/total/max)
0/1280 mbuf+clusters out of packet secondary zone in use (current/cache)
```

```
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
4K/3714K/3719K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/6/6656 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

re0:

```
-----
516/639/1155 mbufs in use (current/cache/total)
515/147/662/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1159K/453K/1612K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```

re1:

```
-----
519/771/1290 mbufs in use (current/cache/total)
518/176/694/30000 mbuf clusters in use (current/cache/total/max)
512/128 mbuf+clusters out of packet secondary zone in use (current/cache)
0/0/0/0 4k (page size) jumbo clusters in use (current/cache/total/max)
0/0/0/0 9k jumbo clusters in use (current/cache/total/max)
0/0/0/0 16k jumbo clusters in use (current/cache/total/max)
1165K/544K/1710K bytes allocated to network (current/cache/total)
0/0/0 requests for mbufs denied (mbufs/clusters/mbuf+clusters)
0/0/0 requests for jumbo clusters denied (4k/9k/16k)
0/4/1024 sbufs in use (current/peak/max)
0 requests for sbufs denied
0 requests for sbufs delayed
0 requests for I/O initiated by sendfile
0 calls to protocol drain routines
```



## show system certificate

<b>Syntax</b>	<code>show system certificate</code> <code>&lt;certificate-id&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series, T Series routers, and QFX Series switches only) Display installed certificates signed by the Juniper Networks certificate authority.
<b>Options</b>	<b>none</b> —Display all installed certificates signed by the Juniper Networks certificate authority.  <b>certificate-id</b> —(Optional) Display the details of a particular certificate.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">show system certificate on page 1000</a> <a href="#">show system certificate (QFX Series) on page 1000</a>
<b>Output Fields</b>	<a href="#">Table 55 on page 999</a> lists the output fields for the <b>show system certificate</b> command. Output fields are listed in the approximate order in which they appear.

**Table 55: show system certificate Output Fields**

Field Name	Field Description
<b>Certificate identifier</b>	Unique identifier associated with a certificate. The certificate identifier is the common name of the subject.
<b>Issuer</b> <b>Subject</b>	Information about the certificate issuer and the distinguished name (DN) of the issuer, respectively: <ul style="list-style-type: none"> <li>• <b>Organization</b>—Name of the owner's organization.</li> <li>• <b>Organizational unit</b>—Name of the owner's department.</li> <li>• <b>Country</b>—Two-character country code in which the owner's system is located.</li> <li>• <b>State</b>—State in the USA in which the owner is using the certificate.</li> <li>• <b>Locality</b>—City in which the owner's system is located.</li> <li>• <b>Common name</b>—Name of the owner of the certificate.</li> <li>• <b>E-mail address</b>—E-mail address of the owner of the certificate.</li> </ul>
<b>Validity</b>	When a certificate is valid.
<b>Signature algorithm</b>	Encryption algorithm applied to the installed certificate.
<b>Public key algorithm</b>	Encryption algorithm applied to the public key.

## Sample Output

### show system certificate

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

### show system certificate (QFX Series)

```
user@host> show system certificate
Certificate identifier: Dallas-v3
Issuer:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
E-mail address:ca@juniper.net
Subject:
Organization: Juniper Networks, Organizational unit: Juniper CA,
Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
E-mail address:ca@juniper.net
Validity:
Not before: Mar 13 03:23:25 2004 GMT
Not after: Mar 24 03:23:25 2014 GMT
Signature algorithm: sha1WithRSAEncryption
Public key algorithm: dsaEncryption
```

## show system commit


<b>Syntax</b>	show system commit <revision> <server>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <b>server</b> introduced in Junos OS Release 12.1 for the PTX Series router. Option <b>revision</b> introduced in Junos OS Release 14.1.
<b>Description</b>	Display the system commit history and any pending commit operation.
<b>Options</b>	<p><b>none</b>—Display the last 50 commit operations listed, most recent to first.</p> <p><b>revision</b>—(Optional) Display the revision number of the active configuration of the Routing Engine(s).</p> <p><b>server</b>—(Optional) Display commit server status.</p>
	<div>  <p><b>NOTE:</b> By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.</p> </div>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear system commit on page 354</a></li> <li><a href="#">show system commit revision</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system commit on page 1003</a> <a href="#">show system commit (At a Particular Time) on page 1003</a> <a href="#">show system commit (At the Next Reboot) on page 1003</a> <a href="#">show system commit (Rollback Pending) on page 1003</a> <a href="#">show system commit (QFX Series) on page 1003</a>
<b>Output Fields</b>	Table 56 on page 1001 describes the output fields for the <b>show system commit</b> command. Output fields are listed in the approximate order in which they appear.

Table 56: show system commit Output Fields

Field Name	Field Description	Level of Output
<number>	Displays the last 50 commit operations listed, most recent to first. The identifier <number> designates a configuration created for recovery using the <b>request system configuration rescue save</b> command.	none

Table 56: show system commit Output Fields (*continued*)

Field Name	Field Description	Level of Output
<code>&lt;time-stamp&gt;</code>	Date and time of the commit operation.	none
<code>&lt;root&gt;/&lt;username&gt;</code>	User who executed the commit operation.	none
<code>&lt;method&gt;</code>	<p>Method used to execute the commit operation:</p> <ul style="list-style-type: none"> <li>• <b>CLI</b>—CLI interactive user performed the commit operation.</li> <li>• <b>Junos XML protocol</b>—Junos XML protocol client performed the commit operation.</li> <li>• <b>synchronize</b>—The <b>commit synchronize</b> command was performed on the other Routing Engine.</li> <li>• <b>snmp</b>—An SNMP <b>set</b> request caused the commit operation.</li> <li>• <b>button</b>—A button on the router or switch was pressed to commit a rescue configuration for recovery.</li> <li>• <b>autoinstall</b>—A configuration obtained through autoinstallation was committed.</li> <li>• <b>other</b>—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit.</li> </ul>	none

## Sample Output

### show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

### show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

### show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

### show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

### show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```

## show system configuration archival

---

**Syntax**    show system configuration archival

**Release Information**    Introduced in Junos OS Release 7.6.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display directory and number of files queued for archival transfer.



**NOTE:** The [edit system configuration] hierarchy is not available on QFabric systems.

---

**Options**    This command has no options.

**Required Privilege Level**    maintenance

**List of Sample Output**    [show system configuration archival on page 1004](#)


### Sample Output

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

## show system configuration rescue

<b>Syntax</b>	show system configuration rescue
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a rescue configuration, if one exists.
<div>  <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems. </div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system configuration archival on page 1004</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system configuration rescue on page 1005</a>

## Sample Output

### show system configuration rescue

```

user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
  services {

```

```
        ftp;  
        rlogin;  
        rsh;  
        telnet;  
    }  
}  
.....
```



## show system connections

<b>List of Syntax</b>	<a href="#">Syntax on page 1007</a> <a href="#">Syntax (EX Series) on page 1007</a> <a href="#">Syntax (TX Matrix Router) on page 1007</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1007</a> <a href="#">Syntax (MX Series Router) on page 1007</a> <a href="#">Syntax (QFX Series) on page 1007</a>
<b>Syntax</b>	<pre>show system connections &lt;extensive&gt; &lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt; &lt;inet   inet6&gt; &lt;show-routing-instances&gt;</pre>
<b>Syntax (EX Series)</b>	<pre>show system connections &lt;extensive&gt; &lt;all-members&gt; &lt;inet   inet6&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;show-routing-instances&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show system connections &lt;extensive&gt; &lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt; &lt;inet   inet6&gt; &lt;show-routing-instances&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show system connections &lt;extensive&gt; &lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt; &lt;inet   inet6&gt; &lt;show-routing-instances&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show system connections &lt;extensive&gt; &lt;all-members&gt; &lt;inet   inet6&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;show-routing-instances&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system connections &lt;extensive&gt; &lt;inet&gt; &lt;infrastructure <i>name</i>&gt; &lt;interconnect-device <i>name</i>&gt; &lt;node-group <i>name</i>&gt; &lt;show-routing-instances&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>

**sfc** option introduced for the TX Matrix Plus router in Junos OS Release 9.6.  
Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.

**Options** **none**—Display information about all active IP sockets on the Routing Engine.

**extensive**—(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block. This option is useful for debugging TCP connections.

**all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system connection activity for all the routers in the chassis.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for all connected T1600 or T4000 LCCs

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for all members of the Virtual Chassis configuration.

**inet | inet6**—(Optional) Display IPv4 connections or IPv6 connections, respectively.

**infrastructure name**—(QFabric systems only) (Optional) Display system connection activity for the fabric control Routing Engines or fabric manager Routing Engines.

**interconnect-device name**—(QFabric systems only) (Optional) Display system connection activity for the Interconnect device.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system connection activity for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system connection activity for a specific router that is connected to the TX Matrix Plus router. Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display system connection activity for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display system connection activity for the Node group.

**scc**—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix routers only) (Optional) Display system connection activity for the TX Matrix Plus router.

**show-routing-instances**—(Optional) Display routing instances.

**Additional Information** By default, when you issue the **show system connections** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system connections on page 1010](#)  
[show system connections extensive on page 1010](#)  
[show system connections lcc \(TX Matrix Router\) on page 1012](#)  
[show system connections show-routing-instances on page 1012](#)  
[show system connections \(TX Matrix Plus Router\) on page 1013](#)  
[show system connections sfc \(TX Matrix Plus Router\) on page 1016](#)  
[show system connections show-routing-instances \(TX Matrix Plus Router\) on page 1018](#)  
[show system connections \(QFX3500 Switch\) on page 1023](#)

**Output Fields** [Table 57 on page 1009](#) describes the output fields for the **show system connections** command. Output fields are listed in the approximate order in which they appear.

**Table 57: show system connections Output Fields**

Field Name	Field Description
<b>Proto</b>	Protocol of the socket: IP, TCP, or UDP for IPv4 or IPv6.
<b>Recv-Q</b>	Number of input packets received by the protocol and waiting to be processed by the application.
<b>Send-Q</b>	Number of output packets sent by the application and waiting to be processed by the protocol.

Table 57: show system connections Output Fields (*continued*)

Field Name	Field Description
Local Address	Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
Foreign Address	Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.
Routing Instance	(Displayed only when the <b>show-routing-instance</b> option is used.) Routing instances associated with active IP sockets on the Routing Engine.
(state)	For TCP, the protocol state of the socket.

## Sample Output

### show system connections

```

user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp      0      2 192.168.4.16.513       208.197.169.254.894    ESTABLISHED
tcp      0      0 192.168.4.16.513       208.197.169.195.945    ESTABLISHED
tcp      0      0 *.23                   *.*                     LISTEN
tcp      0      0 *.22                   *.*                     LISTEN
tcp      0      0 *.513                  *.*                     LISTEN
tcp00 *.514             *.*                     LISTEN
tcp 0 0*.21                   *.*                     LISTEN
tcp00 *.79             *.*                     LISTEN
tcp 00 *.1023                *.*                     LISTEN
tcp 00 *.111                 *.*                     LISTEN
udp00192.168.4.16.1634   208.197.169.249.2049
udp00192.168.4.16.1627   208.197.169.254.2049
udp00192.168.4.16.1371   208.197.169.195.2049
udp00*.*                *.*
udp00*.9999              *.*
udp00 *.161             *.*
udp00192.168.4.16.1039   192.168.4.16.1023
udp00192.168.4.16.1038   192.168.4.16.1023
udp 00 192.168.4.16.1037     192.168.4.16.1023
udp00192.168.4.16.1036   192.168.4.16.1023
udp00*.1022              *.*
udp00*.1023              *.*
udp00*.111               *.*
udp00*.                  *.*

```

### show system connections extensive

```

user@host> show system connections extensive

Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp4      0      6 192.168.187.15.23

```

```

172.27.133.138.3013 ESTABLISHED
sndsbcc: 6 sndsbmbcnt: 256 sndsbmbmax: 272000
sndsblowat: 2048 sndsbhiwat: 34000
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 533120
rcvsblowat: 1 rcvsbhiwat: 66640
proc id: 0 proc name:
iss: 2566994072 sndup: 2566994491
snduna: 2566994491 sndnxt: 2566994494 sndwnd: 64094
sndmax: 2566994494 sndcwnd: 6589 sndsssthresh: 2720
irs: 236981199 rcvup: 236981325
rcvnxt: 236981327 rcvadv: 237046862 rcvwnd: 66640
rtt: 140058623 srtt: 15519 rttv: 908
rtxcur: 1200 rxtshift: 0 rtseq: 2566994491
rttmin: 1000 mss: 1360
flags: SACK_PERMIT [0x2000200]
tcp4 0 0 10.255.165.93.179
10.255.165.203.65141 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 0 proc name:
iss: 2555995917 sndup: 2555995917
snduna: 2555995917 sndnxt: 2555995917 sndwnd: 16384
sndmax: 2555995917 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 2123825753 rcvup: 2123860681
rcvnxt: 2123860681 rcvadv: 2123877065 rcvwnd: 16384
rtt: 0 srtt: 3309 rttv: 72
rtxcur: 1200 rxtshift: 0 rtseq: 2555995898
rttmin: 1000 mss: 500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]
tcp4 0 0 10.255.165.203.65141
10.255.165.93.179 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 5022 proc name: rpd
iss: 2123825753 sndup: 2123860662
snduna: 2123860681 sndnxt: 2123860681 sndwnd: 16384
sndmax: 2123860681 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 2555995917 rcvup: 2555995917
rcvnxt: 2555995917 rcvadv: 2556012301 rcvwnd: 16384
rtt: 0 srtt: 3279 rttv: 22
rtxcur: 1200 rxtshift: 0 rtseq: 2123860662
rttmin: 1000 mss: 500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x100003e0]
tcp4 0 0 10.255.165.203.179
10.255.165.113.52404 ESTABLISHED
sndsbcc: 0 sndsbmbcnt: 0 sndsbmbmax: 131072
sndsblowat: 2048 sndsbhiwat: 16384
rcvsbcc: 0 rcvsbmbcnt: 0 rcvsbmbmax: 131072
rcvsblowat: 1 rcvsbhiwat: 16384
proc id: 0 proc name:
iss: 1109297190 sndup: 1109332099
snduna: 1109332118 sndnxt: 1109332118 sndwnd: 16384
sndmax: 1109332118 sndcwnd: 1000 sndsssthresh: 1073725440
irs: 1476831634 rcvup: 1476866449
rcvnxt: 1476866449 rcvadv: 1476882833 rcvwnd: 16384
rtt: 0 srtt: 3235 rttv: 18
rtxcur: 1200 rxtshift: 0 rtseq: 1109332099

```

```

rttmin:      1000  mss:      500
flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP SACK_PERMIT [0x3e0]

```

### show system connections lcc (TX Matrix Router)

```
user@host> show system connections lcc 2
```

```
lcc2-re0:
```

```
-----
Active Internet connections (including servers)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	(state)
tcp4	0	0	192.168.66.131.1342	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2059	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.4571	192.168.66.130.23	ESTABLISHED
tcp4	0	0	192.168.66.131.2496	192.168.66.130.23	ESTABLISHED
tcp4	0	0	*.3221	*.*	LISTEN
tcp4	0	0	*.23	*.*	LISTEN
tcp4	0	0	*.22	*.*	LISTEN
tcp4	0	0	*.514	*.*	LISTEN
tcp4	0	0	*.513	*.*	LISTEN
tcp4	0	0	*.21	*.*	LISTEN
tcp4	0	0	*.79	*.*	LISTEN
tcp4	0	0	*.6234	*.*	LISTEN
udp4	0	0	*.514	*.*	
udp4	0	0	*.6333	*.*	

### show system connections show-routing-instances

```
user@host> show system connections show-routing-instances
```

```
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address Foreign Address Routing Instance
(state)
```

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	Routing Instance
tcp4	0	0	192.168.69.204.23	172.17.28.19.4267	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.58540	10.209.7.138.23	default
			ESTABLISHED		
tcp4	0	0	192.168.69.204.23	172.17.28.19.1098	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.57668	192.168.9.1.179	default
			ESTABLISHED		
tcp4	0	0	192.168.7.1.179	192.168.8.1.49209	default
			ESTABLISHED		
tcp4	0	0	128.0.0.1.6234	128.0.3.17.1024	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.9000	128.0.0.4.59103	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	128.0.0.4.59103	128.0.0.4.9000	
__juniper_private1__			ESTABLISHED		
tcp4	0	0	*.32012	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.9000	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.33007	*.*	
__juniper_private2__			LISTEN		
tcp46	0	0	*.179	*.*	default
			LISTEN		
tcp4	0	0	*.179	*.*	default
			LISTEN		
tcp4	0	0	*.6154	*.*	
__juniper_private1__			LISTEN		
tcp4	0	0	*.6153	*.*	

```

__juniper_private1__ LISTEN
tcp4      0      0 *.7000    *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.6152    *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.6156    *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.33005   *.*
__juniper_private2__ LISTEN
tcp4      0      0 *.31343   *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.31341   *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.32003   *.*
__juniper_private2__ LISTEN
tcp4      0      0 *.666     *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.38      *.*
__juniper_private1__ LISTEN
tcp4      0      0 *.3221    *.*      default
LISTEN

```

#### show system connections (TX Matrix Plus Router)

```

user@host> show system connections
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
          (state)
tcp4      0      3 192.168.178.11.23
172.17.28.19.3565        ESTABLISHED
tcp4      0      0 192.168.178.11.23
172.17.28.204.62719      ESTABLISHED
tcp4      0      0 192.168.178.11.23
192.168.69.199.51255     ESTABLISHED
tcp4      0      0 192.168.178.11.23
172.24.26.227.42860     ESTABLISHED
tcp4      0      0 *.6156          *.*
LISTEN
tcp4      0      0 162.0.0.4.32012
162.0.0.5.58935         ESTABLISHED
tcp4      0      0 *.32012         *.*
LISTEN
tcp4      0      0 *.33007         *.*
LISTEN
tcp4      0      0 *.666           *.*
LISTEN
tcp4      0      0 162.0.0.4.6161
162.0.0.5.62026         ESTABLISHED
tcp4      0      0 *.33005         *.*
LISTEN
tcp4      0      0 162.0.0.4.9000
162.0.0.4.51611         ESTABLISHED
tcp4      0      0 162.0.0.4.51611
162.0.0.4.9000          ESTABLISHED
tcp4      0      0 *.6151          *.*
LISTEN
tcp4      0      0 *.6154          *.*
LISTEN
tcp4      0      0 *.6153          *.*

```

```

tcp4      0      0 *.31343      LISTEN      *. *
tcp4      0      0 *.31341      LISTEN      *. *
tcp4      0      0 *.9000       LISTEN      *. *
tcp4      0      0 *.6152       LISTEN      *. *
tcp4      0      0 *.32003      LISTEN      *. *
tcp4      0      0 *.33009      LISTEN      *. *
tcp4      0      0 *.3221       LISTEN      *. *
tcp4      0      0 *.23         LISTEN      *. *
tcp4      0      0 *.22         LISTEN      *. *
tcp4      0      0 *.514        LISTEN      *. *
tcp4      0      0 *.513        LISTEN      *. *
tcp4      0      0 *.21         LISTEN      *. *
tcp4      0      0 *.79         LISTEN      *. *
tcp4      0      0 *.514        LISTEN      *. *
tcp4      0      0 *.513        LISTEN      *. *
tcp4      0      0 *.6234       LISTEN      *. *
udp4      0      0 127.0.0.1.123 LISTEN      *. *
udp4      0      0 10.255.178.11.123 LISTEN      *. *
udp4      0      0 *.123        LISTEN      *. *
udp46     0      0 *.514        LISTEN      *. *
udp4      0      0 *.514        LISTEN      *. *
udp46     0      0 *.62027      LISTEN      *. *
udp4      0      0 *.59363      LISTEN      *. *
udp4      0      0 *.31342      LISTEN      *. *
udp46     0      0 *.161        LISTEN      *. *
udp4      0      0 *.161        LISTEN      *. *
udp4      0      0 *.31340      LISTEN      *. *
udp4      0      0 *.31340      LISTEN      *. *
udp46     0      0 *.49152      LISTEN      *. *
udp46     0      0 *.4784       LISTEN      *. *
udp46     0      0 *.3784       LISTEN      *. *
udp4      0      0 *.49152      LISTEN      *. *
udp4      0      0 *.4784       LISTEN      *. *
udp4      0      0 *.3784       LISTEN      *. *
udp4      0      0 *.6333       LISTEN      *. *
ip4       0      0 *. *         LISTEN      *. *
ip4       0      0 *. *         LISTEN      *. *

```

```
lcc0-re0:
```

```
-----
Active Internet connections (including servers)
```

```
Proto Recv-Q Send-Q Local Address
```

```
Foreign Address
```

```
(state)
```

```
tcp4      0      0 192.168.178.3.23
```



```

172.24.26.227.50399
tcp4      0      0 *.6234          ESTABLISHED      *.*
          LISTEN
tcp4      0      0 *.7000          *.*
          LISTEN
tcp4      0      0 *.9000          *.*
          LISTEN
tcp4      0      0 *.33009         *.*
          LISTEN
tcp4      0      0 *.3221          *.*
          LISTEN
tcp4      0      0 *.23            *.*
          LISTEN
tcp4      0      0 *.22            *.*
          LISTEN
tcp4      0      0 *.514           *.*
          LISTEN
tcp4      0      0 *.513           *.*
          LISTEN
tcp4      0      0 *.21            *.*
          LISTEN
tcp4      0      0 *.79            *.*
          LISTEN
tcp4      0      0 *.514           *.*
          LISTEN
tcp4      0      0 *.513           *.*
          LISTEN
udp46     0      0 *.514           *.*
udp4      0      0 *.514           *.*
udp46     0      0 *.59924         *.*
udp4      0      0 *.59412         *.*
udp46     0      0 *.161           *.*
udp4      0      0 *.161           *.*
udp4      0      0 *.31342         *.*
udp4      0      0 *.6333          *.*

```

```

lcc1-re0:
-----

```

```

Active Internet connections (including servers)

```

Proto	Recv-Q	Send-Q	Local Address (state)	Foreign Address
tcp4	0	0	*.6234 LISTEN	*.*
tcp4	0	0	*.7000 LISTEN	*.*
tcp4	0	0	*.9000 LISTEN	*.*
tcp4	0	0	*.3221 LISTEN	*.*
tcp4	0	0	*.23 LISTEN	*.*
tcp4	0	0	*.22 LISTEN	*.*
tcp4	0	0	*.514 LISTEN	*.*
tcp4	0	0	*.513 LISTEN	*.*
tcp4	0	0	*.21 LISTEN	*.*
tcp4	0	0	*.79 LISTEN	*.*

```

tcp4      0      0 *.514          *.*
          LISTEN
tcp4      0      0 *.513          *.*
          LISTEN
tcp4      0      0 *.33009        *.*
          LISTEN
udp46     0      0 *.514          *.*
udp4      0      0 *.514          *.*
udp46     0      0 *.59924        *.*
udp4      0      0 *.59412        *.*
udp4      0      0 *.31342        *.*
udp46     0      0 *.161          *.*
udp4      0      0 *.161          *.*
udp4      0      0 *.6333         *.*

```

lcc2-re0:

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 *.6234            *.*
          LISTEN
tcp4      0      0 *.7000            *.*
          LISTEN
tcp4      0      0 *.9000            *.*
          LISTEN
tcp4      0      0 *.33009           *.*
          LISTEN
tcp4      0      0 *.3221            *.*
          LISTEN
tcp4      0      0 *.23              *.*
          LISTEN
tcp4      0      0 *.22              *.*
          LISTEN
tcp4      0      0 *.514             *.*
...

```

#### show system connections sfc (TX Matrix Plus Router)

```

user@host> show system connections sfc 0
sfc0-re0:

```

```

-----
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
      (state)
tcp4      0      0 162.0.0.4.514        132.0.0.4.952
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        131.0.0.4.694
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        130.0.0.4.860
          TIME_WAIT
tcp4      0      0 162.0.0.4.514        129.0.0.4.716
          TIME_WAIT
tcp4      0      0 162.0.0.4.996        132.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.798        131.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.995        130.0.0.4.514
          TIME_WAIT
tcp4      0      0 162.0.0.4.895        129.0.0.4.514
          TIME_WAIT

```

tcp4	0	0	192.168.178.11.21		
172.17.28.204.64662				TIME_WAIT	
tcp4	0	0	192.168.178.11.21		
172.17.28.204.51612				TIME_WAIT	
tcp4	0	0	*,6156		*,*
			LISTEN		
tcp4	0	0	*,9000		*,*
			LISTEN		
tcp4	0	0	*,666		*,*
			LISTEN		
tcp4	0	2	192.168.178.11.23		
172.17.28.19.3565				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.17.28.204.62719				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
192.168.69.199.51255				ESTABLISHED	
tcp4	0	0	192.168.178.11.23		
172.24.26.227.42860				ESTABLISHED	
tcp4	0	0	162.0.0.4.32012		162.0.0.5.58935
			ESTABLISHED		
tcp4	0	0	*,32012		*,*
			LISTEN		
tcp4	0	0	*,33007		*,*
			LISTEN		
tcp4	0	1432	162.0.0.4.6161		162.0.0.5.62026
			ESTABLISHED		
tcp4	0	0	*,33005		*,*
			LISTEN		
tcp4	0	0	162.0.0.4.9000		162.0.0.4.51611
			FIN_WAIT_2		
tcp4	0	0	162.0.0.4.51611		162.0.0.4.9000
			CLOSE_WAIT		
tcp4	0	0	*,6151		*,*
			LISTEN		
tcp4	0	0	*,6154		*,*
			LISTEN		
tcp4	0	0	*,6153		*,*
			LISTEN		
tcp4	0	0	*,31343		*,*
			LISTEN		
tcp4	0	0	*,31341		*,*
			LISTEN		
tcp4	0	0	*,6152		*,*
			LISTEN		
tcp4	0	0	*,32003		*,*
			LISTEN		
tcp4	0	0	*,33009		*,*
			LISTEN		
tcp4	0	0	*,3221		*,*
			LISTEN		
tcp4	0	0	*,23		*,*
			LISTEN		
tcp4	0	0	*,22		*,*
			LISTEN		
tcp4	0	0	*,514		*,*
			LISTEN		
tcp4	0	0	*,513		*,*
			LISTEN		
tcp4	0	0	*,21		*,*
			LISTEN		
tcp4	0	0	*,79		*,*

```

                                LISTEN
tcp4      0      0 *.514                                *. *
                                LISTEN
tcp4      0      0 *.513                                *. *
                                LISTEN
tcp4      0      0 *.6234                               *. *
                                LISTEN
udp4      0      0 127.0.0.1.123                       *. *
udp4      0      0 10.255.178.11.123                   *. *
udp4      0      0 *.123                                *. *
udp46     0      0 *.514                                *. *
udp4      0      0 *.514                                *. *
udp46     0      0 *.50895                              *. *
udp4      0      0 *.50794                              *. *
udp4      0      0 *.31342                              *. *
udp46     0      0 *.161                                *. *
udp4      0      0 *.161                                *. *
udp4      0      0 *.31340                              *. *
udp4      0      0 *.31340                              *. *
udp46     0      0 *.49152                              *. *
udp46     0      0 *.4784                               *. *
udp46     0      0 *.3784                               *. *
udp4      0      0 *.49152                              *. *
udp4      0      0 *.4784                               *. *
udp4      0      0 *.3784                               *. *
udp4      0      0 *.6333                               *. *
ip4       104    0 *. *                                *. *
ip4       0      0 *. *                                *. *
ip4       0      0 *. *                                *. *

```

#### show system connections show-routing-instances (TX Matrix Plus Router)

```

user@host> show system connections show-routing-instances
sfc0-re0:
-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Foreign Address
-----
                                Routing Instance      (state)
tcp4      0      0 *.6156                   __juniper_private1__    LISTEN      *. *
tcp4      0      0 *.9000                   __juniper_private1__    LISTEN      *. *
tcp4      0      0 *.666                    __juniper_private1__    LISTEN      *. *
tcp4      0      2 192.168.178.11.23        default                  ESTABLISHED 172.17.28.19.3565
tcp4      0      0 192.168.178.11.23        default                  ESTABLISHED 172.17.28.204.62719
tcp4      0      0 192.168.178.11.23        default                  ESTABLISHED 192.168.69.199.51255
tcp4      0      0 192.168.178.11.23        default                  ESTABLISHED 172.24.26.227.42860
tcp4      0      0 162.0.0.4.32012          __juniper_private1__    ESTABLISHED 162.0.0.5.58935
tcp4      0      0 *.32012                  __juniper_private1__    LISTEN      *. *
tcp4      0      0 *.33007                  __juniper_private2__    LISTEN      *. *
tcp4      0      0 162.0.0.4.6161          __juniper_private1__    ESTABLISHED 162.0.0.5.62026
tcp4      0      0 *.33005                  *. *

```

tcp4	0	0	162.0.0.4.9000	__juniper_private2__	LISTEN	162.0.0.4.51611
tcp4	0	0	162.0.0.4.51611	__juniper_private1__	FIN_WAIT_2	162.0.0.4.9000
tcp4	0	0	*.6151	__juniper_private1__	CLOSE_WAIT	*.*
tcp4	0	0	*.6154	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6153	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.31343	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.31341	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6152	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.32003	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
udp4	0	0	127.0.0.1.123	__juniper_private1__	LISTEN	*.*
udp4	0	0	10.255.178.11.123	default		*.*
udp4	0	0	*.123	default		*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514	default		*.*
udp46	0	0	*.50895	default		*.*
udp4	0	0	*.50794	default		*.*
udp4	0	0	*.31342	default		*.*
udp46	0	0	*.161	__juniper_private1__		*.*
udp4	0	0	*.161	default		*.*
				default		

udp4	0	0	*.31340	__juniper_private2__	*.*
udp4	0	0	*.31340	__juniper_private1__	*.*
udp46	0	0	*.49152	default	*.*
udp46	0	0	*.4784	default	*.*
udp46	0	0	*.3784	default	*.*
udp4	0	0	*.49152	default	*.*
udp4	0	0	*.4784	default	*.*
udp4	0	0	*.3784	default	*.*
udp4	0	0	*.6333	__juniper_private1__	*.*
ip4	0	0	*.*	default	*.*
ip4	0	0	*.*	default	*.*
ip4	0	0	*.*	default	*.*

lcc0-re0:

-----

Active Internet connections (including servers) (including routing-instances)

Proto	Recv-Q	Send-Q	Local Address	Routing Instance	(state)	Foreign Address
tcp4	0	0	*.7000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	192.168.178.3.23	default	ESTABLISHED	
tcp4	0	0	*.6234	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.9000	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.33009	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.3221	default	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514			*.*

```

udp46      0      0 *.59924    default      *.
udp4        0      0 *.59412    default      *.
udp46      0      0 *.161      default      *.
udp4        0      0 *.161      default      *.
udp4        0      0 *.31342    default      *.
udp4        0      0 *.6333     __juniper_private1__
   *.
   __juniper_private1__

```

```
lcc1-re0:
```

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000          __juniper_private1__ LISTEN          *.
tcp4      0      0 *.6234          __juniper_private1__ LISTEN          *.
tcp4      0      0 *.9000          __juniper_private1__ LISTEN          *.
tcp4      0      0 *.3221          default             LISTEN          *.
tcp4      0      0 *.23            default             LISTEN          *.
tcp4      0      0 *.22            default             LISTEN          *.
tcp4      0      0 *.514           default             LISTEN          *.
tcp4      0      0 *.513           default             LISTEN          *.
tcp4      0      0 *.21            default             LISTEN          *.
tcp4      0      0 *.79            default             LISTEN          *.
tcp4      0      0 *.514           __juniper_private1__ LISTEN          *.
tcp4      0      0 *.513           __juniper_private1__ LISTEN          *.
tcp4      0      0 *.33009         __juniper_private2__ LISTEN          *.
udp46     0      0 *.514           default             *.
udp4       0      0 *.514           default             *.
udp46     0      0 *.59924         default             *.
udp4       0      0 *.59412         default             *.
udp4       0      0 *.31342         default             *.
udp46     0      0 *.161           __juniper_private1__ *.
udp4       0      0 *.161           default             *.
udp4       0      0 *.6333          default             *.
   __juniper_private1__

```

lcc2-re0:

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.6234                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.9000                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.33009                 __juniper_private2__ LISTEN         *.
tcp4      0      0 *.3221                  default             LISTEN         *.
tcp4      0      0 *.23                    default             LISTEN         *.
tcp4      0      0 *.22                    default             LISTEN         *.
tcp4      0      0 *.514                   default             LISTEN         *.
tcp4      0      0 *.513                   default             LISTEN         *.
tcp4      0      0 *.21                    default             LISTEN         *.
tcp4      0      0 *.79                    default             LISTEN         *.
tcp4      0      0 *.514                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.513                  __juniper_private1__ LISTEN         *.
udp46     0      0 *.514                   default             *.
udp4      0      0 *.514                   default             *.
udp4      0      0 *.31342                 __juniper_private1__ *.
udp46     0      0 *.62103                 default             *.
udp4      0      0 *.59924                 default             *.
udp46     0      0 *.161                   default             *.
udp4      0      0 *.161                   default             *.
udp4      0      0 *.6333                  __juniper_private1__

```

lcc3-re0:

```

-----
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q Local Address           Routing Instance      (state)      Foreign Address
tcp4      0      0 *.7000                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.6234                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.9000                  __juniper_private1__ LISTEN         *.
tcp4      0      0 *.33009                 __juniper_private1__ LISTEN         *.

```



tcp4	0	0	*.3221	__juniper_private2__	LISTEN	*.*
tcp4	0	0	*.23	default	LISTEN	*.*
tcp4	0	0	*.22	default	LISTEN	*.*
tcp4	0	0	*.514	default	LISTEN	*.*
tcp4	0	0	*.513	default	LISTEN	*.*
tcp4	0	0	*.21	default	LISTEN	*.*
tcp4	0	0	*.79	default	LISTEN	*.*
tcp4	0	0	*.514	__juniper_private1__	LISTEN	*.*
tcp4	0	0	*.513	__juniper_private1__	LISTEN	*.*
udp46	0	0	*.514	default		*.*
udp4	0	0	*.514	default		*.*
udp46	0	0	*.62103	default		*.*
udp4	0	0	*.59924	default		*.*
udp4	0	0	*.31342	__juniper_private1__		*.*
udp46	0	0	*.161	default		*.*
udp4	0	0	*.161	default		*.*
udp4	0	0	*.6333	__juniper_private1__		*.*

### show system connections (QFX3500 Switch)

```

user@switch> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address           Foreign Address
              (state)
tcp4          0      0 10.94.204.110.23        172.17.28.19.1308    ESTABLISHED
tcp4          0      0 128.0.0.1.6234          128.0.0.1.65142     ESTABLISHED
tcp4          0      0 128.0.0.1.65142         128.0.0.1.6234     ESTABLISHED
tcp4          0      0 128.0.0.1.33003         128.0.0.1.61441     ESTABLISHED
tcp4          0      0 128.0.0.1.61441         128.0.0.1.33003     ESTABLISHED
tcp46         0      0 *.179                   *.*                  LISTEN
tcp4          0      0 *.179                   *.*                  LISTEN
tcp4          0      0 128.0.0.16.9000         128.0.0.16.50970    ESTABLISHED
tcp4          0      0 128.0.0.16.50970        128.0.0.16.9000     ESTABLISHED
tcp4          0      0 *.38                    *.*                  LISTEN

```

			LISTEN	
tcp4	0	0	*.3491	*.*
			LISTEN	
tcp4	0	0	*.6156	*.*
			LISTEN	
tcp4	0	0	128.0.0.1.33001	128.0.0.1.59437
			ESTABLISHED	
tcp4	0	0	128.0.0.1.59437	128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0	128.0.0.1.33023	128.0.0.1.63605
			ESTABLISHED	
tcp4	0	0	128.0.0.1.63605	128.0.0.1.33023
			ESTABLISHED	
tcp4	0	0	128.0.0.1.33001	128.0.0.1.63830
			ESTABLISHED	
tcp4	0	0	128.0.0.1.63830	128.0.0.1.33001
			ESTABLISHED	
tcp4	0	0	*.667	*.*
			LISTEN	
tcp4	0	0	*.6156	*.*
			LISTEN	
tcp4	0	0	128.0.0.1.7000	128.0.0.1.51580
			ESTABLISHED	
tcp4	0	0	128.0.0.1.51580	128.0.0.1.7000
			ESTABLISHED	
tcp4	0	0	128.0.0.1.6234	128.0.0.1.53646
			ESTABLISHED	
tcp4	0	0	*.33001	*.*
			LISTEN	
tcp4	0	0	*.33003	*.*
			LISTEN	
tcp4	0	0	128.0.0.1.53646	128.0.0.1.6234
			ESTABLISHED	
tcp4	0	0	128.0.0.16.9000	128.0.0.16.63454
			ESTABLISHED	
tcp4	0	0	128.0.0.16.63454	128.0.0.16.9000
			ESTABLISHED	
tcp4	0	0	*.666	*.*
			LISTEN	
tcp4	0	0	*.7000	*.*
			LISTEN	
tcp4	0	0	*.51627	*.*
			LISTEN	
tcp4	0	0	*.3492	*.*
			LISTEN	
tcp4	0	0	*.33023	*.*
			LISTEN	
tcp4	0	0	*.33013	*.*
			LISTEN	
tcp4	0	0	*.7202	*.*
			LISTEN	
tcp4	0	0	*.6151	*.*
			LISTEN	
tcp4	0	0	*.9000	*.*
			LISTEN	
tcp4	0	0	*.6161	*.*
			LISTEN	
tcp4	0	0	*.6011	*.*
			LISTEN	
tcp4	0	0	*.3221	*.*
			LISTEN	

tcp4	0	0 *.23		*. *
			LISTEN	
tcp4	0	0 *.22		*. *
			LISTEN	
tcp4	0	0 *.514		*. *
			LISTEN	
tcp4	0	0 *.513		*. *
			LISTEN	
tcp4	0	0 *.21		*. *
			LISTEN	
tcp4	0	0 *.79		*. *
			LISTEN	
tcp4	0	0 *.514		*. *
			LISTEN	
tcp4	0	0 *.513		*. *
			LISTEN	
tcp4	0	0 *.1127		*. *
			LISTEN	
tcp4	0	0 *.1129		*. *
			LISTEN	
tcp4	0	0 *.1128		*. *
			LISTEN	
tcp4	0	0 *.6234		*. *
			LISTEN	
udp46	0	0 *.514		*. *
udp4	0	0 *.514		*. *
udp4	0	0 128.0.0.1.123		*. *
udp46	0	0 *.53344		*. *
udp4	0	0 *.54261		*. *
udp46	0	0 *.161		*. *
udp4	0	0 *.161		*. *
udp4	0	0 *.31342		*. *
udp4	0	0 *.59137		*. *
udp4	0	0 *. *		*. *
udp46	0	0 *.49152		*. *
udp46	0	0 *.4784		*. *
udp46	0	0 *.3784		*. *
udp4	0	0 *.49152		*. *
udp4	0	0 *.4784		*. *
udp4	0	0 *.3784		*. *
udp4	0	0 10.255.204.110.123		*. *
udp4	0	0 *.123		*. *
udp4	0	0 *.67		*. *
udp4	0	0 *.6333		*. *
udp4	0	0 *.2293		*. *
ip4	0	0 *. *		*. *
ip4	0	0 *. *		*. *
ip4	0	0 *. *		*. *

## show system core-dumps

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1026</a> <a href="#">Syntax (EX Series Switches) on page 1026</a> <a href="#">Syntax (TX Matrix Router) on page 1026</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1026</a> <a href="#">Syntax (QFX Series) on page 1026</a>
<b>Syntax</b>	<code>show system core-dumps</code> <code>&lt;brief   detail&gt;</code> <code>&lt;core-filename&gt;</code> <code>&lt;core-file-info&gt;</code> <code>&lt;re0&gt;</code> <code>&lt;re1&gt;</code> <code>&lt;routing-engine&gt;</code>
<b>Syntax (EX Series Switches)</b>	<code>show system core-dumps</code> <code>&lt;all-members&gt;</code> <code>&lt;brief   detail&gt;</code> <code>&lt;core-filename&gt;</code> <code>&lt;core-file-info&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>show system core-dumps</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;brief   detail&gt;</code> <code>&lt;core-filename&gt;</code> <code>&lt;core-file-info&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>show system core-dumps</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;brief   detail&gt;</code> <code>&lt;core-filename&gt;</code> <code>&lt;core-file-info&gt;</code>
<b>Syntax (QFX Series)</b>	<code>show system core-dumps</code> <code>&lt;brief   detail&gt;</code> <code>&lt;component (<i>UUID</i>   <i>serial number</i>   all)&gt;</code> <code>&lt;core-file-info component (<i>UUID</i>   <i>serial number</i>) <i>core-file-name</i>&gt;</code> <code>&lt;display-period (<i>hours</i>   <i>minutes</i>   <i>seconds</i>)&gt;</code> <code>&lt;display-order&gt;</code> <code>&lt;kernel-crashinfo component (<i>UUID</i>   <i>serial number</i>)&gt;</code> <code>&lt;repository (core   log)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series. <b>re0</b> , <b>re1</b> , and <b>routing-engine</b> options introduced for dual routing engines in Junos OS Release 13.1.

**Description** Show core files on all routers or switches running Junos OS. You can use the **show system core-dumps** command to show a list of system core files created when the router or switch has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and path and filename. If dual routing engines are present, you can view core-dump files for either routing engine or both routing engines together. On a QFabric system, you can view core-dump files on individual QFabric system devices as well as on the entire QFabric system.

You can use the option **core-filename** and its options **core-file-info**, **brief**, and **detail** to display more information about the specified core-dump files.

**Options** **none**—Display a list of all existing core-dump files.



**NOTE:** If dual routing engines are present, lists only the core-dump files for the active routing engine.

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on a TX Matrix router, display system core files for the TX Matrix router switch-card chassis [SCC] and all the T640 routers [LCCs] connected to the TX Matrix router.

On a routing matrix based on a TX Matrix Plus router, display system core files for the TX Matrix Plus router (switch-fabric chassis [SFC]) and all the T1600 routers [LCCs] connected to the TX Matrix Plus router.

**<all-lcc | lcc number>**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a routing matrix based on the TX Matrix router, display core dump files for all T640 routers (line-card chassis [LCCs]) or a specific T640 router [LCC] connected to the TX Matrix router.

On a routing matrix based on the TX Matrix Plus router, display logging information for all T1600 routers (line-card chassis [LCCs]) or a specific T1600 router (LCC) connected to the TX Matrix Plus router. When using the **lcc number** option, replace **number** with a value from 0 through 3.



**NOTE:** The **all-chassis** option displays system core files for the SCC or SFC and the LCCs connected to the SCC or SFC in the routing matrix while the **all-lcc** option only displays system core files for the LCCs in the routing matrix.

**all-members**—(EX4200 switches) (Optional) Display system core files on all members of the Virtual Chassis configuration.

**brief**—(Optional) View details of a binary file.

**component** (*UUID | serial number | all*)—(QFabric systems only) (Optional) Display a list of core-dump files located on individual QFabric system device or on the entire QFabric system.

**core-file-info**—(Optional) Display the stack trace of a core file.

**core-filename**—(Optional) Name of a specific core file to display.

**detail**—(Optional) View stack trace with details of the binary file.

**display-order** (*timestamp-sort | alphanumeric-sort*)—(QFabric systems only) (Optional) Display list of debug artifacts generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds—or according to their filename.

**display-period** (*hours | minutes | seconds*)—(QFabric systems only) (Optional) Display core-dump files generated within the specified period—for example, within the last hour, within the last 20 minutes, or within the last 32 seconds.

**kernel-crashinfo component** (*UUID | serial number*)—(QFabric systems only) (Optional) Display kernel crash information from the EEPROM on a QFabric system device.

**local**—(EX4200 switches only) (Optional) Display system core files on the local Virtual Chassis member.

**member** *member-id*—(EX4200 switches only) (Optional) Display system core files on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**re0**—(Dual routing engines only) Display the core-dump files on re0.

**re1**—(Dual routing engines only) Display the core-dump files on re1.

**repository** (*core | log*)—(QFabric systems only) (Optional) Specify either the core or log repository in which to view core-dump files.

**routing-engine** (*backup | both | local | master | other*)—(Dual routing engines only) Display a list of core-dump files for either the backup, local, master, or other routing engine or both routing engines.

**scc**—(TX Matrix routers only) (Optional) Display system core files on the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display system core files on the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level**

view

**List of Sample Output**

[show system core-dumps on page 1030](#)

[show system core-dumps on page 1031](#)

[show system core-dumps routing-engine both on page 1031](#)

[show system core-dumps \(TX Matrix Plus Router\) on page 1031](#)

[show system core-dumps \(QFX3500 Switch\) on page 1033](#)  
[show system core-dumps \(QFabric Systems\) on page 1033](#)  
[show system core-dumps core-file-info component serial number core-file-name \(QFabric Systems\) on page 1034](#)  
[show system core-dumps component serial number display-order alphanumeric-sort repository core \(QFabric Systems\) on page 1034](#)  
[show system core-dumps display-period \(QFabric Systems\) on page 1034](#)  
[show system core-dumps kernel-crashinfo component serial number \(QFabric Systems\) on page 1036](#)  
[show system core-dumps repository core \(QFabric Systems\) on page 1038](#)  
[show system core-dumps repository log \(QFabric Systems\) on page 1038](#)

**Output Fields** Table 58 on page 1029 describes the output fields for the **show system core-dumps** command. Output fields are listed in the approximate order in which they appear.

**Table 58: show system core-dumps Output Fields**

Field Name	Field Description
<i>Permissions</i>	Read/write permissions for the file named.
<i>Links</i>	Number of links to the file.
<i>Owner</i>	Name of the file owner.
<i>Group</i>	Name of the group with file access.
<i>File size</i>	File size in bytes.
<i>Modified</i>	Last file modification date and time.
<i>Path/filename</i>	File path where the file resides and the filename.
<b>Repository scope:</b>	Repository where core-dump files and log files are stored. The core-dump files are located in the <b>core</b> repository, and the log files are located in the <b>log</b> repository. The default <b>Repository scope</b> is shared since both the <b>core</b> and <b>log</b> repositories are shared by all of the QFabric system devices.
<b>Repository head:</b>	Path to the top-level repository location.
<b>Repository name:</b>	Name of the repository: <b>core</b> or <b>log</b> .
<b>List of nodes for core repository:</b>	List of core-dump files associated with a particular QFabric system device located in the core repository.
<b>Node Group</b>	Name of the QFabric system device.
<b>Node Identifier</b>	UUID or serial number of the QFabric system device.
<b>Num</b>	Number of core-dump and log files.

Table 58: show system core-dumps Output Fields (*continued*)

Field Name	Field Description
<b>Model</b>	Model number of the QFabric system device.
<b>Usage</b>	Usage of the repository in megabytes.
<b>Total usage of core repository:</b>	Total usage of core-dump files associated with a particular QFabric system device located in the core repository. Usage is specified in megabytes and as a percentage.
<b>Total usage of log repository:</b>	Total usage of log files associated with a particular QFabric system device located in the log repository. Usage is specified in megabytes and as a percentage.
<b>List of nodes for core repository:</b>	List of core-dump files associated with a particular QFabric system device located in the core repository.
<b>List of nodes for log repository:</b>	List of log files associated with a particular QFabric system device located in the log repository.
<b>Filename</b>	Name of the core-dump file.
<b>Date</b>	Last core-dump file modification date and time.
<b>Size</b>	Size of the core-dump file.
<b>Core filename</b>	Filename of the core-dump file.
<b>Process name</b>	Name of the process that is generating a core-dump file or log file.
<b>Release</b>	Junos OS release.
<b>Build server</b>	Junos OS build server.
<b>Build date</b>	Junos OS build date.
<b>Stack trace</b>	Stack trace of the core-dump file.

## Sample Output

### show system core-dumps

This example shows the command output if core files exist.

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



### show system core-dumps

This example shows the command output if core files do not exist.

```
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
/var/tmp/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
```

### show system core-dumps routing-engine both

This example shows the command output if dual routing engines are present.

```
user@host> show system core-dumps routing-engine both
re0:
-----
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 496776
-rw-rw---- 1 root field 11910589 Nov 8 13:20 chassisd.core.0.201311081320
...

-rw-rw---- 1 root field 11737227 Oct 28 14:21
rpd.core-tarball.4.tgz.201310281421.3458162
total files: 10

re1:
-----
/var/crash/*core*: No such file or directory
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory

/var/tmp/cores:
total blocks: 3178420
-rw-rw---- 1 root field 19039721 Nov 8 14:29
chassisd.core.0.201311081429.3485600.gz
-rw-rw---- 1 root field 19039793 Nov 8 14:37
chassisd.core.1.201311081437.3485599.gz
..

-rw-rw---- 1 root field 11710113 Oct 17 15:26
rpd.core-tarball.1.1.tgz.201310171526.3430028
```

### show system core-dumps (TX Matrix Plus Router)

```
user@host> show system core-dumps
sfc0-re0:
-----
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory

/var/crash/cores:
total 8

/var/tmp/cores:
total 1627592
-rw-r--r-- 1 root field 535346090 May 15 07:36
rpd.core-tarball.0.090515.0736.tgz
```

```
-rw-r--r-- 1 root field 105632057 May 15 07:37
rpd.core-tarball.1.090515.0737.tgz
-rw-r--r-- 1 root field 101981681 May 15 07:38
rpd.core-tarball.2.090515.0738.tgz
-rw-r--r-- 1 root field 85854573 May 15 07:40
rpd.core-tarball.3.090515.0740.tgz
-rw-r--r-- 1 root field 4157845 May 15 08:18
rpd.core-tarball.4.090515.0818.tgz
```

lcc0-re0:

-----  
/var/crash/kernel.\*: No such file or directory  
/tftpboot/corefiles/\*core\*: No such file or directory

/var/crash/cores:  
total 8

/var/tmp/cores:  
total 12

lcc1-re0:

-----  
/var/crash/kernel.\*: No such file or directory  
/tftpboot/corefiles/\*core\*: No such file or directory

/var/crash/cores:  
total 8

/var/tmp/cores:  
total 10024

```
-rw-r--r-- 1 root field 1875794 Apr 22 15:47
chassisd.core-tarball.0.090422.1547.tgz
-rw-r--r-- 1 root field 1894183 Apr 22 19:02
chassisd.core-tarball.0.090422.1902.tgz
-rw-r--r-- 1 root field 1290240 Apr 26 16:01 ksyncd_1558.core.0.090426.1601
```

lcc2-re0:

-----  
/var/crash/kernel.\*: No such file or directory  
/tftpboot/corefiles/\*core\*: No such file or directory

/var/crash/cores:  
total 21124008

```
-rw-r--r-- 1 root wheel 1022376528 May 2 06:43
core-LCC2-EGFPC7.core.0.090502.0643
-rw-r--r-- 1 root wheel 1022376528 May 2 08:13
core-LCC2-EGFPC7.core.0.090502.0813
-rw-r--r-- 1 root wheel 1022376544 May 5 06:15
core-LCC2-EGFPC7.core.0.090505.0615
-rw-r--r-- 1 root wheel 1022376544 May 6 10:59
core-LCC2-EGFPC7.core.0.090506.1059
-rw-r--r-- 1 root wheel 1022376528 May 2 06:58
core-LCC2-EGFPC7.core.1.090502.0658
-rw-r--r-- 1 root wheel 754271232 May 5 06:33
core-LCC2-EGFPC7.core.1.090505.0633
-rw-r--r-- 1 root wheel 264897536 May 6 11:12
core-LCC2-EGFPC7.core.1.090506.1112
-rw-r--r-- 1 root wheel 1022376528 May 2 07:22
core-LCC2-EGFPC7.core.2.090502.0722
-rw-r--r-- 1 root wheel 163633152 May 5 06:52
core-LCC2-EGFPC7.core.2.090505.0652
```

```

-rw-r--r-- 1 root wheel 171312128 May 6 12:13
core-LCC2-EGFPC7.core.2.090506.1213
-rw-r--r-- 1 root wheel 1022376528 May 2 07:39
core-LCC2-EGFPC7.core.3.090502.0739
-rw-r--r-- 1 root wheel 1022376528 May 2 07:55
core-LCC2-EGFPC7.core.4.090502.0755
-rw-r--r-- 1 root wheel 427277312 May 7 04:47
core-LCC2-STFPC4.core.0.090507.0447
-rw-r--r-- 1 root wheel 419609600 May 7 04:47
core-LCC2-STFPC5.core.0.090507.0447
-rw-r--r-- 1 root wheel 432356352 May 7 04:47
core-LCC2-STFPC6.core.0.090507.0447

/var/tmp/cores:
total 2568
-rw-r--r-- 1 root field 1290240 May 14 14:26 ksyncd_1540.core.0.090514.1426
...

```

### show system core-dumps (QFX3500 Switch)

```

user@switch> show system core-dumps
/var/crash/*core*: No such file or directory
-rw-rw---- 1 root field 1545143 Jun 4 2012 /var/tmp/pafxpc.core.0.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.1.gz
-rw-rw---- 1 root field 1545141 Jun 4 2012 /var/tmp/pafxpc.core.2.gz
-rw-rw---- 1 root field 1545146 Jun 4 2012 /var/tmp/pafxpc.core.3.gz
-rw-rw---- 1 root field 1545142 Jun 5 2012 /var/tmp/pafxpc.core.4.gz
/var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
/tftpboot/corefiles/*core*: No such file or directory
total 5

```

### show system core-dumps (QFabric Systems)

```

user@switch> show system core-dumps
Repository scope: shared
Repository head: /pbdata/export
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	OM
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	OM
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	OM
NW-NG-0	BBAK0394	0	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	OM
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	OM
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	OM
IC-WS001	WS001/YN5999	0	qfxc08-3008	OM
node-device1	BBAK0372	0	qfx3500	OM
node-device1	EE3093	0	qfx3500	OM

Total usage of core repository: 0M of 70000M (0.0%)

```

List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	OM
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	OM
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	OM

NW-NG-0	BBAK0394	1	qfx3500	OM
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	OM
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	OM
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	OM
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	OM
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	OM
IC-WS001	WS001/YW3803	1	qfxc08-3008	OM
node-device1	BBAK0372	1	qfx3500	OM
node-device1	EE3093	1	qfx3500	OM

Total usage of log repository: 0M of 70000M (0.0%)

### show system core-dumps core-file-info component serial number core-file-name (QFabric Systems)

```

user@switch> show system core-dumps core-file-info component
e8ff4b3e-7d92-11e0-be5d-00e081c1fe0e cosd.core.0.1519.05162011131846.gz
Repository scope: shared
Repository head: /pbstorage
Repository name: core
Core filename: /pbstorage/r dumps/e8ff4b3e-7d92-11e0-be5d-
00e081c1fe0e/5658.cosd.core.0.1519.05162011131846
Process name: cosd
Release: 11.3I0
Build server: /c/ssengupta/dfx_ha_v1/obj-i386-dcp/dcp/usr.sbin/cosd
Build date: 2011-05-14 01:11:44 UTC
Stack trace:
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#0 0x8885d183 in select () from /usr/lib/libc.so.6
#1 0x887d4a45 in pselect () from /usr/lib/libc.so.6
#2 0x88774719 in pselect () from /usr/lib/libthr.so.2
#3 0x885de5db in __evGetNext () from /usr/lib/libisc.so.2
#4 0x885debf0 in __evMainLoop () from /usr/lib/libisc.so.2
#5 0x081125b2 in cosd_loop ()
#6 0x0812e19a in main ()

```

### show system core-dumps component serial number display-order alphanumeric-sort repository core (QFabric Systems)

```

user@switch> show system core-dumps component BBAK8891 display-order alphanumeric-sort
repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of core dumps for component BBAK8891
Repository location: /pbdata/export/r dumps/BBAK8891

```

Filename	Date	Size
eswd.core.0.1361.11172011214257.gz	Nov 17 21:43:10 2011	4779553
eswd.core.1.80267.11172011214514.gz	Nov 17 21:45:19 2011	3541648
eswd.core.2.80682.11172011214535.gz	Nov 17 21:45:43 2011	2156683
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375617

Number of core dumps in repository: 4

### show system core-dumps display-period (QFabric Systems)

```

user@switch> show system core-dumps display-period 24h
show system core-dumps display-period 24h
Repository scope: shared
Repository head: /pbdata/export
List of core dumps at repository: /pbdata/export/r dumps
Delta timespec: Last 24h

```

Component: BBAK8273		
Filename	Size	Date
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375794
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990		
Filename	Size	Date
vccpd.core.0.1461.11182011151131.gz	Nov 18 15:11:31 2011	120951
Component: ee19c4f8-0025-11e1-aef6-00e081c52990		
Filename	Size	Date
vccpd.core.0.1462.11182011151131.gz	Nov 18 15:11:31 2011	109420
Component: BBAK8281		
Filename	Size	Date
vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:36 2011	375373
Component: BBAK8891		
Filename	Size	Date
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:35 2011	375617
Component: BBAK8276		
Filename	Size	Date
vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:35 2011	375350
Component: BBAK8868		
Filename	Size	Date
vccpd.core.0.1196.11182011151130.gz	Nov 18 15:11:34 2011	376211
Component: BBAK8835		
Filename	Size	Date
vccpd.core.0.1195.11182011151130.gz	Nov 18 15:11:35 2011	375700
Component: BBAK8283		
Filename	Size	Date
vccpd.core.0.1195.11182011151131.gz	Nov 18 15:11:36 2011	368298
Component: YW3781/YW3781		
Filename	Size	Date
vccpd.core.0.1220.11182011151131.gz	Nov 18 15:11:38 2011	380002
Component: 09726be2-0026-11e1-82d9-00e081c52990		
Filename	Size	Date
vccpd.core.0.1461.11182011151130.gz	Nov 18 15:11:31 2011	119965
Component: BBAK8309		
Filename	Size	Date
vccpd.core.0.1196.11182011151131.gz	Nov 18 15:11:36 2011	378930
Component: 303d476a-0026-11e1-abf4-00e081c52990		
Filename	Size	Date
vccpd.core.0.1460.11182011151131.gz	Nov 18 15:11:31 2011	118385
Component: YW3798/YW3798		
Filename	Size	Date
vccpd.core.0.1219.11182011151131.gz	Nov 18 15:11:36 2011	380455
List of log dumps at repository: /pbdata/export/rlogs		
Delta timespec: Last 24h		
Component: BBAK8273		
Filename	Size	Date

vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:39 2011	20415
Component: cedb7b0e-0025-11e1-9a5f-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1461.11182011151131.tgz	Nov 18 15:11:33 2011	19651
Component: ee19c4f8-0025-11e1-aef6-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1462.11182011151133.tgz	Nov 18 15:11:36 2011	24650
Component: BBAK8281		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:41 2011	19445
Component: BBAK8891		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:41 2011	21916
Component: BBAK8276		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:39 2011	20461
Component: BBAK8868		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151137.tgz	Nov 18 15:11:41 2011	21924
Component: BBAK8835		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151137.tgz	Nov 18 15:11:39 2011	19424
Component: BBAK8283		
Filename	Size	Date
vccpd.tarball.0.1195.11182011151138.tgz	Nov 18 15:11:42 2011	31186
Component: YW3781/YW3781		
Filename	Size	Date
vccpd.tarball.0.1220.11182011151141.tgz	Nov 18 15:11:45 2011	27565
Component: 09726be2-0026-11e1-82d9-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1461.11182011151130.tgz	Nov 18 15:11:34 2011	19613
Component: BBAK8309		
Filename	Size	Date
vccpd.tarball.0.1196.11182011151138.tgz	Nov 18 15:11:46 2011	50362
Component: 303d476a-0026-11e1-abf4-00e081c52990		
Filename	Size	Date
vccpd.tarball.0.1460.11182011151133.tgz	Nov 18 15:11:33 2011	19360
Component: YW3798/YW3798		
Filename	Size	Date
vccpd.tarball.0.1219.11182011151140.tgz	Nov 18 15:11:49 2011	24473

#### show system core-dumps kernel-crashinfo component serial number (QFabric Systems)

```
user@switch> show system core-dumps kernel-crashinfo component A0001/YA0197
Node: A0001/YA0197
```

Information about previous kernel crash:

-- Kernel panic data --

Panic string: kdb\_sysctl\_panic

System uptime: 3 day 20 hr 59 min 40 sec Kernel crash time: 2011-11-15 Wed 15:25:17

Kernel build linkstamp: JUNOS 11.3I #0: 2011-11-10 20:42:27 UTC

-- Stacktrace of panicing context --

Processor 1 (crash monarch):

savectx+0x0 (c9552800,80214efc,802a7fbc,c88ad05c) ra 801b93a8 sz 0

kdm\_kcore\_save\_crashinfo+0x254 (c9552800,0,802a7fbc,c88ad05c) ra 801b9f44 sz 784

kdm\_kcore\_kern\_panic\_event\_handler+0x4b0 (c9552800,0,802a7fbc,c88ad05c) ra 8022a9b8 sz 88

panic+0x1d0 (c9552800,0,4,77fed534) ra 802540c0 sz 56

kdb\_sysctl\_panic+0x70 (c9552800,0,4,77fed534) ra 80237e58 sz 40 sysctl\_root+0x12c (c9552800,0,4,e8bc5cf8) ra 80238e50 sz 48

userland\_sysctl+0x164 (c9552800,0,4,e8bc5cf8) ra 8023956c sz 104

\_\_sysctl+0xe4 (c9552800,0,4,e8bc5cf8) ra 806d62e8 sz 160

trap+0xe1c (c9552800,0,4,e8bc5cf8) ra 80896e68 sz 128

MipsUserGenException+0x1a4 (c9552800,0,4,405cd12c) ra 0 sz 0

pid 82340, process: sysctl

Processor 0:

restoreintr+0x14 (1,81bca820,3,0) ra 806cdc3c sz 0

spinlock\_exit+0x30 (1,81bca820,3,0) ra 8025d354 sz 24

sleepq\_release+0x64 (1,81bca820,3,0) ra 8025e670 sz 24

sleepq\_timeout+0x224 (1,81bca820,3,0) ra 80240294 sz 48

softclock+0x434 (1,81bca820,3,0) ra 802067f8 sz 80

ithread\_loop+0x244 (1,81bca820,3,0) ra 80200e28 sz 64 fork\_exit+0xc0 (1,81bca820,3,0) ra 80897c28 sz 48

MipsNMIException+0x34 (1,81bca820,3,0) ra 0 sz 0

pid 82340, process: sysctl

Processor 2:

cpu\_idle+0x20 (80960000,51bbc,2031df,81bca1b8) ra 80204948 sz 24 idle\_proc+0x130 (80960000,51bbc,2031df,81bca1b8) ra 80200e28 sz 56 fork\_exit+0xc0

(80960000,51bbc,2031df,81bca1b8) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2031df,81bca1b8) ra 0 sz 0

pid 82340, process: sysctl

Processor 3:

cpu\_idle+0x20 (80960000,51bbc,2038df,81bca300) ra 80204948 sz 24 idle\_proc+0x130 (80960000,51bbc,2038df,81bca300) ra 80200e28 sz 56 fork\_exit+0xc0

(80960000,51bbc,2038df,81bca300) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2038df,81bca300) ra 0 sz 0

pid 82340, process: sysctl

Processor 4:

cpu\_idle+0x20 (80960000,51bbc,2037df,81bca448) ra 80204948 sz 24 idle\_proc+0x130 (80960000,51bbc,2037df,81bca448) ra 80200e28 sz 56 fork\_exit+0xc0

(80960000,51bbc,2037df,81bca448) ra 80897c28 sz 48

MipsNMIException+0x34 (80960000,51bbc,2037df,81bca448) ra 0 sz 0

pid 82340, process: sysctl

Processor 5:

restoreintr+0x14 (1,51bbc,203edf,81bca590) ra 806cdc3c sz 0

spinlock\_exit+0x30 (1,51bbc,203edf,81bca590) ra 80204a34 sz 24 idle\_proc+0x21c (1,51bbc,203edf,81bca590) ra 80200e28 sz 56 fork\_exit+0xc0

(1,51bbc,203edf,81bca590) ra 80897c28 sz 48

MipsNMIException+0x34 (1,51bbc,203edf,81bca590) ra 0 sz 0

pid 82340, process: sysctl

```

Processor 6:
cpu_idle+0x20 (80960000,51bbc,205cdf,81bca6d8) ra 80204948 sz 24 idle_proc+0x130
(80960000,51bbc,205cdf,81bca6d8) ra 80200e28 sz 56 fork_exit+0xc0
(80960000,51bbc,205cdf,81bca6d8) ra 80897c28 sz 48
MipsNMIException+0x34 (80960000,51bbc,205cdf,81bca6d8) ra 0 sz 0
pid 82340, process: sysctl

Processor 7:
lockmgr+0x5ac (c97e8484,c8dd9800,0,c8dd9800) ra 8c11c81c sz 48
sal_sem_take+0x134 (c97e8484,c8dd9800,0,c8dd9800) ra 8c351108 sz 56
_bcm_esw_linkscan_thread+0x45c (c97e8484,c8dd9800,0,c8dd9800) ra 8c11cdb4 sz 104
sal_thread_start_wrap+0x74 (c97e8484,c8dd9800,0,c8dd9800) ra 80200e28 sz 32
fork_exit+0xc0 (c97e8484,c8dd9800,0,c8dd9800) ra 80897c28 sz 48
MipsNMIException+0x34 (c97e8484,c8dd9800,0,c8dd9800) ra 0 sz 0
pid 82340, process: sysctl
-- End of stacktrace --

```

### show system core-dumps repository core (QFabric Systems)

```

user@switch> show system core-dumps repository core
Repository scope: shared
Repository head: /pbdata/export
Repository name: core
List of nodes for core repository: /pbdata/export/rdumps/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	0	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	0	fx-jvre	0M
NW-NG-0	BBAK0394	0	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	0	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	0	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	0	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	0	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YW3803	0	qfxc08-3008	0M
IC-WS001	WS001/YN5999	0	qfxc08-3008	0M
node-device1	BBAK0372	0	qfx3500	0M
node-device1	EE3093	0	qfx3500	0M

Total usage of core repository: 0M of 70000M (0.0%)

### show system core-dumps repository log (QFabric Systems)

```

user@switch> show system core-dumps repository log
Repository scope: shared
Repository head: /pbdata/export
Repository name: log
List of nodes for log repository: /pbdata/export/rlogs/

```

Node Group	Node Identifier	Num	Model	Usage
DG-0	BCF7208D-E44F-E011-802F-4171BAAC781D	0	qfx3100	0M
FM-0	73747cd8-0710-11e1-b6a4-00e081c5297e	1	fx-jvre	0M
DRE-0	77116f18-0710-11e1-a2a0-00e081c5297e	1	fx-jvre	0M
NW-NG-0	BBAK0394	1	qfx3500	0M
NW-NG-0	cd78871a-0710-11e1-878e-00e081c5297e	1	fx-jvre	0M
NW-NG-0	d0afda1e-0710-11e1-a1d0-00e081c5297e	3	fx-jvre	0M
FC-0	d31ab7a6-0710-11e1-ad1b-00e081c5297e	1	fx-jvre	0M
FC-1	d4d0f254-0710-11e1-90c3-00e081c5297e	1	fx-jvre	0M
IC-WS001	WS001	0	-	-
IC-WS001	WS001/YN5999	1	qfxc08-3008	0M
IC-WS001	WS001/YW3803	1	qfxc08-3008	0M



node-device1	BBAK0372	1	qfx3500	0M
node-device1	EE3093	1	qfx3500	0M
Total usage of log repository:0M of 70000M (0.0%)				

## show system directory-usage

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1040</a> <a href="#">Syntax (EX Series) on page 1040</a> <a href="#">Syntax (TX Matrix Router) on page 1040</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1040</a> <a href="#">Syntax (MX Series Router) on page 1040</a> <a href="#">Syntax (QFX Series) on page 1040</a>
<b>Syntax</b>	show system directory-usage <depth <i>number</i> > <path>
<b>Syntax (EX Series)</b>	show system directory-usage <all-members> <depth <i>number</i> > <local> <member <i>member-id</i> > <path>
<b>Syntax (TX Matrix Router)</b>	show system directory-usage <all-chassis   all-lcc   lcc <i>number</i>   scc> <depth <i>number</i> > <path>
<b>Syntax (TX Matrix Plus Router)</b>	show system directory-usage <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <depth <i>number</i> > <path>
<b>Syntax (MX Series Router)</b>	show system directory-usage <all-members> <depth <i>number</i> > <local> <member <i>member-id</i> > <path>
<b>Syntax (QFX Series)</b>	show system directory-usage <depth <i>number</i> > <path> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display directory usage information.
<b>Options</b>	none—Display all directory usage information.

**all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display directory usage information about all the T640 routers (in a routing matrix based on a TX Matrix router). Display directory usage information about all the T1600 or T4000 routers (in a routing matrix based on a TX Matrix Plus router) in the chassis.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display directory information for all members of the Virtual Chassis configuration.

**depth *number***—(Optional) Depth of the directory to traverse. This option is useful when you want to limit the output shown for a large file system.

**infrastructure *name***— (QFabric systems only) (Optional) Display directory information for the fabric control Routing Engines and fabric manager Routing Engines.

**interconnect-device *name***— (QFabric systems only) (Optional) Display directory information for the Interconnect device.

**node-group *name***— (QFabric systems only) (Optional) Display directory information for the Node group.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display directory information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display directory information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display directory information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

***path***—(Optional) Path or root directory to traverse.

**scc**—(TX Matrix router only) (Optional) Display directory information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display directory information for the TX Matrix Plus router. Replace *number* with 0.

**Required Privilege Level** view

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system directory-usage scc \(TX Matrix Router\) on page 1043](#)  
[show system directory-usage sfc \(TX Matrix Plus Router\) on page 1043](#)  
[show system directory-usage \(QFX3500 Switch\) on page 1043](#)

**Output Fields** [Table 59 on page 1042](#) describes the output fields for the **show system directory-usage** command. Output fields are listed in the approximate order in which they appear.

**Table 59: show system directory-usage Output Fields**

Field Name	Field Description
<i>bytes</i>	Number of bytes used by files in a directory.
<i>directory-name</i>	Name of the directory.

## Sample Output

### show system directory-usage scc (TX Matrix Router)

```

user@host> show system directory-usage /var/tmp scc
/var/tmp
1.0K    /var/tmp/vi.recover
2.0K    /var/tmp/instmp.tPMk8u
1.0K    /var/tmp/install
        /var/tmp/instmp.GUMpur
4.8M    /var/tmp/instmp.GUMpur/packages
6.4M    /var/tmp/troy1
297M    /var/tmp/dsw
        /var/tmp/pkg_tmp.2073
83K     /var/tmp/pkg_tmp.2073/bin
        /var/tmp/instmp.oMIDb1
89K     /var/tmp/instmp.oMIDb1/bin
        /var/tmp/instmp.byhMjR
4.6M    /var/tmp/instmp.byhMjR/packages
        /var/tmp/instmp.6fqHf3
1.7M    /var/tmp/instmp.6fqHf3/packages
        /var/tmp/instmp.mljECe
4.6M    /var/tmp/instmp.mljECe/packages

```

### show system directory-usage sfc (TX Matrix Plus Router)

```

user@switch> show system directory-usage /var/tmp sfc 0
sfc0-re0:
-----
/var/tmp
46K     /var/tmp/gres-tp
        /var/tmp/sec-download
2.0K    /var/tmp/sec-download/sub-download
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
795M    /var/tmp/cores
766K    /var/tmp/pr440594

```

### show system directory-usage (QFX3500 Switch)

```

user@switch> show system directory-usage
/var/tmp
30K     /var/tmp/gres-tp
2.0K    /var/tmp/rtbdb
2.0K    /var/tmp/vi.recover
2.0K    /var/tmp/install
2.0K    /var/tmp/pics

```

## show system license

<b>Syntax</b>	<code>show system license</code> <code>&lt;installed   keys   usage&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 13.3 for the MX104 3D Universal Edge Routers.
<b>Description</b>	Display licenses and information about how they are used.
<b>Options</b>	<p><b>none</b>—Display all license information.</p> <p><b>installed</b>—(Optional) Display installed licenses only.</p> <p><b>keys</b>—(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.</p> <p><b>usage</b>—(Optional) Display the state of licensed features.</p>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">show system license on page 1045</a> <a href="#">show system license installed on page 1046</a> <a href="#">show system license keys on page 1046</a> <a href="#">show system license usage on page 1046</a> <a href="#">show system license (MX104 Routers) on page 1046</a> <a href="#">show system license installed (MX104 Routers) on page 1047</a> <a href="#">show system license keys (MX104 Routers) on page 1047</a> <a href="#">show system license usage (MX104 Routers) on page 1047</a> <a href="#">show system license (MX104 Routers) on page 1047</a> <a href="#">show system license installed (MX104 Routers) on page 1048</a> <a href="#">show system license keys (MX104 Routers) on page 1048</a> <a href="#">show system license usage (MX104 Routers) on page 1048</a> <a href="#">show system license (MX104 Routers) on page 1049</a> <a href="#">show system license installed (MX104 Routers) on page 1049</a> <a href="#">show system license keys (MX104 Routers) on page 1049</a> <a href="#">show system license usage (MX104 Routers) on page 1050</a> <a href="#">show system license (QFX Series) on page 1050</a>
<b>Output Fields</b>	Table 60 on page 1044 lists the output fields for the <b>show system license</b> command. Output fields are listed in the approximate order in which they appear.

Table 60: show system license Output Fields

Field Name	Field Description
<b>Feature name</b>	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.

Table 60: show system license Output Fields (*continued*)

Field Name	Field Description
<b>Licenses used</b>	<p>Number of licenses used by a router or switch. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.</p> <p><b>NOTE:</b> In Junos OS Release 10.1 and later, the <b>Licenses used</b> column displays the actual usage count based on the number of active sessions or connections as reported by the corresponding feature daemons. This is applicable for scalable license-based features such as Subscriber Access (<b>scale-subscriber</b>), L2TP (<b>scale-l2tp</b>), Mobile IP (<b>scale-mobile-ip</b>), and so on.</p>
<b>Licenses installed</b>	<p>Information about the installed license key:</p> <ul style="list-style-type: none"> <li>• <b>License identifier</b>—Identifier associated with a license key.</li> <li>• <b>State</b>—State of the license key: <b>valid</b> or <b>invalid</b>. An <b>invalid</b> state indicates that the key was entered incorrectly or is not valid for the specific device.</li> <li>• <b>License version</b>—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key.</li> <li>• <b>Valid for device</b>—Device that can use a license key.</li> <li>• <b>Group defined</b>—Group membership of a device.</li> <li>• <b>Features</b>—Feature associated with a license, such as data link switching (DLSw).</li> </ul>
<b>Licenses needed</b>	Number of licenses required for features being used but not yet properly licensed.
<b>Expiry</b>	Amount of time left within the grace period before a license is required for a feature being used.

## Sample Output

### show system license

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

```
License usage:
```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

```
Licenses installed:
```

```
License identifier: XXXXXXXXXX
```

```
License version: 2
```

```
Features:
```

```
subscriber-accounting - Per Subscriber Radius Accounting
permanent
subscriber-authentication - Per Subscriber Radius Authentication
permanent
subscriber-address-assignment - Radius/SRC Address Pool Assignment
permanent
subscriber-vlan - Dynamic Auto-sensed Vlan
```

```

permanent
subscriber-ip    - Dynamic and Static IP
permanent

```

### show system license installed

```

user@host> show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
  subscriber-accounting - Per Subscriber Radius Accounting
  permanent
  subscriber-authentication - Per Subscriber Radius Authentication
  permanent
  subscriber-address-assignment - Radius/SRC Address Pool Assignment
  permanent
  subscriber-vlan - Dynamic Auto-sensed Vlan
  permanent
  subscriber-ip - Dynamic and Static IP
  permanent

```

### show system license keys

```

user@host> show system license keys
XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxx

```

### show system license usage

```

user@host> show system license usage
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
subscriber-accounting	2	2	0	permanent
subscriber-authentication	1	2	0	permanent
subscriber-address-assignment	2	2	0	permanent
subscriber-vlan	2	2	0	permanent
subscriber-ip	0	2	0	permanent
scale-subscriber	2	3	0	permanent
scale-l2tp	4	5	0	permanent
scale-mobile-ip	1	2	0	permanent

### show system license (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host> show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent

```

Licenses installed:
License identifier: XXXXXXXXXX
License version: 2
Features:

```



```

MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
    permanent

```

### show system license installed (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
    permanent

```

### show system license keys (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

```

### show system license usage (MX104 Routers)

In the following output, ports 0 and 1 are activated by installing the license to activate the first two built-in ports.

```

user@host > show system license usage

```

Feature name	Licenses used	Licenses installed	Expiry needed	
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent

### show system license (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license
License usage:

```

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

```

Licenses installed:
License identifier: XXXXXXXXXX
License version: 2

```

```

Features:
  MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent

License identifier: XXXXXXXXXX
License version: 2
Features:
  MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent

```

### show system license installed (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license installed
License identifier: XXXXXXXXXX
License version: 2
Features:
  MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent

License identifier: XXXXXXXXXX
License version: 2
Features:
  MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent

```

### show system license keys (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license keys

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
          xxxxxx xxxx

```

### show system license usage (MX104 Routers)

In the following output, ports 2 and 3 are activated by installing the license to activate the next two built-in ports after installing the license to activate the first two built-in ports.

```

user@host > show system license usage

```

Feature name	Licenses used	Licenses installed	Expiry needed	
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent

scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

### show system license (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

Licenses installed:

License identifier: XXXXXXXXXX

License version: 2

Features:

```
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent
MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent
```

### show system license installed (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

License identifier: XXXXXXXXXX

License version: 2

Features:

```
MX104-2x10Gig-port-0-1 - MX104 2X10Gig Builtin Port(xe-2/0/0 & xe-2/0/1)
upgrade
  permanent
MX104-2x10Gig-port-2-3 - MX104 2X10Gig Builtin Port(xe-2/0/2 & xe-2/0/3)
upgrade
  permanent
```

### show system license keys (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

```
XXXXXXXX XXXXXX XXXXXX XXXXXX XXXXXX XXXXXX
XXXXXXXX XXXXXX XXXXXX XXXXXX XXXXXX
XXXXXXXX XXXXXX X
```

### show system license usage (MX104 Routers)

In the following output, ports 0,1,2, and 3 are activated by installing a single license key to activate all four built-in ports.

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

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
scale-subscriber	0	1000	0	permanent
scale-l2tp	0	1000	0	permanent
scale-mobile-ip	0	1000	0	permanent
MX104-2x10Gig-port-0-1	0	1	0	permanent
MX104-2x10Gig-port-2-3	0	1	0	permanent

### show system license (QFX Series)

```
user@switch> show system license
```

License usage:

Feature name	Licenses used	Licenses installed	Licenses needed	Expiry
qfx-edge-fab	1	1	1	permanent

Licenses installed:  
License identifier: JUNOS417988  
License version: 1  
Features:  
qfx-edge-fab - QFX3000 Series QF/Node feature license  
permanent

## show system processes

<b>List of Syntax</b>	<a href="#">Syntax on page 1051</a> <a href="#">Syntax (EX Series Switches) on page 1051</a> <a href="#">Syntax (MX Series Routers) on page 1051</a> <a href="#">Syntax (QFX Series) on page 1051</a> <a href="#">Syntax (TX Matrix Routers) on page 1051</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1051</a>
<b>Syntax</b>	<pre>show system processes &lt;brief   detail   extensive   summary&gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;providers&gt; &lt;resource-limits (brief   detail) <i>process-name</i>&gt; &lt;wide&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show system processes &lt;all-members&gt; &lt;brief   detail   extensive   summary&gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;providers&gt; &lt;resource-limits (brief   detail) <i>process-name</i>&gt; &lt;wide&gt;</pre>
<b>Syntax (MX Series Routers)</b>	<pre>show system processes &lt;all-members&gt; &lt;brief   detail   extensive   summary&gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;local&gt; &lt;member <i>member-id</i>&gt; &lt;providers&gt; &lt;resource-limits (brief   detail) <i>process-name</i>&gt; &lt;wide&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system processes &lt;brief   detail   extensive   summary &gt; &lt;health (pid <i>process-identifer</i>   process-name <i>process-name</i>)&gt; &lt;interconnect-device <i>name</i>&gt; &lt;node-group <i>name</i>&gt; &lt;providers&gt; &lt;resource-limits&gt; &lt;wide&gt;</pre>
<b>Syntax (TX Matrix Routers)</b>	<pre>show system processes &lt;brief   detail   extensive   summary&gt; &lt;all-chassis  all-lcc   lcc <i>number</i>   scc&gt; &lt;wide&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show system processes &lt;brief   detail   extensive   summary&gt; &lt;all-chassis  all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</pre>

<wide>

<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about software processes that are running on the router or switch and that have controlling terminals.
<b>Options</b>	<p><b>none</b>—Display standard information about system processes.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of detail.</p> <p><b>adaptive-services</b>—(Optional) Display the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.</p> <p><b>alarm-control</b>—(Optional) Display the process to configure the system alarm.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display standard system process information about all the T640 routers (in a routing matrix based on the TX Matrix router) or all the T1600 or T4000 routers (in a routing matrix based on the TX Matrix Plus router) in the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus router only) (Optional) Display standard system process information for all T640 routers (or line-card chassis) connected to the TX Matrix router. Display standard system process information for all connected T1600 or T4000 LCCs.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for all members of the Virtual Chassis configuration.</p> <p><b>ancpd-service</b>—Display the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.</p> <p><b>application-identification</b>—Display the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.</p> <p><b>audit-process</b>—(Optional) Display the RADIUS accounting process.</p> <p><b>auto-configuration</b>—Display the Interface Auto-Configuration process.</p> <p><b>bootp</b>—Display the process that enables a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. DHCP relaying is disabled.</p> <p><b>captive-portal-content-delivery</b>—Display the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.</p>

- ce-l2tp-service**—(Optional) (M10, M10i, M7i, and MX Series routers only) Display the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.
- cfm**—Display Ethernet Operations, Administration, and Maintenance (OAM) connectivity fault management (CFM) process, which can be used to monitor the physical link between two switches.
- chassis-control**—(Optional) Display the chassis management process.
- class-of-service**—(Optional) Display the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.
- clksyncd-service**—Display the external clock synchronization process, which uses synchronous Ethernet (SyncE).
- craft-control**—Display the process for the I/O of the craft interface.
- database-replication**—(EX Series switches and MX Series routers only) (Optional) Display the database replication process.
- datapath-trace-service**—Display the packet path tracing process.
- dhcp-service**—(EX Series switches and MX Series routers only) (Optional) Display the Dynamic Host Configuration Protocol process, which enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
- diameter-service**—(Optional) Display the diameter process.
- disk-monitoring**—(Optional) Display the disk monitoring process, which checks the health of the hard disk drive on the Routing Engine.
- dynamic-flow-capture**—(Optional) Display the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.
- ecc-error-logging**—(Optional) Display the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.
- ethernet-connectivity-fault-management**—Display the process that provides IEEE 802.1ag OAM connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
- ethernet-link-fault-management**—(EX Series switches and MX Series routers only) (Optional) Display the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.
- event-processing**—(Optional) Display the event process (eventd).
- firewall**—(Optional) Display the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

**general-authentication-service**—(EX Series switches and MX Series routers only)  
(Optional) Display the general authentication process.

**health (pid *process-identifier* | process-name *process-name*)**—(Optional) Display process health information, either by process id (PID) or by process name.

**iccp-service**—Display the Inter-Chassis Communication Protocol (ICCP) process.

**idp-policy**—Display the intrusion detection and prevention (IDP) protocol process.

**ilmi**—Display the Integrated Local Management Interface (ILMI) protocol process, which provides bidirectional exchange of management information between two ATM interfaces across a physical connection.

**inet-process**—Display the IP multicast family process.

**init**—Display the process that initializes the USB modem.

**interface-control**—(Optional) Display the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

**kernel-replication**—(Optional) Display the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

**l2-learning**—(Optional) Display the Layer 2 address flooding and learning process.

**l2cpd-service**—Display the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

**lACP**—(Optional) Display the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display standard system process information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display standard system process information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.



- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the local Virtual Chassis member.

**local-policy-decision-function**—Display the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

**logical-system-mux**—Display the logical router multiplexer process (lrmuxd), which manages the multiple instances of the routing protocols process (rpd) on a machine running logical routers.

**mac-validation**—Display the MAC validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display standard system process information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**mib-process**—(Optional) Display the MIB II process, which provides the router's MIB II agent.

**mobile-ip**—(Optional) Display the Mobile IP process, which configures Junos OS Mobile IP features.

**mountd-service**—(EX Series switches and MX Series routers only) (Optional) Display the service for NFS mounts requests.

**mpls-traceroute**—(Optional) Display the MPLS Periodic Traceroute process.

**mspd**—(Optional) Display the Multiservice process.

**multicast-snooping**—(EX Series switches and MX Series routers only) (Optional) Display the multicast snooping process, which makes Layer 2 devices such as VLAN switches aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

**named-service**—(Optional) Display the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

**neighbor-liveness**—Display the process, which specifies the maximum length of time that the router waits for its neighbor to re-establish an LDP session.

**nfsd-service**—(Optional) Display the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

**ntp**—Display the Network Time Protocol (NTP) process, which provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network.

**packet-triggered-subscribers**—Display the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

**peer-selection-service**—(Optional) Display the Peer Selection Service process.

**periodic-packet-services**—Display the Periodic packet management process, which is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

**pfe**—Display the Packet Forwarding Engine management process.

**pgcp-service**—(Optional) Display the pgcpd service process running on the Routing Engine.

**pgm**—Display the Pragmatic General Multicast (PGM) protocol process, which enables a reliable transport layer for multicast applications.

**pic-services-logging**—(Optional) Display the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

**ppp**—(Optional) Display the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

**ppp-service**—Display the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

**pppoe**—(Optional) Display the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

**process-monitor**—Display the process health monitor process (pmond).

**providers**—(Optional) Display provider processes.

**redundancy-interface-process**—(Optional) Display the ASP redundancy process.

**remote-operations**—(Optional) Display the remote operations process, which provides the ping and traceroute MIBs.

**resource-cleanup**—Display the resource cleanup process.

**resource-limits (brief | detail) process-name**—(Optional) Display process resource limits.

**routing**—(Optional) Display the routing protocol process.

**sampling**—(Optional) Display the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

**sbc-configuration-process**—Display the session border controller (SBC) process of the border signaling gateway (BSG).

**scc**—(TX Matrix routers only) (Optional) Display standard system process information for the TX Matrix router (or switch-card chassis).

**sdk-service**—Display the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

**secure-neighbor-discovery**—(EX Series switches and MX Series routers only) (Optional) Display the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

**send**—(Optional) Display the Secure Neighbor Discovery Protocol (SEND) process, which provides support for protecting Neighbor Discovery Protocol (NDP) messages.

**service-deployment**—(Optional) Display the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

**sfc number**—(TX Matrix Plus routers only) (Optional) Display system process information for the TX Matrix Plus router. Replace *number* with 0.

**snmp**—Display the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

**sonet-aps**—Display the SONET Automatic Protection Switching (APS) process, which monitors any SONET interface that participates in APS.

**static-subscribers**—(Optional) Display the Static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

**tunnel-oamd**—(Optional) Display the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

**vrrp**—(EX Series switches and MX Series routers only) (Optional) Display the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**watchdog**—Display the watchdog timer process, which enables the watchdog timer when Junos OS encounters a problem.

**wide**—(Optional) Display process information that might be wider than 80 columns.

**Additional Information** By default, when you issue the **show system processes** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise,

if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level**

view

**Related Documentation**

- [List of Junos OS Processes](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

[show system processes on page 1060](#)  
[show system processes brief on page 1060](#)  
[show system processes detail on page 1061](#)  
[show system processes extensive on page 1061](#)  
[show system processes extensive \(EX9200 Switch\) on page 1062](#)  
[show system processes lcc wide \(TX Matrix Routing Matrix\) on page 1062](#)  
[show system processes summary on page 1063](#)  
[show system processes \(TX Matrix Plus Router\) on page 1063](#)  
[show system processes sfc \(TX Matrix Plus Router\) on page 1070](#)  
[show system processes lcc wide \(TX Matrix Plus Routing Matrix\) on page 1073](#)  
[show system processes \(QFX Series\) on page 1075](#)

**Output Fields**

[Table 61 on page 1058](#) describes the output fields for the **show system processes** command. Output fields are listed in the approximate order in which they appear.

**Table 61: show system processes Output Fields**

Field Name	Field Description	Level of Output
last pid	Last process identifier assigned to the process.	brief extensive summary
load averages	Three load averages followed by the current time.	brief extensive summary
processes	Number of existing processes and the number of processes in each state ( <b>sleeping</b> , <b>running</b> , <b>starting</b> , <b>zombies</b> , and <b>stopped</b> ).	brief extensive summary
Mem	Information about physical and virtual memory allocation.	brief extensive summary
Swap	Information about physical and virtual memory allocation.	brief extensive summary
PID	Process identifier.	detail extensive summary
TT	Control terminal name.	none detail

Table 61: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>STAT</b>	<p>Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process:</p> <ul style="list-style-type: none"> <li>• <b>D</b>—In disk or other short-term, uninterruptible wait</li> <li>• <b>I</b>—Idle (sleeping longer than about 20 seconds)</li> <li>• <b>R</b>—Runnable</li> <li>• <b>S</b>—Sleeping for less than 20 seconds</li> <li>• <b>T</b>—Stopped</li> <li>• <b>Z</b>—Dead (zombie)</li> <li>• <b>+</b> —The process is in the foreground process group of its control terminal.</li> <li>• <b>&lt;</b>—The process has raised CPU scheduling priority.</li> <li>• <b>&gt;</b>—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped.</li> <li>• <b>A</b>—The process requested random page replacement.</li> <li>• <b>E</b>—The process is trying to exit.</li> <li>• <b>L</b>—The process has pages locked in core.</li> <li>• <b>N</b>—The process has reduced CPU scheduling priority.</li> <li>• <b>S</b>—The process requested first-in, first-out (FIFO) page replacement.</li> <li>• <b>s</b>—The process is a session leader.</li> <li>• <b>V</b>—The process is temporarily suspended.</li> <li>• <b>W</b>—The process is swapped out.</li> <li>• <b>X</b>—The process is being traced or debugged.</li> </ul>	none <b>detail</b>
<b>UID</b>	User identifier.	<b>detail</b>
<b>USERNAME</b>	Process owner.	<b>extensive summary</b>
<b>PPID</b>	Parent process identifier.	<b>detail</b>
<b>CPU</b>	<p>(D)—Short-term CPU usage.</p> <p>(E and S)—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.</p>	<b>detail extensive summary</b>
<b>RSS</b>	Resident set size.	<b>detail</b>
<b>WCHAN</b>	Symbolic name of the wait channel.	<b>detail</b>
<b>STARTED</b>	Local time when the process started running.	<b>detail</b>
<b>PRI</b>	Current priority of the process. A lower number indicates a higher priority.	<b>detail extensive summary</b>
<b>NI or NICE</b>	UNIX "niceness" value. A lower number indicates a higher priority.	<b>detail extensive summary</b>
<b>SIZE</b>	Total size of the process (text, data, and stack), in kilobytes.	<b>extensive summary</b>

Table 61: show system processes Output Fields (*continued*)

Field Name	Field Description	Level of Output
RES	Current amount of resident memory, in kilobytes.	extensive summary
STATE	Current state of the process (for example, <b>sleep</b> , <b>wait</b> , <b>run</b> , <b>idle</b> , <b>zombie</b> , or <b>stop</b> ).	extensive summary
TIME	(S)—Number of system and user CPU seconds that the process has used.  (None, D, and E)—Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running.	detail extensive summary
THR	Number of threads in the process	extensive

## Sample Output

### show system processes

```

user@host> show system processes
PID  TT  STAT      TIME COMMAND
   0  ??  DLs      0:00.70 (swapper)
   1  ??  Is       0:00.35 /sbin/init --
   2  ??  DL       0:00.00 (pagedaemon)
   3  ??  DL       0:00.00 (vmdaemon)
   4  ??  DL       0:42.37 (update)
   5  ??  DL       0:00.00 (if_jnx)
  80  ??  Ss       0:14.66 syslogd -s
  96  ??  Is       0:00.01 portmap
 128  ??  Is       0:02.70 cron
 173  ??  Is       0:02.24 /usr/local/sbin/sshd (sshd1)
 189  ??  S        0:03.80 /sbin/watchdog -t180
 190  ??  I        0:00.03 /usr/sbin/timed -N
 191  ??  S        2:24.76 /sbin/ifd -N
 192  ??  S<       0:55.44 /usr/sbin/xntpd -N
 195  ??  S        0:53.11 /usr/sbin/snmpd -N
 196  ??  S        1:15.73 /usr/sbin/mib2d -N
 198  ??  I        0:00.75 /usr/sbin/inetd -N
2677  ??  I        0:00.01 /usr/sbin/mgd -N
2712  ??  Ss       0:00.24 rlogind
2735  ??  R        0:00.00 /bin/ps -ax
1985  p0-  S        0:07.41 ./rpd -N
2713  p0  Is       0:00.24 -tcsh (tcsh)
2726  p0  S+       0:00.07 cli

```

### show system processes brief

```

user@host> show system processes brief
last pid:  543; load averages:  0.00,  0.00,  0.00   18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

```

## show system processes detail

user@host&gt; show system processes detail

PID	UID	PPID	CPU	PRI	NI	RSS	WCHAN	STARTED	TT	STAT	TIME	COMMAND
3151	1049	3129	2	28	0	672	-	1:13PM	p0	R+	0:00.00	ps -ax -r
1	0	0	0	10	0	376	wait	1:51PM	??	Is	0:00.29	/sbin/ini
2	0	0	0	-18	0	12	psleep	1:51PM	??	DL	0:00.00	(pagedae
3	0	0	0	28	0	12	psleep	1:51PM	??	DL	0:00.00	(vmdaemo
4	0	0	0	28	0	12	update	1:51PM	??	DL	0:07.15	(update)
5	0	0	0	2	0	12	pfesel	1:51PM	??	IL	0:02.90	(if_pfe)
27	0	1	0	10	0	17936	mfsidl	1:51PM	??	Is	0:00.46	mfs /dev/
81	0	1	0	2	0	496	select	1:52PM	??	Ss	0:31.21	syslogd -
119	1	1	0	2	0	492	select	1:52PM	??	Is	0:00.00	portmap
134	0	1	0	2	0	580	select	1:52PM	??	S	0:02.95	amd -p -a
151	0	1	0	18	0	532	pause	1:52PM	??	Is	0:00.34	cron
183	0	1	0	2	0	420	select	1:52PM	??	Ss	0:00.07	/usr/loca
206	0	1	0	18	0	72	pause	1:52PM	??	S	0:00.51	/sbin/wat
207	0	1	0	2	0	520	select	1:52PM	??	I	0:00.16	/usr/sbin
208	0	1	0	2	0	536	select	1:52PM	??	S	0:08.21	/sbin/dcd
210	0	1	255	2	-12	740	select	1:52PM	??	S<	0:05.83	/usr/sbin
211	0	1	0	2	0	376	select	1:52PM	??	S	0:00.03	/usr/sbin
215	0	1	0	2	0	548	select	1:52PM	??	I	0:00.50	/usr/sbin
219	0	1	0	3	0	540	ttyin	1:52PM	v0	Is+	0:00.02	/usr/libe
220	0	1	0	3	0	540	ttyin	1:52PM	v1	Is+	0:00.01	/usr/libe
221	0	1	0	3	0	540	ttyin	1:52PM	v2	Is+	0:00.01	/usr/libe
222	0	1	0	3	0	540	ttyin	1:52PM	v3	Is+	0:00.01	/usr/libe
735	0	1	0	2	0	468	select	2:47PM	??	S	0:19.14	/usr/sbin
736	0	1	0	2	0	212	select	2:47PM	??	S	0:14.13	/usr/sbin
1380	0	1	0	3	0	888	ttyin	7:32PM	d0	Is+	0:00.46	bash
3019	0	207	0	2	0	636	select	10:49AM	??	Ss	0:02.93	tnp.chass
3122	0	1380	0	2	0	1764	select	12:33PM	d0	S	0:00.77	./rpd -N
3128	0	215	0	2	0	580	select	12:45PM	??	Ss	0:00.12	rlogind
3129	1049	3128	0	18	0	944	pause	12:45PM	p0	Ss	0:00.14	-tcsh (tc
0	0	0	0	-18	0	0	sched	1:51PM	??	DLs	0:00.10	(swapper

## show system processes extensive

user@host&gt; show system processes extensive

Mem: 241M Active, 99M Inact, 78M Wired, 325M Cache, 69M Buf, 1251M Free  
 Swap: 2048M Total, 2048M Free

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
11	root	1	171	52	OK	12K	RUN	807.5H	98.73%	idle
13	root	1	-20	-139	OK	12K	WAIT	36:17	0.00%	swi7: clock sio
1499	root	1	96	0	7212K	3040K	select	34:01	0.00%	license-check
1621	root	1	96	0	20968K	11216K	select	20:25	0.00%	mib2d
1465	root	2	8	-88	115M	11748K	nanslp	14:32	0.00%	chassisd
1478	root	1	96	0	6336K	3816K	select	11:28	0.00%	ppmd
20	root	1	-68	-187	OK	12K	WAIT	10:28	0.00%	irq10: em0 em1+++*
1490	root	1	96	0	11792K	4336K	select	9:44	0.00%	shm-rtssdbd
1618	root	1	96	0	39584K	7464K	select	8:47	0.00%	pfed
1622	root	1	96	0	15268K	10988K	select	6:16	0.00%	snmpd
1466	root	1	96	0	7408K	2896K	select	5:44	0.00%	alarmd
7	root	1	-16	0	OK	12K	client	5:09	0.00%	ifstate notify
1480	root	1	96	0	5388K	2660K	select	4:29	0.00%	ksyncd
12	root	1	-40	-159	OK	12K	WAIT	4:15	0.00%	swi2: netisr 0
1462	root	1	96	0	1836K	1240K	select	3:57	0.00%	bslockd
55	root	1	-16	0	OK	12K	-	3:44	0.00%	schedcpu
1392	root	1	16	0	OK	12K	bcmsem	3:37	0.00%	bcmLINK.0

```

    47 root      1 -16    0    OK    12K psleep  3:25 0.00% vmkmemdaemon
    36 root      1  20    0    OK    12K syncer  2:46 0.00% syncer
  1484 root      1  96    0 7484K 3428K select  2:38 0.00% clksyncd
  1616 root      1  96    0 4848K 2848K select  2:18 0.00% irsd
  1487 root      1  96    0 32800K 6992K select  2:10 0.00% smid
  1623 root      1  96    0 34616K 5464K select  2:01 0.00% dcd
    15 root      1 -16    0    OK    12K -      1:59 0.00% yarrow
    49 root      1 -16    0    OK    12K .       1:51 0.00% ddostasks

```

### show system processes extensive (EX9200 Switch)

```

user@switch> show system processes extensive
last pid: 3372; load averages: 0.02, 0.02, 0.00 up 0+01:42:22 16:39:57
151 processes: 4 running, 131 sleeping, 1 zombie, 15 waiting

Mem: 935M Active, 122M Inact, 108M Wired, 838M Cache, 214M Buf, 5872M Free
Swap: 8192M Total, 8192M Free

```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
10	root	1	171	52	OK	16K	RUN	96:34	92.19%	idle
3317	root	1	97	0	40412K	30944K	select	0:00	5.13%	mgd
3316	root	1	96	0	26672K	20516K	select	0:00	3.08%	cli
1626	root	2	8	-88	124M	20332K	nanslp	3:19	2.39%	chassisd
260	root	1	-8	0	OK	16K	mdwait	0:16	0.00%	md16
19	root	1	-68	-187	OK	16K	WAIT	0:12	0.00%	irq11: em0 em1
em2*										
1642	root	1	96	0	8052K	3936K	RUN	0:10	0.00%	clksyncd
11	root	1	-20	-139	OK	16K	WAIT	0:07	0.00%	swi7: clock sio
154	root	1	-8	0	OK	16K	mdwait	0:06	0.00%	md8
1784	root	1	96	0	98M	33720K	select	0:05	0.00%	authd
1646	root	1	96	0	7776K	2944K	select	0:03	0.00%	license-check
1807	root	1	96	0	41340K	9944K	select	0:02	0.00%	mib2d

[...Output truncated...]

### show system processes lcc wide (TX Matrix Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

PID	TT	STAT	TIME	COMMAND
0	??	DLs	0:00.00	(swapper)
1	??	ILs	0:00.10	/sbin/preinit -- (init)
2	??	DL	0:00.00	(pagedaemon)
3	??	DL	0:00.00	(vmdaemon)
4	??	DL	0:00.00	(bufdaemon)
5	??	DL	0:00.04	(syncer)
6	??	DL	0:00.00	(netdaemon)
7	??	IL	0:00.00	(if_pic_listen)
8	??	IL	0:00.00	(scs_housekeeping)
9	??	IL	0:00.00	(if_pfe_listen)
10	??	DL	0:00.00	(vmuncachedaemon)
11	??	SL	0:00.02	(cb_poll)
172	??	ILs	0:00.21	mfs -o noauto /dev/ad1s1b /tmp (newfs)
2909	??	Is	0:00.00	pccardd
2932	??	Ss	0:00.07	syslogd -r -s
3039	??	Is	0:00.00	cron
3217	??	I	0:00.00	/sbin/watchdog -d
3218	??	I	0:00.02	/usr/sbin/tnetd -N
3221	??	S	0:00.11	/usr/sbin/alarmd -N
3222	??	S	0:00.85	/usr/sbin/craftd -N



```

3223 ?? S      0:00.05 /usr/sbin/mgd -N
3224 ?? I      0:00.02 /usr/sbin/inetd -N
3225 ?? I      0:00.00 /usr/sbin/tnp.sntpd -N
3226 ?? I      0:00.01 /usr/sbin/tnp.sntpc -N
3228 ?? I      0:00.01 /usr/sbin/smartd -N
3231 ?? I      0:00.01 /usr/sbin/eccd -N
3425 ?? S      0:00.09 /usr/sbin/dfwd -N
3426 ?? S      0:00.19 /sbin/dcd -N
3427 ?? I      0:00.04 /usr/sbin/pfed -N
3430 ?? S      0:00.10 /usr/sbin/ksyncd -N
3482 ?? S      1:53.63 /usr/sbin/chassisd -N
4285 ?? SL     0:00.01 (peer proxy)
4286 ?? SL     0:00.00 (peer proxy)
4303 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
4304 ?? R      0:00.00 /bin/ps -ax -ww
3270 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0

```

### show system processes summary

```
user@host> show system processes summary
```

```
last pid: 543; load averages: 0.00, 0.00, 0.00 18:29:47
37 processes: 1 running, 36 sleeping
```

```
Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
```

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
527	root	2	0	176K	580K	select	0:00	0.04%	0.04%	rlogind
543	root	30	0	604K	768K	RUN	0:00	0.00%	0.00%	top

### show system processes (TX Matrix Plus Router)

```
user@host> show system processes
```

```
sfc0-re0:
```

```

-----
PID  TT  STAT      TIME COMMAND
 0  ??  Wls      0:00.00 [swapper]
 1  ??  ILs      0:00.18 /packages/mnt/jbase/sbin/init --
 2  ??  DL       0:00.20 [g_event]
 3  ??  DL       0:00.39 [g_up]
 4  ??  DL       0:00.32 [g_down]
 5  ??  DL       0:00.00 [thread taskq]
 6  ??  DL       0:00.09 [kqueue taskq]
 7  ??  DL       0:00.01 [pagedaemon]
 8  ??  DL       0:00.00 [vmdaemon]
 9  ??  DL       0:06.63 [pagezero]
10  ??  DL       0:00.00 [ktrace]
11  ??  RL      310:52.98 [idle]
12  ??  WL       0:11.03 [swi2: net]
13  ??  WL       0:27.58 [swi7: clock sio]
14  ??  WL       0:00.00 [swi6: vm]
15  ??  DL       0:03.02 [yarrow]
16  ??  WL       0:00.00 [swi9: +]
17  ??  WL       0:00.00 [swi8: +]
18  ??  WL       0:00.00 [swi5: cambio]
19  ??  WL       0:00.00 [swi9: task queue]
20  ??  WL       0:11.41 [irq16: uhci0 uhci*]
21  ??  DL       0:00.00 [usb0]
22  ??  DL       0:00.00 [usbtask]
23  ??  WL       0:39.51 [irq17: uhci1 uhci*]
24  ??  DL       0:00.00 [usb1]

```

```
25 ?? WL 0:00.00 [irq18: uhci2 uhci*]
26 ?? DL 0:00.83 [usb2]
27 ?? DL 0:00.00 [usb3]
28 ?? DL 0:00.00 [usb4]
29 ?? DL 0:00.00 [usb5]
30 ?? DL 0:00.73 [usb6]
31 ?? DL 0:00.00 [usb7]
32 ?? WL 0:00.00 [irq14: ata0]
33 ?? WL 0:00.00 [irq15: ata1]
34 ?? WL 0:00.00 [irq1: atkbd0]
35 ?? WL 0:00.00 [swi0: sio]
36 ?? WL 0:00.00 [irq11: isab0]
37 ?? WL 0:00.00 [swi3: ip6opt ipopt]
38 ?? WL 0:00.00 [swi4: ip6mismatch+]
39 ?? WL 0:00.00 [swi1: ipfwd]
40 ?? DL 0:00.02 [bufdaemon]
41 ?? DL 0:00.02 [vn1ru]
42 ?? DL 0:00.39 [syncer]
43 ?? DL 0:00.05 [softdepflush]
44 ?? DL 0:00.00 [netdaemon]
45 ?? DL 0:00.02 [vmuncachedaemon]
46 ?? DL 0:00.00 [if_pic_listen]
47 ?? DL 0:00.35 [vmkmemdaemon]
48 ?? DL 0:00.00 [cb_poll]
49 ?? DL 0:00.06 [if_pfe_listen]
50 ?? DL 0:00.00 [scs_housekeeping]
51 ?? IL 0:00.00 [kern_dump_proc]
52 ?? IL 0:00.00 [nfsiod 0]
53 ?? IL 0:00.00 [nfsiod 1]
54 ?? IL 0:00.00 [nfsiod 2]
55 ?? IL 0:00.00 [nfsiod 3]
56 ?? DL 0:00.37 [schedcpu]
57 ?? DL 0:00.56 [md0]
79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.34 [bcmTX]
1342 ?? SL 0:01.68 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.40 [bcmLINK.0]
1345 ?? SL 0:33.83 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? S 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.19 /usr/sbin/mgd -N
1512 ?? I 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.10 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
```

```

1619 ?? DL      0:00.30 [peer proxy]
2391 ?? Is      0:00.01 telnetd
7331 ?? Ss      0:00.03 telnetd
9538 ?? DL      0:01.16 [jsr_kkcm]
9613 ?? DL      0:00.18 [peer proxy]
23781 ?? Ss      0:00.01 telnetd
23926 ?? Ss      0:00.01 mgd: (mgd) (regress)/dev/ttyp2 (mgd)
36867 ?? S       0:03.14 /usr/sbin/rpd -N
36874 ?? S       0:00.08 /usr/sbin/lmpd
36876 ?? S       0:00.17 /usr/sbin/lacpd -N
36877 ?? S       0:00.15 /usr/sbin/bfdd -N
36878 ?? S       0:05.05 /usr/sbin/ppmd -N
36907 ?? S       0:25.07 /usr/sbin/chassisd -N
37775 ?? S       0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S       0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S       0:00.38 /usr/sbin/l2ald -N
45730 ?? S<      0:00.12 /usr/sbin/apspd -N
45731 ?? SN      0:00.10 /usr/sbin/sampled -N
45732 ?? S       0:00.03 /usr/sbin/ilmid -N
45733 ?? S       0:00.09 /usr/sbin/rmopd -N
45734 ?? S       0:00.30 /usr/sbin/cosd
45735 ?? I       0:00.00 /usr/sbin/rtspd -N
45736 ?? S       0:00.06 /usr/sbin/fsad -N
45737 ?? S       0:00.05 /usr/sbin/rdd -N
45738 ?? S       0:00.10 /usr/sbin/pppd -N
45739 ?? S       0:00.05 /usr/sbin/dfcd -N
45740 ?? S       0:00.07 /usr/sbin/lfmd -N
45741 ?? S       0:00.01 /usr/sbin/mpiisoamd -N
45742 ?? I       0:00.01 /usr/sbin/sendd -N
45743 ?? S       0:00.08 /usr/sbin/appidd -N
45744 ?? S       0:00.05 /usr/sbin/mspd -N
45745 ?? S       0:00.25 /usr/sbin/jdiameterd -N
45746 ?? S       0:00.10 /usr/sbin/pfed -N
45747 ?? S       0:00.19 /usr/sbin/lpdfd -N
45748 ?? S       0:00.63 /sbin/dcd -N
45750 ?? S       0:00.45 /usr/sbin/mib2d -N
45751 ?? S       0:00.15 /usr/sbin/dfwd -N
45752 ?? S       0:00.15 /usr/sbin/irsd -N
45764 ?? S       0:20.59 /usr/sbin/snmpd -N
56479 ?? Ss      0:00.00 mgd: (mgd) (root) (mgd)
56480 ?? R       0:00.00 /bin/ps -ax
1142 d0- I       0:00.01 /usr/sbin/usbd -N
1160 d0- S       0:29.17 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+      0:00.00 /usr/libexec/getty std.9600 ttyd0
2392 p1 Is       0:00.00 login [pam] (login)
2393 p1 I        0:00.00 -csh (csh)
2394 p1 I        0:00.00 su -
2395 p1 I+       0:00.01 -su (csh)
23782 p2 Is      0:00.00 login [pam] (login)
23881 p2 I        0:00.00 -csh (csh)
23925 p2 S+      0:00.03 cli
7332 p3 Is       0:00.00 login [pam] (login)
7333 p3 I        0:00.00 -csh (csh)
23780 p3 S+      0:00.02 telnet aj

```

```
1cc0-re0:
```

```

-----
PID  TT  STAT  TIME  COMMAND
  0  ??  Wls   0:00.00 [swapper]
  1  ??  ILs   0:00.16 /packages/mnt/jbase/sbin/init --
  2  ??  DL    0:00.01 [g_event]

```

```

3  ??  DL    0:00.16 [g_up]
4  ??  DL    0:00.11 [g_down]
5  ??  DL    0:00.00 [thread taskq]
6  ??  DL    0:00.00 [kqueue taskq]
7  ??  DL    0:00.00 [pagedaemon]
8  ??  DL    0:00.00 [vmdaemon]
9  ??  DL    0:01.77 [pagezero]
10 ??  DL    0:00.00 [ktrace]
11 ??  RL    17:22.31 [idle]
12 ??  WL    0:00.32 [swi2: net]
13 ??  WL    0:01.21 [swi7: clock sio]
14 ??  WL    0:00.00 [swi6: vm]
15 ??  DL    0:00.10 [yarrow]
16 ??  WL    0:00.00 [swi9: +]
17 ??  WL    0:00.00 [swi8: +]
18 ??  WL    0:00.00 [swi5: cambio]
19 ??  WL    0:00.00 [swi9: task queue]
20 ??  WL    0:02.73 [irq10: bcm0 uhci1*]
21 ??  WL    0:00.02 [irq11: cb0 uhci0+*]
22 ??  DL    0:00.00 [usb0]
23 ??  DL    0:00.00 [usbtask]
24 ??  DL    0:00.00 [usb1]
25 ??  DL    0:00.05 [usb2]
26 ??  DL    0:00.00 [usb3]
27 ??  DL    0:00.00 [usb4]
28 ??  DL    0:00.00 [usb5]
29 ??  DL    0:00.04 [usb6]
30 ??  DL    0:00.00 [usb7]
31 ??  WL    0:00.00 [irq14: ata0]
32 ??  WL    0:00.00 [irq15: ata1]
33 ??  WL    0:00.00 [irq1: atkbd0]
34 ??  WL    0:00.00 [swi0: sio]
35 ??  WL    0:00.00 [swi3: ip6opt ipopt]
36 ??  WL    0:00.00 [swi4: ip6mismatch+]
37 ??  WL    0:00.00 [swi1: ipfwd]
38 ??  DL    0:00.00 [bufdaemon]
39 ??  DL    0:00.00 [vn1ru]
40 ??  DL    0:00.01 [syncer]
41 ??  DL    0:00.00 [softdepflush]
42 ??  DL    0:00.00 [netdaemon]
43 ??  DL    0:00.00 [vmuncachedaemon]
44 ??  DL    0:00.00 [if_pic_listen]
45 ??  DL    0:00.02 [vmknemdaemon]
46 ??  DL    0:00.01 [cb_poll]
47 ??  DL    0:00.00 [if_pfe_listen]
48 ??  DL    0:00.00 [scs_housekeeping]
49 ??  IL    0:00.00 [kern_dump_proc]
50 ??  IL    0:00.00 [nfsiod 0]
51 ??  IL    0:00.00 [nfsiod 1]
52 ??  IL    0:00.00 [nfsiod 2]
53 ??  IL    0:00.00 [nfsiod 3]
54 ??  DL    0:00.01 [schedcpu]
55 ??  DL    0:00.73 [md0]
77 ??  DL    0:03.54 [md1]
98 ??  DL    0:00.37 [md2]
116 ?? DL    0:00.02 [md3]
137 ?? DL    0:00.56 [md4]
158 ?? DL    0:00.15 [md5]
179 ?? DL    0:00.00 [md6]
215 ?? DL    0:00.03 [md7]
225 ?? DL    0:00.03 [md8]

```

```

1078 ?? DL      0:00.00 [jsr_kkcm]
1363 ?? SL      0:00.09 [bcmTX]
1364 ?? SL      0:00.10 [bcmXGS3AsyncTX]
1365 ?? SL      0:03.08 [bcmLINK.0]
1370 ?? Is      0:00.00 /usr/sbin/cron
1522 ?? S       0:00.00 /sbin/watchdog -t-1
1523 ?? S       0:00.05 /usr/libexec/bslockd -mp -N
1524 ?? I       0:00.01 /usr/sbin/tnetd -N
1526 ?? S       0:04.98 /usr/sbin/chassisd -N
1527 ?? S       0:00.04 /usr/sbin/alarmd -N
1528 ?? I       0:00.40 /usr/sbin/craftd -N
1529 ?? S       0:00.08 /usr/sbin/mgd -N
1532 ?? I       0:00.04 /usr/sbin/inetd -N
1533 ?? I       0:00.00 /usr/sbin/tnp.snptd -N
1534 ?? I       0:00.00 /usr/sbin/tnp.snptc -N
1536 ?? S       0:00.01 /usr/sbin/smartd -N
1540 ?? I       0:00.07 /usr/sbin/jcsd -N
1541 ?? S       0:00.11 /usr/sbin/idpd -N
1542 ?? I       0:00.00 /usr/libexec/getty Pc ttyv0
2089 ?? DL      0:00.01 [peer proxy]
2090 ?? DL      0:00.01 [peer proxy]
2091 ?? DL      0:00.01 [peer proxy]
2657 ?? S       0:00.02 /usr/sbin/dfwd -N
2658 ?? S       0:00.02 /sbin/dcd -N
2659 ?? S       0:00.05 /usr/sbin/snmpd -N
2660 ?? S       0:00.01 /usr/sbin/mib2d -N
2661 ?? S       0:00.01 /usr/sbin/pfed -N
2662 ?? S       0:00.01 /usr/sbin/irsd -N
2667 ?? S       0:00.13 /usr/sbin/ksyncd -N
2690 ?? Ss      0:00.00 mgd: (mgd) (root) (mgd)
2691 ?? R       0:00.00 /bin/ps -ax
1164 d0- S      0:00.00 /usr/sbin/usbd -N
1182 d0- S      0:00.34 /usr/sbin/eventd -N -r -s -A
1543 d0 Is+     0:00.00 /usr/libexec/getty std.9600 ttyd0

```

```
lcc1-re0:
```

```

-----
PID  TT  STAT    TIME COMMAND
  0  ??  WLS     0:00.00 [swapper]
  1  ??  ILs     0:00.17 /packages/mnt/jbase/sbin/init --
  2  ??  DL      0:00.01 [g_event]
  3  ??  DL      0:00.16 [g_up]
  4  ??  DL      0:00.11 [g_down]
  5  ??  DL      0:00.00 [thread taskq]
  6  ??  DL      0:00.00 [kqueue taskq]
  7  ??  DL      0:00.00 [pagedaemon]
  8  ??  DL      0:00.00 [vmdaemon]
  9  ??  DL      0:01.77 [pagezero]
 10  ??  DL      0:00.00 [ktrace]
 11  ??  RL     17:22.83 [idle]
 12  ??  WL      0:00.35 [swi2: net]
 13  ??  WL      0:01.20 [swi7: clock sio]
 14  ??  WL      0:00.00 [swi6: vm]
 15  ??  DL      0:00.10 [yarrow]
 16  ??  WL      0:00.00 [swi9: +]
 17  ??  WL      0:00.00 [swi8: +]
 18  ??  WL      0:00.00 [swi5: cambio]
 19  ??  WL      0:00.00 [swi9: task queue]
 20  ??  WL      0:02.87 [irq10: bcm0 uhci1*]
 21  ??  WL      0:00.02 [irq11: cb0 uhci0+*]
 22  ??  DL      0:00.00 [usb0]

```

```

23 ?? DL 0:00.00 [usbtask]
24 ?? DL 0:00.00 [usb1]
25 ?? DL 0:00.05 [usb2]
26 ?? DL 0:00.00 [usb3]
27 ?? DL 0:00.00 [usb4]
28 ?? DL 0:00.00 [usb5]
29 ?? DL 0:00.04 [usb6]
30 ?? DL 0:00.00 [usb7]
31 ?? WL 0:00.00 [irq14: ata0]
32 ?? WL 0:00.00 [irq15: ata1]
33 ?? WL 0:00.00 [irq1: atkbd0]
34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vn1ru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]
43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.40 [md1]
98 ?? DL 0:00.37 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.10 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? I 0:00.01 /usr/sbin/tnetd -N
1500 ?? S 0:04.97 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.40 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.04 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpc -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.18 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2068 ?? DL 0:00.01 [peer proxy]

```

```

2069 ?? DL 0:00.01 [peer proxy]
2070 ?? DL 0:00.01 [peer proxy]
2666 ?? S 0:00.02 /sbin/dcd -N
2667 ?? S 0:00.01 /usr/sbin/irsd -N
2668 ?? S 0:00.01 /usr/sbin/pfed -N
2669 ?? S 0:00.05 /usr/sbin/snmpd -N
2670 ?? S 0:00.01 /usr/sbin/mib2d -N
2671 ?? S 0:00.02 /usr/sbin/dfwd -N
2675 ?? S 0:00.13 /usr/sbin/ksyncd -N
2699 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
2700 ?? R 0:00.00 /bin/ps -ax
1138 d0- S 0:00.00 /usr/sbin/usbd -N
1156 d0- S 0:00.37 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0

```

lcc2-re0:

```

-----
PID TT STAT TIME COMMAND
 0 ?? Wls 0:00.00 [swapper]
 1 ?? ILs 0:00.18 /packages/mnt/jbase/sbin/init --
 2 ?? DL 0:00.01 [g_event]
 3 ?? DL 0:00.17 [g_up]
 4 ?? DL 0:00.12 [g_down]
 5 ?? DL 0:00.00 [thread taskq]
 6 ?? DL 0:00.00 [kqueue taskq]
 7 ?? DL 0:00.00 [pagedaemon]
 8 ?? DL 0:00.00 [vmdaemon]
 9 ?? DL 0:01.77 [pagezero]
10 ?? DL 0:00.00 [ktrace]
11 ?? RL 17:19.13 [idle]
12 ?? WL 0:00.36 [swi2: net]
13 ?? WL 0:01.20 [swi7: clock sio]
14 ?? WL 0:00.00 [swi6: vm]
15 ?? DL 0:00.13 [yarrow]
16 ?? WL 0:00.00 [swi9: +]
17 ?? WL 0:00.00 [swi8: +]
18 ?? WL 0:00.00 [swi5: cambio]
19 ?? WL 0:00.00 [swi9: task queue]
20 ?? WL 0:03.03 [irq10: bcm0 uhci1*]
21 ?? WL 0:00.02 [irq11: cb0 uhci0+*]
22 ?? DL 0:00.00 [usb0]
23 ?? DL 0:00.00 [usbtask]
24 ?? DL 0:00.00 [usb1]
25 ?? DL 0:00.05 [usb2]
26 ?? DL 0:00.00 [usb3]
27 ?? DL 0:00.00 [usb4]
28 ?? DL 0:00.00 [usb5]
29 ?? DL 0:00.04 [usb6]
30 ?? DL 0:00.00 [usb7]
31 ?? WL 0:00.00 [irq14: ata0]
32 ?? WL 0:00.00 [irq15: ata1]
33 ?? WL 0:00.00 [irq1: atkbd0]
34 ?? WL 0:00.00 [swi0: sio]
35 ?? WL 0:00.00 [swi3: ip6opt ipopt]
36 ?? WL 0:00.00 [swi4: ip6mismatch+]
37 ?? WL 0:00.00 [swi1: ipfwd]
38 ?? DL 0:00.00 [bufdaemon]
39 ?? DL 0:00.00 [vn1ru]
40 ?? DL 0:00.01 [syncer]
41 ?? DL 0:00.00 [softdepflush]
42 ?? DL 0:00.00 [netdaemon]

```

```

43 ?? DL 0:00.00 [vmuncachedaemon]
44 ?? DL 0:00.00 [if_pic_listen]
45 ?? DL 0:00.02 [vmkmemdaemon]
46 ?? DL 0:00.01 [cb_poll]
47 ?? DL 0:00.00 [if_pfe_listen]
48 ?? DL 0:00.00 [scs_housekeeping]
49 ?? IL 0:00.00 [kern_dump_proc]
50 ?? IL 0:00.00 [nfsiod 0]
51 ?? IL 0:00.00 [nfsiod 1]
52 ?? IL 0:00.00 [nfsiod 2]
53 ?? IL 0:00.00 [nfsiod 3]
54 ?? DL 0:00.02 [schedcpu]
55 ?? DL 0:00.75 [md0]
77 ?? DL 0:03.48 [md1]
98 ?? DL 0:00.59 [md2]
116 ?? DL 0:00.02 [md3]
137 ?? DL 0:00.56 [md4]
158 ?? DL 0:00.15 [md5]
179 ?? DL 0:00.00 [md6]
215 ?? DL 0:00.03 [md7]
225 ?? DL 0:00.03 [md8]
1052 ?? DL 0:00.00 [jsr_kkcm]
1337 ?? SL 0:00.09 [bcmTX]
1338 ?? SL 0:00.10 [bcmXGS3AsyncTX]
1339 ?? SL 0:03.22 [bcmLINK.0]
1344 ?? Is 0:00.00 /usr/sbin/cron
1496 ?? S 0:00.00 /sbin/watchdog -t-1
1497 ?? S 0:00.05 /usr/libexec/bslockd -mp -N
1498 ?? S 0:00.01 /usr/sbin/tnetd -N
1500 ?? R 0:05.17 /usr/sbin/chassisd -N
1501 ?? S 0:00.04 /usr/sbin/alarmd -N
1502 ?? I 0:00.39 /usr/sbin/craftd -N
1503 ?? S 0:00.08 /usr/sbin/mgd -N
1506 ?? I 0:00.05 /usr/sbin/inetd -N
1507 ?? I 0:00.00 /usr/sbin/tnp.snmpd -N
1508 ?? I 0:00.00 /usr/sbin/tnp.sntpd -N
1510 ?? S 0:00.01 /usr/sbin/smartd -N
1514 ?? I 0:00.07 /usr/sbin/jcsd -N
1515 ?? S 0:00.17 /usr/sbin/idpd -N
1516 ?? I 0:00.00 /usr/libexec/getty Pc ttyv0
2591 ?? DL 0:00.01 [peer proxy]
2592 ?? DL 0:00.01 [peer proxy]
2593 ?? DL 0:00.01 [peer proxy]
2597 ?? DL 0:00.00 [peer proxy]
3192 ?? S 0:00.01 /usr/sbin/irsd -N
3193 ?? S 0:00.05 /usr/sbin/snmpd -N
3194 ?? S 0:00.02 /sbin/dcd -N
3195 ?? S 0:00.01 /usr/sbin/pfed -N
3196 ?? S 0:00.01 /usr/sbin/mib2d -N
3197 ?? S 0:00.02 /usr/sbin/dfwd -N
3198 ?? S 0:00.13 /usr/sbin/ksyncd -N
3228 ?? Ss 0:00.00 mgd: (mgd) (root) (mgd)
3229 ?? R 0:00.00 /bin/ps -ax
1138 d0- S 0:00.00 /usr/sbin/usbd -N
1156 d0- S 0:00.42 /usr/sbin/eventd -N -r -s -A
1517 d0 Is+ 0:00.00 /usr/libexec/getty std.9600 ttyd0
...

```

#### show system processes sfc (TX Matrix Plus Router)

```
user@host> show system processes sfc 0
```



sfc0-re0:

```

-----
PID  TT  STAT      TIME COMMAND
  0  ??  Wls      0:00.00 [swapper]
  1  ??  SLs      0:00.18 /packages/mnt/jbase/sbin/init --
  2  ??  DL       0:00.20 [g_event]
  3  ??  DL       0:00.39 [g_up]
  4  ??  DL       0:00.32 [g_down]
  5  ??  DL       0:00.00 [thread taskq]
  6  ??  DL       0:00.09 [kqueue taskq]
  7  ??  DL       0:00.01 [pagedaemon]
  8  ??  DL       0:00.00 [vmdaemon]
  9  ??  DL       0:06.63 [pagezero]
10  ??  DL       0:00.00 [ktrace]
11  ??  RL      312:09.00 [idle]
12  ??  WL       0:11.07 [swi2: net]
13  ??  WL       0:27.70 [swi7: clock sio]
14  ??  WL       0:00.00 [swi6: vm]
15  ??  DL       0:03.03 [yarrow]
16  ??  WL       0:00.00 [swi9: +]
17  ??  WL       0:00.00 [swi8: +]
18  ??  WL       0:00.00 [swi5: cambio]
19  ??  WL       0:00.00 [swi9: task queue]
20  ??  WL       0:11.46 [irq16: uhci0 uhci*]
21  ??  DL       0:00.00 [usb0]
22  ??  DL       0:00.00 [usbtask]
23  ??  WL       0:39.63 [irq17: uhci1 uhci*]
24  ??  DL       0:00.00 [usb1]
25  ??  WL       0:00.00 [irq18: uhci2 uhci*]
26  ??  DL       0:00.84 [usb2]
27  ??  DL       0:00.00 [usb3]
28  ??  DL       0:00.00 [usb4]
29  ??  DL       0:00.00 [usb5]
30  ??  DL       0:00.73 [usb6]
31  ??  DL       0:00.00 [usb7]
32  ??  WL       0:00.00 [irq14: ata0]
33  ??  WL       0:00.00 [irq15: ata1]
34  ??  WL       0:00.00 [irq1: atkbd0]
35  ??  WL       0:00.00 [swi0: sio]
36  ??  WL       0:00.00 [irq11: isab0]
37  ??  WL       0:00.00 [swi3: ip6opt ipopt]
38  ??  WL       0:00.00 [swi4: ip6mismatch+]
39  ??  WL       0:00.00 [swi1: ipfwd]
40  ??  DL       0:00.02 [bufdaemon]
41  ??  DL       0:00.02 [vnlr]
42  ??  DL       0:00.39 [syncer]
43  ??  DL       0:00.05 [softdepflush]
44  ??  DL       0:00.00 [netdaemon]
45  ??  DL       0:00.02 [vmuncachedaemon]
46  ??  DL       0:00.00 [if_pic_listen]
47  ??  DL       0:00.35 [vmkmemdaemon]
48  ??  DL       0:00.00 [cb_poll]
49  ??  DL       0:00.06 [if_pfe_listen]
50  ??  DL       0:00.00 [scs_housekeeping]
51  ??  IL       0:00.00 [kern_dump_proc]
52  ??  IL       0:00.00 [nfsiod 0]
53  ??  IL       0:00.00 [nfsiod 1]
54  ??  IL       0:00.00 [nfsiod 2]
55  ??  IL       0:00.00 [nfsiod 3]
56  ??  DL       0:00.37 [schedcpu]
57  ??  DL       0:00.56 [md0]

```

```

79 ?? DL 0:02.58 [md1]
100 ?? DL 0:00.03 [md2]
118 ?? DL 0:00.01 [md3]
139 ?? DL 0:00.95 [md4]
160 ?? DL 0:00.12 [md5]
181 ?? DL 0:00.00 [md6]
217 ?? DL 0:00.02 [md7]
227 ?? DL 0:00.05 [md8]
1341 ?? SL 0:01.35 [bcmTX]
1342 ?? SL 0:01.69 [bcmXGS3AsyncTX]
1343 ?? SL 0:41.57 [bcmLINK.0]
1345 ?? SL 0:33.97 [bcmLINK.1]
1350 ?? Is 0:00.01 /usr/sbin/cron
1502 ?? S 0:00.01 /sbin/watchdog -t-1
1503 ?? S 0:00.86 /usr/libexec/bslockd -mp -N
1504 ?? I 0:00.01 /usr/sbin/tnetd -N
1507 ?? S 0:01.32 /usr/sbin/alarmd -N
1508 ?? S 0:14.54 /usr/sbin/craftd -N
1509 ?? S 0:01.20 /usr/sbin/mgd -N
1512 ?? S 0:00.05 /usr/sbin/inetd -N
1513 ?? S 0:00.10 /usr/sbin/tnp.sntpd -N
1517 ?? S 0:00.11 /usr/sbin/smartd -N
1525 ?? S 0:01.11 /usr/sbin/idpd -N
1526 ?? S 0:01.43 /usr/sbin/license-check -U -M -p 10 -i 10
1527 ?? I 0:00.01 /usr/libexec/getty Pc ttyv0
1616 ?? DL 0:00.30 [peer proxy]
1617 ?? DL 0:00.32 [peer proxy]
1618 ?? DL 0:00.34 [peer proxy]
1619 ?? DL 0:00.30 [peer proxy]
2391 ?? Is 0:00.01 telnetd
7331 ?? Ss 0:00.03 telnetd
9538 ?? DL 0:01.16 [jsr_kkcm]
9613 ?? DL 0:00.18 [peer proxy]
23781 ?? Ss 0:00.01 telnetd
23926 ?? Ss 0:00.03 mgd: (mgd) (regress)/dev/ttyp2 (mgd)
36867 ?? S 0:03.14 /usr/sbin/rpd -N
36874 ?? S 0:00.08 /usr/sbin/lmpd
36876 ?? S 0:00.17 /usr/sbin/lacpd -N
36877 ?? S 0:00.15 /usr/sbin/bfdd -N
36878 ?? S 0:05.05 /usr/sbin/ppmd -N
36907 ?? S 0:26.63 /usr/sbin/chassisd -N
37775 ?? S 0:00.01 /usr/sbin/bdbrepd -N
45727 ?? S 0:00.02 /usr/sbin/xntpd -j -N -g (ntpd)
45729 ?? S 0:00.40 /usr/sbin/l2ald -N
45730 ?? S< 0:00.13 /usr/sbin/apds -N
45731 ?? SN 0:00.10 /usr/sbin/sampled -N
45732 ?? S 0:00.03 /usr/sbin/ilmid -N
45733 ?? S 0:00.09 /usr/sbin/rmopd -N
45734 ?? S 0:00.31 /usr/sbin/cosd
45735 ?? I 0:00.00 /usr/sbin/rtspd -N
45736 ?? S 0:00.06 /usr/sbin/fsad -N
45737 ?? S 0:00.05 /usr/sbin/rdd -N
45738 ?? S 0:00.10 /usr/sbin/pppd -N
45739 ?? S 0:00.05 /usr/sbin/dfcd -N
45740 ?? S 0:00.08 /usr/sbin/lfmd -N
45741 ?? S 0:00.01 /usr/sbin/mpiisoamd -N
45742 ?? I 0:00.01 /usr/sbin/sendd -N
45743 ?? S 0:00.08 /usr/sbin/appidd -N
45744 ?? S 0:00.05 /usr/sbin/mspd -N
45745 ?? S 0:00.27 /usr/sbin/jdiameterd -N
45746 ?? S 0:00.10 /usr/sbin/pfed -N

```

```

45747 ?? S      0:00.19 /usr/sbin/lpd -N
45748 ?? S      0:00.64 /sbin/dcd -N
45750 ?? S      0:00.46 /usr/sbin/mib2d -N
45751 ?? S      0:00.16 /usr/sbin/dfwd -N
45752 ?? S      0:00.15 /usr/sbin/irsd -N
45764 ?? S      0:20.60 /usr/sbin/snmpd -N
56481 ?? Ss     0:00.02 telnetd
56548 ?? Rs     0:00.19 mgd: (mgd) (regress)/dev/tty0 (mgd)
56577 ?? Ss     0:00.00 mgd: (mgd) (root) (mgd)
56578 ?? R      0:00.00 /bin/ps -ax
1142 d0- S      0:00.01 /usr/sbin/usbd -N
1160 d0- S      0:29.71 /usr/sbin/eventd -N -r -s -A
6527 d0 Is+    0:00.00 /usr/libexec/getty std.9600 ttyd0
56482 p0 Is     0:00.00 login [pam] (login)
56483 p0 S       0:00.01 -csh (csh)
56547 p0 S+     0:00.02 cli
2392 p1 Is     0:00.00 login [pam] (login)
2393 p1 I       0:00.00 -csh (csh)
2394 p1 I       0:00.00 su -
2395 p1 I+     0:00.01 -su (csh)
23782 p2 Is     0:00.00 login [pam] (login)
23881 p2 I       0:00.00 -csh (csh)
23925 p2 S+     0:00.03 cli
7332 p3 Is     0:00.00 login [pam] (login)
7333 p3 I       0:00.00 -csh (csh)
23780 p3 S+     0:00.02 telnet aj

```

#### show system processes lcc wide (TX Matrix Plus Routing Matrix)

```

user@host> show system processes lcc 2 wide
lcc2-re0:

```

PID	TT	STAT	TIME	PROVIDER	COMMAND
0	??	Wls	0:00.00	(null)	[swapper]
1	??	ILs	0:00.19		/packages/mnt/jbase/sbin/init --
2	??	DL	0:00.02		[g_event]
3	??	DL	0:00.19		[g_up]
4	??	DL	0:00.13		[g_down]
5	??	DL	0:00.00		[thread taskq]
6	??	DL	0:00.00		[kqueue taskq]
7	??	DL	0:00.00		[pagedaemon]
8	??	DL	0:00.00		[vmdaemon]
9	??	DL	0:01.77		[pagezero]
10	??	DL	0:00.00		[ktrace]
11	??	RL	20:33.81		[idle]
12	??	WL	0:00.38		[swi2: net]
13	??	WL	0:01.43		[swi7: clock sio]
14	??	WL	0:00.00		[swi6: vm]
15	??	DL	0:00.14		[yarrow]
16	??	WL	0:00.00		[swi9: +]
17	??	WL	0:00.00		[swi8: +]
18	??	WL	0:00.00		[swi5: cambio]
19	??	WL	0:00.00		[swi9: task queue]
20	??	WL	0:03.18		[irq10: bcm0 uhci1*]
21	??	WL	0:00.03		[irq11: cb0 uhci0+*]
22	??	DL	0:00.00		[usb0]
23	??	DL	0:00.00		[usbtask]
24	??	DL	0:00.00		[usb1]
25	??	DL	0:00.06		[usb2]
26	??	DL	0:00.00		[usb3]
27	??	DL	0:00.00		[usb4]

28	??	DL	0:00.00	[usb5]
29	??	DL	0:00.05	[usb6]
30	??	DL	0:00.00	[usb7]
31	??	WL	0:00.00	[irq14: ata0]
32	??	WL	0:00.00	[irq15: ata1]
33	??	WL	0:00.00	[irq1: atkbd0]
34	??	WL	0:00.00	[swi0: sio]
35	??	WL	0:00.00	[swi3: ip6opt ipopt]
36	??	WL	0:00.00	[swi4: ip6mismatch+]
37	??	WL	0:00.00	[swi1: ipfwd]
38	??	DL	0:00.00	[bufdaemon]
39	??	DL	0:00.00	[vn1ru]
40	??	DL	0:00.02	[syncer]
41	??	DL	0:00.01	[softdepflush]
42	??	DL	0:00.00	[netdaemon]
43	??	DL	0:00.00	[vmuncachedaemon]
44	??	DL	0:00.00	[if_pic_listen]
45	??	DL	0:00.03	[vmkmemdaemon]
46	??	DL	0:00.01	[cb_poll]
47	??	DL	0:00.00	[if_pfe_listen]
48	??	DL	0:00.00	[scs_housekeeping]
49	??	IL	0:00.00	[kern_dump_proc]
50	??	IL	0:00.00	[nfsiod 0]
51	??	IL	0:00.00	[nfsiod 1]
52	??	IL	0:00.00	[nfsiod 2]
53	??	IL	0:00.00	[nfsiod 3]
54	??	DL	0:00.02	[schedcpu]
55	??	DL	0:00.75	[md0]
77	??	DL	0:03.84	[md1]
98	??	DL	0:00.59	[md2]
116	??	DL	0:00.02	[md3]
137	??	DL	0:00.72	[md4]
158	??	DL	0:00.15	[md5]
179	??	DL	0:00.00	[md6]
215	??	DL	0:00.03	[md7]
225	??	DL	0:00.03	[md8]
1052	??	DL	0:00.00	[jsr_kkcm]
1337	??	SL	0:00.11	[bcmTX]
1338	??	SL	0:00.12	[bcmXGS3AsyncTX]
1339	??	SL	0:03.82	[bcmLINK.0]
1344	??	Is	0:00.00	/usr/sbin/cron
1496	??	I	0:00.00	/sbin/watchdog -t-1
1497	??	S	0:00.06	/usr/libexec/bslockd -mp -N
1498	??	I	0:00.01	/usr/sbin/tnetd -N
1500	??	S	0:09.93	/usr/sbin/chassisd -N
1501	??	S	0:00.05	/usr/sbin/alarmd -N
1502	??	I	0:00.39	/usr/sbin/craftd -N
1503	??	S	0:00.09	/usr/sbin/mgd -N
1506	??	I	0:00.05	/usr/sbin/inetd -N
1507	??	I	0:00.00	/usr/sbin/tnp.sntpd -N
1508	??	I	0:00.00	/usr/sbin/tnp.sntpc -N
1510	??	S	0:00.01	/usr/sbin/smartd -N
1514	??	I	0:00.07	/usr/sbin/jcsd -N
1515	??	S	0:00.17	/usr/sbin/idpd -N
1516	??	I	0:00.00	/usr/libexec/getty Pc ttyv0
2591	??	DL	0:00.01	[peer proxy]
2592	??	DL	0:00.01	[peer proxy]
2593	??	DL	0:00.01	[peer proxy]
2597	??	DL	0:00.01	[peer proxy]
3192	??	S	0:00.02	/usr/sbin/irsd -N
3193	??	S	0:00.05	/usr/sbin/snmpd -N

```

3194 ?? S      0:00.04      /sbin/dcd -N
3195 ?? I      0:00.01      /usr/sbin/pfed -N
3196 ?? S      0:00.02      /usr/sbin/mib2d -N
3197 ?? I      0:00.03      /usr/sbin/dfwd -N
3198 ?? S      0:00.15      /usr/sbin/ksyncd -N
3559 ?? Ss    0:00.00      mgd: (mgd) (root) (mgd)
3560 ?? R      0:00.00      /bin/ps -ax -Jpww
1138 d0- S      0:00.00      /usr/sbin/usbd -N
1156 d0- S      0:00.50      /usr/sbin/eventd -N -r -s -A
1517 d0 Is+    0:00.00      /usr/libexec/getty std.9600 ttyd0

```

### show system processes (QFX Series)

```

user@switch> show system processes
PID TT STAT      TIME COMMAND
  0 ?? WLS -2341043:-31.01 [swapper]
  1 ?? SLs  0:01.34 /packages/mnt/jbase/sbin/init --
  2 ?? DL   2:48.31 [g_event]
  3 ?? DL   1:47.44 [g_up]
  4 ?? DL   1:37.82 [g_down]
  5 ?? DL   0:00.00 [kdm_tcp_poller]
  6 ?? DL   0:00.00 [thread taskq]
  7 ?? DL   0:04.86 [kqueue taskq]
  9 ?? DL   0:03.94 [pagedaemon]
 10 ?? DL   0:00.00 [ktrace]
 11 ?? RL   0:00.00 [idle: cpu31]
 12 ?? RL   0:00.00 [idle: cpu30]
 13 ?? RL   0:00.00 [idle: cpu29]
 14 ?? RL   0:00.00 [idle: cpu28]
 15 ?? RL   0:00.00 [idle: cpu27]
 16 ?? RL   0:00.00 [idle: cpu26]
 17 ?? RL   0:00.00 [idle: cpu25]
 18 ?? RL   0:00.00 [idle: cpu24]
 19 ?? RL   0:00.00 [idle: cpu23]
 20 ?? RL   0:00.00 [idle: cpu22]
 21 ?? RL   0:00.00 [idle: cpu21]
 22 ?? RL   0:00.00 [idle: cpu20]
 23 ?? RL   0:00.00 [idle: cpu19]
 24 ?? RL   0:00.00 [idle: cpu18]
 25 ?? RL   0:00.00 [idle: cpu17]
 26 ?? RL   0:00.00 [idle: cpu16]
 27 ?? RL   0:00.00 [idle: cpu15]
 28 ?? RL   0:00.00 [idle: cpu14]
 29 ?? RL   0:00.00 [idle: cpu13]
 30 ?? RL   0:00.00 [idle: cpu12]
 31 ?? RL   0:00.00 [idle: cpu11]
 32 ?? RL   0:00.00 [idle: cpu10]
 33 ?? RL   0:00.00 [idle: cpu9]
 34 ?? RL 18184:07.25 [idle: cpu8]
 35 ?? RL   0:00.00 [idle: cpu7]
 36 ?? RL 17862:11.31 [idle: cpu6]
 37 ?? RL 19343:45.16 [idle: cpu5]
 38 ?? RL 5192:38.30 [idle: cpu4]
 39 ?? RL   0:00.00 [idle: cpu3]
 40 ?? RL 19278:02.24 [idle: cpu2]
 41 ?? RL 19291:00.72 [idle: cpu1]
 42 ?? RL 18910:31.21 [idle: cpu0]
 43 ?? WL   19:03.74 [swi2: net]
 44 ?? WL 261:43.82 [swi7: clock sio]
 45 ?? WL   0:00.00 [swi6: vm]
 46 ?? DL   2:18.57 [yarrow]

```

```

47 ?? WL 0:00.00 [swi9: +]
48 ?? WL 0:00.00 [swi8: +]
49 ?? WL 0:12.36 [swi5: cambio]
50 ?? WL 0:00.00 [swi9: task queue]
51 ?? WL 0:00.00 [swi0: sio]
52 ?? WL 0:32.40 [irq39: ehci0]
53 ?? DL 0:00.21 [usb0]
54 ?? DL 0:00.00 [usbtask]
55 ?? WL 0:00.00 [irq22: xlr_lbus0]
56 ?? WL 0:00.00 [irq38: xlr_lbus0]
57 ?? WL 0:00.00 [swi3: ip6opt ipopt]
58 ?? WL 0:00.00 [swi4: ip6mismatch+]
59 ?? WL 0:00.00 [swi1: ipfwd]
60 ?? DL 0:18.65 [pagezero]
61 ?? DL 0:18.59 [bufdaemon]
62 ?? DL 1:10.44 [vnlr_u_mem]
63 ?? DL 1:51.66 [syncer]
64 ?? DL 0:20.22 [vnlr_u]
65 ?? DL 0:40.48 [softdepflush]
66 ?? DL 0:00.00 [netdaemon]
67 ?? DL 20:47.67 [vmkmemdaemon]
68 ?? DL 0:00.00 [if_pfe_listen]
69 ?? SL 0:02.80 [kdm_checkkcore]
70 ?? SL 0:03.34 [kdm_savekcore]
71 ?? SL 0:04.31 [kdm_livekcore]
72 ?? SL 0:06.14 [kdm_logger]
73 ?? SL 0:04.31 [kdm_kdb]
74 ?? SL 0:00.02 [devrt_kernel_thread]
75 ?? DL 0:21.54 [vmuncachedaemon]
76 ?? DL 0:00.00 [if_pic_listen0]
77 ?? SL 0:00.00 [nfsiod 0]
78 ?? SL 0:00.00 [nfsiod 1]
79 ?? SL 0:00.00 [nfsiod 2]
80 ?? SL 0:00.00 [nfsiod 3]
81 ?? WL 5:59.98 [irq13: +]
82 ?? RL 105:06.81 [pkt_sender: cpu0]
83 ?? DL 0:03.62 [md0]
95 ?? DL 0:37.04 [md1]
115 ?? DL 0:06.01 [md2]
135 ?? DL 0:00.75 [md3]
155 ?? DL 0:21.17 [md4]
175 ?? DL 0:01.90 [md5]
195 ?? DL 0:06.26 [md6]
231 ?? DL 0:00.01 [md7]
755 ?? Ss 0:04.17 /usr/sbin/cron
847 ?? S 0:00.10 /usr/sbin/tinetd -N
849 ?? S 0:06.82 /usr/sbin/mgd -N
850 ?? S 0:00.32 /usr/sbin/inetd -N
852 ?? S 1:05.34 /usr/sbin/dhcpd -N
853 ?? S 0:00.18 /usr/sbin/inetd -p /var/run/inetd_4.pid -N -JU __juni
855 ?? L 1181:02.21 /usr/sbin/dc-pfe -N (pafxpc)
857 ?? S 17:55.86 /usr/sbin/vccpd -N
896 ?? S 93:43.45 /usr/sbin/chassism -N
953 ?? S 0:02.89 /sbin/watchdog -t-1
954 ?? S 3:34.00 /sbin/dcd -N
955 ?? S 10:30.13 /usr/sbin/chassisd -N
956 ?? DL 0:00.21 [peer proxy]
957 ?? S 4:07.43 /usr/sbin/alarmd -N
958 ?? S 0:31.69 /usr/sbin/craftd -N
959 ?? S 0:55.16 /usr/sbin/mib2d -N
960 ?? S 3:40.64 /usr/sbin/rpd -N

```

```

961 ?? S      0:00.03 /usr/sbin/tnp.sntpd -N
962 ?? S      0:51.94 /usr/sbin/pfed -N
963 ?? S      0:47.31 /usr/sbin/rmopd -N
964 ?? S      0:33.65 /usr/sbin/cosd
965 ?? S      1:48.41 /usr/sbin/ppmd -N
966 ?? S      0:07.18 /usr/sbin/dfwd -N
967 ?? S      1:02.56 /usr/sbin/bfdd -N
968 ?? S      0:00.63 /usr/sbin/rdd -N
969 ?? S      0:40.61 /usr/sbin/dfcd -N
971 ?? S      0:07.81 /usr/sbin/bdbrepd -N
972 ?? S      0:00.28 /usr/sbin/sendd -N
973 ?? S      1:37.69 /usr/sbin/xntpd -j -N -g -JU __juniper_private4__ (nt
974 ?? S      5:56.28 /usr/sbin/snmpd -N -JU __juniper_private4__
975 ?? S      16:46.82 /usr/sbin/jdiameterd -N
976 ?? S      2:34.13 /usr/sbin/eswd -N
977 ?? S      1:03.05 /usr/sbin/sflowd -N
978 ?? S      0:22.30 /usr/sbin/fcd -N
979 ?? S      1:07.01 /usr/sbin/vccpdf -N
982 ?? S      0:25.25 /usr/sbin/mcsnoopd -N
983 ?? S      3:45.68 /usr/sbin/rpdf -N
1043 ?? S      0:37.87 /usr/sbin/lacpd -N
1048 ?? DL     0:01.29 [peer proxy]
1111 ?? WL     0:00.00 [swi2: FMNITHRD+]
1112 ?? DL     0:00.03 [peer proxy]
12816 ?? S     15:35.32 /usr/sbin/sfid -N
30893 ?? Ss    0:00.65 sshd: tlewis@tty0 (sshd)
30897 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty0 (mgd)
30905 ?? Ss    0:00.64 sshd: tlewis@tty1 (sshd)
30909 ?? Ss    0:00.15 mgd: (mgd) (tlewis)/dev/tty1 (mgd)
30910 ?? Ss    0:01.26 sshd: tcheng@tty2 (sshd)
30914 ?? Ss    0:00.80 mgd: (mgd) (tcheng)/dev/tty2 (mgd)
30937 ?? R      0:00.03 /bin/ps -ax
    661 d0- S    0:21.24 /usr/sbin/eventd -N -r -s -A
    860 d0 Ss+   0:00.07 /usr/libexec/getty std.9600 ttyd0
30896 p0 Ss+   0:00.55 -cli (cli)
30908 p1 Ss+   0:00.50 -cli (cli)
30913 p2 Ss+   0:00.85 -cli (cli)

```

## show system reboot

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1078</a> <a href="#">Syntax (EX Series Switches) on page 1078</a> <a href="#">Syntax (TX Matrix Router) on page 1078</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1078</a> <a href="#">Syntax (MX Series Router) on page 1078</a> <a href="#">Syntax (QFX Series) on page 1078</a>
<b>Syntax</b>	show system reboot <both-routing-engines>
<b>Syntax (EX Series Switches)</b>	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system reboot <all-chassis   all-lcc   lcc <i>number</i>   scc> <both-routing-engines>
<b>Syntax (TX Matrix Plus Router)</b>	show system reboot <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> > <both-routing-engines>
<b>Syntax (MX Series Router)</b>	show system reboot <all-members> <both-routing-engines> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system reboot <both-routing-engines> <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-device <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display pending system reboots or halts.
<b>Options</b>	<b>none</b> —Display pending reboots or halts on the active Routing Engine.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for all the T640 routers in the chassis that are connected to the TX Matrix router. On a TX Matrix router, display halt or reboot request information for all the T1600 or T4000 routers in the chassis that are connected to the TX Matrix Plus router.



**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for all members of the Virtual Chassis configuration.

**all-lcc**—(TX Matrix routers and TX Matrix Plus router only) (Optional) On a TX Matrix router, display system halt or reboot request information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for all connected T1600 or T4000 LCCs.

**both-routing-engines**—(Systems with multiple Routing Engines) (Optional) Display halt or reboot request information on both Routing Engines.

**infrastructure *name***—(QFabric systems only) (Optional) Display reboot request information on the fabric manager Routing Engines and fabric control Routing Engines.

**interconnect-device *name***—(QFabric systems only) (Optional) Display reboot request information on the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display halt or reboot request information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display halt or reboot request information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display halt or reboot request information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display reboot request information on the Node group.

**scc**—(TX Matrix router only) (Optional) Display halt or reboot request information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus router only) (Optional) Display halt or reboot request information for the TX Matrix Plus router.

**Additional Information** By default, when you issue the **show system reboot** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

**Required Privilege Level** maintenance

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system reboot on page 1080](#)  
[show system reboot all-lcc \(TX Matrix Router\) on page 1080](#)  
[show system reboot sfc \(TX Matrix Plus Router\) on page 1080](#)  
[show system reboot \(QFX3500 Switch\) on page 1080](#)

## Sample Output

### [show system reboot](#)

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

### [show system reboot all-lcc \(TX Matrix Router\)](#)

```
user@host> show system reboot all-lcc
lcc0-re0:
```

```
-----
No shutdown/reboot scheduled.
```

```
lcc2-re0:
```

```
-----
No shutdown/reboot scheduled.
```

### [show system reboot sfc \(TX Matrix Plus Router\)](#)

```
user@host> show system sfc 0
No shutdown/reboot scheduled.
```

### [show system reboot \(QFX3500 Switch\)](#)

```
user@switch> show system reboot
No shutdown/reboot scheduled.
```

## show system resource-cleanup processes

<b>Syntax</b>	show system resource-cleanup processes <detail> <pid <i>number</i> > <process-name <i>name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.3. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the list of processes that have been registered for resource cleanup services.
<b>Options</b>	<p><b>detail</b>—(Optional) Display the list of processes that have been registered for resource cleanup services, along with the resources that have been requested for cleanup.</p> <p><b>pid <i>number</i></b>—(Optional) Display a process that has been registered for resource cleanup services by specifying the Process Identifier number.</p> <p><b>process-name <i>name</i></b>—(Optional) Display a process that has been registered for resource cleanup services by name of the process.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>resource-cleanup</i></li> <li><i>traceoptions (Resource Cleanup)</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system resource-cleanup processes on page 1081</a> <a href="#">show system resource-cleanup processes detail on page 1082</a>
<b>Output Fields</b>	For a description of the output fields, see <a href="#">Table 62 on page 1081</a> . Output fields are listed in the approximate order in which they appear.

**Table 62: show system resource-cleanup processes Output Fields**

Field Name	Field Description
<b>PID</b>	Process ID, a number that identifies a process.
<b>Process name</b>	String that identifies the process.
<b>Resources to clean</b>	Resources that have been registered to be cleaned up.

## Sample Output


### show system resource-cleanup processes

```
user@host> show system resource-cleanup processes
PID      Process name      Resources to clean
420      jnx-exampld       GENCFG, SYSV shared memory
```

### show system resource-cleanup processes detail

```
user@host> show system resource-cleanup processes detail
PID      Process name      Resources to clean
420      jnx-exampld        GENCFG blob major ID 0x8000, minor ID 0x0000
          SYSV shared memory ID 65536, key 1108955839
          SYSV shared memory ID 65537, key 1108955837
```

## show system rollback

<b>Syntax</b>	<code>show system rollback <i>number</i></code> <code>&lt;compare <i>number</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of a previously committed configuration, or the differences between two previously committed configurations.
<div>  <b>NOTE:</b> The <code>show system rollback</code> command is a purely operational mode command and cannot be issued with <code>run</code> from the configuration mode.         </div>	
<b>Options</b>	<p><b><i>number</i></b>—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.</p> <p><b><code>compare <i>number</i></code></b>—(Optional) Number of another previously committed (rollback) configuration to compare to rollback <b><i>number</i></b>. The output displays the differences between the two configurations. The range of values is 0 through 49.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show system rollback compare on page 1083</a>

## Sample Output

### show system rollback compare

```

user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;
+       }
+     }
+   }
+ }

```

```
+      }
+    }
+    ge-1/3/0 {
+      unit 0 {
+        family inet {
+          filter {
+            input mf_plp;
+          }
+          address 12.1.1.1/30;
+        }
+      }
+    }
+  }
+}
```

## show system services service-deployment

<b>Syntax</b>	show system services service-deployment
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about a Session and Resource Control (SRC) client.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	system view
<b>List of Sample Output</b>	<a href="#">show system services service-deployment on page 1085</a>
<b>Output Fields</b>	<a href="#">Table 63 on page 1085</a> lists the output fields for the <b>show system services service-deployment</b> command. Output fields are listed in the approximate order in which they appear.

**Table 63: show system services service-deployment Output Fields**

Field Name	Field Description
PDT Keepalive settings	Configured PDT keepalive interval, in seconds.
Keepalives sent	Number of keepalives sent.
Notifications sent	Number of notifications sent.
Last update from peer	Time at which the last update from a peer was received.

## Sample Output

### show system services service-deployment

```

user@host> show system services service-deployment
Connected to 192.4.4.4 port 10288 since 2004-05-03 11:04:34 PDT Keepalive settings:
Interval 15 seconds Keepalives sent: 750 Notifications sent: 0 Last update from
peer: 00:00:06 ago

```

## show system software

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1086</a> <a href="#">Syntax (EX Series Switches) on page 1086</a> <a href="#">Syntax (TX Matrix Router) on page 1086</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1086</a> <a href="#">Syntax (J Series Routers) on page 1086</a> <a href="#">Syntax (QFX Series) on page 1086</a>
<b>Syntax</b>	<code>show system software</code> <code>&lt;detail&gt;</code>
<b>Syntax (EX Series Switches)</b>	<code>show system software</code> <code>&lt;all-members&gt;</code> <code>&lt;detail&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>show system software</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;detail&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>show system software</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;detail&gt;</code>
<b>Syntax (J Series Routers)</b>	<code>show system software</code> <code>&lt;backup&gt;</code> <code>&lt;detail&gt;</code>
<b>Syntax (QFX Series)</b>	<code>show system software</code> <code>&lt;detail&gt;</code> <code>&lt;infrastructure <i>name</i>&gt;</code> <code>&lt;interconnect-device <i>name</i>&gt;</code> <code>&lt;node-group <i>name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the Junos OS extensions loaded on your router or switch.
<b>Options</b>	<b>none</b> —Display standard information about all loaded Junos OS extensions.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system software information for all the T640 routers (TX Matrix Router) or all the routers (TX Matrix Plus Router) in the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for all T640 routers connected to the



TX Matrix router. On a TX Matrix Plus router, display system software information for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches only) (Optional) Display the system software running on all members of the Virtual Chassis configuration.

**backup**—(J Series routers only) (Optional) Display the status of old system software packages only.

**detail**—(Optional) Display detailed information about available Junos OS extensions.

**infrastructure name**—(QFabric systems only) (Optional) Display the system software running on the fabric control Routing Engine and the fabric manager Routing Engine.

**interconnect-device name**—(QFabric systems only) (Optional) Display the system software running on the Interconnect device.

**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system software information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system software information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Display the system software running on the local Virtual Chassis member.

**member member-id**—(EX4200 switches only) (Optional) Display the system software running on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**node-group name**—(QFabric systems only) (Optional) Display the system software running on the Node group.

**scc**—(Routing matrix only) (Optional) Display the system software running on a TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display system software information for the TX Matrix Plus router.

**Required Privilege Level** maintenance

Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li></ul>
List of Sample Output	<a href="#">show system software on page 1088</a> <a href="#">show system software (TX Matrix Plus Router) on page 1088</a> <a href="#">show system software (QFX Series) on page 1092</a>
Output Fields	When you enter this command, you are provided a list of Junos OS packages installed on the router and their corresponding Junos OS release number.

## Sample Output

### [show system software](#)

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]

Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]
Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]

Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]

Information for jroute:

Comment:
JUNOS Routing Software Suite [7.2R1.7]

Information for junos:

Comment:
JUNOS Base OS boot [7.2R1.7]
```

### [show system software \(TX Matrix Plus Router\)](#)

```
user@host> show system software
sfc0-re0:
-----
Information for jbase:
```

Comment:  
JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:  
JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:  
JUNOS Online Documentation [9.6-20090515.0]  
Information for jkernel:

Comment:  
JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:  
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:  
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:Comment:  
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aac1:

Comment:  
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:  
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:  
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:

JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:

JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:

JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:

JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:

JUNOS Base OS boot [9.6-20090515.0]

...

lcc0-re0:

-----  
Information for jbase:

Comment:

JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [9.6-20090515.0]

Information for jdocs:

Comment:

JUNOS Online Documentation [9.6-20090515.0]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [9.6-20090515.0]

Information for jpfe:

Comment:  
JUNOS Packet Forwarding Engine Support (T-Series) [9.6-20090515.0]

Information for jpfe-common:

Comment:  
JUNOS Packet Forwarding Engine Support (M/T Common) [9.6-20090515.0]

Information for jroute:

Comment:  
JUNOS Routing Software Suite [9.6-20090515.0]

Information for jservices-aacl:

Comment:  
JUNOS Services ACL Container package [9.6-20090515.0]

Information for jservices-appid:

Comment:  
JUNOS AppId Services [9.6-20090515.0]

Information for jservices-bgf:

Comment:  
JUNOS Border Gateway Function package [9.6-20090515.0]

Information for jservices-idp:

Comment:  
JUNOS IDP Services [9.6-20090515.0]

Information for jservices-llpdf:

Comment:  
JUNOS Services LL-PDF Container package [9.6-20090515.0]

Information for jservices-sfw:

Comment:  
JUNOS Services Stateful Firewall [9.6-20090515.0]

Information for jservices-voice:

Comment:

JUNOS Voice Services Container package [9.6-20090515.0]

Information for junos:

Comment:

JUNOS Base OS boot [9.6-20090515.0]

lcc1-re0:

-----  
Information for jbase:

Comment:

JUNOS Base OS Software Suite [9.6-20090515.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [9.6-20090515.0]

...

### show system software (QFX Series)

user@switch> **show system software**

Information for jbase:

Comment:

JUNOS Base OS Software Suite [11.3-20110730.0]

Information for jcrypto:

Comment:

JUNOS Crypto Software Suite [11.3-20110730.0]

Information for jdocs:

Comment:

JUNOS Online Documentation [11.3-20110730.0]

Information for jkernel:

Comment:

JUNOS Kernel Software Suite [11.3-20110730.0]

Information for jpfe:

Comment:

JUNOS Packet Forwarding Engine Support (QFX) [11.3-20110730.0]

Information for jroute:

Comment:

JUNOS Routing Software Suite [11.3-20110730.0]

Information for jswitch:

Comment:

JUNOS Enterprise Software Suite [11.3-20110730.0]

Information for junos:

Comment:

JUNOS Base OS boot [11.3-20110730.0]

Information for jweb:

Comment:

JUNOS Web Management [11.3-20110730.0]

## show system statistics

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1094</a> <a href="#">Syntax (EX Series Switches) on page 1094</a> <a href="#">Syntax (TX Matrix Router) on page 1094</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1094</a> <a href="#">Syntax (MX Series Router) on page 1094</a> <a href="#">Syntax (QFX Series) on page 1094</a>
<b>Syntax</b>	show system statistics
<b>Syntax (EX Series Switches)</b>	show system statistics <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system statistics <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system statistics <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system statistics <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system statistics
<b>Release Information</b>	Command introduced before JUNOS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display system-wide protocol-related statistics.
<b>Options</b>	<b>none</b> —Display system statistics for all the following protocols: <ul style="list-style-type: none"><li>• <b>arp</b>—Address Resolution Protocol</li><li>• <b>bridge</b>—IEEE 802.1 Bridging</li><li>• <b>clns</b>—Connectionless Network Service</li><li>• <b>esis</b>—End System-to-Intermediate System</li><li>• <b>ethoamcfm</b>—Ethernet OAM protocol for connectivity fault management</li><li>• <b>ethoamlfm</b>—Ethernet OAM protocol for link fault management</li><li>• <b>icmp</b>—Internet Control Message Protocol</li><li>• <b>icmp6</b>—Internet Control Message Protocol version 6</li><li>• <b>igmp</b>—Internet Group Management Protocol</li></ul>



- **ip**—Internet Protocol version 4
- **ip6**—Internet Protocol version 6
- **mpls**—Multiprotocol Label Switching
- **rdp**—Reliable Datagram Protocol
- **tcp**—Transmission Control Protocol
- **tnp**—Trivial Network Protocol
- **ttp**—TNP Tunneling Protocol
- **tudp**—Trivial User Datagram Protocol
- **udp**—User Datagram Protocol
- **vpls**—Virtual Private LAN Service

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for all the routers in the chassis.

**all-lcc**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for all routers (line-card chassis) connected to the TX Matrix Plus router

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for all members of the Virtual Chassis configuration.

**lcc number**—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for a protocol for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for a protocol for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the local Virtual Chassis member.

**member member-id**—(EX4200 switches and MX Series routers only) (Optional) Display system statistics for a protocol for the specified member of the Virtual Chassis

configuration. For EX4200 switches, replace **member-id** with a value from 0 through 9. For an MX Series Virtual Chassis, replace **member-id** with a value of 0 or 1.

**scc**—(TX Matrix routers only) (Optional) Display system statistics for a protocol for the TX Matrix router (or switch-card chassis).

**sfc number**—(TX Matrix Plus routers only) (Optional) Display system statistics for a protocol for the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with 0.

**Additional Information** By default, when you issue the **show system statistics** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on the TX Matrix router) or T1600 (in a routing matrix based on the TX Matrix Plus router) backup Routing Engines that are connected to it.

**Required Privilege Level** view

**List of Sample Output** [show system statistics on page 1096](#)  
[show system statistics \(EX Series Switches\) on page 1103](#)  
[show system statistics \(TX Matrix Router\) on page 1112](#)  
[show system statistics \(QFX Series\) on page 1119](#)

## Sample Output

### show system statistics

```
user@host> show system statistics
ip:
    3682087 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped (queue overflow)
    0 fragments dropped after timeout
    0 fragments dropped due to over limit
    0 packets reassembled ok
    3664774 packets for this host
    17316 packets for unknown/unsupported protocol
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    6528 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs
    0 output packets discarded due to no route
```

```

0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
1123 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
1123 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
    echo reply: 75
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 75
    router advertisement: 130
75 message responses generated
tcp:
3844 packets sent
    3618 data packets (1055596 bytes)
    0 data packets (0 bytes) retransmitted
    0 resends initiated by MTU discovery
    205 ack-only packets (148 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    1079 control packets
5815 packets received
    3377 acks (for 1055657 bytes)
    24 duplicate acks
    0 acks for unsent data
    2655 packets (15004 bytes) received in-sequence
    1 completely duplicate packet (0 bytes)
    0 old duplicate packets
    0 packets with some dup. data (0 bytes duped)
    0 out-of-order packets (0 bytes)
    0 packets (0 bytes) of data after window
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
1 connection request
32 connection accepts
0 bad connection attempts

```

```
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
    27 connections updated cached RTT on close
    27 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
    0 connections dropped by rexmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
344 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    32 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
1058 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors

udp:
3658884 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
3657342 dropped due to no socket
3657342 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
4291311496 delivered
1551 datagrams output

ipsec:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
```

```
0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route

igmp:
17186 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

arp:
44181302 datagrams received
2 ARP requests received
2028 ARP replies received
3156 resolution requests received
0 unrestricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 with bogus interface
787 with incorrect length
712 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
7611 with multicast target address
0 with my own hardware address
14241699 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29929250 which were not for me
0 packets discarded waiting for resolution
6 packets sent after waiting for resolution
17812 ARP requests sent
2 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry

ip6:
0 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
```

```
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
Mbuf statistics:
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too may headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol

icmp6:
0 calls to icmp_error
0 errors not generated because old message was icmp error or so
0 errors not generated because rate limitation
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length
Histogram of error messages to be generated:
    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options

ipsec6:
0 inbound packets processed successfully
0 inbound packets violated process security policy
0 inbound packets with no SA available
0 invalid inbound packets
0 inbound packets failed due to insufficient memory
0 inbound packets failed getting SPI
0 inbound packets failed on AH replay check
0 inbound packets failed on ESP replay check
0 inbound AH packets considered authentic
0 inbound AH packets failed on authentication
0 inbound ESP packets considered authentic
0 inbound ESP packets failed on authentication
0 outbound packets processed successfully
0 outbound packets violated process security policy
0 outbound packets with no SA available
```

```
0 invalid outbound packets
0 outbound packets failed due to insufficient memory
0 outbound packets with no route

c1n1:
0 total packets received
0 packets delivered
0 too small
0 bad header length
0 bad checksum
0 bad version
0 unknown or unsupported protocol
0 bogus sdl size
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 address fields were not reasonable
0 segment information forgotten
0 forwarded packets
0 total packets sent
0 output packets discarded
0 non-forwarded packets
0 packets fragmented
0 fragments sent
0 fragments discarded
0 fragments timed out
0 fragmentation prohibited
0 packets reconstructed
0 packets destined to dead nexthop
0 packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure

esis:
0 total pkts received
0 total packets consumed by protocol
0 pdus received with bad checksum
0 pdus received with bad version number
0 pdus received with bad type field
0 short pdus received
0 bogus sdl size
0 bad header length
0 unknown or unsupported protocol
0 no free memory in socket buffer
0 send packets discarded
0 sbappend failure
0 mcopy failure
0 ISO family not configured

tnp:
146776365 unicast packets received
0 broadcast packets received
0 fragmented packets received
0 hello packets dropped
0 fragments dropped
0 fragment reassembly queue flushes
0 hello packets received
0 control packets received
49681642 rdp packets received
337175 udp packets received
96757548 tunnel packets received
0 input packets discarded with no protocol
98397591 unicast packets sent
```

```
0 broadcast packets sent
0 fragmented packets sent
0 hello packets dropped
0 fragments dropped
0 hello packets sent
0 control packets sent
49681642 rdp packets sent
337175 udp packets sent
48378774 tunnel packets sent
0 packets sent with unknown protocol

rdp:
49681642 input packets
0 discards for bad checksum
0 discards bad sequence number
0 refused connections
2031964 acks received
0 dropped due to full socket buffers
49692 retransmits
49681642 output packets
24815968 acks sent
28 connects
0 closes
22783990 keepalives received
22783990 keepalives sent

tudp:
337175 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
0 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
337175 delivered
337175 datagrams output

ttp:
398749 packets sent
0 packets sent while unconnected
0 packets sent while interface down
0 packets sent couldn't get buffer
0 packets sent couldn't find neighbor
44696687 L2 packets received
0 unknown L3 packets received
3682087 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 cyclotron cycle L3 packets received
0 cyclotron send L3 packets received
0 packets received while unconnected
0 packets received from unknown ifl
0 input packets couldn't get buffer
0 input packets with bad type
0 input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
```



```

68877 Input packets dropped based on tlv result
0 input packets for which rt lookup is bypassed

mpls:
0 total mpls packets received
0 packets forwarded
0 packets dropped
0 with header too small
0 after tagging, can't fit link MTU
0 with IPv4 explicit NULL tag
0 with IPv4 explicit NULL cksum errors
0 with router alert tag
0 lsp ping packets (ttl-expired/router alert)
0 with ttl expired
0 with tag encoding error
0 packets discarded, no route

vpls:
0 total packets received
0 with size smaller than minimum
0 with incorrect version number
0 packets for this host
0 packets with no logical interface
0 packets with no family
0 packets with no route table
0 packets with no auxiliary table
0 packets with no corefacing entry
0 packets with no CE-facing entry
0 mac route learning requests
0 mac routes learnt
0 requests to learn an existing route
0 learning requests while learning disabled on interface
0 learning requests over capacity
0 mac routes moved
0 requests to move static route
0 mac route aging requests
0 mac routes aged
0 bogus address in aging requests
0 requests to age static route
0 requests to re-ageout aged route
0 requests involving multiple peer FEs
0 aging acks from PFE
0 aging non-acks from PFE
0 aging requests timed out waiting on FEs
0 aging requests over max-rate
0 errors finding peer FEs

```

### show system statistics (EX Series Switches)

```

user@host> show system statistics
Tcp:
571779 packets sent
21517 data packets (1797102 bytes)
2 data packets retransmitted (20 bytes)
0 resends initiated by MTU discovery
3708 ack only packets (531 packets delayed)
0 URG only packets
1 window probe packets
1 window update packets
1093063 control packets
1132541 packets received
20961 acks(for 1796102 bytes)
5861 duplicate acks

```

```
0 acks for unsent data
19556 packets received in-sequence(232079 bytes)
3018 completely duplicate packets(0 bytes)
0 old duplicate packets
4 packets with some duplicate data(4 bytes duped)
2 out-of-order packets(2 bytes)
0 packets of data after window(0 bytes)
0 window probes
39 window update packets
0 packets received after close
0 discarded for bad checksums
0 discarded for bad header offset fields
0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)
546596 connections closed (including 6 drops)
    47 connections updated cached RTT on close
    47 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
3028 keepalive timeouts
    3027 keepalive probes sent
    1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
    0 retransmitted
    0 dupsyn
    4 dropped
    78 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
```

```

0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output
ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
icmp:
0 drops due to rate limit
9 calls to icmp_error

```

```
0 errors not generated because old message was icmp
Output histogram:
    295 echo reply
    9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    295 echo
295 message responses generated

igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent

raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMPv2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMPv2 packets received
0 Input drops due to bogus protocol
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket

arp:
186413 datagrams received
88 ARP requests received
88 ARP replies received
0 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
```

```

0 datagrams with multicast source address
0 datagrams with multicast source address
0 datagrams with my own hardware address
164 datagrams for an address not on the interface
0 datagrams with a broadcast source address
0 datagrams with source address duplicate to mine
186065 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
50 ARP requests sent
88 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

ip6:
0 total packets received
0 packets with size smaller than minimum
0 packets with data size < data length
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f

icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums

```

```
0 Messages with bad length
    0 No route
    0 Administratively prohibited
    0 Beyond scope
    0 Address unreachable
    0 Port unreachable
    0 packet too big
    0 Time exceed transit
    0 Time exceed reassembly
    0 Erroneous header field
    0 Unrecognized next header
    0 Unrecognized option
    0 redirect
    0 Unknown
0 Message responses generated
0 Messages with too many ND options
pfkey:
0 Requests sent from userland
0 Bytes sent from userland
histogram by message type:
    0 reserved
    0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
    0 reserved
    0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
c1n1:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
```

```

0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
  0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdu with unknown or unsupported protocol
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections

```

```
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result
0 Input packets for which rt lookup is bypassed
mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
```



```

0 Packets used first nexthop in ecmp unilist
vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket
bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop

```

```
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
0 packets with no nexthop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
  0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table
```

#### show system statistics (TX Matrix Router)

```
user@host> show system statistics
sfc0-re0:
```

-----

## Tcp:

```

361694 packets sent
    326507 data packets (103237236 bytes)
    2343 data packets retransmitted (2673324 bytes)
    0 resends initiated by MTU discovery
    33857 ack only packets (31613 packets delayed)
    0 URG only packets
    14 window probe packets
    387 window update packets
    1108 control packets
345879 packets received
    298207 acks(for 103141728 bytes)
    438 duplicate acks
    0 acks for unsent data
    204578 packets received in-sequence(13820995 bytes)
    6 completely duplicate packets(18 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    899 window update packets
    166 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
406 connection requests
233 connection accepts
0 bad connection attempts
0 listen queue overflows
616 connections established (including accepts)
911 connections closed (including 41 drops)
    346 connections updated cached RTT on close
    346 connections updated cached RTT variance on close
    200 connections updated cached ssthresh on close
23 embryonic connections dropped
298155 segments updated rtt(of 287216 attempts)
1163 retransmit timeouts
    27 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
5 keepalive timeouts
    5 keepalive probes sent
    0 connections dropped by keepalive
69922 correct ACK header predictions
34993 correct data packet header predictions
233 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    233 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received

```

- 23 SACK recovery episodes
- 68 segment retransmits in SACK recovery episodes
- 71542 byte retransmits in SACK recovery episodes
- 158 SACK options (SACK blocks) received
- 0 SACK options (SACK blocks) sent
- 0 SACK scoreboard overflow
- 0 ACKs sent in response to in-window but not exact RSTs
- 0 ACKs sent in response to in-window SYNs on established connections
- 0 rcv packets dropped by TCP due to bad address
- 0 out-of-sequence segment drops due to insufficient memory
- 259 RST packets
- 0 ICMP packets ignored by TCP
- 0 send packets dropped by TCP due to auth errors
- 0 rcv packets dropped by TCP due to auth errors
- 0 outgoing segments dropped due to policing

1cc0-re0:

-----  
Tcp:

- 346 packets sent
  - 222 data packets (22894 bytes)
  - 0 data packets retransmitted (0 bytes)
  - 0 resends initiated by MTU discovery
  - 80 ack only packets (12 packets delayed)
  - 0 URG only packets
  - 0 window probe packets
  - 5 window update packets
  - 42 control packets
- 358 packets received
  - 268 acks(for 22939 bytes)
  - 9 duplicate acks
  - 0 acks for unsent data
  - 203 packets received in-sequence(33820 bytes)
  - 0 completely duplicate packets(0 bytes)
  - 0 old duplicate packets
  - 0 packets with some duplicate data(0 bytes duped)
  - 0 out-of-order packets(0 bytes)
  - 0 packets of data after window(0 bytes)
  - 0 window probes
  - 6 window update packets
  - 0 packets received after close
  - 0 discarded for bad checksums
  - 0 discarded for bad header offset fields
  - 0 discarded because packet too short
- 13 connection requests
- 18 connection accepts
- 0 bad connection attempts
- 0 listen queue overflows
- 31 connections established (including accepts)
- 35 connections closed (including 2 drops)
  - 3 connections updated cached RTT on close
  - 3 connections updated cached RTT variance on close
  - 0 connections updated cached ssthresh on close
- 0 embryonic connections dropped
- 268 segments updated rtt(of 247 attempts)
- 0 retransmit timeouts
  - 0 connections dropped by retransmit timeout
- 0 persist timeouts
  - 0 connections dropped by persist timeout
- 0 keepalive timeouts
  - 0 keepalive probes sent

```

    0 connections dropped by keepalive
0 correct ACK header predictions
42 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

1cc1-re0:

-----  
 Tcp:

```

348 packets sent
    223 data packets (22895 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    81 ack only packets (13 packets delayed)
    0 URG only packets
    0 window probe packets
    5 window update packets
    42 control packets
360 packets received
    269 acks(for 22940 bytes)
    9 duplicate acks
    0 acks for unsent data
    203 packets received in-sequence(33820 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    6 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields

```

```
0 discarded because packet too short
13 connection requests
18 connection accepts
0 bad connection attempts
0 listen queue overflows
31 connections established (including accepts)
36 connections closed (including 2 drops)
    3 connections updated cached RTT on close
    3 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
269 segments updated rtt(of 248 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
43 correct data packet header predictions
18 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    18 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
```

1cc2-re0:

-----  
Tcp:

```
405 packets sent
    271 data packets (23926 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    86 ack only packets (13 packets delayed)
    0 URG only packets
```

```

    0 window probe packets
    5 window update packets
    46 control packets
418 packets received
    321 acks(for 23975 bytes)
    9 duplicate acks
    0 acks for unsent data
    234 packets received in-sequence(34403 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
15 connection requests
19 connection accepts
0 bad connection attempts
0 listen queue overflows
34 connections established (including accepts)
39 connections closed (including 2 drops)
    4 connections updated cached RTT on close
    4 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
321 segments updated rtt(of 299 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
0 correct ACK header predictions
48 correct data packet header predictions
19 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    19 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs

```

- 0 ACKs sent in response to in-window SYNs on established connections
- 0 rcv packets dropped by TCP due to bad address
- 0 out-of-sequence segment drops due to insufficient memory
- 5 RST packets
- 0 ICMP packets ignored by TCP
- 0 send packets dropped by TCP due to auth errors
- 0 rcv packets dropped by TCP due to auth errors
- 0 outgoing segments dropped due to policing

lcc3-re0:

-----  
Tcp:

- 346 packets sent
  - 221 data packets (22895 bytes)
  - 0 data packets retransmitted (0 bytes)
  - 0 resends initiated by MTU discovery
  - 81 ack only packets (13 packets delayed)
  - 0 URG only packets
  - 0 window probe packets
  - 5 window update packets
  - 42 control packets
- 360 packets received
  - 267 acks(for 22940 bytes)
  - 9 duplicate acks
  - 0 acks for unsent data
  - 203 packets received in-sequence(33820 bytes)
  - 0 completely duplicate packets(0 bytes)
  - 0 old duplicate packets
  - 0 packets with some duplicate data(0 bytes duped)
  - 0 out-of-order packets(0 bytes)
  - 0 packets of data after window(0 bytes)
  - 0 window probes
  - 6 window update packets
  - 0 packets received after close
  - 0 discarded for bad checksums
  - 0 discarded for bad header offset fields
  - 0 discarded because packet too short
- 13 connection requests
- 18 connection accepts
- 0 bad connection attempts
- 0 listen queue overflows
- 31 connections established (including accepts)
- 35 connections closed (including 2 drops)
  - 3 connections updated cached RTT on close
  - 3 connections updated cached RTT variance on close
  - 0 connections updated cached ssthresh on close
- 0 embryonic connections dropped
- 267 segments updated rtt(of 246 attempts)
- 0 retransmit timeouts
  - 0 connections dropped by retransmit timeout
- 0 persist timeouts
  - 0 connections dropped by persist timeout
- 0 keepalive timeouts
  - 0 keepalive probes sent
  - 0 connections dropped by keepalive
- 0 correct ACK header predictions
- 43 correct data packet header predictions
- 18 syncache entries added
  - 0 retransmitted
  - 0 dupsyn
  - 0 dropped



```

18 completed
0 bucket overflow
0 cache overflow
0 reset
0 stale
0 aborted
0 badack
0 unreach
0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
5 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

#### show system statistics (QFX Series)

```

user@switch> show system statistics
Tcp:
571779 packets sent
21517 data packets (1797102 bytes)
2 data packets retransmitted (20 bytes)
0 resends initiated by MTU discovery
3708 ack only packets (531 packets delayed)
0 URG only packets
1 window probe packets
1 window update packets
1093063 control packets
1132541 packets received
20961 acks(for 1796102 bytes)
5861 duplicate acks
0 acks for unsent data
19556 packets received in-sequence(232079 bytes)
3018 completely duplicate packets(0 bytes)
0 old duplicate packets
4 packets with some duplicate data(4 bytes duped)
2 out-of-order packets(2 bytes)
0 packets of data after window(0 bytes)
0 window probes
39 window update packets
0 packets received after close
0 discarded for bad checksums
0 discarded for bad header offset fields
0 discarded because packet too short
546519 connection requests
78 connection accepts
0 bad connection attempts
0 listen queue overflows
100 connections established (including accepts)

```

```
546596 connections closed (including 6 drops)
47 connections updated cached RTT on close
47 connections updated cached RTT variance on close
0 connections updated cached ssthresh on close
546497 embryonic connections dropped
20453 segments updated rtt(of 566914 attempts)
2 retransmit timeouts
0 connections dropped by retransmit timeout
0 persist timeouts
0 connections dropped by persist timeout
3028 keepalive timeouts
3027 keepalive probes sent
1 connections dropped by keepalive
7515 correct ACK header predictions
12258 correct data packet header predictions
78 syncache entries added
0 retransmitted
0 dupsyn
4 dropped
78 completed
0 bucket overflow
0 cache overflow
0 reset
0 stale
0 aborted
0 badack
0 unreach
0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
546544 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing
udp:
147 datagrams received
0 with incomplete header
0 with bad data length field
0 with bad checksum
9 dropped due to no socket
0 broadcast/multicast datagrams dropped due to no socket
0 dropped due to full socket buffers
0 not for hashed pcb
138 delivered
0 datagrams output
ip:
73704 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
```

```

0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped (queue overflow)
0 fragments dropped after timeout
0 fragments dropped due to over limit
0 packets reassembled ok
1133057 packets for this host
0 packets for unknown/unsupported protocol
40146 packets forwarded
0 packets not forwardable
40146 redirects sent
1121700 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
icmp:
0 drops due to rate limit
9 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
295 echo reply
9 destination unreachable
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
295 echo
295 message responses generated
igmp:
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum

```

```
0 membership queries received
0 membership queries received with invalid fields
0 membership reports received
0 membership reports received with invalid fields
0 membership reports received for groups to which we belong
0 Membership reports sent
raw_if:
0 RAW packets transmitted
0 PPPOE packets transmitted
0 ISDN packets transmitted
0 DIALER packets transmitted
0 PPP packets transmitted to pppd
0 PPP packets transmitted to jppd
0 IGMP2 packets transmitted
13 output drops due to tx error
0 MPU packets transmitted
0 PPPOE packets received
0 ISDN packets received
0 DIALER packets received
0 PPP packets received from pppd
0 MPU packets received
0 PPP packets received from jppd
0 IGMP2 packets received
0 Input drops due to bogus protocol
0 input drops due to no mbufs available
0 input drops due to no space in socket
0 input drops due to no socket
arp:
186413 datagrams received
88 ARP requests received
88 ARP replies received
0 resolution request received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requests not proxied
0 restricted proxy requests not proxied
0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast source address
0 datagrams with my own hardware address
164 datagrams for an address not on the interface
0 datagrams with a broadcast source address
0 datagrams with source address duplicate to mine
186065 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
50 ARP requests sent
88 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
```

```
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
ip6:
0 total packets received
0 packets with size smaller than minimum
0 packets with data size < data length
0 packets with bad options
0 packets with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragments that exceeded limit
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join
0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
0 forward cache hit
0 forward cache miss
0 Packets destined to dead next hop
0 option packets dropped due to rate limit
0 Packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
icmp6:
0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
0 No route
0 Administratively prohibited
0 Beyond scope
0 Address unreachable
0 Port unreachable
0 packet too big
0 Time exceed transit
0 Time exceed reassembly
0 Erroneous header field
0 Unrecognized next header
0 Unrecognized option
0 redirect
0 Unknown
0 Message responses generated
0 Messages with too many ND options
pfkey:
0 Requests sent from userland
```

```
0 Bytes sent from userland
histogram by message type:
0 reserved
0 dump
0 Messages with invalid length field
0 Messages with invalid version field
0 Messages with invalid message type field
0 Messages too short
0 Messages with memory allocation failure
0 Messages with duplicate extension
0 Messages with invalid extension type
0 Messages with invalid sa type
0 Messages with invalid address extension
0 Requests sent to userland
0 Bytes sent to userland
histogram by message type:
0 reserved
0 dump
0 Messages toward single socket
0 Messages toward all sockets
0 Messages toward registered sockets
0 Messages with memory allocation failure
c1n1:
0 Total packets received
0 Packets delivered
0 Too small packets
0 Packets with bad header length
0 Packets with bad checksum
0 Bad version packets
0 Unknown or unsupported protocol packets
0 Packets with bogus sdl size
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 Address fields were not reasonable
0 Segment information forgotten
0 Forwarded packets
0 Total packets sent
0 Output packets discarded
0 Non-forwarded packets
0 Packets fragmented
0 Fragments sent
0 Fragments discarded
0 Fragments timed out
0 Fragmentation prohibited
0 Packets reconstructed
0 Packets destined to dead nexthop
0 Packets discarded due to no route
0 Error pdu rate drops
0 ER pdu generation failure
esis:
0 Total pkts received
0 Total packets consumed by protocol
0 Pdus received with bad checksum
0 Pdus received with bad version number
0 Pdus received with bad type field
0 Short pdus received
0 Pdus with bogus sdl size
0 Pdus with bad header length
0 Pdus with unknown or unsupported protocol
```

```
0 No free memory in socket buffer
0 Send packets discarded
0 Sbappend failure
0 Mcopy failure
0 ISO family not configured
tnp:
0 Unicast packets received
0 Broadcast packets received
0 Fragmented packets received
0 Hello packets dropped
0 Fragments dropped
0 Fragment reassembly queue flushes
0 Packets with tnp src address collision received
0 Hello packets received
0 Control packets received
0 Rdp packets received
0 Udp packets received
0 Tunnel packets received
0 Input packets discarded with no protocol
0 Packets of version unspecified received
0 Packets of version 1 received
0 Packets of version 2 received
0 Packets of version 3 received
0 Unicast packets sent
0 Broadcast packets sent
0 Fragmented packets sent
0 Hello packets dropped
0 Fragments dropped
0 Hello packets sent
0 Control packets sent
0 Rdp packets sent
0 Udp packets sent
0 Tunnel packets sent
0 Packets sent with unknown protocol
0 Packets of version unspecified sent
0 Packets of version 1 sent
0 Packets of version 2 sent
0 Packets of version 3 sent
rdp:
0 Input packets
0 Packets discarded for bad checksum
0 Packets discarded due to bad sequence number
0 Refused connections
0 Acks received
0 Packets dropped due to full socket buffers
0 Retransmits
0 Output packets
0 Acks sent
0 Connects
0 Closes
0 Keepalives received
0 Keepalives sent
tudp:
67 Datagrams received
0 Datagrams with incomplete header
0 Datagrams with bad data length field
0 Datagrams with bad checksum
0 Datagrams dropped due to no socket
0 Broadcast/multicast datagrams dropped due to no socket
0 Datagrams dropped due to full socket buffers
67 Delivered
```

```
68 Datagrams output
ttp:
0 Packets sent
0 Packets sent while unconnected
0 Packets sent while interface down
0 Packets sent couldn't get buffer
0 Packets sent couldn't find neighbor
0 L2 packets received
0 Unknown L3 packets received
0 IPv4 L3 packets received
0 MPLS L3 packets received
0 MPLS->IPv4 L3 packets received
0 IPv4->MPLS L3 packets received
0 IPv6 L3 packets received
0 ARP L3 packets received
0 CLNP L3 packets received
0 TNP L3 packets received
0 NULL L3 packets received
0 Cyclotron cycle L3 packets received
0 Cyclotron send L3 packets received
0 Packets received while unconnected
0 Packets received from unknown ifl
0 Input packets couldn't get buffer
0 Input packets with bad type
0 Input packets with discard type
0 Input packets with too many tlvs
0 Input packets with bad tlv header
70633 Input packets with bad tlv type
68877 Input packets dropped based on tlv result0 Input packets for which rt lookup
  is bypassed
mpls:
0 Total MPLS packets received
0 Packets forwarded
0 Packets dropped
0 Packets with header too small
0 After tagging, packets can't fit link MTU
0 Packets with IPv4 explicit NULL tag
0 Packets with IPv4 explicit NULL cksum errors
0 Packets with router alert tag
0 LSP ping packets (ttl-expired/router alert)
0 Packets with ttl expired
0 Packets with tag encoding error
0 Packets discarded due to no route
0 Packets used first nexthop in ecmp unilist
vpls:
0 Total packets received
0 Packets with size smaller than minimum
0 Packets with incorrect version number
0 Packets for this host
0 Packets with no logical interface
0 Packets with no family
0 Packets with no route table
582 Copyright © 2010, Juniper Networks, Inc.
0 Packets with no auxiliary table
0 Packets with no corefacing entry
0 packets with no CE-facing entry
0 MAC route learning requests
0 MAC routes learnt
0 Requests to learn an existing route
0 Learning requests while learning disabled on interface
0 Learning requests over capacity
```



```
0 MAC routes moved
0 Requests to move static route
0 MAC route aging requests
0 MAC routes aged
0 Bogus address in aging requests
0 Requests to age static route
0 Requests to re-ageout aged route
0 Requests involving multiple peer FEs
0 Aging acks from PFE
0 Aging non-acks from PFE
0 Aging requests timed out waiting on FEs
0 Aging requests over max-rate
0 Errors finding peer FEs
0 Unsupported platform
0 Packets dropped due to no l3 route table
0 Packets dropped due to no local ifl
0 Packets punted
0 Packets dropped due to no socket
bridge:
Input:
0 packets received
0 packets forwarded
0 packets failed to forward
0 packets dropped
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with stp state lookup failures
0 packets dropped due to stp blocked/listening
0 packets dropped due to stp learning
0 packets with src MAC learning failures
0 packets with input control processing failures
Forward:
0 packets sent successfully
0 packets with send failures
0 packets forwarded to l3 interface
0 packets with l3 send failures
0 packets discarded
0 packets with l2ifl store failures
0 packets with ifl mismatch failures
0 packets with packet duplication failures
0 packets with tag lookup failures
0 packets with no route for DMAC
0 packets with no route table
0 packets with no nexthop
0 packets with dead nexthop
0 packets with eof reached error
Learning:
0 MACs learned
0 packets sent to l3 interface
0 packets with l3 send failures
0 packets hit holdq while learning
0 MAC moves
0 packets discarded
0 packets with no route for SMAC
0 packets with no nexthop
0 packets with dead nexthop
0 packets dropped due to no resolve route
0 packets with l3 ifd lookup failures
0 packets with l3 ifl lookup failures
0 packets with l3 invalid rnh
0 packets with no route for SMAC in clone learning
```

```
0 packets with no nexhop in clone learning
0 packets with dead nexthop in clone learning
0 packets dropped due to no resolve nh in clone learning
Output:
0 packets forwarded
0 packets failed to forward
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with input control processing failures
Send:
0 packets sent successfully
0 packets with send failures
0 packets dropped due to interface down
0 packets with dev output failures
0 blocked ifl discards
0 packets with tag lookup failures
0 packets with stp state lookup failures
0 packets with tag insertion failures
0 packets with tag removal failures
Flood:
0 packets flooded
0 flood failures
IGMP:
0 packets sent successfully
0 packets with send failures
0 packets forwarded
0 packets failed to forward
0 packets with mpull failures
0 packets with vmember lookup failures
0 packets with vlan lookup failures
0 packets with ifl lookup failures
0 packets with tag lookup failures
Misc:
0 packets with size smaller than minimum
0 packets with double tags
0 packets with no ifl
0 packets with no family
0 packets with no route table
```

## show system storage

<b>List of Syntax</b>	<a href="#">Syntax on page 1129</a> <a href="#">Syntax (EX Series Switches) on page 1129</a> <a href="#">Syntax (MX Series Router) on page 1129</a> <a href="#">Syntax (QFX Series) on page 1129</a> <a href="#">Syntax (SRX Series) on page 1129</a> <a href="#">Syntax (TX Matrix Router) on page 1129</a> <a href="#">Syntax (TX Matrix Plus Router and TX Matrix Plus Router with 3D SIBs) on page 1129</a>
<b>Syntax</b>	<pre>show system storage &lt;detail&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show system storage &lt;detail&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show system storage &lt;detail&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show system storage &lt;detail&gt; &lt;infrastructure name&gt; &lt;interconnect-device name&gt; &lt;node-group name&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (SRX Series)</b>	<pre>show system storage &lt;detail&gt; &lt;partitions&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show system storage &lt;detail&gt; &lt;all-chassis   all-lcc   lcc number   scc&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Syntax (TX Matrix Plus Router and TX Matrix Plus Router with 3D SIBs)</b>	<pre>show system storage &lt;detail&gt; &lt;all-chassis   all-lcc   lcc number   sfc number&gt; &lt;invoke-on (all-routing-engines   other-routing-engine)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>

**sfc** option introduced for the TX Matrix Plus router in JUNOS Release 9.6.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Option **invoke-on (all-routing-engines | other-routing-engine)** introduced in Junos OS Release 14.1

<b>Description</b>	Display statistics about the amount of free disk space in the router's or switch's file systems.
<b>Options</b>	<p><b>none</b>—Display standard information about the amount of free disk space in the router's or switch's file systems.</p> <p><b>detail</b>—(Optional) Display detailed output.</p> <p><b>invoke-on all-routing-engines</b>—(Optional) Display the system storage information on all master and backup Routing Engines on a routing matrix based on the TX Matrix or TX Matrix Plus router or on a router that has dual Routing Engines.</p> <p><b>invoke-on other-routing-engines</b>—(Optional) Display the system storage information on the other Routing Engine. For example, if you issue this command on the master Routing Engine on an M320 router, the JUNOS Software displays the system storage information on the backup Routing Engine. On a routing matrix based on the TX Matrix or TX Matrix Plus router, if you issue this command on the TX Matrix or TX Matrix Plus router's master Routing Engine, the JUNOS Software displays all the system storage information on all the backup Routing Engines.</p> <p><b>all-chassis</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system storage statistics for all the routers in the chassis.</p> <p><b>all-lcc</b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for all routers connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for all members of the Virtual Chassis configuration.</p> <p><b>infrastructure <i>name</i></b>—(QFabric systems only) (Optional) Display system storage statistics for the fabric control Routing Engines or fabric manager Routing Engines.</p> <p><b>interconnect-device <i>name</i></b>—(QFabric systems only) (Optional) Display system storage statistics for the Interconnect device.</p> <p><b>lcc <i>number</i></b>—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system storage statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system storage statistics for a specific router that is connected to the TX Matrix Plus router.</p>

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display system storage statistics for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display system storage statistics for the Node group.

**scc**—(TX Matrix routers only) (Optional) Display system storage statistics for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display system storage statistics for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system storage** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)
- [show system storage partitions \(View SRX Series\)](#)

**List of Sample Output**

- [show system storage on page 1132](#)
- [show system storage \(TX Matrix Plus Router\) on page 1132](#)
- [show system storage \(QFX3500 Switch\) on page 1134](#)
- [show system storage invoke-on all-routing-engines on page 1135](#)
- [show system storage invoke-on other-routing-engine on page 1136](#)

**Output Fields** Table 64 on page 1132 describes the output fields for the **show system storage** command. Output fields are listed in the approximate order in which they appear.

**Table 64: show system storage Output Fields**

Field Name	Field Description
<b>Filesystem</b>	Name of the filesystem.
<b>Size</b>	Size of the filesystem.
<b>Used</b>	Amount of space used in the filesystem.
<b>Avail</b>	Amount of space available in the filesystem.
<b>Capacity</b>	Percentage of the filesystem space that is being used.
<b>Mounted on</b>	Directory in which the filesystem is mounted.

## Sample Output

### show system storage

```

user@host> show system storage
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      77M       37M       34M      52%      /
devfs           16K       16K        0B     100%    /dev/
/dev/vn0        12M       12M        0B     100%    /packages/mnt/jbase
/dev/vn1        39M       39M        0B     100%
/packages/mnt/jkernel-7.2R1.7
/dev/vn2        12M       12M        0B     100%
/packages/mnt/jpfe-M40-7.2R1.7
/dev/vn3        2.3M      2.3M        0B     100%
/packages/mnt/jdocs-7.2R1.7
/dev/vn4        14M       14M        0B     100%
/packages/mnt/jroute-7.2R1.7
/dev/vn5        4.5M      4.5M        0B     100%
/packages/mnt/jcrypto-7.2R1.7
mfs:172         1.5G      4.0K       1.3G      0%      /tmp
/dev/ad0s1e      12M       20K        11M      0%      /config
procfs          4.0K      4.0K        0B     100%    /proc
/dev/ad1s1f      9.4G      4.9G       3.7G      57%     /var

```

### show system storage (TX Matrix Plus Router)

```

user@host> show system storage
sfc0-re0:
-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a      3.4G      178M       2.9G      6%      /
devfs           1.0K      1.0K        0B     100%    /dev
devfs           1.0K      1.0K        0B     100%    /dev/
/dev/md0         33M       33M        0B     100%    /packages/mnt/jbase
/dev/md1        216M      216M        0B     100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2         66M       66M        0B     100%
/packages/mnt/jpfe-T-9.6-20090519.0

```

/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	1.0M	1.8G	0%	/mfs
/dev/ad0s1e	383M	82K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	52G	7.5G	40G	16%	/var

lcc0-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config
procfs	4.0K	4.0K	0B	100%	/proc
/dev/ad1s1f	52G	6.3G	41G	13%	/var

lcc1-re0:

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/ad0s1a	3.4G	178M	2.9G	6%	/
devfs	1.0K	1.0K	0B	100%	/dev
devfs	1.0K	1.0K	0B	100%	/dev/
/dev/md0	33M	33M	0B	100%	/packages/mnt/jbase
/dev/md1	216M	216M	0B	100%	
/packages/mnt/jkernel-9.6-20090519.0					
/dev/md2	66M	66M	0B	100%	
/packages/mnt/jpfe-T-9.6-20090519.0					
/dev/md3	4.1M	4.1M	0B	100%	
/packages/mnt/jdocs-9.6-20090519.0					
/dev/md4	57M	57M	0B	100%	
/packages/mnt/jroute-9.6-20090519.0					
/dev/md5	15M	15M	0B	100%	
/packages/mnt/jcrypto-9.6-20090519.0					
/dev/md6	34M	34M	0B	100%	
/packages/mnt/jpfe-common-9.6-20090519.0					
/dev/md7	2.0G	10.0K	1.8G	0%	/tmp
/dev/md8	2.0G	540K	1.8G	0%	/mfs
/dev/ad0s1e	383M	88K	352M	0%	/config

```

procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      13G      7.7G     64%  /var

lcc2-re0:
-----
Filesystem            Size      Used      Avail Capacity  Mounted on
/dev/ad0s1a           3.4G      178M      2.9G        6% /
devfs                 1.0K      1.0K      0B      100% /dev
devfs                 1.0K      1.0K      0B      100% /dev/
/dev/md0              33M       33M       0B      100% /packages/mnt/jbase
/dev/md1             216M      216M       0B      100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2              66M       66M       0B      100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3              4.1M      4.1M       0B      100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4              57M       57M       0B      100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5              15M       15M       0B      100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6              34M       34M       0B      100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7              2.0G      10.0K      1.8G        0% /tmp
/dev/md8              2.0G      540K      1.8G        0% /mfs
/dev/ad0s1e           383M       64K      352M        0% /config
procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      3.7G      17G       18%  /var

lcc3-re0:
-----
Filesystem            Size      Used      Avail Capacity  Mounted on
/dev/ad0s1a           3.4G      178M      2.9G        6% /
devfs                 1.0K      1.0K      0B      100% /dev
devfs                 1.0K      1.0K      0B      100% /dev/
/dev/md0              33M       33M       0B      100% /packages/mnt/jbase
/dev/md1             216M      216M       0B      100%
/packages/mnt/jkernel-9.6-20090519.0
/dev/md2              66M       66M       0B      100%
/packages/mnt/jpfe-T-9.6-20090519.0
/dev/md3              4.1M      4.1M       0B      100%
/packages/mnt/jdocs-9.6-20090519.0
/dev/md4              57M       57M       0B      100%
/packages/mnt/jroute-9.6-20090519.0
/dev/md5              15M       15M       0B      100%
/packages/mnt/jcrypto-9.6-20090519.0
/dev/md6              34M       34M       0B      100%
/packages/mnt/jpfe-common-9.6-20090519.0
/dev/md7              2.0G      10.0K      1.8G        0% /tmp
/dev/md8              2.0G      540K      1.8G        0% /mfs
/dev/ad0s1e           383M       34K      352M        0% /config
procfs                4.0K      4.0K      0B      100% /proc
/dev/ad1s1f           23G      18G      3.5G       84%  /var

```

#### show system storage (QFX3500 Switch)

```

user@switch> show system storage
Filesystem            Size      Used      Avail Capacity  Mounted on
/dev/da0s2a           343M      192M      123M       61% /
devfs                 1.0K      1.0K      0B      100% /dev
/dev/md0             119M      119M       0B      100% /packages/mnt/jbase
/dev/md1             513M      513M       0B      100%

```



```

/packages/mnt/jkernel-qfx-11.1R1.5
/dev/md2          37M          37M          0B          100%
/packages/mnt/jpfe-qfx-e9xxx-11.1R1.5
/dev/md3          6.0M          6.0M          0B          100%
/packages/mnt/jdocs-qfx-11.1R1.5
/dev/md4          216M         216M          0B          100%
/packages/mnt/jroute-qfx-11.1R1.5
/dev/md5          59M          59M          0B          100%
/packages/mnt/jcrypto-qfx-11.1R1.5
/dev/md6          85M          85M          0B          100%
/packages/mnt/jswitch-qfx-11.1R1.5
/dev/md7          63M          8.0K          58M          0% /tmp
/dev/da0s2f       228M         14M          196M          7% /var
/dev/da0s3d       590M         3.0M         540M          1% /var/tmp
/dev/da0s3e       104M         162K          95M          0% /config
procfs           4.0K          4.0K          0B          100% /proc

```

### show system storage invoke-on all-routing-engines

```
user@host> show system storage invoke-on all-routing-engines
```

```
re0:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B        100%     /dev
/dev/md0        76M       76M       0B        100%     /packages/mnt/jbase
/dev/md1        40M       40M       0B        100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M      219M      0B        100%
/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3        5.4M      5.4M      0B        100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4        116M      116M      0B        100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5        44M       44M       0B        100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6        70M       70M       0B        100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7        182K      182K      0B        100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8        499M      499M      0B        100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9        41M       41M       0B        100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10       12M       12M       0B        100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11       3.2G      8.0K      2.9G       0% /tmp
/dev/md12       3.2G      1.1M      2.9G       0% /mfs
/dev/ad0s1e     376M      220K      346M       0% /config
procfs         4.0K      4.0K      0B        100% /proc
/dev/ad1s1f     50G       43G      3.2G      93% /var

```

```
re1:
```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B        100%     /dev
/dev/md0        76M       76M       0B        100%     /packages/mnt/jbase
/dev/md1        40M       40M       0B        100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M      219M      0B        100%

```

```

/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3          5.4M      5.4M      0B      100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4          116M     116M      0B      100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5          44M      44M      0B      100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6          70M      70M      0B      100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7          182K     182K      0B      100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8          499M     499M      0B      100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9          41M      41M      0B      100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10         12M      12M      0B      100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11         3.2G      8.0K      2.9G      0% /tmp
/dev/md12         3.2G     662K      2.9G      0% /mfs
/dev/ad0s1e       375M     230K     344M      0% /config
procfs           4.0K      4.0K      0B      100% /proc
/dev/ad1s1f       52G      46G      2.2G      95% /var

```

#### show system storage invoke-on other-routing-engine

```

user@host> show system storage invoke-on other-routing-engine
rel:

```

```

-----
Filesystem      Size      Used      Avail  Capacity  Mounted on
/dev/ad0s1a     3.3G      440M      2.6G      14%      /
devfs           1.0K      1.0K      0B      100%     /dev
/dev/md0        76M      76M      0B      100%     /packages/mnt/jbase
/dev/md1        40M      40M      0B      100%
/packages/mnt/jkernel64-14.1-20140407.1
/dev/md2        219M     219M      0B      100%
/packages/mnt/jpfe-T-14.1-20140407.1
/dev/md3        5.4M      5.4M      0B      100%
/packages/mnt/jdocs-14.1-20140407.1
/dev/md4        116M     116M      0B      100%
/packages/mnt/jroute-14.1-20140407.1
/dev/md5        44M      44M      0B      100%
/packages/mnt/jcrypto64-14.1-20140407.1
/dev/md6        70M      70M      0B      100%
/packages/mnt/jpfe-common-14.1-20140407.1
/dev/md7        182K     182K      0B      100%
/packages/mnt/jplatform-14.1-20140407.1
/dev/md8        499M     499M      0B      100%
/packages/mnt/jruntime-14.1-20140407.1
/dev/md9        41M      41M      0B      100%
/packages/mnt/jruntime64-14.1-20140407.1
/dev/md10       12M      12M      0B      100%
/packages/mnt/py-base-i386-14.1-20140407.1
/dev/md11       3.2G      8.0K      2.9G      0% /tmp
/dev/md12       3.2G     662K      2.9G      0% /mfs
/dev/ad0s1e     375M     230K     344M      0% /config
procfs         4.0K      4.0K      0B      100% /proc
/dev/ad1s1f     52G      46G      2.2G      95% /var

```

## show system uptime

<b>List of Syntax</b>	<a href="#">Syntax on page 1137</a> <a href="#">Syntax (EX Series Switches) on page 1137</a> <a href="#">Syntax (QFX Series) on page 1137</a> <a href="#">Syntax (TX Matrix Router) on page 1137</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1137</a> <a href="#">Syntax (MX Series Router) on page 1137</a>
<b>Syntax</b>	show system uptime
<b>Syntax (EX Series Switches)</b>	show system uptime <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system uptime <director-group <i>name</i> > <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Syntax (TX Matrix Router)</b>	show system uptime <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system uptime <detail> <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system uptime <all-members> <invoke-on> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in JUNOS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.
<b>Options</b>	<b>none</b> —Show time since the system rebooted and processes started.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started on all the routers in the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus

router, show time since the system rebooted and processes started for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

**director-group *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Director group.

**infrastructure *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Interconnect device.

**invoke-on**—(MX Series routers only) (Optional) Display the time since the system rebooted and processes started on the master Routing Engine, backup Routing Engine, or both, on a router with two Routing Engines.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show time since the system rebooted and processes started for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show time since the system rebooted and processes started for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Show time since the system rebooted and processes started on the Node group.

**scc**—(TX Matrix routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Show time since the system rebooted and processes started for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system uptime** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation**

- [Monitoring System Process Information on page 333](#)
- [Monitoring System Properties on page 334](#)
- [10-Gigabit Ethernet LAN/WAN PIC with XFP \(T640 Router\)](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output**

[show system uptime on page 1140](#)  
[show system uptime all-lcc \(TX Matrix Router\) on page 1140](#)  
[show system uptime all-lcc \(TX Matrix Plus Router\) on page 1140](#)  
[show system uptime \(EX Series\) on page 1141](#)  
[show system uptime \(QFX Series\) on page 1141](#)

**Output Fields** [Table 65 on page 1139](#) describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

**Table 65: show system uptime Output Fields**

Field Name	Field Description
<b>Current time</b>	Current system time in UTC.
<b>System booted</b>	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
<b>Protocols started</b>	Date and time when the routing protocols were last started and how long they have been running.
<b>Last configured</b>	Date and time when a configuration was last committed. Also shows the name of the user who issued the last <b>commit</b> command.
<b>time and up</b>	Current time, in the local time zone, and how long the router or switch has been operational.
<b>users</b>	Number of users logged in to the router or switch.
<b>load averages</b>	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

## Sample Output

### show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

### show system uptime all-lcc (TX Matrix Router)

```
user@host> show system uptime all-lcc
lcc0-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-13 03:13:55 PDT (06:41:40 ago)
Last configured: 2004-09-13 03:17:48 PDT (06:37:47 ago) by root
9:55AM PDT up 6:42, 1 user, load averages: 0.02, 0.03, 0.00
lcc2-re0:
-----
Current time: 2004-09-13 09:55:35 PDT
System booted: 2004-09-12 03:23:43 PDT (1d 06:31 ago)
Last configured: 2004-09-13 03:05:36 PDT (06:49:59 ago) by root
9:55AM PDT up 1 day, 6:32, 1 user, load averages: 0.02, 0.01, 0.00
```

### show system uptime all-lcc (TX Matrix Plus Router)

```
user@host> show system uptime all-lcc
sfc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:33 PDT (17:44:57 ago)
Protocols started: 2009-05-24 06:40:30 PDT (17:44:00 ago)
Last configured: 2009-05-24 06:33:27 PDT (17:51:03 ago) by gregdo
12:24AM up 17:45, 2 users, load averages: 0.07, 0.05, 0.01

lcc0-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:46 PDT (17:44:44 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:47 PDT (17:43:43 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc1-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:38 PDT (17:44:52 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:18 PDT (17:44:12 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

lcc2-re0:
-----
Current time: 2009-05-25 00:24:30 PDT
System booted: 2009-05-24 06:39:48 PDT (17:44:42 ago)
error: the routing subsystem is not running
Last configured: 2009-05-24 06:40:44 PDT (17:43:46 ago) by root
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00
```

lcc3-re0:

-----  
Current time: 2009-05-25 00:24:30 PDT  
System booted: 2009-05-24 06:39:44 PDT (17:44:46 ago)  
error: the routing subsystem is not running  
Last configured: 2009-05-24 06:40:08 PDT (17:44:22 ago) by root  
12:24AM up 17:45, 0 users, load averages: 0.00, 0.00, 0.00

#### show system uptime (EX Series)

```
user@switch> show system uptime
Current time: 2014-03-12 16:39:56 UTC
System booted: 2014-03-12 14:58:05 UTC (01:41:51 ago)
Protocols started: 2014-03-12 14:59:48 UTC (01:40:08 ago)
Last configured: 2014-03-12 14:58:58 UTC (01:40:58 ago) by root
4:39PM up 1:42, 4 users, load averages: 0.02, 0.02, 0.00
```

#### show system uptime (QFX Series)

```
user@switch> show system uptime
Current time: 2010-08-27 03:12:30 PDT
System booted: 2010-08-13 17:11:54 PDT (1w6d 10:00 ago)
Protocols started: 2010-08-13 17:13:56 PDT (1w6d 09:58 ago)
Last configured: 2010-08-26 05:54:00 PDT (21:18:30 ago) by regress
3:12AM up 13 days, 10:01, 3 users, load averages: 0.00, 0.00, 0.00
```

## show system users

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1142</a> <a href="#">Syntax (TX Matrix Router) on page 1142</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1142</a> <a href="#">Syntax (MX Series Router) on page 1142</a>
<b>Syntax</b>	show system users <no-resolve>
<b>Syntax (TX Matrix Router)</b>	show system users <all-chassis   all-lcc   lccnumber   scc> <no-resolve>
<b>Syntax (TX Matrix Plus Router)</b>	show system users <detail> <all-chassis   all-lcc   lcc number   sfc number> <no-resolve>
<b>Syntax (MX Series Router)</b>	show system users <all-members> <local> <member member-id> <no-resolve>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in JUNOS OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	List information about the users who are currently logged in to the router or switch.



**NOTE:** The `show system users` command lists the information about administrative users that are logged in to a router or switch using the CLI, J-Web, or an SSH client. The output does not list information about web users or automated users that are logged in from a remote client application using Junos XML APIs, such as NETCONF.

---

- Options**    **none**—List information about the users who are currently logged in to the router or switch.
- all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) Show users currently logged in to all the routers in the chassis.
- all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to all connected T1600 or T4000 LCCs.
- all-members**—(MX Series routers only) (Optional) Display users currently logged in to all members of the Virtual Chassis configuration.



**lcc number**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, show users currently logged in to a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, show users currently logged in to a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display users currently logged in to the local Virtual Chassis member.

**member member-id**—(MX Series routers only) (Optional) Display users currently logged in to the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**no-resolve**—(Optional) Do not attempt to resolve IP addresses to hostnames.

**scc**—(TX Matrix routers only) (Optional) Show users currently logged in to the TX Matrix router (or switch-card chassis).

**sfc number**—(TX Matrix Plus routers only) (Optional) Show users currently logged in to the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system users** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** [• Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system users on page 1144](#)  
[show system users lcc no-resolve \(TX Matrix, TX Matrix Plus Router\) on page 1144](#)  
[show system users \(TX Matrix Plus Router\) on page 1144](#)  
[show system users \(QFX Series\) on page 1145](#)  
[show system users no-resolve \(QFX Series\) on page 1145](#)

**Output Fields** Table 66 on page 1144 describes the output fields for the **show system users** command. Output fields are listed in the approximate order in which they appear.

**Table 66: show system users Output Fields**

Field Name	Field Description
<b>time and up</b>	Current time, in the local time zone, and how long the router or switch has been operational.
<b>users</b>	Number of users logged in to the router or switch.
<b>load averages</b>	Load averages for the last 1 minute, 5 minutes, and 15 minutes.
<b>USER</b>	Username.
<b>TTY</b>	Terminal through which the user is logged in.
<b>FROM</b>	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.
<b>LOGIN@</b>	Time when the user logged in.
<b>IDLE</b>	How long the user has been idle.
<b>WHAT</b>	Processes that the user is running.

## Sample Output

### show system users

```
user@host> show system users
 7:30PM up 4 days, 2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER   TTY FROM          LOGIN@  IDLE WHAT
root   d0  -              Fri05PM 4days -csh (csh)
blue   p0  level5.company.net 7:30PM  - cli
```

### show system users lcc no-resolve (TX Matrix, TX Matrix Plus Router)

```
user@host> show system users lcc 2 no-resolve

lcc2-re0:
-----
10:34AM PDT up 1 day, 7:11, 5 users, load averages: 0.03, 0.01, 0.00
USER   TTY   FROM          LOGIN@  IDLE WHAT
root   d0    -              3:21AM  7:12 /bin/csh
user1  p0    scc-re0        10:15AM  - telnet hostA
user1  p1    scc-re0        10:16AM  - telnet hostA
user1  p2    scc-re0        10:19AM  - telnet hostA
user1  p3    scc-re0        10:24AM  - telnet hostA
```

### show system users (TX Matrix Plus Router)

```
user@host> show system users
sfc0-re0:
-----
1:41AM up 26 mins, 3 users, load averages: 0.08, 0.04, 0.03
```

```

USER      TTY      FROM                                LOGIN@  IDLE WHAT
user2     p0       10.209.208.123                    1:18AM  21 cli
user2     p1       172.17.29.207                    1:37AM   2 cli
user2     p2       172.17.28.19                     1:40AM   - cli

lcc0-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.00, 0.03

lcc1-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.00, 0.02, 0.03

lcc2-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.16, 0.06, 0.02

lcc3-re0:
-----
1:41AM up 26 mins, 0 users, load averages: 0.12, 0.04, 0.04

user3@aj> show system users
sfc0-re0:
-----
1:42AM up 28 mins, 4 users, load averages: 0.02, 0.03, 0.02
USER      TTY      FROM                                LOGIN@  IDLE WHAT
user3     p0       pssraj-t61.jnpr.net              1:18AM  22 cli
user3     p1       eng-shell14.juniper.net          1:37AM   - cli
user3     p2       bigpink.juniper.net              1:40AM   - cli
user3     p3       sv-cutty-01.englab.juniper.net    1:42AM   - csh (csh)

lcc0-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.02, 0.01, 0.03

lcc1-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.07, 0.04, 0.03

lcc2-re0:
-----
1:42AM up 27 mins, 0 users, load averages: 0.07, 0.06, 0.02

lcc3-re0:
-----
1:42AM up 28 mins, 0 users, load averages: 0.05, 0.04, 0.04

```

### show system users (QFX Series)

```

user@switch> show system users
USER      TTY      FROM                                LOGIN@  IDLE WHAT
tlewis    p0       172.22.18.117                    2:54AM  39 -cli (cli)
tlewis    p1       172.22.18.117                    3:01AM   - -cli (cli)
tcheng    p2       172.22.17.197                    3:08AM  11 -cli (cli)

```

### show system users no-resolve (QFX Series)

```

user@switch> show system users no-resolve
USER      TTY      FROM                                LOGIN@  IDLE WHAT
tlewis    p0       172.22.18.117                    2:54AM  39 -cli (cli)

```

tLewis	p1	172.22.18.117	3:01AM	- -cli (cli)
tcheng	p2	172.22.17.197	3:08AM	11 -cli (cli)

## show system virtual-memory

<b>List of Syntax</b>	<a href="#">Syntax on page 1147</a> <a href="#">Syntax (EX Series) on page 1147</a> <a href="#">Syntax (TX Matrix Router) on page 1147</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1147</a> <a href="#">Syntax (MX Series Router) on page 1147</a> <a href="#">Syntax (QFX Series) on page 1147</a>
<b>Syntax</b>	show system virtual-memory
<b>Syntax (EX Series)</b>	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system virtual-memory <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system virtual-memory <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system virtual-memory <all-members> <local> <member <i>member-id</i> >
<b>Syntax (QFX Series)</b>	show system virtual-memory <infrastructure <i>name</i> > <interconnect-device <i>name</i> > <node-group <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the usage of Junos OS kernel memory listed first by size of allocation and then by type of usage. Use the <b>show system virtual-memory</b> command for troubleshooting with Juniper Networks Customer Support.
<b>Options</b>	<b>none</b> —Display kernel dynamic memory usage information.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for all chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for all connected T1600 or T4000 LCCs.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for all members of the Virtual Chassis configuration.

**infrastructure *name***—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the fabric control Routing Engine and fabric manager Routing Engine.

**interconnect-device *name***—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Interconnect device.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display kernel dynamic memory usage information for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display kernel dynamic memory usage information for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display kernel dynamic memory usage information for the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**node-group *name***—(QFabric systems only) (Optional) Display kernel dynamic memory usage information for the Node group.

**scc**—(TX Matrix routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display kernel dynamic memory usage information for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system virtual-memory** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix

or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.



**NOTE:** The `show system virtual-memory` command with the `| display XML` pipe option now displays XML output for the command in the parent tags: `<vmstat-memstat-malloc>`, `<vmstat-memstat-zone>`, `<vmstat-sumstat>`, `<vmstat-intr>`, and `<vmstat-kernel-state>` with each child element as a separate XML tag. In Junos OS Releases 10.1 and earlier, the `| display XML` option for this command does not have an XML API element and the entire output is displayed in a single `<output>` tag element.

Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
List of Sample Output	<a href="#">show system virtual-memory on page 1151</a> <a href="#">show system virtual-memory scc (TX Matrix Router) on page 1155</a> <a href="#">show system virtual-memory sfc (TX Matrix Plus Router) on page 1156</a> <a href="#">show system virtual-memory   display xml on page 1159</a> <a href="#">show system virtual-memory (QFX Series) on page 1182</a>
Output Fields	<p><a href="#">Table 67 on page 1150</a> lists the output fields for the <code>show system virtual-memory</code> command. Output fields are listed in the approximate order in which they appear.</p>

Table 67: show system virtual-memory Output Fields

Field Name	Field Description
<b>Memory statistics by bucket size</b>	
<b>Size</b>	Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
<b>In Use</b>	Number of memory blocks of this size that are in use (bytes).
<b>Free</b>	Number of memory blocks of this size that are free (bytes).
<b>Requests</b>	Number of memory allocation requests made.
<b>HighWater</b>	Maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.
<b>Couldfree</b>	Total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.
<b>Memory usage type by bucket size</b>	
<b>Size</b>	Memory block size (bytes).
<b>Type(s)</b>	Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.
<b>Memory statistics by type</b>	
<b>Type</b>	Kernel module that is using dynamic memory.
<b>InUse</b>	Number of memory blocks used by this type. The number is rounded up.
<b>MemUse</b>	Amount of memory in use, in kilobytes (KB).
<b>HighUse</b>	Maximum memory ever used by this type.
<b>Limit</b>	Maximum memory that can be allocated to this type.
<b>Requests</b>	Total number of dynamic memory allocation requests this type has made.
<b>Type Limit</b>	Number of times requests were blocked for reaching the maximum limit.
<b>Kern Limit</b>	Number of times requests were blocked for the kernel map.
<b>Size(s)</b>	Memory block sizes this type is using.
<b>Memory Totals</b>	
<b>In Use</b>	Total kernel dynamic memory in use (bytes, rounded up).
<b>Free</b>	Total kernel dynamic memory free (bytes, rounded up).



Table 67: show system virtual-memory Output Fields (*continued*)

Field Name	Field Description
<b>Requests</b>	Total number of memory allocation requests.
<b>ITEM</b>	Kernel module that is using memory.
<b>Size</b>	Memory block size (bytes).
<b>Limit</b>	Maximum memory that can be allocated to this type.
<b>Used</b>	Number of memory blocks used by this type. The number is rounded up.
<b>Free</b>	Number of memory blocks available to this type.
<b>Requests</b>	Total number of memory allocation requests this type has made.
<b>interrupt</b>	Timer events and scheduling interruptions.
<b>total</b>	Total number of interruptions for each type.
<b>rate</b>	Interruption rate.
<b>Total</b>	Total for all interruptions.

## Sample Output

### show system virtual-memory

```

user@host> show system virtual-memory
Memory statistics by bucket size
Size    In Use    Free    Requests  HighWater  Couldfree
16      906      118     154876    1280       0
32      455      313     209956    640        0
64      4412     260     75380     320        20
128     3200     32       19361     160        81
256     1510     10       8844      80         4
512     446      2        5085      40         0
1K      18       2        5901      20         0
2K     1128     2        4445      10        1368
4K      185     1         456       5          0
8K       5      1        2653      5          0
16K     181     0         233       5          0
32K       2     0        1848      5          0
64K      20     0         22        5          0
128K      5     0          5        5          0
256K      2     0          2        5          0
512K      1     0          1        5          0

Memory usage type by bucket size
Size    Type(s)
16    uc_devlist, nexusdev, iftable, temp, devbuf, atexit, COS, BPF,
      DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, proc-args, kld,
      MD disk, rman, ATA generic, bus, sysctl, ippool, pfestat, ifstate,

```

```

pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode
32 atkbddev, dirrem, mkdir, diradd, freefile, freefrag, indirdep,
bmsafemap, newblk, temp, devbuf, COS, vnodes, cluster_save buffer,
pcb, soname, proc-args, sigio, kld, Gzip trees, taskqueue, SWAP,
eventhandler, bus, sysctl, uidinfo, subproc, pgrp, pfestat, itable32,
ifstate, pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode, rtnexthop
64 isadev, iftable, MFS node, allocindir, allocdirect, pagedep, temp,
devbuf, lockf, COS, NULLFS hash, DEVFS name, vnodes,
cluster_save buffer, vfscache, pcb, soname, proc-args, file,
AR driver, AD driver, Gzip trees, rman, eventhandler, bus, sysctl,
subproc, pfestat, pic, ifstate, pfe_ipc, mkey, ifaddr, rtable, ipfw
128 ZONE, freeblks, inodedep, temp, devbuf, zombie, COS, DEVFS node,
vnodes, mount, vfscache, pcb, soname, proc-args, ttys, dev_t,
timecounter, kld, Gzip trees, ISOFS node, bus, uidinfo, cred,
session, pic, itable16, ifstate, pfe_ipc, rtable, ifstat, metrics,
rtnexthop, iffamily
256 iflogical, iftable, MFS node, FFS node, newblk, temp, devbuf,
NFS daemon, vnodes, proc-args, kqueue, file desc, Gzip trees, bus,
subproc, itable16, ifstate, pfe_ipc, sysctl, rtnexthop
512 UFS mount, temp, devbuf, mount, BIO buffer, ptys, ttys, AR driver,
Gzip trees, ISOFS mount, msg, iocltops, ATA generic, bus, proc,
pfestat, lr, ifstate, pfe_ipc, rtable, ipfw, ifstat, rtnexthop
1K iftable, temp, devbuf, NQ NFS Lease, kqueue, kld, AD driver,
Gzip trees, sem, MD disk, bus, ifstate, pfe_ipc, ipfw
2K uc_devlist, UFS mount, temp, devbuf, BIO buffer, pcb, AR driver,
Gzip trees, iocltops, bus, ipfw, ifstat, rcache
4K memdesc, iftable, UFS mount, temp, devbuf, kld, Gzip trees, sem, msg
8K temp, devbuf, syncache, Gzip trees
16K indirdep, temp, devbuf, shm, msg
32K pagedep, kld, Gzip trees
64K VM pgdata, devbuf, MSDOSFS mount
128K UFS ihash, inodedep, NFS hash, kld, ISOFS mount
256K mbuf, vfscache
512K SWAP

```

Memory statistics by type					Type	Kern		
Type	InUse	MemUse	HighUse	Limit	Requests	Limit	Limit	Size(s)
isadev	13	1K	1K127753K	13	0	0	0	64
atkbddev	2	1K	1K127753K	2	0	0	0	32
uc_devlist	24	3K	3K127753K	24	0	0	0	16,2K
nexusdev	3	1K	1K127753K	3	0	0	0	16
memdesc	1	4K	4K127753K	1	0	0	0	4K
mbuf	1	152K	152K127753K	1	0	0	0	256K
iflogical	6	2K	2K127753K	6	0	0	0	256
iftable	17	9K	9K127753K	18	0	0	0	16,64,256,1K,4K
ZONE	15	2K	2K127753K	15	0	0	0	128
VM pgdata	1	64K	64K127753K	1	0	0	0	64K
UFS mount	12	26K	26K127753K	12	0	0	0	512,2K,4K
UFS ihash	1	128K	128K127753K	1	0	0	0	128K
MFS node	6	2K	3K127753K	35	0	0	0	64,256
FFS node	906	227K	227K127753K	1352	0	0	0	256
dirrem	0	0K	4K127753K	500	0	0	0	32
mkdir	0	0K	1K127753K	38	0	0	0	32
diradd	0	0K	6K127753K	521	0	0	0	32
freefile	0	0K	4K127753K	374	0	0	0	32
freeblks	0	0K	8K127753K	219	0	0	0	128
freefrag	0	0K	1K127753K	193	0	0	0	32
allocindir	0	0K	25K127753K	1518	0	0	0	64
indirdep	0	0K	17K127753K	76	0	0	0	32,16K
allocdirect	0	0K	10K127753K	760	0	0	0	64
bmsafemap	0	0K	1K127753K	72	0	0	0	32

newblk	1	1K	1K127753K	2279	0	0	32,256
inodedep	1	128K	175K127753K	2367	0	0	128,128K
pagedep	1	32K	33K127753K	47	0	0	64,32K
temp	1239	92K	96K127753K	8364	0	0	16,32,64K
devbuf	1413	5527K	5527K127753K	1535	0	0	16,32,64,128,256
lockf	38	3K	3K127753K	2906	0	0	64
atexit	1	1K	1K127753K	1	0	0	16
zombie	0	0K	2K127753K	3850	0	0	128
NFS hash	1	128K	128K127753K	1	0	0	128K
NQNFS Lease	1	1K	1K127753K	1	0	0	1K
NFS daemon	1	1K	1K127753K	1	0	0	256
syncache	1	8K	8K127753K	1	0	0	8K
COS	353	44K	44K127753K	353	0	0	16,32,64,128
BPF	189	3K	3K127753K	189	0	0	16
MSDOSFS mount	1	64K	64K127753K	1	0	0	64K
NULLFS hash	1	1K	1K127753K	1	0	0	64
DEVFS mount	2	1K	1K127753K	2	0	0	16
DEVFS name	487	31K	31K127753K	487	0	0	64
DEVFS node	471	58K	58K127753K	479	0	0	16,128
vnodes	28	7K	7K127753K	429	0	0	16,32,64,128,256
mount	15	8K	8K127753K	18	0	0	16,128,512
cluster_save buffer	0	0K	1K127753K	55	0	0	32,64
vfscache	1898	376K	376K127753K	3228	0	0	64,128,256K
BIO buffer	49	98K	398K127753K	495	0	0	512,2K
pcb	159	16K	17K127753K	399	0	0	16,32,64,128,2K
soname	82	10K	10K127753K	42847	0	0	16,32,64,128
proc-args	57	2K	3K127753K	2105	0	0	16,32,64,128,256
ptys	32	16K	16K127753K	32	0	0	512
ttys	254	33K	33K127753K	522	0	0	128,512
kqueue	5	3K	4K127753K	23	0	0	256,1K
sigio	1	1K	1K127753K	27	0	0	32
file	383	24K	24K127753K	16060	0	0	64
file desc	76	19K	20K127753K	3968	0	0	256
shm	1	12K	12K127753K	1	0	0	16K
dev_t	286	36K	36K127753K	286	0	0	128
timecounter	10	2K	2K127753K	10	0	0	128
kld	11	117K	122K127753K	34	0	0	16,32,128,1K,4K
AR driver	1	1K	3K127753K	5	0	0	64,512,2K
AD driver	2	2K	3K127753K	2755	0	0	64,1K
Gzip trees	0	0K	46K127753K	133848	0	0	32,64,128,256
ISOFS node	1136	142K	142K127753K	1189	0	0	128
ISOFS mount	9	132K	132K127753K	10	0	0	512,128K
sem	3	6K	6K127753K	3	0	0	1K,4K
MD disk	2	2K	2K127753K	2	0	0	16,1K
msg	4	25K	25K127753K	4	0	0	512,4K,16K
rman	59	4K	4K127753K	461	0	0	16,64
ioctlops	0	0K	2K127753K	992	0	0	512,2K
taskqueue	2	1K	1K127753K	2	0	0	32
SWAP	2	413K	413K127753K	2	0	0	32,512K
ATA generic	6	3K	3K127753K	6	0	0	16,512
eventhandler	17	1K	1K127753K	17	0	0	32,64
bus	340	30K	31K127753K	794	0	0	16,32,64,128,256
sysctl	0	0K	1K127753K	130262	0	0	16,32,64
uidinfo	4	1K	1K127753K	10	0	0	32,128
cred	22	3K	3K127753K	3450	0	0	128
subproc	156	10K	10K127753K	7882	0	0	32,64,256
proc	2	1K	1K127753K	2	0	0	512
session	12	2K	2K127753K	34	0	0	128
pgrp	16	1K	1K127753K	45	0	0	32
ippool	1	1K	1K127753K	1	0	0	16
pfestat	0	0K	1K127753K	47349	0	0	16,32,64,512

pic	5	1K	1K127753K	5	0	0	64,128
lr	1	1K	1K127753K	1	0	0	512
itable32	110	4K	4K127753K	110	0	0	32
itable16	161	26K	26K127753K	161	0	0	128,256
ifstate	694	159K	160K127753K	1735	0	0	16,32,64,128,1K
pfe_ipc	0	0K	1K127753K	56218	0	0	16,32,64,128,1K
mkey	250	4K	4K127753K	824	0	0	16,32,64
ifaddr	9	1K	1K127753K	9	0	0	64
sysctl	0	0K	1K127753K	30	0	0	256
rtable	49	6K	6K127753K	307	0	0	16,32,64,128,512
ifmaddr	22	1K	1K127753K	22	0	0	16,32
ipfw	23	10K	10K127753K	48	0	0	16,32,64,512,2K
ifstat	698	805K	805K127753K	698	0	0	128,512,2K
rcache	4	8K	8K127753K	4	0	0	2K
rnode	27	1K	1K127753K	285	0	0	16,32
metrics	1	1K	1K127753K	3	0	0	128
rtnexthop	57	9K	9K127753K	312	0	0	32,128,256,512
iffamily	12	2K	2K127753K	12	0	0	128

Memory Totals:	In Use	Free	Requests
	9311K	54K	489068

ITEM	SIZE	LIMIT	USED	FREE	REQUESTS
PIPE:	192,	0,	4,	81,	4422
SWAPMETA:	160,	95814,	0,	0,	0
unpcb:	160,	0,	114,	36,	279
ripcb:	192,	25330,	5,	37,	5
syncache:	128,	15359,	0,	64,	5
tcpcb:	576,	25330,	23,	12,	32
udpcb:	192,	25330,	14,	28,	255
socket:	256,	25330,	246,	26,	819
KNOTE:	96,	0,	27,	57,	71
NFSNODE:	352,	0,	0,	0,	0
NFSMOUNT:	544,	0,	0,	0,	0
VNODE:	224,	0,	2778,	43,	2778
NAMEI:	1024,	0,	0,	8,	40725
VMSPACE:	192,	0,	57,	71,	3906
PROC:	448,	0,	73,	17,	3923
DP fakepg:	64,	0,	0,	0,	0
PV ENTRY:	28,	499566,	44530,	152053,	1525141
MAP ENTRY:	48,	0,	1439,	134,	351075
KMAP ENTRY:	48,	35645,	179,	119,	10904
MAP:	108,	0,	7,	3,	7
VM OBJECT:	92,	0,	2575,	109,	66912

```

792644 cpu context switches
9863474 device interrupts
286510 software interrupts
390851 traps
3596829 system calls
  16 kernel threads created
 3880 fork() calls
   27 vfork() calls
    0 rfork() calls
    0 swap pager pageins
    0 swap pager pages paged in
    0 swap pager pageouts
    0 swap pager pages paged out
  380 vnode pager pageins
  395 vnode pager pages paged in
  122 vnode pager pageouts

```

```

1476 vnode pager pages paged out
    0 page daemon wakeups
    0 pages examined by the page daemon
101 pages reactivated
161722 copy-on-write faults
    0 copy-on-write optimized faults
84623 zero fill pages zeroed
83063 zero fill pages prezeroed
    7 intransit blocking page faults
535606 total VM faults taken
    0 pages affected by kernel thread creation
238254 pages affected by fork()
    2535 pages affected by vfork()
    0 pages affected by rfork()
283379 pages freed
    0 pages freed by daemon
190091 pages freed by exiting processes
17458 pages active
29166 pages inactive
    0 pages in VM cache
10395 pages wired down
134610 pages free
    4096 bytes per page
183419 total name lookups
    cache hits (90% pos + 7% neg) system 0% per-directory
    deletions 0%, falsehits 0%, toolong 0%

interrupt          total          rate
ata0 irq14         113338           3
mux irq7           727643          21
fxp1 irq10         1178671          34
sio0 irq4           833              0
clk irq0           3439769          99
rtc irq8           4403221          127
Total              9863475          286

Kernel direct memory map:
    4423 pages used
    4057340 pages maximum

```

*Note: Kernel direct memory map only displays for 64 bit platform.*

### show system virtual-memory scc (TX Matrix Router)

```
user@host> show system virtual-memory scc
```

```

Memory statistics by bucket size
Size  In Use  Free  Requests  HighWater  Couldfree
16    898    126    749493    1280        0
32    2018   1310   980643    640        632
64    3490   13342  935420    320        5365
...

Memory usage type by bucket size
Size  Type(s)
16    uc_devlist, COS, BPF, DEVFS mount, DEVFS node, vnodes, mount, pcb,
      soname, rman, bus, sysctl, ifstate, pfe_ipc, mkey, socket, rtable,
      ifmaddr, ipfw, rnode, iftable, temp, devbuf, atexit, proc-args, kld,
      MD disk
32    atkbddev, Gzip trees, dirrem, mkdir, diradd, freefile, freefrag,
      indirdep, bmsafemap, newblk, tseg_qent, COS, vnodes,

```

...

```

Memory statistics by type
      Type InUse MemUse HighUse Limit Requests Limit Limit Size(s)
      isadev 12 1K 1K166400K 12 0 0 64
      atkbdddev 2 1K 1K166400K 2 0 0 32
      uc_devlist 24 3K 3K166400K 24 0 0 16,2K
      ....

```

```

Memory Totals: In Use Free Requests
                6091K 1554K 2897122

```

### show system virtual-memory sfc (TX Matrix Plus Router)

```

user@host> show system virtual-memory sfc 0
sfc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue 1 1K - 1 64
      entropy 1024 64K - 1024 64
      linker 487 6272K - 1163 16,32,64,4096,32768,131072
      USB 127 10K - 127 16,32,64,128,256,1024,2048
      lockf 46 3K - 98418 64
      USBdev 10 2K - 34 16,128,2048,16384
ifstateSLLNode 0 0K - 1096 16
      devbuf 21243 15683K - 21810
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
      temp 1283 151K - 2483472
16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072
      ip6ndp 0 0K - 4 64
      in6ifmulti 1 1K - 1 64
      in6grentry 1 1K - 1 64
      iflogical 20 5K - 29 2048
      iffamily 45 6K - 69 32,1024,2048
      rtnexthop 266 46K - 608013 32,256,512,1024,2048,4096
      metrics 31 4K - 54 256
      rnode 212 4K - 607848 16,32
      rcache 4 8K - 4 65536
      iflist 0 0K - 6 16,64
      ifdevice 11 8K - 17 16,32768
      ifstat 424 472K - 427 512,16384,65536
      ipfw 42 23K - 145
16,32,64,128,256,512,1024,16384,32768,65536,131072
      ifmaddr 415 11K - 415 16,32
      rtable 329 28K - 608066 16,32,64,128,1024,16384
      sysctl 0 0K - 887976 16,32,64,4096,16384,32768
      ifaddr 64 5K - 70 32,64,128
      mkey 331 6K - 12528 16,128
      pfe_ipc 0 0K - 7299115
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768,65536,131072
      ifstate 1245054 70088K - 3040437
16,32,64,128,256,512,1024,2048,4096,8192,16384,32768
      idxbucket 1 1K - 1 16
      itable16 5069 1250K - 5103 1024,4096
      itable32 157 10K - 157 64
      itable64 2 1K - 2 128
      lr 1 1K - 4 16384
      pic 37 6K - 37 64,16384
      pffestat 0 0K - 6220 32,64,128,256,131072
      gencfg 1486 424K - 2614 16,32,64,256,512,16384,32768,65536

```

```

        jsr      2      1K      -      22  16
        idl      1      4K      -      165
32, 64, 128, 256, 512, 1024, 2048, 8192, 16384, 32768, 65536, 131072
        rtmsg    0      0K      -      16  131072
        module  250     16K      -      250  64, 128
        mtx_pool 1       8K      -       1  64, 128
        DEVFS3   113    13K      -      114  256
        DEVFS1   106    24K      -      106  2048
        pgrp     15     1K      -      8600 64
        session  11     2K      -      2829 512
        proc      2     1K      -       2  16384
        subproc  296    572K     -     24689 2048, 131072
        cred      38     5K      -    619244 256
        plimit    18     4K      -     21311 2048
        uidinfo   3     1K      -       10  32, 512
        sysctluid 2701    82K     -     2701 16, 32, 64
        sysctltmp 0       0K      -     15572 16, 32, 64, 1024
        umtx     171    11K      -      171  64
        SWAP      2    277K      -       2  64
        bus      779   125K     -     3072 16, 32, 64, 128, 32768
        bus-sc    67     62K     -     1477
16, 32, 64, 512, 1024, 2048, 8192, 16384, 65536, 131072
        devstat   8     17K      -       8  16, 131072
        eventhandler 46     2K      -      47  32, 128
        kobj      93   186K      -     111  65536
        DEVFS      8     1K      -       9  16, 64
        rman     106     7K      -      490 16, 32, 64
        sbuf       0     0K      -     28234 16, 32, 32768, 131072
...
lcc0-re0:

```

```

-----
      Type InUse MemUse HighUse Requests Size(s)
CAM dev queue    1     1K      -       1  64
      entropy  1024    64K      -    1024  64
      linker   487   6272K     -    1163 16, 32, 64, 4096, 32768, 131072
      USB     127    10K      -      127 16, 32, 64, 128, 256, 1024, 2048
      lockf    23     2K      -   169585  64
      USBdev   10     2K      -       34 16, 128, 2048, 16384
      devbuf  5128  10760K     -     5310
16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, 65536, 131072
      temp   1285    151K     -    10770
16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072
      ip6ndp    0     0K      -       4  64
      iflogical 20     5K      -      29  2048
      iffamilly 45     6K      -      69  32, 1024, 2048
      rtnexthop 189    29K      -   1211988 32, 256, 512, 1024, 2048, 4096
      metrics   11     2K      -       16  256
      rnode    135     3K      -    606391 16, 32
      rcache     4     8K      -       4  65536
      iflist     0     0K      -       6  16, 64
      ifdevice  11     8K      -      17  16, 32768
      ifstat   412   471K      -     415  512, 16384, 65536
      ipfw      42    23K      -       91
16, 32, 64, 128, 256, 512, 1024, 16384, 32768, 65536, 131072
      ifmaddr  415    11K      -      415 16, 32
      rtable   225    20K      -    606584 16, 32, 64, 128, 1024, 16384
      sysctl    0     0K      -   2302479 16, 32, 64
      ifaddr    53     4K      -       69  32, 64, 128
      mkey     133     3K      -     8974 16, 128
      pfe_ipc    0     0K      -   19035108
16, 32, 64, 128, 512, 1024, 2048, 8192, 16384, 32768, 65536, 131072

```

ifstate	710270	42176K	-	9583703	
16,32,64,128,256,512,1024,2048,8192,16384,32768					
idxbucket	1	1K	-	1	16
itable16	5045	1245K	-	1825178	1024,4096
itable32	157	10K	-	157	64
itable64	2	1K	-	2	128
lr	1	1K	-	4	16384
pic	37	6K	-	37	64,16384
pfestat	0	0K	-	1682	32,64,128,256,131072
gencfg	1486	424K	-	2812	16,32,64,256,512,16384,32768,65536
jsr	0	0K	-	22	16
idl	0	0K	-	4	32768,131072
rtmsg	0	0K	-	3	131072
module	250	16K	-	250	64,128
mtx_pool	1	8K	-	1	64,128
DEVFS3	108	12K	-	109	256
DEVFS1	101	23K	-	101	2048
pgrp	5	1K	-	917	64
session	5	1K	-	917	512
proc	2	1K	-	2	16384
subproc	217	441K	-	4867	2048,131072
cred	21	3K	-	48719	256
plimit	9	2K	-	5255	2048
uidinfo	2	1K	-	2	32,512
sysctluid	2786	85K	-	2786	16,32,64
sysctltmp	0	0K	-	1833	16,32,64,1024
umtx	126	8K	-	126	64
SWAP	2	277K	-	2	64
bus	780	125K	-	2734	16,32,64,128,32768
bus-sc	69	69K	-	1194	
16,32,64,512,1024,2048,8192,16384,65536,131072					
devstat	8	17K	-	8	16,131072
eventhandler	45	2K	-	46	32,128
kobj	93	186K	-	111	65536
DEVFS	8	1K	-	9	16,64
rman	94	6K	-	477	16,32,64
sbuf	0	0K	-	532	16,32,32768,131072
NULLFS hash	1	1K	-	1	64
taskqueue	5	1K	-	5	64
turnstiles	127	8K	-	127	64
Unitno	6	1K	-	44	16,64
ioctlops	0	0K	-	1771718	16,32,64,128,8192,16384,65536,131072
iov	0	0K	-	79425	16,64,128,256,512,1024,2048,131072
msg	4	25K	-	4	32768,131072
sem	4	7K	-	4	16384,32768,131072
shm	2	13K	-	4	32768
ttys	93	16K	-	195	512,32768
soname	31	3K	-	389284	16,32,64,256
pcb	101	16K	-	4374	
16,32,64,128,1024,2048,4096,16384,65536					
BIO buffer	40	80K	-	750	65536
vfscache	1	512K	-	1	65536
cluster_save buffer	0	OK	-	55	32,64
VFS hash	1	256K	-	1	32,64
vnodes	1	1K	-	1	512
mount	266	21K	-	481	16,32,64,128,256,4096,32768
vnodemarker	0	0K	-	2497	16384
pfs_nodes	25	3K	-	25	128
pfs_vncache	144	5K	-	386	32
STP	1	1K	-	1	64



```

      GEOM      173      15K      -      1068
16,32,64,128,256,512,2048,16384,32768,131072
      syncache      1      8K      -      1
16,32,64,128,256,512,2048,16384,32768,131072
      tlv_stat      0      0K      -      223
16,32,64,128,256,512,2048,16384,32768,131072
      NFS daemon      1      8K      -      1
16,32,64,128,256,512,2048,16384,32768,131072
      p1003.1b      1      1K      -      1 16
      MD disk      9      18K      -      9 65536
      ata_generic      2      2K      -      25 16,16384,32768
      ISOFS mount      7      1K      -      13 512
      ISOFS node 1439      135K      -      1453 128
      CAM SIM      1      1K      -      1 64
      CAM XPT      6      1K      -      9 16,64,16384
      CAM periph      1      1K      -      1 128
      ad_driver      2      1K      -      2 256
      pagedep      1      64K      -      105 64
      inodedep      1      256K      -      552 256
      newblk      1      1K      -      327 64,4096
      bmsafemap      0      0K      -      19 64
      allocdirect      0      0K      -      326 128
      freefrag      0      0K      -      31 32
      freeblks      0      0K      -      103 2048
      freefile      0      0K      -      175 32
      diradd      0      0K      -      590 64
      mkdir      0      0K      -      166 32
      dirrem      0      0K      -      382 32
      savedino      0      0K      -      283 512
      UFS mount      15      36K      -      15 2048,65536,131072
      ata_dma      6      1K      -      6 256
      UMAHash      1      4K      -      5 4096,16384,32768,65536,131072
      cdev      26      3K      -      26 256
      file desc 111      25K      -      5199 16,1024,2048,16384
      VM pgdata      2      65K      -      2 64
      sigio      1      1K      -      27 32
      kenv      30      5K      -      33 16,32,64,131072
      atkbddev      2      1K      -      2 32
      kqueue      0      0K      -      88 1024,4096,32768
      proc-args      28      2K      -      3970 32,64,128,256,512,1024
      isadev      23      2K      -      23 64
      zombie      1      1K      -      4651 128
      ithread      92      7K      -      92 16,64,256
      legacydrv      3      1K      -      3 16
      memdesc      1      4K      -      1 131072
      nexusdev      2      1K      -      2 16
      CAM queue      3      1K      -      3 16
      KTRACE      100      10K      -      100 128
      kbdmux      5      9K      -      5 128,2048,65536,131072
ITEM      SIZE      LIMIT      USED      FREE      REQUESTS
UMA Kegs:      136,      0,      71,      1,      71
...
```

### show system virtual-memory | display xml

```

user@host> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.2R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
    </vmstat-memstat-malloc>
  </system-virtual-memory-information>
</rpc-reply>
```

```
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>entropy</memstat-name>
<inuse>1024</inuse>
<memuse>64</memuse>
<high-use>--</high-use>
<memstat-req>1024</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>linker</memstat-name>
<inuse>481</inuse>
<memuse>1871</memuse>
<high-use>--</high-use>
<memstat-req>1145</memstat-req>
<memstat-size>16,32,64,4096,32768,131072</memstat-size>
<memstat-name>lockf</memstat-name>
<inuse>56</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>5998</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>devbuf</memstat-name>
<inuse>2094</inuse>
<memuse>3877</memuse>
<high-use>--</high-use>
<memstat-req>2099</memstat-req>

<memstat-size>16,32,64,128,512,1024,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>temp</memstat-name>
<inuse>21</inuse>
<memuse>66</memuse>
<high-use>--</high-use>
<memstat-req>3127</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ip6ndp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6ifmulti</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>in6grentry</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
```

```

<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

```

```
<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifmaddr</memstat-name>
  <inuse>103</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>103</memstat-req>
  <memstat-size>16,32</memstat-size>
  <memstat-name>rtable</memstat-name>
  <inuse>129</inuse>
  <memuse>14</memuse>
  <high-use>--</high-use>
  <memstat-req>139</memstat-req>
  <memstat-size>16,32,64,128,1024,16384</memstat-size>
  <memstat-name>sysctl</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>14847</memstat-req>
  <memstat-size>16,32,64,4096,16384,32768</memstat-size>
  <memstat-name>ifaddr</memstat-name>
  <inuse>29</inuse>
  <memuse>3</memuse>
  <high-use>--</high-use>
  <memstat-req>29</memstat-req>
  <memstat-size>64,128</memstat-size>
  <memstat-name>mkey</memstat-name>
  <inuse>345</inuse>
  <memuse>6</memuse>
  <high-use>--</high-use>
  <memstat-req>2527</memstat-req>
  <memstat-size>16,128</memstat-size>
  <memstat-name>pfe_ipc</memstat-name>
  <inuse>0</inuse>
  <memuse>0</memuse>
  <high-use>--</high-use>
  <memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>
  <memstat-name>ifstate</memstat-name>
  <inuse>594</inuse>
  <memuse>51</memuse>
  <high-use>--</high-use>
  <memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
  <memstat-name>itable16</memstat-name>
  <inuse>276</inuse>
  <memuse>52</memuse>
  <high-use>--</high-use>
  <memstat-req>294</memstat-req>
  <memstat-size>1024,4096</memstat-size>
  <memstat-name>itable32</memstat-name>
  <inuse>160</inuse>
  <memuse>10</memuse>
  <high-use>--</high-use>
  <memstat-req>160</memstat-req>
  <memstat-size>64</memstat-size>
  <memstat-name>itable64</memstat-name>
  <inuse>2</inuse>
  <memuse>1</memuse>
```

```

<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>

```

```
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctlpid</memstat-name>
<inuse>1117</inuse>
<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```

```

<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>

```

```
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>iocltops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
```



```

<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>

<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>

```

```
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD_disk</memstat-name>
<inuse>10</inuse>
```

```

<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISOFs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISOFs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
```

```

<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>

```

```
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
```

```

<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>
  <used>30</used>
  <free>20</free>
  <zone-req>30</zone-req>
  <zone-name>32 Bucket:</zone-name>
  <zone-size>140</zone-size>
  <count-limit>0</count-limit>
  <used>33</used>
  <free>23</free>

```

```
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
```



```
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
```

```
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
```

```
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL UMA zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
```

```
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
```

```
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
```

```

<zone-req>0</zone-req>
<zone-name>ripcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>5</used>
<free>29</free>
<zone-req>5</zone-req>
<zone-name>SWAPMETA:</zone-name>
<zone-size>276</zone-size>
<count-limit>94948</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>FFS inode:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>72</free>
<zone-req>1306</zone-req>
<zone-name>FFS1 dinode:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>24</free>
<zone-req>1306</zone-req>
<zone-name>FFS2 dinode:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
</vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>
  <copy-on-write-faults>48364</copy-on-write-faults>
  <copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
  <zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
  <zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
  <transit-blocking-page-faults>85</transit-blocking-page-faults>
  <total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>

```

```

    <pages-affected-by-fork>95343</pages-affected-by-fork>
    <pages-affected-by-vfork>3526</pages-affected-by-vfork>
    <pages-affected-by-rfork>0</pages-affected-by-rfork>
    <pages-freed>221502</pages-freed>
    <pages-freed-by-daemon>0</pages-freed-by-daemon>
    <pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
    <pages-active>45826</pages-active>
    <pages-inactive>13227</pages-inactive>
    <pages-in-vm-cache>49278</pages-in-vm-cache>
    <pages-wired-down>10640</pages-wired-down>
    <pages-free>70706</pages-free>
    <bytes-per-page>4096</bytes-per-page>
    <swap-pages-used>0</swap-pages-used>
    <peak-swap-pages-used>0</peak-swap-pages-used>
    <total-name-lookups>214496</total-name-lookups>
    <positive-cache-hits>92</positive-cache-hits>
    <negative-cache-hits>5</negative-cache-hits>
    <pass2>0</pass2>
    <cache-deletions>0</cache-deletions>
    <cache-falsehits>0</cache-falsehits>
    <toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk          </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0        </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc         </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0   </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1       </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0       </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
<kernel-direct-mm-size-information>
  <vm-directmm-size-used>4644</vm-directmm-size-used>
  <vm-directmm-size-max>4057334</vm-directmm-size-max>
</kernel-direct-mm-size-information>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```

Note: <kernel-direct-mm-size-information> only displays for 64 bit platform.

## show system virtual-memory (QFX Series)

```
user@switch> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/11.1R1/junos">
  <system-virtual-memory-information>
    <vmstat-memstat-malloc>
      <memstat-name>CAM dev queue</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>entropy</memstat-name>
      <inuse>1024</inuse>
      <memuse>64</memuse>
      <high-use>-</high-use>
      <memstat-req>1024</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>linker</memstat-name>
      <inuse>481</inuse>
      <memuse>1871</memuse>
      <high-use>-</high-use>
      <memstat-req>1145</memstat-req>
      <memstat-size>16, 32, 64, 4096, 32768, 131072</memstat-size>
      <memstat-name>lockf</memstat-name>
      <inuse>56</inuse>
      <memuse>4</memuse>
      <high-use>-</high-use>
      <memstat-req>5998</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>devbuf</memstat-name>
      <inuse>2094</inuse>
      <memuse>3877</memuse>
      <high-use>-</high-use>
      <memstat-req>2099</memstat-req>

      <memstat-size>16, 32, 64, 128, 512, 1024, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>temp</memstat-name>
      <inuse>21</inuse>
      <memuse>66</memuse>
      <high-use>-</high-use>
      <memstat-req>3127</memstat-req>

      <memstat-size>16, 32, 64, 128, 256, 512, 2048, 4096, 8192, 16384, 32768, 65536, 131072</memstat-size>

      <memstat-name>ip6ndp</memstat-name>
      <inuse>0</inuse>
      <memuse>0</memuse>
      <high-use>-</high-use>
      <memstat-req>4</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6ifmulti</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
      <high-use>-</high-use>
      <memstat-req>1</memstat-req>
      <memstat-size>64</memstat-size>
      <memstat-name>in6grentry</memstat-name>
      <inuse>1</inuse>
      <memuse>1</memuse>
```



```

<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>iflogical</memstat-name>
<inuse>13</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>64,2048</memstat-size>
<memstat-name>iffamily</memstat-name>
<inuse>28</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>28</memstat-req>
<memstat-size>32,1024,2048</memstat-size>
<memstat-name>rtnexthop</memstat-name>
<inuse>127</inuse>
<memuse>18</memuse>
<high-use>--</high-use>
<memstat-req>129</memstat-req>
<memstat-size>32,256,512,1024,2048,4096</memstat-size>
<memstat-name>metrics</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>inifmulti</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>ingrentry</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>rnode</memstat-name>
<inuse>68</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>76</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rcache</memstat-name>
<inuse>4</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ifdevice</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>ifstat</memstat-name>
<inuse>40</inuse>
<memuse>22</memuse>
<high-use>--</high-use>

```

```
<memstat-req>40</memstat-req>
<memstat-size>512,16384,32768</memstat-size>
<memstat-name>ipfw</memstat-name>
<inuse>42</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>91</memstat-req>

<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
<memstat-name>ifmaddr</memstat-name>
<inuse>103</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>103</memstat-req>
<memstat-size>16,32</memstat-size>
<memstat-name>rtable</memstat-name>
<inuse>129</inuse>
<memuse>14</memuse>
<high-use>--</high-use>
<memstat-req>139</memstat-req>
<memstat-size>16,32,64,128,1024,16384</memstat-size>
<memstat-name>sysctl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>14847</memstat-req>
<memstat-size>16,32,64,4096,16384,32768</memstat-size>
<memstat-name>ifaddr</memstat-name>
<inuse>29</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>29</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mkey</memstat-name>
<inuse>345</inuse>
<memuse>6</memuse>
<high-use>--</high-use>
<memstat-req>2527</memstat-req>
<memstat-size>16,128</memstat-size>
<memstat-name>pfe_ipc</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>

<memstat-name>ifstate</memstat-name>
<inuse>594</inuse>
<memuse>51</memuse>
<high-use>--</high-use>
<memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
<memstat-name>itable16</memstat-name>
<inuse>276</inuse>
<memuse>52</memuse>
<high-use>--</high-use>
<memstat-req>294</memstat-req>
<memstat-size>1024,4096</memstat-size>
<memstat-name>itable32</memstat-name>
```

```

<inuse>160</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>160</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>itable64</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>lr</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pic</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64,512</memstat-size>
<memstat-name>pfestat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>162</memstat-req>
<memstat-size>16,32,128,256,16384</memstat-size>
<memstat-name>gencfg</memstat-name>
<inuse>224</inuse>
<memuse>56</memuse>
<high-use>--</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>--</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>

```

```
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>--</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctluid</memstat-name>
<inuse>1117</inuse>
```

```

<memuse>34</memuse>
<high-use>--</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>--</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
<memstat-name>devstat</memstat-name>
<inuse>10</inuse>
<memuse>21</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>16,131072</memstat-size>
<memstat-name>eventhandler</memstat-name>
<inuse>35</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>36</memstat-req>
<memstat-size>32,128</memstat-size>
<memstat-name>kobj</memstat-name>
<inuse>93</inuse>
<memuse>186</memuse>
<high-use>--</high-use>
<memstat-req>111</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>DEVFS</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>rman</memstat-name>
<inuse>71</inuse>

```

```
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>433</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sbuf</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>522</memstat-req>
<memstat-size>16,32,32768,131072</memstat-size>
<memstat-name>NULLFS hash</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>taskqueue</memstat-name>
<inuse>5</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>turnstiles</memstat-name>
<inuse>145</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>145</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>Unitno</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>iocltops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
```

```

<high-use>--</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>--</high-use>
<memstat-req>528</memstat-req>
<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>--</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>--</high-use>
<memstat-req>484</memstat-req>
<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
<memstat-name>BIO buffer</memstat-name>
<inuse>43</inuse>
<memuse>86</memuse>
<high-use>--</high-use>
<memstat-req>405</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>vfscache</memstat-name>
<inuse>1</inuse>
<memuse>256</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>cluster_save buffer</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>VFS hash</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32,64</memstat-size>
<memstat-name>vnodes</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>mount</memstat-name>
<inuse>290</inuse>
<memuse>23</memuse>
<high-use>--</high-use>
<memstat-req>535</memstat-req>
<memstat-size>16,32,64,128,256,4096,32768</memstat-size>
<memstat-name>vnodemarker</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>498</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>pfs_nodes</memstat-name>
<inuse>25</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>25</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>pfs_vncache</memstat-name>
<inuse>27</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>53</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>STP</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>GEOM</memstat-name>
<inuse>146</inuse>
<memuse>11</memuse>
<high-use>--</high-use>
<memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>syncache</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>tlv_stat</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
<memstat-name>NFS_daemon</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
```



```
<memstat-name>p1003.1b</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>MD disk</memstat-name>
<inuse>10</inuse>
<memuse>20</memuse>
<high-use>--</high-use>
<memstat-req>10</memstat-req>
<memstat-size>65536</memstat-size>
<memstat-name>ata_generic</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>6</memstat-req>
<memstat-size>16,16384,32768</memstat-size>
<memstat-name>ISofs mount</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>ISofs node</memstat-name>
<inuse>1440</inuse>
<memuse>135</memuse>
<high-use>--</high-use>
<memstat-req>1457</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>CAM SIM</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>
```

```
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>--</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
```

```

<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>--</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>--</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>

```

```
<high-use>--</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>27</memstat-req>
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>--</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>--</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>zombie</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>--</high-use>
<memstat-req>1278</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ithread</memstat-name>
<inuse>69</inuse>
<memuse>5</memuse>
<high-use>--</high-use>
<memstat-req>69</memstat-req>
<memstat-size>16,64,256</memstat-size>
<memstat-name>legacydrv</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>memdesc</memstat-name>
<inuse>1</inuse>
<memuse>4</memuse>
<high-use>--</high-use>
<memstat-req>1</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>nexusdev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
```

```

<memstat-req>2</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>CAM queue</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>3</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>$PIR</memstat-name>
<inuse>4</inuse>
<memuse>1</memuse>
<high-use>--</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>KTRACE</memstat-name>
<inuse>100</inuse>
<memuse>10</memuse>
<high-use>--</high-use>
<memstat-req>100</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>kbdmux</memstat-name>
<inuse>5</inuse>
<memuse>9</memuse>
<high-use>--</high-use>
<memstat-req>5</memstat-req>
<memstat-size>128,2048,65536,131072</memstat-size>
</vmstat-memstat-malloc>
<vmstat-memstat-zone>
  <zone-name>UMA Kegs:</zone-name>
  <zone-size>136</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>1</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Zones:</zone-name>
  <zone-size>120</zone-size>
  <count-limit>0</count-limit>
  <used>71</used>
  <free>19</free>
  <zone-req>71</zone-req>
  <zone-name>UMA Slabs:</zone-name>
  <zone-size>64</zone-size>
  <count-limit>0</count-limit>
  <used>490</used>
  <free>41</free>
  <zone-req>579</zone-req>
  <zone-name>UMA RCntSlabs:</zone-name>
  <zone-size>104</zone-size>
  <count-limit>0</count-limit>
  <used>276</used>
  <free>20</free>
  <zone-req>276</zone-req>
  <zone-name>UMA Hash:</zone-name>
  <zone-size>128</zone-size>
  <count-limit>0</count-limit>
  <used>4</used>
  <free>26</free>
  <zone-req>5</zone-req>
  <zone-name>16 Bucket:</zone-name>
  <zone-size>76</zone-size>
  <count-limit>0</count-limit>

```

```
<used>30</used>
<free>20</free>
<zone-req>30</zone-req>
<zone-name>32 Bucket:</zone-name>
<zone-size>140</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>23</free>
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
```

```
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
```

```
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
```



```

<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL_UMA_zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>

```

```
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
```

```

<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>

```

```

<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ripcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>5</used>
<free>29</free>
<zone-req>5</zone-req>
<zone-name>SWAPMETA:</zone-name>
<zone-size>276</zone-size>
<count-limit>94948</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>FFS inode:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>72</free>
<zone-req>1306</zone-req>
<zone-name>FFS1 dinode:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>24</free>
<zone-req>1306</zone-req>
<zone-name>FFS2 dinode:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
</vmstat-memstat-zone>
<vmstat-sumstat>
  <cpu-context-switch>934906</cpu-context-switch>
  <dev-intr>1707986</dev-intr>
  <soft-intr>33819</soft-intr>
  <traps>203604</traps>
  <sys-calls>1200636</sys-calls>
  <kernel-thrds>60</kernel-thrds>
  <fork-calls>1313</fork-calls>
  <vfork-calls>21</vfork-calls>
  <rfork-calls>0</rfork-calls>
  <swap-pageins>0</swap-pageins>
  <swap-pagedin>0</swap-pagedin>
  <swap-pageouts>0</swap-pageouts>
  <swap-pagedout>0</swap-pagedout>
  <vnode-pageins>23094</vnode-pageins>
  <vnode-pagedin>23119</vnode-pagedin>
  <vnode-pageouts>226</vnode-pageouts>
  <vnode-pagedout>3143</vnode-pagedout>
  <page-daemon-wakeup>0</page-daemon-wakeup>
  <page-daemon-examined-pages>0</page-daemon-examined-pages>
  <pages-reactivated>8821</pages-reactivated>

```

```

<copy-on-write-faults>48364</copy-on-write-faults>
<copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
<zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
<zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
<transit-blocking-page-faults>85</transit-blocking-page-faults>
<total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>
<pages-affected-by-fork>95343</pages-affected-by-fork>
<pages-affected-by-vfork>3526</pages-affected-by-vfork>
<pages-affected-by-rfork>0</pages-affected-by-rfork>
<pages-freed>221502</pages-freed>
<pages-freed-by-daemon>0</pages-freed-by-daemon>
<pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
<pages-active>45826</pages-active>
<pages-inactive>13227</pages-inactive>
<pages-in-vm-cache>49278</pages-in-vm-cache>
<pages-wired-down>10640</pages-wired-down>
<pages-free>70706</pages-free>
<bytes-per-page>4096</bytes-per-page>
<swap-pages-used>0</swap-pages-used>
<peak-swap-pages-used>0</peak-swap-pages-used>
<total-name-lookups>214496</total-name-lookups>
<positive-cache-hits>92</positive-cache-hits>
<negative-cache-hits>5</negative-cache-hits>
<pass2>0</pass2>
<cache-deletions>0</cache-deletions>
<cache-falsehits>0</cache-falsehits>
<toolong>0</toolong>
</vmstat-sumstat>
<vmstat-intr>
  <intr-name>irq0: clk      </intr-name>
  <intr-cnt>1243455</intr-cnt>
  <intr-rate>999</intr-rate>
  <intr-name>irq4: sio0     </intr-name>
  <intr-cnt>1140</intr-cnt>
  <intr-rate>0</intr-rate>
  <intr-name>irq8: rtc      </intr-name>
  <intr-cnt>159164</intr-cnt>
  <intr-rate>127</intr-rate>
  <intr-name>irq9: cbb1 fxp0 </intr-name>
  <intr-cnt>28490</intr-cnt>
  <intr-rate>22</intr-rate>
  <intr-name>irq10: fxp1    </intr-name>
  <intr-cnt>20593</intr-cnt>
  <intr-rate>16</intr-rate>
  <intr-name>irq14: ata0    </intr-name>
  <intr-cnt>5031</intr-cnt>
  <intr-rate>4</intr-rate>
  <intr-name>Total</intr-name>
  <intr-cnt>1457873</intr-cnt>
  <intr-rate>1171</intr-rate>
</vmstat-intr>
<vm-kernel-state>
  <vm-kmem-map-free>248524800</vm-kmem-map-free>
</vm-kernel-state>
</system-virtual-memory-information>
<cli>
  <banner></banner>
</cli>
</rpc-reply>

```



## show version

<b>List of Syntax</b>	<a href="#">Syntax on page 1205</a> <a href="#">Syntax (EX Series Switches) on page 1205</a> <a href="#">Syntax (TX Matrix Router) on page 1205</a> <a href="#">Syntax (TX Matrix Plus Router) on page 1205</a> <a href="#">Syntax (MX Series Router) on page 1205</a> <a href="#">Syntax (QFX Series) on page 1205</a>
<b>Syntax</b>	<pre>show version &lt;brief   detail&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show version &lt;all-members&gt; &lt;brief   detail&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Syntax (TX Matrix Router)</b>	<pre>show version &lt;brief   detail&gt; &lt;all-chassis   all-lcc   lcc number   scc&gt;</pre>
<b>Syntax (TX Matrix Plus Router)</b>	<pre>show version &lt;all-chassis   all-lcc   lcc number   sfc number&gt; &lt;brief   detail&gt;</pre>
<b>Syntax (MX Series Router)</b>	<pre>show version &lt;brief   detail&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show version &lt;brief   detail&gt; &lt;component component-name   all&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Display the hostname and version information about the software running on the router or switch.</p> <p>Beginning in Junos OS Release 13.3, the <b>show version</b> command output includes the <b>Junos</b> field that displays the Junos OS version running on the device. This field provides a consistent means of identifying the Junos OS version, rather than extracting that information from the list of installed sub-packages.</p>
<b>Options</b>	<p><b>none</b>—Display standard information about the hostname and version of the software running on the router or switch.</p>

**brief | detail**—(Optional) Display the specified level of output.

**all-members**—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on all members of the Virtual Chassis configuration.

**component all**—(QFabric systems only) (Optional) Display the host name and version information about the software running on all the components on the QFabric system.

**component *component-name***—(QFabric systems only) (Optional) Display the host name and version information about the software running on a specific QFabric system component. Replace *component-name* with the name of the QFabric system component. The *component-name* can be the name of a diagnostics Routing Engine, Director group, fabric control Routing Engine, fabric manager Routing Engine, Interconnect device, or Node group.

**local**—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the local Virtual Chassis member.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Display standard information about the hostname and version of the software running on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace *member-id* with a value from 0 through 9. For an MX Series Virtual Chassis, replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix router (or switch-card chassis).

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display the host name and version information about the software running on for a specified T640 router (line-card chassis or LCC) that is connected to the TX Matrix router. On a TX Matrix Plus router, display the host name and version information about the software running for a specified T1600 or T4000 router (LCC) that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display the hostname and version information about the software running on the TX Matrix Plus router (or switch-fabric chassis). Replace *number* with 0.



**Additional Information** By default, when you issue the **show version** command on a TX Matrix or TX Matrix Plus master Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) master Routing Engines connected to it. Likewise, if you issue the same command on the TX Matrix or TX Matrix Plus backup Routing Engine, the command is broadcast to all the T640 (in a routing matrix based on a TX Matrix router) or T1600 or T4000 (in a routing matrix based on a TX Matrix Plus router) backup Routing Engines that are connected to it.

**Required Privilege Level** view

**List of Sample Output** [show version \(Devices Running Junos OS Release 13.3 and Later\) on page 1208](#)  
[show version on page 1208](#)  
[show version \(TX Matrix Plus Router\) on page 1209](#)  
[show version \(TX Matrix Plus Router with 3D SIBs\) on page 1211](#)  
[show version \(MX Series Router\) on page 1215](#)  
[show version \(QFX3500 Switch\) on page 1215](#)  
[show version \(QFabric System\) on page 1215](#)  
[show version component all \(QFabric System\) on page 1216](#)

## Sample Output

### show version (Devices Running Junos OS Release 13.3 and Later)

The following output is from the MX240 Router and shows the **Junos** field introduced in Junos OS 13.3. Depending on the platform running Junos OS 13.3, you might see different installed sub-packages, but the **Junos** field is common across all platforms that run Junos OS 13.3 and later.

```
user@host > show version
Hostname: lab
Model: mx240
Junos: 13.3R1.4
JUNOS Base OS boot [13.3R1.4]
JUNOS Base OS Software Suite [13.3R1.4]
JUNOS Kernel Software Suite [13.3R1.4]
JUNOS Crypto Software Suite [13.3R1.4]
JUNOS Packet Forwarding Engine Support (M/T/EX Common) [13.3R1.4]
JUNOS Packet Forwarding Engine Support (MX Common) [13.3R1.4]
JUNOS Online Documentation [13.3R1.4]
JUNOS Services ACL Container package [13.3R1.4]
JUNOS Services Application Level Gateways [13.3R1.4]
JUNOS AppId Services [13.3R1.4]
JUNOS Border Gateway Function package [13.3R1.4]
JUNOS Services Captive Portal and Content Delivery Container package [13.3R1.4]
JUNOS Services HTTP Content Management package [13.3R1.4]
JUNOS IDP Services [13.3R1.4]
JUNOS Services Jflow Container package [13.3R1.4]
JUNOS Services LL-PDF Container package [13.3R1.4]
JUNOS Services MobileNext Software package [13.3R1.4]
JUNOS Services Mobile Subscriber Service Container package [13.3R1.4]
JUNOS Services NAT [13.3R1.4]
JUNOS Services PTSP Container package [13.3R1.4]
JUNOS Services RPM [13.3R1.4]
JUNOS Services Stateful Firewall [13.3R1.4]
JUNOS Voice Services Container package [13.3R1.4]
JUNOS Services Crypto [13.3R1.4]
JUNOS Services SSL [13.3R1.4]
JUNOS Services IPSec [13.3R1.4]
JUNOS platform Software Suite [13.3R1.4]
JUNOS Runtime Software Suite [13.3R1.4]
JUNOS Routing Software Suite [13.3R1.4]
JUNOS py-base-i386 [13.3R1.4]
```

### show version

```
user@host> show version
Hostname: router1
Model: m20
JUNOS Base OS boot [7.2-20050312.0]
JUNOS Base OS Software Suite [7.2-20050312.0]
JUNOS Kernel Software Suite [7.2R1.7]
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]
JUNOS Routing Software Suite [7.2R1.7]
JUNOS Online Documentation [7.2R1.7]
JUNOS Crypto Software Suite [7.2R1.7]

{master}

user@host> show version psd 1
```

```
psd1-re0:
```

```
-----
Hostname: china
Model: t640
JUNOS Base OS boot [9.1I20080311_1959_builder]
JUNOS Base OS Software Suite [9.1-20080321.0]
JUNOS Kernel Software Suite [9.1-20080321.0]
JUNOS Crypto Software Suite [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [9.1-20080321.0]
JUNOS Packet Forwarding Engine Support (T-series) [9.1-20080321.0]
JUNOS Online Documentation [9.1-20080321.0]
JUNOS Routing Software Suite [9.1-20080321.0]
labpkg [7.0]
```

### show version (TX Matrix Plus Router)

```
user@host> show version
```

```
sfc0-re0:
```

```
-----
Hostname: host
Model: txp
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]
```

```
lcc0-re0:
```

```
-----
Hostname: host1
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
```

JUNOS Services Application Level Gateways [12.3-20121019.0]  
JUNOS AppId Services [12.3-20121019.0]  
JUNOS Border Gateway Function package [12.3-20121019.0]  
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]  
JUNOS Services HTTP Content Management package [12.3-20121019.0]  
JUNOS IDP Services [12.3-20121019.0]  
JUNOS Services LL-PDF Container package [12.3-20121019.0]  
JUNOS Services NAT [12.3-20121019.0]  
JUNOS Services PTSP Container package [12.3-20121019.0]  
JUNOS Services RPM [12.3-20121019.0]  
JUNOS Services Stateful Firewall [12.3-20121019.0]  
JUNOS Voice Services Container package [12.3-20121019.0]  
JUNOS Services Example Container package [12.3-20121019.0]  
JUNOS Services Crypto [12.3-20121019.0]  
JUNOS Services SSL [12.3-20121019.0]  
JUNOS Services IPSec [12.3-20121019.0]  
JUNOS Runtime Software Suite [12.3-20121019.0]  
JUNOS Routing Software Suite [12.3-20121019.0]

lcc1-re0:

-----  
Hostname: host2  
Model: t1600  
JUNOS Base OS boot [12.3-20121019.0]  
JUNOS Base OS Software Suite [12.3-20121019.0]  
JUNOS Kernel Software Suite [12.3-20121019.0]  
JUNOS Crypto Software Suite [12.3-20121019.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]  
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]  
JUNOS Online Documentation [12.3-20121019.0]  
JUNOS Services ACL Container package [12.3-20121019.0]  
JUNOS Services Application Level Gateways [12.3-20121019.0]  
JUNOS AppId Services [12.3-20121019.0]  
JUNOS Border Gateway Function package [12.3-20121019.0]  
JUNOS Services Captive Portal and Content Delivery Container package [12.3-20121019.0]  
JUNOS Services HTTP Content Management package [12.3-20121019.0]  
JUNOS IDP Services [12.3-20121019.0]  
JUNOS Services LL-PDF Container package [12.3-20121019.0]  
JUNOS Services NAT [12.3-20121019.0]  
JUNOS Services PTSP Container package [12.3-20121019.0]  
JUNOS Services RPM [12.3-20121019.0]  
JUNOS Services Stateful Firewall [12.3-20121019.0]  
JUNOS Voice Services Container package [12.3-20121019.0]  
JUNOS Services Example Container package [12.3-20121019.0]  
JUNOS Services Crypto [12.3-20121019.0]  
JUNOS Services SSL [12.3-20121019.0]  
JUNOS Services IPSec [12.3-20121019.0]  
JUNOS Runtime Software Suite [12.3-20121019.0]  
JUNOS Routing Software Suite [12.3-20121019.0]

lcc2-re0:

-----  
Hostname: host3  
Model: t1600  
JUNOS Base OS boot [12.3-20121019.0]  
JUNOS Base OS Software Suite [12.3-20121019.0]  
JUNOS Kernel Software Suite [12.3-20121019.0]  
JUNOS Crypto Software Suite [12.3-20121019.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]

```

JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

```
lcc3-re0:
```

```

-----
Hostname: host4
Model: t1600
JUNOS Base OS boot [12.3-20121019.0]
JUNOS Base OS Software Suite [12.3-20121019.0]
JUNOS Kernel Software Suite [12.3-20121019.0]
JUNOS Crypto Software Suite [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [12.3-20121019.0]
JUNOS Packet Forwarding Engine Support (T-Series) [12.3-20121019.0]
JUNOS Online Documentation [12.3-20121019.0]
JUNOS Services AACL Container package [12.3-20121019.0]
JUNOS Services Application Level Gateways [12.3-20121019.0]
JUNOS AppId Services [12.3-20121019.0]
JUNOS Border Gateway Function package [12.3-20121019.0]
JUNOS Services Captive Portal and Content Delivery Container package
[12.3-20121019.0]
JUNOS Services HTTP Content Management package [12.3-20121019.0]
JUNOS IDP Services [12.3-20121019.0]
JUNOS Services LL-PDF Container package [12.3-20121019.0]
JUNOS Services NAT [12.3-20121019.0]
JUNOS Services PTSP Container package [12.3-20121019.0]
JUNOS Services RPM [12.3-20121019.0]
JUNOS Services Stateful Firewall [12.3-20121019.0]
JUNOS Voice Services Container package [12.3-20121019.0]
JUNOS Services Example Container package [12.3-20121019.0]
JUNOS Services Crypto [12.3-20121019.0]
JUNOS Services SSL [12.3-20121019.0]
JUNOS Services IPSec [12.3-20121019.0]
JUNOS Runtime Software Suite [12.3-20121019.0]
JUNOS Routing Software Suite [12.3-20121019.0]

```

#### show version (TX Matrix Plus Router with 3D SIBs)

```

user@host>show version
sfc0-re0:

```

```

-----
Hostname: sfc0

```

```
Model: txp
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc0-re0:
```

```
-----
Hostname: lcc0
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services ACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
```

```
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc2-re0:
```

```
-----
Hostname: lcc2
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]
```

```
lcc4-re0:
```

```
-----
Hostname: lcc4
Model: t4000
JUNOS Base OS boot [13.1-20130306.0]
JUNOS Base OS Software Suite [13.1-20130306.0]
JUNOS Kernel Software Suite [13.1-20130306.0]
JUNOS Crypto Software Suite [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
```

JUNOS Services Jflow Container package [13.1-20130306.0]  
JUNOS Services LL-PDF Container package [13.1-20130306.0]  
JUNOS Services MobileNext Software package [13.1-20130306.0]  
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]  
JUNOS Services NAT [13.1-20130306.0]  
JUNOS Services PTSP Container package [13.1-20130306.0]  
JUNOS Services RPM [13.1-20130306.0]  
JUNOS Services Stateful Firewall [13.1-20130306.0]  
JUNOS Voice Services Container package [13.1-20130306.0]  
JUNOS Services Example Container package [13.1-20130306.0]  
JUNOS Services Crypto [13.1-20130306.0]  
JUNOS Services SSL [13.1-20130306.0]  
JUNOS Services IPSec [13.1-20130306.0]  
JUNOS Runtime Software Suite [13.1-20130306.0]  
JUNOS Routing Software Suite [13.1-20130306.0]

lcc6-re0:

-----  
Hostname: lcc6  
Model: t1600  
JUNOS Base OS boot [13.1-20130306.0]  
JUNOS Base OS Software Suite [13.1-20130306.0]  
JUNOS Kernel Software Suite [13.1-20130306.0]  
JUNOS Crypto Software Suite [13.1-20130306.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]  
JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]  
JUNOS Online Documentation [13.1-20130306.0]  
JUNOS Services AACL Container package [13.1-20130306.0]  
JUNOS Services Application Level Gateways [13.1-20130306.0]  
JUNOS AppId Services [13.1-20130306.0]  
JUNOS Border Gateway Function package [13.1-20130306.0]  
JUNOS Services Captive Portal and Content Delivery Container package [13.1-20130306.0]  
JUNOS Services HTTP Content Management package [13.1-20130306.0]  
JUNOS IDP Services [13.1-20130306.0]  
JUNOS Services Jflow Container package [13.1-20130306.0]  
JUNOS Services LL-PDF Container package [13.1-20130306.0]  
JUNOS Services MobileNext Software package [13.1-20130306.0]  
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]  
JUNOS Services NAT [13.1-20130306.0]  
JUNOS Services PTSP Container package [13.1-20130306.0]  
JUNOS Services RPM [13.1-20130306.0]  
JUNOS Services Stateful Firewall [13.1-20130306.0]  
JUNOS Voice Services Container package [13.1-20130306.0]  
JUNOS Services Example Container package [13.1-20130306.0]  
JUNOS Services Crypto [13.1-20130306.0]  
JUNOS Services SSL [13.1-20130306.0]  
JUNOS Services IPSec [13.1-20130306.0]  
JUNOS Runtime Software Suite [13.1-20130306.0]  
JUNOS Routing Software Suite [13.1-20130306.0]

lcc7-re0:

-----  
Hostname: lcc7  
Model: t1600  
JUNOS Base OS boot [13.1-20130306.0]  
JUNOS Base OS Software Suite [13.1-20130306.0]  
JUNOS Kernel Software Suite [13.1-20130306.0]  
JUNOS Crypto Software Suite [13.1-20130306.0]  
JUNOS Packet Forwarding Engine Support (M/T Common) [13.1-20130306.0]



```

JUNOS Packet Forwarding Engine Support (T-Series) [13.1-20130306.0]
JUNOS Online Documentation [13.1-20130306.0]
JUNOS Services AACL Container package [13.1-20130306.0]
JUNOS Services Application Level Gateways [13.1-20130306.0]
JUNOS AppId Services [13.1-20130306.0]
JUNOS Border Gateway Function package [13.1-20130306.0]
JUNOS Services Captive Portal and Content Delivery Container package
[13.1-20130306.0]
JUNOS Services HTTP Content Management package [13.1-20130306.0]
JUNOS IDP Services [13.1-20130306.0]
JUNOS Services Jflow Container package [13.1-20130306.0]
JUNOS Services LL-PDF Container package [13.1-20130306.0]
JUNOS Services MobileNext Software package [13.1-20130306.0]
JUNOS Services Mobile Subscriber Service Container package [13.1-20130306.0]
JUNOS Services NAT [13.1-20130306.0]
JUNOS Services PTSP Container package [13.1-20130306.0]
JUNOS Services RPM [13.1-20130306.0]
JUNOS Services Stateful Firewall [13.1-20130306.0]
JUNOS Voice Services Container package [13.1-20130306.0]
JUNOS Services Example Container package [13.1-20130306.0]
JUNOS Services Crypto [13.1-20130306.0]
JUNOS Services SSL [13.1-20130306.0]
JUNOS Services IPSec [13.1-20130306.0]
JUNOS Runtime Software Suite [13.1-20130306.0]
JUNOS Routing Software Suite [13.1-20130306.0]

```

#### show version (MX Series Router)

```

user@host5> show version
Hostname: host5
Model: mx80
JUNOS Base OS boot [11.3-20110717.0]
JUNOS Base OS Software Suite [11.3-20110717.0]
JUNOS Kernel Software Suite [11.3-20110717.0]
JUNOS Crypto Software Suite [11.3-20110717.0]
JUNOS Packet Forwarding Engine Support (MX80) [11.3-20110717.0]
JUNOS Online Documentation [11.3-20110717.0]
JUNOS Routing Software Suite [11.3-20110717.0]

```

#### show version (QFX3500 Switch)

```

user@switch> show version
Hostname: switch
Model: qfx_s3500
JUNOS Base OS boot [11.1R1]
JUNOS Base OS Software Suite [11.1R1]
JUNOS Kernel Software Suite [11.1R1]
JUNOS Crypto Software Suite [11.1R1]
JUNOS Online Documentation [11.1R1]
JUNOS Enterprise Software Suite [11.1R1]
JUNOS Packet Forwarding Engine Support (QFX) [11.1R1]
JUNOS Routing Software Suite [11.1R1]

```

#### show version (QFabric System)

```

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

```

**show version component all (QFabric System)**

```
user@switch> show version component all
dg1:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

dg0:
-
Hostname: qfabric
Model: qfx3100
JUNOS Base Version [11.3R1.6]

NW-NG-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-0:
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

FC-1:
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]

DRE-0:
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
```

```
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
FM-0:
```

```
-
```

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```

```
nodedevice1:
```

```
-
```

```
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
```


```
interconnectdevice1:
```

```
-
```

```
Hostname: qfabric
Model: QFX3108
JUNOS Base OS boot [11.3R1.6]
JUNOS Base OS Software Suite [11.3R1.6]
JUNOS Kernel Software Suite [11.3R1.6]
JUNOS Crypto Software Suite [11.3R1.6]
JUNOS Online Documentation [11.3R1.6]
JUNOS Enterprise Software Suite [11.3R1.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3R1.6]
JUNOS Routing Software Suite [11.3R1.6]
warning: from interconnectdevice0: Disconnected
```

## start shell

---

<b>Syntax</b>	<code>start shell (csh   sh)</code> <code>&lt;user username&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type <b>exit</b> from the shell.
<div> <b>NOTE:</b><ul style="list-style-type: none"><li>To issue this command, the user must have the required login access privileges configured by including the <b>permissions</b> statement at the <b>[edit system login class <i>class-name</i>]</b> hierarchy level.</li><li>UNIX wheel group membership or permissions are no longer required to issue this command.</li></ul></div>	
<b>Options</b>	<b>csh</b> —Create a UNIX C shell.  <b>sh</b> —Create a UNIX Bourne shell.  <b>user <i>username</i></b> —(Optional) Start the shell as another user.
<b>Additional Information</b>	When you are in the shell, the shell prompt has the following format:  <code>username@hostname%</code> An example of the prompt is:  <code>root@host%</code>
<b>Required Privilege Level</b>	shell and maintenance
<b>List of Sample Output</b>	<a href="#">start shell csh on page 1218</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### start shell csh

```
user@host> start shell csh
%
exit
%
```

```
username@hostname% start shell sh
%

exit
user@host>
```

## test configuration

---

<b>Syntax</b>	<code>test configuration filename</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found. This command only accepts text files.
<b>Options</b>	<b>filename</b> —Name of the configuration file.  <b>syntax-only</b> —Check the syntax of a partial configuration file, without checking for commit errors. This option introduced in Junos OS Release 12.1.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">test configuration on page 1220</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
    'paris-23;'
      syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
    'login ;'
      statement must contain additional statements
configuration syntax failed
```

## traceroute

**List of Syntax**   [Syntax on page 1221](#)  
                               [Syntax \(QFX Series\) on page 1221](#)

**Syntax**   `traceroute host`  
                   `<as-number-lookup>`  
                   `<bypass-routing>`  
                   `<clns>`  
                   `<gateway address>`  
                   `<inet | inet6>`  
                   `<interface interface-name>`  
                   `<logical system logical-system-name>`  
                   `<monitor host>`  
                   `<mpls (ldp FEC address | rsvp label-switched-path-name)>`  
                   `<no-resolve>`  
                   `<propagate-ttl>`  
                   `<routing-instance routing-instance-name>`  
                   `<source source-address>`  
                   `<tos value>`  
                   `<ttl value>`  
                   `<wait seconds>`

**Syntax (QFX Series)**   `traceroute host`  
                               `<as-number-lookup>`  
                               `<bypass-routing>`  
                               `<gateway address>`  
                               `<inet>`  
                               `<interface interface-name>`  
                               `<monitor host>`  
                               `<no-resolve>`  
                               `<routing-instance routing-instance-name>`  
                               `<source source-address>`  
                               `<tos value>`  
                               `<ttl value>`  
                               `<wait seconds>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                               Command introduced in Junos OS Release 9.0 for EX Series switches.  
                               **mpls** option introduced in Junos OS Release 9.2.  
                               Command introduced in Junos OS Release 11.1 for the QFX Series.  
                               **propagate-ttl** option introduced in Junos OS Release 12.1.

**Description**   Display the route that packets take to a specified network host. Use **traceroute** as a debugging tool to locate points of failure in a network.

**Options**   **host**—IP address or name of remote host.

**as-number-lookup**—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

**bypass-routing**—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached

network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

**clns**—(Optional) Trace the route belonging to the Connectionless Network Service (CLNS).

**gateway address**—(Optional) Address of a router or switch through which the route transits.

**inet | inet6**—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

**interface interface-name**—(Optional) Name of the interface over which to send packets.

**logical-system logical-system-name**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**monitor host**—(Optional) Display real-time monitoring information for the specified host.

**mpls (ldp FEC address | rsvp label-switched-path name)**—(Optional) See [traceroute mpls ldp](#) and [traceroute mpls rsvp](#).

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**propagate-ttl**—(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-propagate-ttl** configuration statement.



**NOTE:** Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

---

**routing-instance routing-instance-name**—(Optional) Name of the routing instance for the traceroute attempt.

**source source-address**—(Optional) Source address of the outgoing traceroute packets.

**tos value**—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is 0 through 255.

**ttl value**—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is 0 through 128.

**wait seconds**—(Optional) Maximum time to wait for a response to the traceroute request.

**Required Privilege Level**    network



- Related Documentation**
- [traceroute monitor on page 1225](#)
- List of Sample Output**
- [traceroute on page 1223](#)
  - [traceroute as-number-lookup host on page 1223](#)
  - [traceroute no-resolve on page 1223](#)
  - [traceroute propagate-ttl on page 1224](#)
  - [traceroute \(Between CE Routers, Layer 3 VPN\) on page 1224](#)
  - [traceroute \(Through an MPLS LSP\) on page 1224](#)
- Output Fields**
- [Table 68 on page 1223](#) describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 68: traceroute Output Fields

Field Name	Field Description
<b>traceroute to</b>	IP address of the receiver.
<b>hops max</b>	Maximum number of hops allowed.
<b>byte packets</b>	Size of packets being sent.
<i>number-of-hops</i>	Number of hops from the source to the named router or switch.
<i>router-name</i>	Name of the router or switch for this hop.
<i>address</i>	Address of the router or switch for this hop.
<b>Round trip time</b>	Average round-trip time, in milliseconds (ms).

## Sample Output

### traceroute

```

user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1 blue23 (10.168.1.254)    2.370 ms  2.853 ms  0.367 ms
 2 red14 (10.168.255.250)  0.778 ms  2.937 ms  0.446 ms
 3 yellow (10.156.169.254) 7.737 ms  89.905 ms 0.834 ms

```

### traceroute as-number-lookup host

```

user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1 10.39.1.1 (10.39.1.1) 0.779 ms 0.728 ms 0.562 ms
 2 10.39.1.6 (10.39.1.6) [AS 32] 0.657 ms 0.611 ms 0.617 ms
 3 10.100.1.1 (10.100.1.1) [AS 10, 40, 50] 0.880 ms 0.808 ms 0.774 ms

```

### traceroute no-resolve

```

user@host> traceroute santacruz no-resolve

```

```
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

### traceroute propagate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
    MPLS Label=299776 CoS=0 TTL=1 S=0
    MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
    MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

### traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to vpn09.skybank.net (10.255.14.179), 30 hops max, 40
byte packets
 1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
 2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
    MPLS Label=100006 CoS=0 TTL=1 S=1
 3  vpn09.skybank.net (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

### traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
    MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

## traceroute monitor

<b>List of Syntax</b>	<a href="#">Syntax on page 1225</a> <a href="#">Syntax (QFX Series) on page 1225</a>
<b>Syntax</b>	<pre>traceroute monitor <i>host</i> &lt;count <i>value</i>&gt; &lt;inet   inet6&gt; &lt;interval <i>seconds</i>&gt; &lt;no resolve&gt; &lt;size <i>value</i>&gt; &lt;source <i>source-address</i>&gt; &lt;summary&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>traceroute monitor <i>host</i> &lt;count <i>value</i>&gt; &lt;inet&gt; &lt;interval <i>seconds</i>&gt; &lt;no resolve&gt; &lt;size <i>value</i>&gt; &lt;source <i>source-address</i>&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.0</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display live monitoring of each hop in the route that packets take to a specified network host. Use as a debugging tool to locate points of failure in a network.
<b>Options</b>	<p><b><i>host</i></b>—IP address or name of remote host.</p> <p><b><i>count value</i></b>—Number of ping requests, in packets, to send in summary mode. The default value is <b>10</b>.</p> <p><b><i>inet   inet6</i></b>—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.</p> <p><b><i>interval seconds</i></b>—(Optional) Number of seconds to wait before sending ping requests. The default value is <b>1</b>.</p> <p><b><i>no resolve</i></b>—(Optional) Do not attempt to display addresses symbolically.</p> <p><b><i>size value</i></b>—(Optional) Receive the specified number of bytes for each packet. The range is <b>0</b> through <b>65468</b> bytes. The default value is <b>64</b>.</p> <p><b><i>source source-address</i></b>—(Optional) Source address of the outgoing ping packets.</p> <p><b><i>summary</i></b>—(Optional) Generate and display a summary of live monitoring of each hop on the route that packets take to a specified network host.</p>
<b>Required Privilege Level</b>	network
<b>List of Sample Output</b>	<a href="#">traceroute monitor on page 1226</a>

**Output Fields** Table 69 on page 1226 describes the output fields for the **traceroute monitor** command. Output fields are listed in the approximate order in which they appear.

**Table 69: traceroute monitor Output Fields**

Field Name	Field Description
<b>Host</b>	Hostname or IP address of the router at each hop.
<b>Loss%</b>	Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage.
<b>Snt</b>	Number of ping requests sent to the router at this hop.
<b>Last</b>	Most recent round-trip time, in milliseconds, to the router at this hop.
<b>Avg</b>	Average round-trip time, in milliseconds, to the router at this hop.
<b>Best</b>	Shortest round-trip time, in milliseconds, to the router at this hop.
<b>Wrst</b>	Longest round-trip time, in milliseconds, to the router at this hop.
<b>StDev</b>	Standard deviation of round-trip times, in milliseconds, to the router at this hop.

## Sample Output

### traceroute monitor

```
user@host> traceroute monitor 10.16.0.1
```

	Loss%	Snt	Last	Avg	Best	Wrst	StDev
Host							
1. 10.17.41.254	0.0%	17	0.7	1.0	0.6	5.4	1.2
2. secret.net	0.0%	17	0.6	1.0	0.6	6.6	1.4
3. top-secret.net	0.0%	17	0.6	0.6	0.6	0.6	0.0

## CHAPTER 6

# Troubleshooting

- Troubleshooting Procedures on page 1227

### Troubleshooting Procedures

---

- Creating an Emergency Boot Device on page 1227
- Performing a Recovery Installation on page 1229
- Rebooting and Halting a Device on page 1230
- Recovering from a Failed Software Installation on page 1232
- Recovering the Root Password on page 1233
- Troubleshooting Network Interfaces on page 1234
- Troubleshooting an Aggregated Ethernet Interface on page 1234

### Creating an Emergency Boot Device

If Junos OS on the device is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.



**NOTE:** In the following procedure, we assume that you are creating the emergency boot device on a QFX device or EX4600 device. You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device:

1. Use FTP to copy the installation media image into the **/var/tmp** directory on the device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
```

```
%
```

4. Switch to the root account using the **su** command:

```
% su
```

```
Password: password
```



**NOTE:** The password is the root password for the device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command on the QFX3500, QFX3600, and QFX3600-I devices:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/install-media-qfx3500.junos_11.1 of=/dev/da1 bs=16k
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

6. Enter the following command on the QFX5100 and EX4600 devices:

```
root@device% dd if=/var/tmp/filename of=/dev/da0 bs=1048576
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/jinstall-vjunos-usb-13.2.img of=/dev/da0 bs=1048576
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

7. Log out of the shell:

```
root@device% exit
```

```
% exit
```

```
user@device>
```

#### Related Documentation

- [USB Port Specifications for the QFX Series](#)
- [Performing a Recovery Installation on page 116](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 7233](#)
- [Performing a Recovery Installation on page 118](#)

## Performing a Recovery Installation

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

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

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device” on page 176](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



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

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

To reinstall Junos OS:

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



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

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

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

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

3. The software prompts you with the following options:

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

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

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

The following messages are displayed:

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

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

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

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

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

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

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

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

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

**Related Documentation**

- [Creating an Emergency Boot Device on page 176](#)

## Rebooting and Halting a Device

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

```
user@switch> request system reboot ?
Possible completions:
<[Enter]>          Execute this command
```



```

all-members      Reboot all virtual chassis members
at               Time at which to perform the operation
both-routing-engines Reboot both the Routing Engines
fast-boot        Enable fast reboot
in               Number of minutes to delay before operation
local            Reboot local virtual chassis member
member           Reboot specific virtual chassis member (0..9)
message          Message to display to all users
other-routing-engine Reboot the other Routing Engine
|               Pipe through a command
{master:0}

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

```



**NOTE:** Not all options shown in the preceding command output are available on all QFX Series and EX4600 devices. For example, the `fast-boot` option is available only on QFX5100. See the documentation for the [request system reboot](#) command for details about options.

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



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

```

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

```



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

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

- Related Documentation**
- [clear system reboot on page 355](#)
  - [request system reboot on page 415](#)
  - [request system halt on page 400](#)
  - [request system power-off on page 410](#)
  - *Connecting a QFX Series Device to a Management Console*

## Recovering from a Failed Software Installation

**Problem Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution** If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [- --format] [- --external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



**NOTE:** The root password cannot be recovered on a QFabric system.



**NOTE:** You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.  

```
ok boot -s
```
12. At the following prompt, enter **recovery** to start the root password recovery procedure.  

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: recovery
```
13. Enter configuration mode in the CLI.
14. Set the root password. For example:  

```
user@switch# set system root-authentication plain-text-password
```
15. At the following prompt, enter the new root password. For example:  

```
New password: juniper1
Retype new password:
```
16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.  

```
root@host# commit
commit complete
```
18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the device.  

```
Reboot the system? [y/n] y
```

**Related Documentation** • [Configuring the Root Password on page 1354](#)

## Troubleshooting Network Interfaces

### The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

**Problem** **Description:** The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command **show interfaces *interface-name***, the disabled port is not listed.

**Cause** By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution** Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

## Troubleshooting an Aggregated Ethernet Interface

**Problem** **Description:** The **show interfaces terse** command shows that the LAG is down.

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related  
Documentation**

- [Verifying the Status of a LAG Interface on page 2750](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)



## PART 3

# Configuration and File Management

- [Overview on page 1239](#)
- [Configuration on page 1245](#)
- [Administration on page 1281](#)
- [Troubleshooting on page 1311](#)





## CHAPTER 7

# Overview

- [Configuration Files Overview on page 1239](#)
- [Software Overview on page 1240](#)

## Configuration Files Overview

---

- [Configuration File Terms on page 1239](#)

## Configuration File Terms

[Table 3 on page 11](#) lists the various configuration file terms and their definitions.

**Table 70: Configuration File Terms**

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

- Related Documentation**
- [Loading a Previous Configuration File on page 1252](#)
  - [Reverting to the Rescue Configuration on page 189](#)
  - [Understanding Configuration Files on page 1242](#)

## Software Overview

---

- [Forms of the configure Command on page 1240](#)
- [Junos OS Commit Model for Router or Switch Configuration on page 1241](#)
- [Understanding Configuration Files on page 1242](#)
- [Understanding How the Junos OS Configuration Is Stored on page 1243](#)

### Forms of the configure Command

The Junos OS supports three forms of the **configure** command: **configure**, **configure private**, and **configure exclusive**. These forms control how users edit and commit configurations and can be useful when multiple users configure the software. See [Table 71 on page 1240](#).

**Table 71: Forms of the configure Command**

Command	Edit Access	Commit Access
<b>configure</b>	<ul style="list-style-type: none"><li>• No one can lock the configuration. All users can make configuration changes.</li></ul> <p>When you enter configuration mode, the CLI displays the following information:</p> <ul style="list-style-type: none"><li>• A list of other users editing the configuration.</li><li>• Hierarchy levels the users are viewing or editing.</li><li>• Whether the configuration has been changed, but not committed.</li><li>• When multiple users enter conflicting configurations, the most recent change to be entered takes precedence.</li></ul>	<ul style="list-style-type: none"><li>• No one can lock the configuration. All users can commit all changes to the configuration.</li><li>• If you and another user make changes and the other user commits changes, your changes are committed as well.</li></ul>

---

Table 71: Forms of the configure Command (*continued*)

Command	Edit Access	Commit Access
<b>configure exclusive</b>	<ul style="list-style-type: none"> <li>One user locks the configuration and makes changes without interference from other users.</li> <li>Other users can enter and exit configuration mode, but they cannot commit the configuration.</li> <li>If you enter configuration mode while another user has locked the configuration (with the <b>configure exclusive</b> command), the CLI displays the user and the hierarchy level the user is viewing or editing.</li> <li>If you enter configuration mode while another user has locked the configuration, you can forcibly log out that user with the <b>request system logout</b> operational mode command. For details, see the <a href="#">CLI Explorer</a>.</li> </ul>	
<b>configure private</b>	<ul style="list-style-type: none"> <li>Multiple users can edit the configuration at the same time.</li> <li>Each user has a private candidate configuration to edit independently of other users.</li> <li>When multiple users enter conflicting configurations, the first commit operation takes precedence over subsequent commit operations.</li> </ul>	<ul style="list-style-type: none"> <li>When you commit the configuration, the router verifies that the operational (running) configuration has not been modified by another user before accepting your private candidate configuration as the new operational configuration.</li> <li>If the configuration has been modified by another user, you can merge the modifications into your private candidate configuration and attempt to commit again.</li> </ul>

#### Related Documentation

- *Committing a Junos OS Configuration*
- *Example: Using the configure Command*
- *Displaying Users Currently Editing the Junos OS Configuration*
- *Using the configure exclusive Command*
- *Updating the configure private Configuration*
- *Displaying set Commands from the Junos OS Configuration*

## Junos OS Commit Model for Router or Switch Configuration

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

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

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

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



**NOTE:** The QFX3500 switch has only one Routing Engine.

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

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



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

**Related Documentation**

- *Configuring Junos OS for the First Time on a Router or Switch with a Single Routing Engine*
- [commit](#) on page 345

## Understanding Configuration Files

A configuration file stores the complete configuration of a switch. The current configuration of a switch is called the active configuration. You can alter this current configuration and you can also return to a previous configuration or to a rescue configuration.

Juniper Networks Junos OS saves the 50 most recently committed configuration files on a switch so that you can return to a previous configuration. The configuration files are named:

- **juniper.conf.gz**—The current active configuration.
- **juniper.conf.1.gz** to **juniper.conf.49.gz**—Rollback configurations.

To make changes to the configuration file, you have to work in the configuration mode in the CLI. When making changes to a configuration file, you are viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the active configuration or causing potential damage to your current network operations. Once you commit the changes made to the candidate configuration, the system updates the active configuration.

#### Related Documentation

- [Uploading a Configuration File on page 1261](#)
- [Loading a Previous Configuration File on page 1252](#)
- [Reverting to the Rescue Configuration on page 189](#)
- [Configuration File Terms on page 11](#)

## Understanding How the Junos OS Configuration Is Stored

When you edit a configuration, you work in a copy of the current configuration to create a candidate configuration. The changes you make to the candidate configuration are visible in the CLI immediately, so if multiple users are editing the configuration at the same time, all users can see all changes.

To have a candidate configuration take effect, you *commit* the changes. At this point, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration, when you commit the candidate configuration, all changes made by all the users take effect.

In addition to saving the current configuration, the CLI saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0, which is the current operational version and the default configuration that the system returns to if you roll back to a previous configuration. The oldest saved configuration is version 49.

By default, the Junos OS saves the current configuration and three previous versions of the committed configuration on the CompactFlash card. The currently operational Junos OS configuration is stored in the file **juniper.conf.gz**, and the last three committed configurations are stored in the files **juniper.conf.1.gz**, **juniper.conf.2.gz**, and **juniper.conf.3.gz**. These four files are located in the router or switch's CompactFlash card in the directory **/config**.

The remaining 46 previous versions of committed configurations, the files **juniper.conf.4** through **juniper.conf.49**, are stored in the directory **/var/db/config** on the hard disk.

#### Related Documentation

- [Using Junos OS to Specify the Number of Configurations Stored on the CompactFlash Card](#)
- [Returning to the Most Recently Committed Junos OS Configuration on page 1252](#)

- [Returning to a Previously Committed Junos OS Configuration on page 1253](#)
- [Loading a Configuration from a File on page 1249](#)

## CHAPTER 8

# Configuration

- [Configuration Tasks on page 1245](#)
- [Configuration Statements on page 1265](#)
- [Default Configurations on page 1271](#)
- [Configuration Examples on page 1277](#)

### Configuration Tasks

---

- [Comparing Configuration Changes with a Prior Version on page 1245](#)
- [Compressing the Current Configuration File on page 1247](#)
- [Creating and Returning to a Rescue Configuration on page 1248](#)
- [Loading a Configuration from a File on page 1249](#)
- [Loading a Previous Configuration File on page 1252](#)
- [Returning to the Most Recently Committed Junos OS Configuration on page 1252](#)
- [Returning to a Previously Committed Junos OS Configuration on page 1253](#)
- [Reverting to the Default Factory Configuration on page 1258](#)
- [Reverting to the Rescue Configuration on page 1258](#)
- [Rolling Back Junos OS Configuration Changes on page 1259](#)
- [Saving a Configuration to a File on page 1260](#)
- [Setting or Deleting the Rescue Configuration on page 1261](#)
- [Uploading a Configuration File on page 1261](#)
- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1263](#)

### Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

[edit]

```
user@host# show | compare (filename| rollback n)
```

**filename** is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.

**n** is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (`/config/juniper.conf`).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ( ).

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
  type internal;
  hold-time 60;
  advertise-inactive;
  allow 1.1.1.1/32;
}
group fred {
  type external;
  peer-as 33333;
  allow 2.2.2.2/32;
}
group test-peers {
  type external;
  allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
```



```
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
  -type external;
  -allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# show
group my-group {
  type internal;
  hold-time 90;
  allow 1.1.1.1/32;
}
group fred {
  type external;
  advertise-inactive;
  peer-as 3333;
  allow 2.2.2.2/32;
}
```

#### Related Documentation

- [Creating and Returning to a Rescue Configuration on page 1248](#)

## Compressing the Current Configuration File

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

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



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

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

```
[edit system]
compress-configuration-files;
```

Commit the current configuration file to include the **compression-configuration-files** statement. Commit the configuration again to compress the current configuration file:

```
[edit system]
user@host# set compress-configuration-files
user@host# commit
```

```
commit complete
user@host# commit
commit complete
```

- If you do not want to compress the current operational configuration file, include the **no-compress-configuration-files** statement at the **[edit system]** hierarchy level:

```
[edit system]
no-compression-configuration-files;
```

Commit the current configuration file to include the **no-compress-configuration-files** statement. Commit the configuration again to uncompress the current configuration file:

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

#### Related Documentation

- [Junos OS Commit Model for Router or Switch Configuration on page 14](#)
- [compress-configuration-files on page 265](#)

## Creating and Returning to a Rescue Configuration

A *rescue configuration* allows you to define a known working configuration or a configuration with a known state that you can roll back to at any time. This alleviates the necessity of having to remember the rollback number with the **rollback** command. You use the rescue configuration when you need to roll back to a known configuration or as a last resort if your router or switch configuration and the backup configuration files become damaged beyond repair.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the **request system configuration rescue save** command:

```
user@host> request system configuration rescue save
```

To return to the rescue configuration, use the **rollback rescue** configuration mode command:

```
[edit]
user@host# rollback rescue
load complete
```



**NOTE:** If the rescue configuration does not exist, or if the rescue configuration is not a complete, viable configuration, the **rollback** command fails, an error message appears, and the current configuration remains active.

To activate the rescue configuration that you have loaded, use the **commit** command:

```
[edit]
user@host# rollback rescue
```

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

To delete an existing rescue configuration, issue the **request system configuration rescue delete** command:

```
user@host> request system configuration rescue delete
user@host>
```

For more information about the **request system configuration rescue delete** and **request system configuration rescue save** commands, see the [CLI Explorer](#).

- Related Documentation**
- [Comparing Configuration Changes with a Prior Version on page 1245](#)
  - [Saving a Configuration to a File on page 1257](#)

## Loading a Configuration from a File

You can create a file, copy the file to the local router, and then load the file into the CLI. After you have loaded the file, you can commit it to activate the configuration on the router, or you can edit the configuration interactively using the CLI and commit it at a later time.

You can also create a configuration while typing at the terminal and then load it. Loading a configuration from the terminal is generally useful when you are cutting existing portions of the configuration and pasting them elsewhere in the configuration.

To load an existing configuration file that is located on the router, use the **load** configuration mode command:

```
[edit]
user@host# load (factory-default | merge | override | patch | replace | set | update)
filename <relative>
```

For information about specifying the filename, see *Viewing Files and Directories on a Device Running Junos OS*.

To load a configuration from the terminal, use the following version of the **load** configuration mode command. Press Ctrl-d to end input.

```
[edit]
user@host# load (factory-default | merge | override | patch | replace | set | update)
terminal <relative>
```

To replace an entire configuration, specify the **override** option at any level of the hierarchy. A **load override** operation completely replaces the current candidate configuration with the file you are loading. Thus, if you saved a complete configuration, use this option.

An **override** operation discards the current candidate configuration and loads the configuration in **filename** or the configuration that you type at the terminal. When you use the **override** option and commit the configuration, all system processes reparse the configuration. For an example, see [Figure 13 on page 1277](#).

To replace portions of a configuration, specify the **replace** option. The **load replace** operation looks for **replace:** tags that you added to the loaded file, and replaces the parts of the candidate configuration with whatever is specified after the tag. This is useful when you want more control over exactly what is being changed. For this operation to work, you must include **replace:** tags in the file or configuration you type at the terminal. The software searches for the **replace:** tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the **replace** operation adds to the configuration the statements marked with the **replace:** tag. For an example, see [Figure 14 on page 1278](#).

If, in an **override** or **merge** operation, you specify a file or type text that contains **replace:** tags, the **replace:** tags are ignored and the **override** or **merge** operation is performed.

If you are performing a **replace** operation and the file you specify or text you type does not contain any **replace:** tags, the **replace** operation is effectively equivalent to a **merge** operation. This might be useful if you are running automated scripts and cannot know in advance whether the scripts need to perform a **replace** or a **merge** operation. The scripts can use the **replace** operation to cover either case.

The **load merge** operation adds the saved file to the existing candidate configuration. This is useful if you are adding new configuration sections. For example, suppose that you are adding a BGP configuration to the **[edit protocols]** hierarchy level, where there was no BGP configuration before, you can use the **load merge** operation to combine the saved file configuration to the existing candidate configuration. If the existing configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the existing configuration.

To replace only the configuration that has changed, specify the **update** option at any level of the hierarchy. The **load update** operation compares the candidate configuration and the file you are loading, and only changes the parts of the candidate configuration that are different from the new configuration. You would use this, for example, if there is an existing BGP configuration and the file you are loading changes it in some way.

To change part of the configuration with a patch file, specify the **patch** option. The **load patch** operation loads a file or terminal input that contains configuration changes. First, on a device that already has the configuration changes, you type the **show | compare** command to output the differences between two configurations. Then you can load the differences on another router. The advantage of the **load patch** command is that it saves you from having to copy snippets from different hierarchy levels into a text file prior to loading them into the target device. This might be a useful time saver if you are configuring several devices with the same options. For example, suppose that you configure a routing policy on Device router1 and you want to replicate the policy configuration on Device router2, router3, and router4, you can use the **load patch** operation.

First, run the **show | compare** command.

```
user@router1# show | compare rollback 3
[edit protocols ospf]
+ export default-static;
- export static-default
[edit policy-options]
+ policy-statement default-static {
```

```
+      from protocol static;
+      then accept;
+  }
```

Copy the output of the **show | compare** command to the clipboard, making sure to include the hierarchy levels. On Device router2, router3, and router4, type **load patch terminal** and paste the output. Press Enter and then press Ctrl-d to end the operation. If the patch input specifies different values for an existing statement, the patch input overrides the existing statement.

To use the **merge**, **replace**, **set**, or **update** option without specifying the full hierarchy level, specify the **relative** option. For example:

```
[edit system]
user@host# show static-host-mapping
bob sysid 987.654.321ab
[edit system]
user@host# load replace terminal relative
[Type ^D at a new line to end input]
replace: static-host-mapping {
  bob sysid 0123.456.789bc;
}
load complete
[edit system]
user@host# show static-host-mapping
bob sysid 0123.456.789bc;
```

To load a configuration that contains the **set** configuration mode command, specify the **set** option. This option executes the configuration instructions line by line as they are stored in a file or from a terminal. The instructions can contain any configuration mode command, such as **set**, **edit**, **exit**, and **top**. For an example, see [Figure 17 on page 1279](#).

To copy a configuration file from another network system to the local router, you can use the SSH and Telnet utilities, as described in the [CLI Explorer](#).



**NOTE:** If you are using Junos OS in a Common Criteria environment, system log messages are created whenever a secret attribute is changed (for example, password changes or changes to the RADIUS shared secret). These changes are logged during the following configuration load operations:

```
load merge
load replace
load override
load update
```

For more information, see the *Secure Configuration Guide for Common Criteria and Junos-FIPS*.

#### Related Documentation

- [Examples: Loading a Configuration from a File on page 1277](#)

## Loading a Previous Configuration File

You can use the **rollback** <*number*> command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

### Syntax

**rollback** <*number*>

### Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.
  - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
  - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

**Related Documentation**

- [Configuration File Terms on page 11](#)

## Returning to the Most Recently Committed Junos OS Configuration

To return to the most recently committed configuration and load it into configuration mode without activating it, use the **rollback** configuration mode command:

```
[edit]
user@host# rollback

load complete
```

To activate the configuration to which you rolled back, use the **commit** command:

```
[edit]
user@host# rollback
load complete
[edit]
user@host# commit
```

- Related Documentation**
- [Rolling Back Junos OS Configuration Changes on page 1259](#)
  - [Returning to a Previously Committed Junos OS Configuration on page 1253](#)
  - [Understanding How the Junos OS Configuration Is Stored on page 1243](#)

## Returning to a Previously Committed Junos OS Configuration

This topic explains how you can return to a configuration prior to the most recently committed one, and contains the following sections:

- [Returning to a Configuration Prior to the One Most Recently Committed on page 1253](#)
- [Displaying Previous Configurations on page 1253](#)
- [Comparing Configuration Changes with a Prior Version on page 1254](#)
- [Creating and Returning to a Rescue Configuration on page 1256](#)
- [Saving a Configuration to a File on page 1257](#)

### Returning to a Configuration Prior to the One Most Recently Committed

To return to a configuration prior to the most recently committed one, include the configuration number, 0 through 49, in the **rollback** command. The most recently saved configuration is number 0 (which is the default configuration to which the system returns), and the oldest saved configuration is number 49.

```
[edit]
user@host# rollback number
load complete
```

### Displaying Previous Configurations

To display previous configurations, including the rollback number, date, time, the name of the user who committed changes, and the method of commit, use the **rollback ?** command.

```
[edit]
user@host# rollback ?
Possible completions:
<[Enter]> Execute this command
<number> Numeric argument
0      2005-02-27 12:52:10 PST by abc via cli
1      2005-02-26 14:47:42 PST by def via cli
2      2005-02-14 21:55:45 PST by ghi via cli
3      2005-02-10 16:11:30 PST by jkl via cli
4      2005-02-10 16:02:35 PST by mno via cli
5      2005-03-16 15:10:41 PST by pqr via cli
6      2005-03-16 14:54:21 PST by stu via cli
7      2005-03-16 14:51:38 PST by vwx via cli
8      2005-03-16 14:43:29 PST by yzz via cli
9      2005-03-16 14:15:37 PST by abc via cli
10     2005-03-16 14:13:57 PST by def via cli
11     2005-03-16 12:57:19 PST by root via other
12     2005-03-16 10:45:23 PST by root via other
13     2005-03-16 10:08:13 PST by root via other
```

```
14      2005-03-16 01:20:56 PST by root via other
15      2005-03-16 00:40:37 PST by ghi via cli
16      2005-03-16 00:39:29 PST by jkl via cli
17      2005-03-16 00:32:36 PST by mno via cli
18      2005-03-16 00:31:17 PST by pqr via cli
19      2005-03-15 19:59:00 PST by stu via cli
20      2005-03-15 19:53:39 PST by vwx via cli
21      2005-03-15 18:07:19 PST by yzz via cli
22      2005-03-15 17:59:03 PST by abc via cli
23      2005-03-15 15:05:14 PST by def via cli
24      2005-03-15 15:04:51 PST by ghi via cli
25      2005-03-15 15:03:42 PST by jkl via cli
26      2005-03-15 15:01:52 PST by mno via cli
27      2005-03-15 14:58:34 PST by pqr via cli
28      2005-03-15 13:09:37 PST by root via other
29      2005-03-12 11:01:20 PST by stu via cli
30      2005-03-12 10:57:35 PST by vwx via cli
31      2005-03-11 10:25:07 PST by yzz via cli
32      2005-03-10 23:40:58 PST by abc via cli
33      2005-03-10 23:40:38 PST by def via cli
34      2005-03-10 23:14:27 PST by ghi via cli
35      2005-03-10 23:10:16 PST by jkl via cli
36      2005-03-10 23:01:51 PST by mno via cli
37      2005-03-10 22:49:57 PST by pqr via cli
38      2005-03-10 22:24:07 PST by stu via cli
39      2005-03-10 22:20:14 PST by vwx via cli
40      2005-03-10 22:16:56 PST by yzz via cli
41      2005-03-10 22:16:41 PST by abc via cli
42      2005-03-10 20:44:00 PST by def via cli
43      2005-03-10 20:43:29 PST by ghi via cli
44      2005-03-10 20:39:14 PST by jkl via cli
45      2005-03-10 20:31:30 PST by root via other
46      2005-03-10 18:57:01 PST by mno via cli
47      2005-03-10 18:56:18 PST by pqr via cli
48      2005-03-10 18:47:49 PST by stu via cli
49      2005-03-10 18:47:34 PST by vw via cli
|Pipe through a command
[edit]
```

---

### Comparing Configuration Changes with a Prior Version

In configuration mode only, when you have made changes to the configuration and want to compare the candidate configuration with a prior version, you can use the **compare** command to display the configuration. The **compare** command compares the candidate configuration with either the current committed configuration or a configuration file and displays the differences between the two configurations. To compare configurations, specify the **compare** command after the pipe:

```
[edit]
user@host# show | compare (filename) rollback n)
```

***filename*** is the full path to a configuration file. The file must be in the proper format: a hierarchy of statements.



*n* is the index into the list of previously committed configurations. The most recently saved configuration is number 0, and the oldest saved configuration is number 49. If you do not specify arguments, the candidate configuration is compared against the active configuration file (`/config/juniper.conf`).

The comparison output uses the following conventions:

- Statements that are only in the candidate configuration are prefixed with a plus sign (+).
- Statements that are only in the comparison file are prefixed with a minus sign (-).
- Statements that are unchanged are prefixed with a single blank space ( ).

The following example shows various changes, then a comparison of the candidate configuration with the active configuration, showing only the changes made at the **[edit protocols bgp]** hierarchy level:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# show
group my-group {
    type internal;
    hold-time 60;
    advertise-inactive;
    allow 1.1.1.1/32;
}
group fred {
    type external;
    peer-as 33333;
    allow 2.2.2.2/32;
}
group test-peers {
    type external;
    allow 3.3.3.3/32;
}
[edit protocols bgp]
user@host# set group my-group hold-time 90
[edit protocols bgp]
user@host# delete group my-group advertise-inactive
[edit protocols bgp]
user@host# set group fred advertise-inactive
[edit protocols bgp]
user@host# delete group test-peers
[edit protocols bgp]
user@host# show | compare
[edit protocols bgp group my-group]
-hold-time 60;
+hold-time 90;
-advertise-inactive;
[edit protocols bgp group fred]
+advertise-inactive;
[edit protocols bgp]
-group test-peers {
```

```
-type external;  
-allow 3.3.3.3/32;  
}  
[edit protocols bgp]  
user@host# show  
group my-group {  
  type internal;  
  hold-time 90;  
  allow 1.1.1.1/32;  
}  
group fred {  
  type external;  
  advertise-inactive;  
  peer-as 3333;  
  allow 2.2.2.2/32;  
}
```

---

### Creating and Returning to a Rescue Configuration

A *rescue configuration* allows you to define a known working configuration or a configuration with a known state that you can roll back to at any time. This alleviates the necessity of having to remember the rollback number with the **rollback** command. You use the rescue configuration when you need to roll back to a known configuration or as a last resort if your router or switch configuration and the backup configuration files become damaged beyond repair.

To save the most recently committed configuration as the rescue configuration so that you can return to it at any time, issue the **request system configuration rescue save** command:

```
user@host> request system configuration rescue save
```

To return to the rescue configuration, use the **rollback rescue** configuration mode command:

```
[edit]  
user@host# rollback rescue  
load complete
```



**NOTE:** If the rescue configuration does not exist, or if the rescue configuration is not a complete, viable configuration, the rollback command fails, an error message appears, and the current configuration remains active.

---

To activate the rescue configuration that you have loaded, use the **commit** command:

```
[edit]  
user@host# rollback rescue  
load complete  
[edit]  
user@host# commit
```

To delete an existing rescue configuration, issue the **request system configuration rescue delete** command:

```
user@host> request system configuration rescue delete
user@host>
```

For more information about the **request system configuration rescue delete** and **request system configuration rescue save** commands, see the [CLI Explorer](#).

### Saving a Configuration to a File

Save the Junos OS configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the **save** configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

By default, the configuration is saved to a file in your home directory, which is on the flash drive.

When you issue this command from anywhere in the hierarchy (except the top level), a **replace** tag is automatically included at the beginning of the file. You can use the **replace** tag to control how a configuration is loaded from a file.

```
user@host> file show /var/home/user/myconf
replace:
protocols {
  bgp {
    disable;
    group int {
      type internal;
    }
  }
  isis {
    disable;
    interface all {
      level 1 disable;
    }
    interface fxp0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    ...
  }
}
```

- Related Documentation**
- [Returning to the Most Recently Committed Junos OS Configuration on page 1252](#)
  - [Loading a Configuration from a File on page 1249](#)
  - [Viewing Files and Directories on a Device Running Junos OS](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

- Related Documentation**
- [Understanding Configuration Files on page 1242](#)
  - [Loading a Previous Configuration File on page 1252](#)
  - [Reverting to the Rescue Configuration on page 189](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```
2. Commit your changes.

```
[edit]
user@switch# commit filename
```

- Related Documentation**
- [Setting or Deleting the Rescue Configuration on page 1261](#)
  - [Reverting to the Default Factory Configuration on page 188](#)

- [Configuration File Terms on page 11](#)

## Rolling Back Junos OS Configuration Changes

This topic shows how to use the **rollback** command to return to the most recently committed Junos OS configuration. The **rollback** command is useful if you make configuration changes and then decide not to keep the changes.

The following procedure shows how to configure an SNMP health monitor on a device running Junos OS and then return to the most recently committed configuration that does not include the health monitor. When configured, the SNMP health monitor provides the network management system (NMS) with predefined monitoring for file system usage, CPU usage, and memory usage on the device.

1. Enter configuration mode:

```
user@host> configure
entering configuration mode
[edit]
user@host#
```

2. Show the current configuration (if any) for SNMP:

```
[edit]
user@host# show snmp
```

No **snmp** statements appear because SNMP has not been configured on the device.

3. Configure the health monitor:

```
[edit]
user@host# set snmp health-monitor
```

4. Show the new configuration:

```
[edit]
user@host# show snmp
health-monitor;
```

The **health-monitor** statement indicates that SNMP health monitoring is configured on the device.

5. Enter the **rollback** configuration mode command to return to the most recently committed configuration:

```
[edit]
user@host# rollback
load complete
```

6. Show the configuration again to make sure your change is no longer present:

```
[edit]
user@host# show snmp
```

No **snmp** configuration statements appear. The health monitor is no longer configured.

7. Enter the **commit** command to activate the configuration to which you rolled back:

```
[edit]
```

```
user@host# commit
```

8. Exit configuration mode:

```
[edit]
user@host# exit
Exiting configuration mode
```

You can also use the **rollback** command to return to earlier configurations.

**Related  
Documentation**

- [Returning to the Most Recently Committed Junos OS Configuration on page 1252](#)

## Saving a Configuration to a File

Save the Junos OS configuration to a file so that you can edit it with a text editor of your choice. You can save your current configuration to an ASCII file, which saves the configuration in its current form, including any uncommitted changes. If more than one user is modifying the configuration, all changes made by all users are saved.

To save software configuration changes to an ASCII file, use the **save** configuration mode command:

```
[edit]
user@host# save filename
[edit]
user@host#
```

The contents of the current level of the statement hierarchy (and below) are saved, along with the statement hierarchy containing it. This allows a section of the configuration to be saved, while fully specifying the statement hierarchy.

By default, the configuration is saved to a file in your home directory, which is on the flash drive.

When you issue this command from anywhere in the hierarchy (except the top level), a **replace** tag is automatically included at the beginning of the file. You can use the **replace** tag to control how a configuration is loaded from a file.

```
user@host> file show /var/home/user/myconf
replace:
protocols {
  bgp {
    disable;
    group int {
      type internal;
    }
  }
  isis {
    disable;
    interface all {
      level 1 disable;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

```

    }
    ospf {
        traffic-engineering;
        reference-bandwidth 4g;
        ...
    }
}

```

## Setting or Deleting the Rescue Configuration

A rescue configuration is user-defined configuration that restores connectivity to the device. You set a current committed configuration to be the rescue configuration through the CLI. If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To set the current active configuration as the rescue configuration:

```
user@switch> request system configuration rescue save
```

To delete an existing rescue configuration:

```
user@switch> request system configuration rescue delete
```

### Related Documentation

- [Reverting to the Default Factory Configuration on page 188](#)
- [Loading a Previous Configuration File on page 1252](#)
- [Configuration File Terms on page 11](#)
- [CLI Explorer](#)

## Uploading a Configuration File

You can create a configuration file on your local system, copy the file to the switch, and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Create the configuration file using a text editor such as Notepad, making sure that the syntax of the configuration file is correct. For more information about testing the syntax of a configuration file see the *Junos OS System Basics and Services Command Reference* at <http://www.juniper.net/techpubs/software/junos/index.html>.
2. In the configuration text file, use an option to perform the required action when the file is loaded. [Table 72 on page 1262](#) lists and describes some options for the **load** command.

Table 72: Options for the load Command

Options	Description
<b>merge</b>	Combines the current active configuration and the configuration in the filename you specify or the one that you type at the terminal. A <b>merge</b> operation is useful when you are adding a new section to an existing configuration. If the active configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the active configuration.
<b>override</b>	Discards the current candidate configuration and loads the configuration in the filename you specify or the one that you type at the terminal. When you use the <b>override</b> option and commit the configuration, all system processes reparse the configuration. You can use the <b>override</b> option at any level of the hierarchy.
<b>replace</b>	Searches for the <b>replace</b> tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the <b>replace</b> operation adds the statements marked with the <b>replace</b> tag to the active configuration.  <b>NOTE:</b> For this operation to work, you must include <b>replace</b> tags in the text file or in the configuration you type at the terminal.

- Press Ctrl+a to select all the text in the configuration file.
- Press Ctrl+c to copy the contents of the configuration text file to the Clipboard.
- Log in to the switch using your username and password.
- To enter configuration mode:  
user@switch> **configure**  
  
You will see this output, with the hash or pound mark indicating configuration mode.  
Entering configuration mode  
[edit]  
user@switch#
- Load the configuration file:  
[edit]  
user@switch# **load merge terminal**
- At the cursor, paste the contents of the Clipboard using the mouse and the Paste icon:  
[edit]  
user@switch# **load merge terminal**  
[Type ^D at a new line to end input]  
>Cursor is here. Paste the contents of the clipboard here<
- Press Enter.
- Press Ctrl+d to set the end-of-file marker.

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt. You can also edit the configuration interactively using the CLI and commit it at a later time.



**Related Documentation**

- [Understanding Configuration Files on page 1242](#)

## Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site

You can configure a router or switch to transfer its configuration to an archive file periodically. The following tasks describe how to transfer the configuration to an archive site:

1. [Configuring the Router or Switch to Transfer Its Currently Active Configuration to an Archive on page 1263](#)
2. [Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site on page 1263](#)
3. [Configuring Transfer of the Current Active Configuration When a Configuration Is Committed on page 1264](#)
4. [Configuring Archive Sites for Transfer of Active Configuration Files on page 1264](#)

### Configuring the Router or Switch to Transfer Its Currently Active Configuration to an Archive

If you want to back up your device's current configuration to an archive site, you can configure the router or switch to transfer its currently active configuration by FTP or secure copy (SCP) periodically or after each commit.

To configure the router or switch to transfer its currently active configuration to an archive site, include statements at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
archive-sites {
  ftp://username<:password>@host-address<:port>/url-path;
  scp://username<:password>@host-address<:port>/url-path;
}
transfer-interval interval;
transfer-on-commit;
```



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([ ]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path"

### Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site

To configure the router or switch to periodically transfer its currently active configuration to an archive site, include the **transfer-interval** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
transfer-interval interval;
```

The *interval* is a period of time ranging from 15 through 2880 minutes.

## Configuring Transfer of the Current Active Configuration When a Configuration Is Committed

To configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration, include the **transfer-on-commit** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
transfer-on-commit;
```



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ( [ ]). For example,  
`"scp://username<:password>@[ipv6-host-address]<:port>/url-path"`

## Configuring Archive Sites for Transfer of Active Configuration Files

When you configure the router or switch to transfer its configuration files, you specify an archive site to which the files are transferred. If you specify more than one archive site, the router or switch attempts to transfer files to the first archive site in the list, moving to the next site only if the transfer fails.

When you use the **archive-sites** statement, you can specify a destination as an FTP URL, or SCP-style remote file specification. The URL type **file://** is also supported.

To configure the archive site, include the **archive-sites** statement at the **[edit system archival configuration]** hierarchy level:

```
[edit system archival configuration]
archive-sites {
  ftp://username@host:<port>url-path password password;
  scp://username@host:<port>url-path password password;
  file://<path>/<filename>;
}
```



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ( [ ]). For example,  
`"scp://username<:password>@[ipv6-host-address]<:port>/url-path"`

When you specify the archive site, do not add a forward slash (/) to the end of the URL.

The destination filename is saved in the following format, where *n* corresponds to the number of the compressed configuration rollback file that has been archived:

```
<router-name>_juniper.conf.n.gz_YYYYMMDD_HHMMSS
```



.....

**NOTE:** The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.



.....

## Configuration Statements



---

- [archival on page 1266](#)
- [archive-sites \(Configuration File\) on page 1267](#)
- [configuration on page 1269](#)
- [transfer-interval \(Configuration\) on page 1270](#)
- [transfer-on-commit on page 1271](#)

## archival

<b>Syntax</b>	<pre> archival {   configuration {     archive-sites {       file://&lt;path&gt;/&lt;filename&gt;;       ftp://username@host:&lt;port&gt;url-path password password;       http://username@host:&lt;port&gt;url-path password password;       pasvftp://username@host:&lt;port&gt;url-path password password;       scp://username@host:&lt;port&gt;url-path password password;     }     transfer-interval interval;     transfer-on-commit;   } } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.
<div>  <b>NOTE:</b> The <code>edit system archival</code> hierarchy is not available on QFabric systems. </div>	
<b>Options</b>	The remaining statements are explained separately.
<div>  <b>NOTE:</b> The <code>[edit system archival]</code> hierarchy is not available on QFabric systems. </div>	
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1263</li> </ul>

## archive-sites (Configuration File)

<b>Syntax</b>	<pre>archive-sites {     file://&lt;path&gt;/&lt;filename&gt;;     ftp://username@host:&lt;port&gt;url-path password password;     http://username@host:&lt;port&gt;url-path password password;     pasvftp://username@host:&lt;port&gt;url-path password password;     scp://username@host:&lt;port&gt;url-path password password; }</pre>
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks ( " ") and enclose the IPv6 host address in brackets ( [ ] ). For example,</p> <pre>"scp://username&lt;:password&gt;@[ipv6-host-address]&lt;:port&gt;/url-path"</pre> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <pre>router-name_juniper.conf.n.gz_YYYYMMDD_HHMMSS.</pre> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
<b>Options</b>	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p><b>file://</b> —transfer on a path to a named file</p> <p><b>ftp://</b> —transfer using active FTP server</p> <p><b>pasvftp://</b> —transfer to a device that only accepts passive FTP services</p>

**scp://** —transfer to a known host using background SCP file transfers

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Archive Sites for Transfer of Active Configuration Files on page 1264</a></li><li>• <a href="#">Junos OS Commit Model for Router or Switch Configuration on page 14</a></li><li>• <a href="#">configuration on page 1269</a></li><li>• <a href="#">transfer-on-commit on page 1271</a></li></ul>
------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## configuration

**Syntax**

```
configuration {
  transfer-interval interval;
  transfer-on-commit;
  archive-sites {
    file://<path>/<filename>;
    ftp://username@host:<port>url-path password password;
    http://username@host:<port>url-path password password;
    pasvftp://username@host:<port>url-path password password;
    scp://username@host:<port>url-path password password;
  }
}
```

**Hierarchy Level** [edit system archival]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the router or switch to periodically transfer its currently active configuration (or after each commit).



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.

**Options** The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1263](#)
- [archive on page 6780](#)
- [archive-sites on page 1267](#)
- [transfer-interval on page 1270](#)
- [transfer-on-commit on page 1271](#)

## transfer-interval (Configuration)

---

<b>Syntax</b>	<code>transfer-interval <i>interval</i>;</code>
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the router or switch to periodically transfer its currently active configuration to an archive site.



**NOTE:** The `edit system archival` hierarchy is not available on QFabric systems.

---

**Options** *interval*—Interval at which to transfer the current configuration to an archive site.  
**Range:** 15 through 2880 minutes



**NOTE:** The `[edit system archival]` hierarchy is not available on QFabric systems.

---

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site on page 1263](#)
- [archive on page 6780](#)
- [configuration on page 1269](#)
- [transfer-on-commit on page 1271](#)



## transfer-on-commit

<b>Syntax</b>	transfer-on-commit;
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([ ]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path".



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Transfer of the Current Active Configuration When a Configuration Is Committed on page 1264</a></li> <li>• <a href="#">archive on page 6780</a></li> <li>• <a href="#">configuration on page 1269</a></li> <li>• <a href="#">transfer-interval on page 1270</a></li> </ul>

## Default Configurations

- [QFX3500 Switch Default Configuration on page 1271](#)

### QFX3500 Switch Default Configuration

Each QFX Series product is programmed with a factory default configuration that contains the values set for each configuration parameter when a switch is shipped. The default configuration file sets values for system parameters such as **syslog** and **commit**, configures storm control and Ethernet switching on all interfaces, and enables IGMP snooping, RSTP, and LLDP protocols.

When you commit changes to the configuration, a new configuration file is created, which becomes the active configuration. You can always revert to the factory default configuration if you need to.

The following factory default configuration file is for a QFX3500 switch with 48 ports:



**NOTE:** In this example, xe-0/0/0 through xe-0/0/47 are the network interface ports.

```
protocols {
  igmp-snooping {
    vlan all;
  }
  rstp;
  lldp {
    interface all;
  }
}
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/3 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/4 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/5 {
    unit 0 {
      family ethernet-switching;
    }
  }
  xe-0/0/6 {
    unit 0 {
      family ethernet-switching;
    }
  }
```

```
}
xe-0/0/7 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/8 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/9 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/10 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/11 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/12 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/13 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/14 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/15 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/16 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/17 {
  unit 0 {
    family ethernet-switching;
  }
}
```

```
xe-0/0/18 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/19 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/20 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/21 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/22 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/23 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/24 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/25 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/26 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/27 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/28 {  
  unit 0 {  
    family ethernet-switching;  
  }  
}  
xe-0/0/29 {
```

```
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/30 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/31 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/32 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/33 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/34 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/35 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/36 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/37 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/38 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/39 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/40 {
    unit 0 {
```

```
        family ethernet-switching;
    }
}
xe-0/0/41 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/42 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/43 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/44 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/45 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/46 {
    unit 0 {
        family ethernet-switching;
    }
}
xe-0/0/47 {
    unit 0 {
        family ethernet-switching;
    }
}
}
ethernet-switching-options {
    storm-control {
        interface all;
    }
}
system {
    syslog {
        archive size 256k;
        file default-log-messages {
            structured-data;
        }
    }
    user * {
        any emergency;
    }
    file messages {
        any notice;
        authorization info;
    }
}
```

```
    }
    file interactive-commands {
        interactive-commands any;
    }
}
ports {
    console type vt100;
}
compress-configuration-files;
login {
    password {
        minimum-length 6;
        minimum-changes 1;
        change-type set transitions;
        format md5;
    }
}
commit {
    factory-settings {
        reset-chassis-lcd-menu;
    }
}
}
```

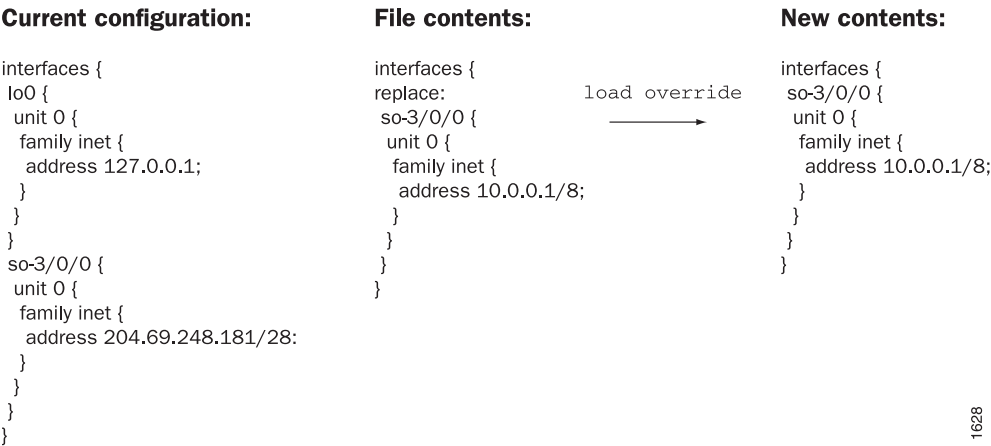
- Related Documentation
- [Reverting to the Default Factory Configuration on page 188](#)
  - [Configuring a QFX3500 Device as a Standalone Switch on page 175](#)
  - [Understanding Configuration Files on page 1242](#)
  - [Interfaces Overview on page 2389](#)

## Configuration Examples

- [Examples: Loading a Configuration from a File on page 1277](#)

### Examples: Loading a Configuration from a File

Figure 13: Overriding the Current Configuration



1628

Figure 14: Using the replace Option




Figure 15: Using the merge Option






Figure 16: Using a Patch File

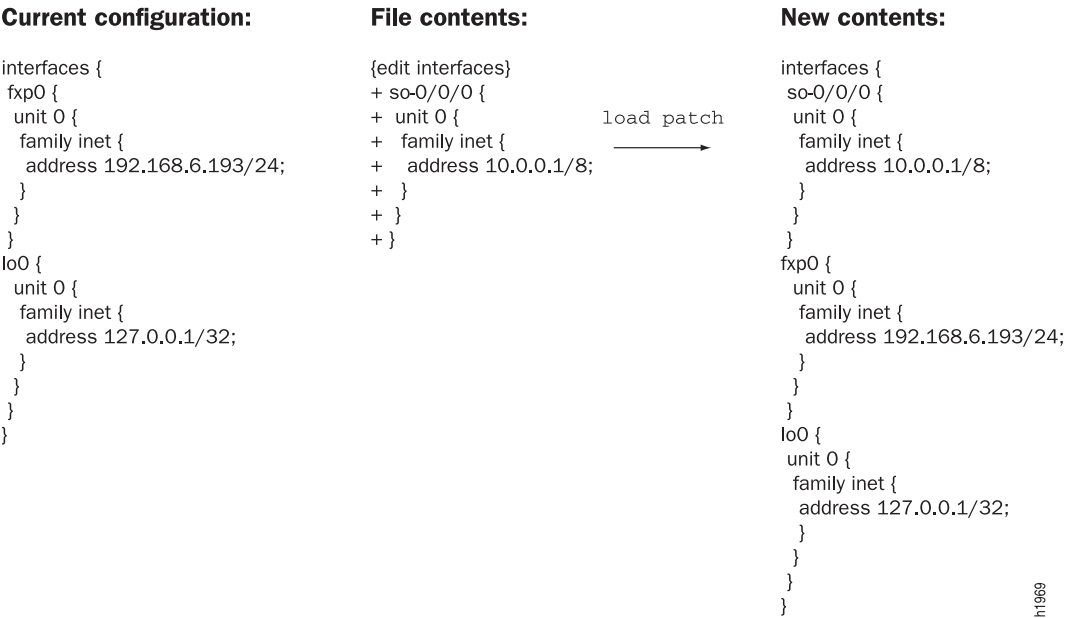
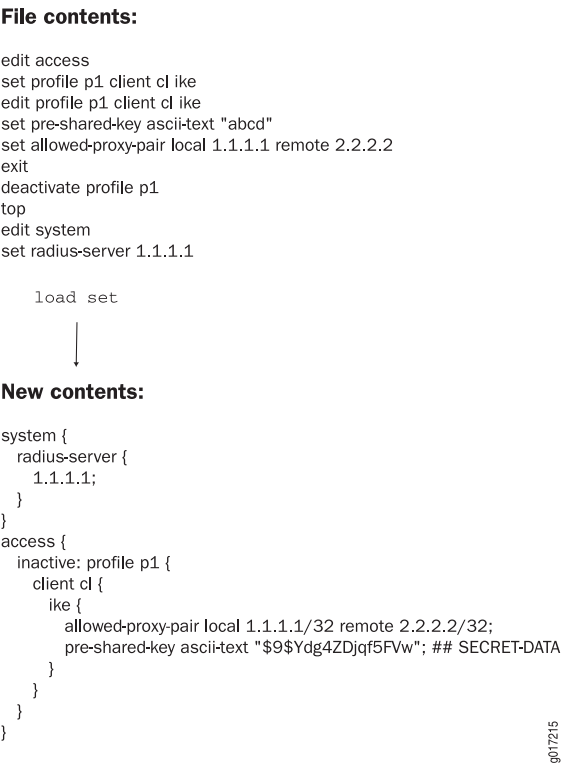


Figure 17: Using the set Option



Related Documentation

- [Loading a Configuration from a File on page 1249](#)



## CHAPTER 9

# Administration

- [Operational Commands on page 1281](#)

### Operational Commands

---

- [clear log](#)
- [clear system commit](#)
- [file archive](#)
- [file checksum md5](#)
- [file checksum sha1](#)
- [file checksum sha-256](#)
- [file compare](#)
- [file delete](#)
- [file list](#)
- [file rename](#)
- [file show](#)
- [request system configuration rescue delete](#)
- [request system configuration rescue save](#)
- [show system commit](#)
- [show system configuration archival](#)
- [show system configuration rescue](#)
- [show system rollback](#)
- [test configuration](#)

## clear log

---

<b>Syntax</b>	<code>clear log <i>filename</i></code> <code>&lt;all&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Remove contents of a log file.
<b>Options</b>	<i>filename</i> —Name of the specific log file to delete.  <code>all</code> —(Optional) Delete the specified log file and all archived versions of it.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show log on page 948</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear log on page 1282</a>
<b>Output Fields</b>	See <a href="#">file list</a> for an explanation of output fields.

## Sample Output

### clear log

The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel          26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
-----

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
-----
-rw-r-----  1 root  wheel           57 Sep 15 03:44 /var/log/sampled
total 1
```

## clear system commit

---

<b>Syntax</b>	clear system commit
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear any pending commit operation.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance (or the actual user who scheduled the commit)
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system commit on page 1001</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear system commit on page 1283</a> <a href="#">clear system commit (None Pending) on page 1283</a> <a href="#">clear system commit (User Does Not Have Required Privilege Level) on page 1283</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

#### clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

#### clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

## file archive

---

<b>Syntax</b>	<code>file archive destination <i>destination</i> source <i>source</i> &lt;compress&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.
<b>Options</b>	<p><b>destination <i>destination</i></b>—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:</p> <ul style="list-style-type: none"><li>• For archived files—The suffix <b>.tar</b></li><li>• For archived and compressed files—The suffix <b>.tgz</b></li></ul> <p><b>source <i>source</i></b>—Source of the original file or files. Specify the source as a URL or filename.</p> <p><b>compress</b>—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix <b>.tgz</b>.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 42</a></li></ul>
<b>List of Sample Output</b>	<a href="#">file archive (Multiple Files) on page 1284</a> <a href="#">file archive (Single File) on page 1284</a> <a href="#">file archive (with Compression) on page 1285</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file archive (Multiple Files)

The following sample command archives all message files in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host>
```

### file archive (Single File)

The following sample command archives one message file in the local directory `/var/log/messages` as the single file `messages-archive.tar`.

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

### file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz**.

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.
```

## file checksum md5

---

<b>Syntax</b>	<code>file checksum md5 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Calculate the Message Digest 5 (MD5) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename.  <b>filename</b> —Name of a local file for which to calculate the MD5 checksum.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum sha-256 on page 365</a></li><li>• <a href="#">file checksum sha1 on page 364</a></li><li>• <i>op</i></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum md5 on page 1286</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```



## file checksum sha1

<b>Syntax</b>	<code>file checksum sha1 &lt;pathname&gt; filename</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.
<b>Options</b>	<p><b>pathname</b>—(Optional) Path to a filename.</p> <p><b>filename</b>—Name of a local file for which to calculate the SHA-1 checksum.</p>
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li> <li>• <a href="#">file checksum md5 on page 363</a></li> <li>• <a href="#">file checksum sha-256 on page 365</a></li> <li>• <i>op</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha1 on page 1287</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax
```

```
SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

## file checksum sha-256

---

<b>Syntax</b>	<code>file checksum sha-256 &lt;pathname&gt; filename</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.
<b>Options</b>	<b>pathname</b> —(Optional) Path to a filename.  <b>filename</b> —Name of a local file for which to calculate the SHA-256 checksum.
<b>Required Privilege Level</b>	maintenance view view-configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <i>Executing an Op Script from a Remote Site</i> in the <i>Junos OS Configuration and Operations Automation Guide</i></li><li>• <a href="#">file checksum md5 on page 363</a></li><li>• <a href="#">file checksum sha1 on page 364</a></li><li>• <i>op</i></li></ul>
<b>List of Sample Output</b>	<a href="#">file checksum sha-256 on page 1288</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

## file compare

<b>Syntax</b>	<pre>file compare (files <i>filename filename</i>) &lt;context   unified&gt; &lt;ignore-white-space&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Compare two local files and describe the differences between them in default, context, or unified output styles:</p> <ul style="list-style-type: none"> <li>• <b>Default</b>—In the first line of output, <b>c</b> means lines were changed between the two files, <b>d</b> means lines were deleted between the two files, and <b>a</b> means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (&lt;) in front of output lines refers to the first file. A right angle bracket (&gt;) in front of output lines refers to the second file.</li> <li>• <b>Context</b>—The display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (!) have changed. Additions are marked with a plus sign (+), and deletions are marked with a minus sign (-).</li> <li>• <b>Unified</b>—The display is preceded by the line number from the first and the second file (xx,xxx,x). Before the line number, additions to the file are marked with a plus sign (+), and deletions to the file are marked with a minus sign (-). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.</li> </ul>
<b>Options</b>	<p><b>files <i>filename</i></b>—Names of two local files to compare.</p> <p><b>context</b>—(Optional) Display output in context format.</p> <p><b>ignore-white-space</b>—(Optional) Ignore changes in the amount of white space.</p> <p><b>unified</b>—(Optional) Display output in unified format.</p>
<b>Required Privilege Level</b>	none
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Format for Specifying Filenames and URLs in Junos OS CLI Commands on page 42</a></li> <li>• <a href="#">Viewing Core Files from Junos OS Processes on page 196</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">file compare files on page 1290</a></p> <p><a href="#">file compare files context on page 1290</a></p> <p><a href="#">file compare files unified on page 1290</a></p> <p><a href="#">file compare files unified ignore-white-space on page 1290</a></p>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<          full-name "File 1";
---
>          full-name "File 2";
102c102
<          class foo; # 'foo' is not defined
---
>          class super-user;
```

### file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one   Wed Dec  3 17:12:50 2003
--- /tmp/two   Wed Dec  3 09:13:14 2003
*****
*** 97,104 ****
        }
    }
    user bill {
!         full-name "Bill Smith";
!         class foo; # 'foo' is not defined
        authentication {
            encrypted-password SECRET;
        }
--- 97,105 ----
    }
    user bill {
!         full-name "Bill Smith";
!         uid 1089;
!         class super-user;
        authentication {
            encrypted-password SECRET;
        }
    }
```

### file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one   Wed Dec  3 17:12:50 2003
+++ /tmp/two   Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
    }
}
user bill {
-     full-name "Bill Smith";
-     class foo; # 'foo' is not defined
+     full-name "Bill Smith";
+     uid 1089;
+     class super-user;
    authentication {
        encrypted-passwordSECRET;
    }
}
```

### file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
```

```
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
     user bill {
         full-name "Bill Smith";
         uid 1089;
-        class foo; # 'foo' is not defined
+        class super-user;
         authentication {
             encrypted-password <SECRET>; # SECRET-DATA
         }
     }
```

## file delete

---

<b>Syntax</b>	<code>file delete <i>filename</i></code> <code>&lt;purge&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Delete a file on the local router or switch.
<b>Options</b>	<b><i>filename</i></b> —Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.  <b><i>purge</i></b> —(Optional) Overwrite regular files before deleting them.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file delete on page 1292</a> <a href="#">file delete (Routing Matrix) on page 1292</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

### file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

## file list

<b>Syntax</b>	file list <detail   recursive> <filename>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a list of files on the local router or switch.
<b>Options</b>	<p><b>none</b>—Display a list of all files for the current directory.</p> <p><b>detail   recursive</b>—(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.</p> <p><b>filename</b>—(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.</p>
<b>Additional Information</b>	The default directory is the home directory of the user logged in to the router or switch. To view available directories, enter a space and then a backslash (/) after the <b>file list</b> command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the <b>file list</b> command.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file list on page 1293</a> <a href="#">file list (Routing Matrix) on page 1293</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

### file list (Routing Matrix)

```
user@host> file list lcc0-re0:var/tmp
lcc0-re0:
-----
/var/tmp/:
.gdbinit
.pccardd
Test/
chassisd*
chassisd.nathan*
check_time*
```

```
cores/  
diagTestPrep*  
diagtest*  
diagtest.regress*  
do_switchovers*  
dump_test*  
err.manoj.log  
esw_clearstats*  
esw_counter*  
esw_debug*  
esw_debug_ge*  
esw_filt_test*  
esw_filter_tnp_addr*  
esw_getstats*  
esw_phy*  
esw_stats*
```



## file rename

<b>Syntax</b>	<code>file rename <i>source destination</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Rename a file on the local router or switch.
<b>Options</b>	<i>destination</i> —New name for the file.  <i>source</i> —Original name of the file. For a routing matrix, the filename must include the chassis information.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file rename on page 1295</a> <a href="#">file rename (Routing Matrix) on page 1295</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file rename

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

### file rename (Routing Matrix)

The following example lists the files in `/var/tmp`, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list lcc0-re1:/var/tmp
lcc0-re1:
-----

/var/tmp:
.pccardd
sartre.conf
snmpd
syslogd.core-tarball.0.tgz
```

```
user@host> file rename lcc0-re0:/var/tmp/snmpd /var/tmp/snmpd.rr
```

```
user@host> file list lcc0-re1:/var/tmp
```

```
lcc0-re1:
```

```
-----
```

```
/var/tmp:
```

```
.pccardd
```

```
sartre.conf
```

```
snmpd.rr
```

```
syslogd.core-tarball.0.tgz
```

## file show

<b>Syntax</b>	<code>file show <i>filename</i></code> <code>&lt;encoding (base64   raw)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of a file.
<b>Options</b>	<b><i>filename</i></b> —Name of a file. For a routing matrix, the filename must include the chassis information.  <b><code>encoding (base64   raw)</code></b> —(Optional) Encode file contents with base64 encoding or show raw text.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">file show on page 1297</a> <a href="#">file show (Routing Matrix) on page 1297</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

### file show (Routing Matrix)

```
user@host> file show lcc0-re0:/var/tmp/gdbinit
lcc0-re0:
-----
#####
# Settings
#####


set print pretty

#####
# Basic stuff
#####

define msgbuf
    printf "%s", msgbuf->msg_ptr
end
```

```
# hex dump of a block of memory
# usage: dump address length
define dump
  p $arg0, $arg1
  set $ch = $arg0
  set $j = 0
  set $n = $arg1
  while ($j < $n)
    #printf "%x %x ",&$ch[$j],$ch[$j]
    printf "%x ",$ch[$j]
    set $j = $j + 1
    if (!($j % 16))
      printf "\n"
    end
  end
end
end
```

## request system configuration rescue delete

<b>Syntax</b>	request system configuration rescue delete
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Delete an existing rescue configuration.
<div>  <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems.         </div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request system configuration rescue save on page 399</a></li> <li>• <a href="#">request system software rollback on page 459</a></li> <li>• <a href="#">show system commit on page 1001</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request system configuration rescue delete on page 1299</a>
<b>Output Fields</b>	This command produces no output.


### Sample Output

#### request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

## request system configuration rescue save

---


<b>Syntax</b>	request system configuration rescue save
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the <b>rollback</b> command.
<div> <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems.</div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">request system software delete on page 430</a></li><li>• <a href="#">request system software rollback on page 459</a></li><li>• <a href="#">show system commit on page 1001</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request system configuration rescue save on page 1300</a>
<b>Output Fields</b>	This command produces no output.

### Sample Output

#### request system configuration rescue save

```
user@host> request system configuration rescue save
```

## show system commit

<b>Syntax</b>	<pre>show system commit &lt;revision&gt; &lt;server&gt;</pre>	
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>server</b> introduced in Junos OS Release 12.1 for the PTX Series router.</p> <p>Option <b>revision</b> introduced in Junos OS Release 14.1.</p>	
<b>Description</b>	Display the system commit history and any pending commit operation.	
<b>Options</b>	<p><b>none</b>—Display the last 50 commit operations listed, most recent to first.</p> <p><b>revision</b>—(Optional) Display the revision number of the active configuration of the Routing Engine(s).</p> <p><b>server</b>—(Optional) Display commit server status.</p>	
	<div>  <p><b>NOTE:</b> By default, the status of the commit server is “Not running”. The commit server starts running only when a commit job is added to the batch.</p> </div>	
<b>Required Privilege Level</b>	view	
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear system commit on page 354</a></li> <li><a href="#">show system commit revision</a></li> </ul>	
<b>List of Sample Output</b>	<a href="#">show system commit on page 1303</a> <a href="#">show system commit (At a Particular Time) on page 1303</a> <a href="#">show system commit (At the Next Reboot) on page 1303</a> <a href="#">show system commit (Rollback Pending) on page 1303</a> <a href="#">show system commit (QFX Series) on page 1303</a>	
<b>Output Fields</b>	<p><a href="#">Table 56 on page 1001</a> describes the output fields for the <b>show system commit</b> command. Output fields are listed in the approximate order in which they appear.</p>	

**Table 73: show system commit Output Fields**

Field Name	Field Description	Level of Output
<b>&lt;number&gt;</b>	Displays the last 50 commit operations listed, most recent to first. The identifier <b>&lt;number&gt;</b> designates a configuration created for recovery using the <b>request system configuration rescue save</b> command.	<b>none</b>

Table 73: show system commit Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>&lt;time-stamp&gt;</b>	Date and time of the commit operation.	<b>none</b>
<b>&lt;root&gt;/&lt;username&gt;</b>	User who executed the commit operation.	<b>none</b>
<b>&lt;method&gt;</b>	<p>Method used to execute the commit operation:</p> <ul style="list-style-type: none"> <li>• <b>CLI</b>—CLI interactive user performed the commit operation.</li> <li>• <b>Junos XML protocol</b>—Junos XML protocol client performed the commit operation.</li> <li>• <b>synchronize</b>—The <b>commit synchronize</b> command was performed on the other Routing Engine.</li> <li>• <b>snmp</b>—An SNMP <b>set</b> request caused the commit operation.</li> <li>• <b>button</b>—A button on the router or switch was pressed to commit a rescue configuration for recovery.</li> <li>• <b>autoinstall</b>—A configuration obtained through autoinstallation was committed.</li> <li>• <b>other</b>—When there is no login name associated with the session, the values for user and client default to root and other. For example, during a reboot after package installation, mgd commits the configuration as a system commit, and there is no login associated with the commit.</li> </ul>	<b>none</b>



## Sample Output

### show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue 2002-05-10 15:32:03 PDT by root via other
```

### show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

### show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

### show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

### show system commit (QFX Series)

```
user@switch> show system commit
0 2011-11-25 19:17:49 PST by root via cli
```

## show system configuration archival

---

**Syntax**    show system configuration archival

**Release Information**    Introduced in Junos OS Release 7.6.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display directory and number of files queued for archival transfer.



**NOTE:** The [edit system configuration] hierarchy is not available on QFabric systems.

---

**Options**    This command has no options.

**Required Privilege Level**    maintenance

**List of Sample Output**    [show system configuration archival on page 1304](#)


### Sample Output

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

## show system configuration rescue

<b>Syntax</b>	show system configuration rescue
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display a rescue configuration, if one exists.
<div>  <b>NOTE:</b> The [edit system configuration] hierarchy is not available on QFabric systems. </div>	
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show system configuration archival on page 1004</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show system configuration rescue on page 1305</a>

## Sample Output

### show system configuration rescue


```

user@switch> show system configuration rescue
version "7.3"; groups {
  global {
    system {
      host-name router1;
      domain-name customer.net;
      domain-search [ customer.net ];
      backup-router 192.168.124.254;
      name-server {
        172.17.28.11;
        172.17.28.101;
        172.17.28.100;
        172.17.28.10;
      }
      login {
        user regress {
          uid 928;
          class ;
          shell csh;
          authentication {
            encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
SECRET-DATA
          }
        }
      }
    }
  }
  services {

```

```
        ftp;  
        rlogin;  
        rsh;  
        telnet;  
    }  
}  
.....
```

## show system rollback

<b>Syntax</b>	<code>show system rollback <i>number</i></code> <code>&lt;compare <i>number</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of a previously committed configuration, or the differences between two previously committed configurations.
<div>  <b>NOTE:</b> The <code>show system rollback</code> command is a purely operational mode command and cannot be issued with <code>run</code> from the configuration mode. </div>	
<b>Options</b>	<p><b><i>number</i></b>—Number of a configuration to view. The output displays the configuration. The range of values is 0 through 49.</p> <p><b><code>compare <i>number</i></code></b>—(Optional) Number of another previously committed (rollback) configuration to compare to rollback <b><i>number</i></b>. The output displays the differences between the two configurations. The range of values is 0 through 49.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show system rollback compare on page 1307</a>

## Sample Output

### show system rollback compare

```

user@host> show system rollback 3 compare 1
[edit]
+ interfaces {
+   ge-1/1/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 14.1.1.1/30;
+       }
+     }
+   }
+   ge-1/2/1 {
+     unit 0 {
+       family inet {
+         filter {
+           input mf_plp;
+         }
+         address 13.1.1.1/30;
+       }
+     }
+   }
+ }

```

```
+      }
+    }
+    ge-1/3/0 {
+      unit 0 {
+        family inet {
+          filter {
+            input mf_plp;
+          }
+          address 12.1.1.1/30;
+        }
+      }
+    }
+  }
+}
```

## test configuration

<b>Syntax</b>	<code>test configuration <i>filename</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Verify that the syntax of a configuration file is correct. If the configuration contains any syntax or commit check errors, a message is displayed to indicate the line number and column number in which the error was found. This command only accepts text files.
<b>Options</b>	<b><i>filename</i></b> —Name of the configuration file.  <b>syntax-only</b> —Check the syntax of a partial configuration file, without checking for commit errors. This option introduced in Junos OS Release 12.1.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">test configuration on page 1309</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### test configuration

```

user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
[edit system]
    'paris-23;'
      syntax error
terminal:4:(11) statement must contain additional statements: ;
[edit system login]
    'login ;'
      statement must contain additional statements
configuration syntax failed

```





## CHAPTER 10

# Troubleshooting

- [Troubleshooting Procedures on page 1311](#)

## Troubleshooting Procedures

---

- [Loading a Previous Configuration File on page 1311](#)
- [Reverting to the Default Factory Configuration on page 1312](#)
- [Reverting to the Rescue Configuration on page 1312](#)

### Loading a Previous Configuration File

You can use the **rollback** *<number>* command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

#### Syntax

**rollback** *<number>*

#### Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.
  - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
  - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
```

```
user@switch# commit
```

- Related Documentation**
- [Configuration File Terms on page 11](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

- Related Documentation**
- [Understanding Configuration Files on page 1242](#)
  - [Loading a Previous Configuration File on page 1252](#)
  - [Reverting to the Rescue Configuration on page 189](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```
2. Commit your changes.

```
[edit]
user@switch# commit filename
```

- Related Documentation**
- [Setting or Deleting the Rescue Configuration on page 1261](#)
  - [Reverting to the Default Factory Configuration on page 188](#)
  - [Configuration File Terms on page 11](#)

## PART 4

# User and Access Management

- [Overview on page 1315](#)
- [Configuration on page 1343](#)
- [Administration on page 1475](#)



## CHAPTER 11

# Overview

- [Software Overview on page 1315](#)
- [Access Control Overview on page 1317](#)

## Software Overview

---

- [Understanding Software Infrastructure and Processes on page 1315](#)

## Understanding Software Infrastructure and Processes

Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the switch.

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

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

- [Routing Engine and Packet Forwarding Engine on page 1315](#)
- [Junos OS Processes on page 1316](#)

## Routing Engine and Packet Forwarding Engine

---

A switch has two primary software processing components:

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

- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

### Junos OS Processes

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

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

Table 7 on page 30 describes the primary Junos OS processes.

**Table 74: Junos OS Processes**

Process	Name	Description
Chassis process	chassisd	<p>Detects hardware on the system that is used to configure network interfaces.</p> <p>Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered.</p> <p>Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.</p>
DNS Server process	named-service	Resolves hostnames into addresses.
Dynamic Host Configuration Protocol (DHCP) process	dhcp-service	Enables a DHCP server to allocate network IP addresses and deliver configuration settings to client hosts without user intervention.
Ethernet switching process	eswd	<p>Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree Protocol, and access port security.</p> <p>Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.</p>
Firewall management process	firewall	Manages the firewall configuration and helps accept or reject packets that are transiting an interface on a switch.
Forwarding process	pfem	Defines how routing protocols operate on the partition. The overall performance of the partition is largely determined by the effectiveness of the forwarding process.
Interface process	dcd	Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.
Integrated Local Management Interface (ILMI) process	ilmi	Provides bidirectional exchange of management information between two ATM interfaces across a physical connection.
Link Management Protocol (LMP) process	link-management	Establishes and maintains LMP control channels.

Table 74: Junos OS Processes (*continued*)

Process	Name	Description
Management process	mgd	<p>Provides communication between the other processes and an interface to the configuration database.</p> <p>Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured.</p> <p>Interacts with the other processes when commands are issued through one of the user interfaces on the partition.</p> <p>If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.</p>
Multicast snooping process	<del>multicast</del> snooping	Makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.
Secure Neighbor Discovery (SEND) Protocol process	send	Protects Neighbor Discovery Protocol (NDP) messages.
Simple Network Management Protocol (SNMP) process	snmp	Enables the monitoring of network devices from a central location and provides the switch's SNMP master agent.
Tunnel OAM process	tunnel-oamd	Enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.
Virtual Router Redundancy Protocol (VRRP) process	vrrp	Enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**Related Documentation**

- [Junos OS Baseline Network Operations Guide](#)
- [Junos OS Administration Library for Routing Devices](#)

## Access Control Overview

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)
- [Understanding Login Authentication on page 1318](#)
- [Understanding LLDP on page 1319](#)
- [Understanding RADIUS Accounting on page 1320](#)
- [Understanding VSAs on page 1321](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1321](#)
- [Juniper Networks Vendor-Specific TACACS+ Attributes on page 1324](#)

- [Understanding Junos OS Access Privilege Levels on page 1325](#)
- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1330](#)
- [Junos OS User Authentication Methods on page 1334](#)
- [Junos OS User Accounts Overview on page 1335](#)
- [Junos OS Login Classes Overview on page 1337](#)
- [Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies on page 1338](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1339](#)
- [Special Requirements for Junos OS Plain-Text Passwords on page 1339](#)

## Overview of Template Accounts for RADIUS and TACACS+ Authentication

When you use local password authentication, you must create a local user account for every user who wants to access the system. However, when you are using RADIUS or TACACS+ authentication, you can create single accounts (for authorization purposes) that are shared by a set of users. You create these accounts using the remote and local user template accounts. When a user is using a template account, the command-line interface (CLI) username is the login name; however, the privileges, file ownership, and effective user ID are inherited from the template account.

### Related Documentation

- [Understanding Remote Authentication Servers](#)
- [Configuring Remote Template Accounts for User Authentication on page 1354](#)
- [Configuring Local User Template Accounts for User Authentication on page 1347](#)

## Understanding Login Authentication

You can control access to your network using several different authentication methods—media access control (MAC) RADIUS, for example. Authentication prevents unauthorized devices and users from gaining access to your LAN. For MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a DHCP server.

You can enable end devices to access the network without authenticating on the RADIUS server by configuring the MAC address of the end device in the static MAC bypass list by configuring the MAC address using the **authentication-whitelist** statement.

You can configure one or more authentication methods on a single interface and thereby enable fallback to the next method if the first or second method is unsuccessful.

On a single interface you can configure one or a combination of several authentication methods.

This topic covers:

- [MAC RADIUS Authentication on page 1319](#)



## MAC RADIUS Authentication

You can configure MAC RADIUS authentication on interfaces that are connected to end devices.

The EAP method supported for MAC RADIUS authentication is EAP-MD5.

When you configure the **mac-radius restrict** option, the switch immediately attempts a MAC- RADIUS authentication by sending a request to the RADIUS server for authentication of the MAC address of the end device. If MAC address of the end device is configured for RADIUS authentication, LAN access between the two switches is created.

### Related Documentation

- [Configuring RADIUS Authentication \(QFX Series\) on page 1351](#)

## Understanding LLDP

The device uses Link Layer Discovery Protocol (LLDP) to learn and distribute device information on network links. The information enables the switch to identify a variety of devices quickly. This quick identification results in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include specifics, such as chassis and port identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in Junos OS.

The device supports the following basic TLVs:

- **Chassis Identifier**—The MAC address associated with the local system.
- **Port Identifier**—The port identification for the specified port in the local system.
- **Port Description**—The user-configured port description. The port description can be a maximum of 256 characters.
- **System Name**—The user-configured name of the local system. The system name can be a maximum of 256 characters.
- **System Description**—The system description containing information about the software and current image running on the system. This information cannot be configured, but is taken from the software.
- **System Capabilities**—The primary function performed by the system. The capabilities that system supports are defined; for example, bridge or router. This information cannot be configured, but is based on the model of the product.
- **Management Address**—The IP management address of the local system.

The device supports the following 802.3 TLVs:

- **Power via MDI**—A TLV that advertises media dependent interface (MDI) power support, power source equipment (PSE) power pair, and power class information.
- **MAC/PHY Configuration Status**—A TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information cannot be configured, but is based on the physical interface structure.
- **Link Aggregation**—A TLV that advertises whether the port is aggregated and its aggregated port ID.
- **Maximum Frame Size**—A TLV that advertises the Maximum Transmission Unit (MTU) of the interface sending LLDP frames.
- **Port Vlan**—A TLV that advertises the VLAN name configured on the interface.

**Related  
Documentation**

- [Configuring LLDP on page 1345](#)

## Understanding RADIUS Accounting

Devices support IETF RFC 2866, *RADIUS Accounting*. Configuring RADIUS accounting on the device supports collecting statistical data about users logging in to or out from a LAN and sending the data to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, analyzing and tracking usage patterns, or billing a user based upon the amount of time or type of services accessed.

To configure RADIUS accounting, specify one or more RADIUS accounting servers to receive the statistical data from the device, and select the type of accounting data to be collected.

The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify a list of RADIUS accounting servers. If the primary server (the first one configured) is unavailable, each RADIUS server in the list is tried in the order in which they are configured in the Junos OS.

The RADIUS accounting process between the device and a RADIUS server works like this:

1. A RADIUS accounting server listens for User Datagram Protocol (UDP) packets on a specific port. For example, on FreeRADIUS, the default port is 1813.
2. The device forwards an *accounting-request* packet containing an event record to the accounting server. The event record associated with this supplicant contains an *Acct-Status-Type* attribute whose value indicates the beginning of user service for this supplicant. When the supplicant's session ends, the accounting request contains an *Acct-Status-Type* attribute value indicating the end of user service. The RADIUS accounting server records this as a stop-accounting record containing session information and the length of the session.
3. The RADIUS accounting server logs these events in a file as start-accounting or stop-accounting records. On FreeRADIUS, the filename is the server's address; for example, 122.69.1.250.

4. The accounting server sends an *accounting-response* packet back to the device confirming it has received the accounting request.
5. If the device does not receive a response from the server, it continues to send accounting requests until an accounting response is returned from the accounting server.

The statistics collected through this process can be displayed from the RADIUS server; to see those statistics, the user accesses the log file configured to receive them.

**Related Documentation**

- [Configuring RADIUS System Accounting on page 1349](#)

## Understanding VSAs

Devices support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs) and are described in RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*.

Through VSAs, you can configure port-filtering attributes on the RADIUS server. VSAs are cleartext fields sent from the RADIUS server to the device as a result of authentication success or failure. Authentication prevents unauthorized user access by blocking a supplicant at the port until the device is authenticated by the RADIUS server. The VSA attributes are interpreted by the device during authentication, and the device takes appropriate actions. Implementing port-filtering attributes with authentication on the RADIUS server provides a central location for controlling LAN access for supplicants.

These port-filtering attributes specific to Juniper Networks are encapsulated in a RADIUS server VSA with the vendor ID set to the Juniper Networks ID number, 2636.

As well as configuring port-filtering attributes through VSAs, you can apply a port firewall filter that has already been configured on the device directly to the RADIUS server. Like port-filtering attributes, the filter is applied during the authentication process, and its actions are applied at the device port. Adding a port firewall filter to a RADIUS server eliminates the need to add the filter to multiple ports and devices.

**Related Documentation**

- [Configuring Firewall Filters on page 5290](#)
- [Configuring RADIUS Authentication \(QFX Series\) on page 1351](#)
- [VSA Match Conditions and Actions on page 1376](#)

## Juniper Networks Vendor-Specific RADIUS Attributes

Junos OS supports the configuration of Juniper Networks RADIUS vendor-specific attributes (VSAs). These VSAs are encapsulated in a RADIUS vendor-specific attribute with the vendor ID set to the Juniper Networks ID number, 2636. [Table 75 on page 1322](#) lists the Juniper Networks VSAs you can configure.

Table 75: Juniper Networks Vendor-Specific RADIUS Attributes

Name	Description	Type	Length	String
Juniper-Local-User-Name	Indicates the name of the user template used by this user when logging in to a device. This attribute is used only in Access-Accept packets.	1	≥3	One or more octets containing printable ASCII characters.
Juniper-Allow-Commands	Contains an extended regular expression that enables the user to run operational mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	2	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">“Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands”</a> on page 1339.
Juniper-Deny-Commands	Contains an extended regular expression that denies the user permission to run operation mode commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	3	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">“Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands”</a> on page 1339.
Juniper-Allow-Configuration	Contains an extended regular expression that enables the user to run configuration mode commands in addition to the commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	4	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">“Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies”</a> on page 1338.
Juniper-Deny-Configuration	Contains an extended regular expression that denies the user permission to run configuration commands authorized by the user's login class permission bits. This attribute is used only in Access-Accept packets.	5	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">“Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies”</a> on page 1338.
Juniper-Interactive-Command	Indicates the interactive command entered by the user. This attribute is used only in Accounting-Request packets.	8	≥3	One or more octets containing printable ASCII characters.

Table 75: Juniper Networks Vendor-Specific RADIUS Attributes (*continued*)

Name	Description	Type	Length	String
Juniper-Configuration-Change	Indicates the interactive command that results in a configuration (database) change. This attribute is used only in Accounting-Request packets.	9	≥3	One or more octets containing printable ASCII characters.
Juniper-User-Permissions	<p>Contains information the server uses to specify user permissions. This attribute is used only in Access-Accept packets.</p> <p><b>NOTE:</b> When the <b>Juniper-User-Permissions</b> attribute is configured to grant the Junos OS <b>maintenance</b> or <b>all</b> permissions on a RADIUS server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the <b>su root</b> command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions <b>maintenance</b> or <b>all</b>, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	10	≥3	<p>One or more octets containing printable ASCII characters.</p> <p>The string is a list of permission flags separated by a space. The exact name of each flag must be specified in its entirety. See <a href="#">Table 77 on page 1326</a>.</p>
Juniper-Authentication-Type	Indicates the authentication method (local database, or RADIUS server) used to authenticate a user. If the user is authenticated using a local database, the attribute value shows 'local'. If the user is authenticated using RADIUS server, the attribute value shows 'remote'.	11	≥5	One or more octets containing printable ASCII characters.
Juniper-Session-Port	Indicates the source port number of the established session.	12	size of integer	Integer

For more information about the VSAs, see RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*.

- Related Documentation**
- [Configuring RADIUS Authentication](#)
  - [Configuring RADIUS Authentication \(QFX Series\) on page 1351](#)

## Juniper Networks Vendor-Specific TACACS+ Attributes

Junos OS supports the configuration of Juniper Networks TACACS+ vendor-specific attributes (VSAs). These VSAs are encapsulated in a TACACS+ vendor-specific attribute with the vendor ID set to the Juniper Networks ID number, 2636. [Table 76 on page 1324](#) lists the Juniper Networks VSAs you can configure.

**Table 76: Juniper Networks Vendor-Specific TACACS+ Attributes**

Name	Description	Length	String
<b>local-user-name</b>	Indicates the name of the user template used by this user when logging in to a device.	≥3	One or more octets containing printable ASCII characters.
<b>allow-commands</b>	Contains an extended regular expression that enables the user to run operational mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">Table 81 on page 1339</a> .
<b>allow-configuration</b>	Contains an extended regular expression that enables the user to run configuration mode commands in addition to those commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">"Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies" on page 1338</a> .
<b>deny-commands</b>	Contains an extended regular expression that denies the user permission to run operational mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">Table 81 on page 1339</a> .
<b>deny-configuration</b>	Contains an extended regular expression that denies the user permission to run configuration mode commands authorized by the user's login class permission bits.	≥3	One or more octets containing printable ASCII characters, in the form of an extended regular expression. See <a href="#">Table 80 on page 1338</a> .

Table 76: Juniper Networks Vendor-Specific TACACS+ Attributes (*continued*)

Name	Description	Length	String
<b>user-permissions</b>	<p>Contains information the server uses to specify user permissions.</p> <p><b>NOTE:</b> When the <b>user-permissions</b> attribute is configured to grant the Junos OS <b>maintenance</b> or <b>all</b> permissions on an IPv4 or IPv6 TACACS+ server, the UNIX wheel group membership is not automatically added to a user's list of group memberships. Some operations such as running the <b>su root</b> command from a local shell require wheel group membership permissions. However, when a user is configured locally with the permissions <b>maintenance</b> or <b>all</b>, the user is automatically granted membership to the UNIX wheel group. Therefore, we recommend that you create a template user account with the required permissions and associate individual user accounts with the template user account.</p>	≥3	One or more octets containing printable ASCII characters. See <a href="#">Table 77 on page 1326</a> .
<b>authentication-type</b>	Indicates the authentication method (local database, or TACACS+ server) used to authenticate a user. If the user is authenticated using a local database, the attribute value shows 'local'. If the user is authenticated using TACACS+ server, the attribute value shows 'remote'.	≥5	One or more octets containing printable ASCII characters.
<b>session-port</b>	Indicates the source port number of the established session.	size of integer	Integer

- Related Documentation**
- [Configuring TACACS+ Authentication](#)
  - [Configuring TACACS+ Authentication \(QFX Series\) on page 1364](#)

## Understanding Junos OS Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with them. Users can execute only those commands and configure and view only those statements for which they have access privileges. The access privileges for each login class are defined by one or more *permission flags*.

For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

The following sections provide additional information about permissions:

- [Junos OS Login Class Permission Flags on page 1326](#)
- [Allowing or Denying Individual Commands for Junos OS Login Classes on page 1329](#)

## Junos OS Login Class Permission Flags

The **permissions** statement specifies one or more of the permission flags listed in [Table 77 on page 1326](#). Permission flags are not cumulative, so for each class you must list all the permission flags needed, including **view** to display information and **configure** to enter configuration mode. Two forms of permissions control for individual parts of the configuration are:

- "Plain" form—Provides read-only capability for that permission type. An example is **interface**.
- Form that ends in **-control**—Provides read and write capability for that permission type. An example is **interface-control**.

[Table 77 on page 1326](#) lists the Junos® operating system (Junos OS) login class permission flags that you can configure by including the **permissions** statement at the **[edit system login class *class-name*]** hierarchy level.

**Table 77: Login Class Permission Flags**

Permission Flag	Description
<b>access</b>	Can view the access configuration in configuration mode and with the <b>show configuration</b> operational mode command.
<b>access-control</b>	Can view and configure access information at the <b>[edit access]</b> hierarchy level.
<b>admin</b>	Can view user account information in configuration mode and with the <b>show configuration</b> operational mode command.
<b>admin-control</b>	Can view user accounts and configure them at the <b>[edit system login]</b> hierarchy level.
<b>all-control</b>	Can access all operational mode commands and configuration mode commands. Can modify configuration in all the configuration hierarchy levels.
<b>clear</b>	Can clear (delete) information learned from the network that is stored in various network databases by using the <b>clear</b> commands.
<b>configure</b>	Can enter configuration mode by using the <b>configure</b> command.
<b>control</b>	Can perform all control-level operations—all operations configured with the <b>-control</b> permission flags.
<b>field</b>	Can view field debug commands. Reserved for debugging support.
<b>firewall</b>	Can view the firewall filter configuration in configuration mode.



Table 77: Login Class Permission Flags (*continued*)

Permission Flag	Description
<b>firewall-control</b>	Can view and configure firewall filter information at the <b>[edit firewall]</b> hierarchy level.
<b>floppy</b>	Can read from and write to the removable media.
<b>flow-tap</b>	Can view the flow-tap configuration in configuration mode.
<b>flow-tap-control</b>	Can view the flow-tap configuration in configuration mode and can configure flow-tap configuration information at the <b>[edit services flow-tap]</b> hierarchy level.
<b>flow-tap-operation</b>	<p>Can make flow-tap requests to the router or switch. For example, a Dynamic Tasking Control Protocol (DTCP) client must have <b>flow-tap-operation</b> permission to authenticate itself to the Junos OS as an administrative user.</p> <p><b>NOTE:</b> The <b>flow-tap-operation</b> option is not included in the <b>all-control</b> permissions flag.</p>
<b>idp-profiler-operation</b>	Can view profiler data.
<b>interface</b>	Can view the interface configuration in configuration mode and with the <b>show configuration</b> operational mode command.
<b>interface-control</b>	<p>Can view chassis, class of service (CoS), groups, forwarding options, and interfaces configuration information. Can edit configuration at the following hierarchy levels:</p> <ul style="list-style-type: none"> <li>• <b>[edit chassis]</b></li> <li>• <b>[edit class-of-service]</b></li> <li>• <b>[edit groups]</b></li> <li>• <b>[edit forwarding-options]</b></li> <li>• <b>[edit interfaces]</b></li> </ul>
<b>maintenance</b>	Can perform system maintenance, including starting a local shell on the router and becoming the superuser in the shell by using the <b>su root</b> command, and can halt and reboot the router by using the <b>request system</b> commands.
<b>network</b>	Can access the network by using the <b>ping</b> , <b>ssh</b> , <b>telnet</b> , and <b>traceroute</b> commands.
<b>pgcp-session-mirroring</b>	Can view the <b>pgcp</b> session mirroring configuration.
<b>pgcp-session-mirroring-control</b>	Can modify the <b>pgcp</b> session mirroring configuration.
<b>reset</b>	Can restart software processes by using the <b>restart</b> command and can configure whether software processes are enabled or disabled at the <b>[edit system processes]</b> hierarchy level.

Table 77: Login Class Permission Flags (*continued*)

Permission Flag	Description
<b>rollback</b>	Can use the <b>rollback</b> command to return to a previously committed configuration other than the most recently committed one.
<b>routing</b>	Can view general routing, routing protocol, and routing policy configuration information in configuration and operational modes.
<b>routing-control</b>	Can view general routing, routing protocol, and routing policy configuration information and can configure general routing at the <b>[edit routing-options]</b> hierarchy level, routing protocols at the <b>[edit protocols]</b> hierarchy level, and routing policy at the <b>[edit policy-options]</b> hierarchy level.
<b>secret</b>	Can view passwords and other authentication keys in the configuration.
<b>secret-control</b>	Can view passwords and other authentication keys in the configuration and can modify them in configuration mode.
<b>security</b>	Can view security configuration in configuration mode and with the <b>show configuration</b> operational mode command.
<b>security-control</b>	Can view and configure security information at the <b>[edit security]</b> hierarchy level.
<b>shell</b>	Can start a local shell on the router or switch by using the <b>start shell</b> command.
<b>snmp</b>	Can view Simple Network Management Protocol (SNMP) configuration information in configuration and operational modes.
<b>snmp-control</b>	Can view SNMP configuration information and can modify SNMP configuration at the <b>[edit snmp]</b> hierarchy level.
<b>system</b>	Can view system-level information in configuration and operational modes.
<b>system-control</b>	Can view system-level configuration information and configure it at the <b>[edit system]</b> hierarchy level.
<b>trace</b>	Can view trace file settings and configure trace file properties.
<b>trace-control</b>	Can modify trace file settings and configure trace file properties.
<b>view</b>	Can use various commands to display current system-wide, routing table, and protocol-specific values and statistics. Cannot view the secret configuration.

Table 77: Login Class Permission Flags (*continued*)

Permission Flag	Description
<b>view-configuration</b>	Can view all of the configuration excluding secrets, system scripts, and event options.  <b>NOTE:</b> Only users with the <b>maintenance</b> permission can view commit script, op script, or event script configuration.

### Allowing or Denying Individual Commands for Junos OS Login Classes

By default, all top-level CLI commands have associated access privilege levels. Users can execute only those commands and view only those statements for which they have access privileges. For each login class, you can explicitly deny or allow the use of operational and configuration mode commands that would otherwise be permitted or not allowed by a privilege level specified in the **permissions** statement.

Permission flags are used to grant a user access to operational mode commands and configuration hierarchy levels and statements. By specifying a specific permission flag on the user's login class at the **[edit system login class]** hierarchy level, you grant the user access to the corresponding commands and configuration hierarchy levels and statements. To grant access to all commands and configuration statements, use the **all** permissions flag. For permission flags that grant access to configuration hierarchy levels and statements, the flags grant read-only privilege to that configuration. For example, the **interface** permissions flag grants read-only access to the **[edit interfaces]** hierarchy level. The **-control** form of the flag grants read-write access to that configuration. Using the preceding example, **interface-control** grants read-write access to the **[edit interfaces]** hierarchy level.

- The **all** login class permission bits take precedence over extended regular expressions when a user issues **rollback** command with **rollback** permission flag enabled.
- Expressions used to allow and deny commands for users on RADIUS and TACACS+ servers have been simplified. Instead of a single, long expression with multiple commands (**allow-commands=cmd1 cmd2 ... cmdn**), you can specify each command as a separate expression. This new syntax is valid for **allow-configuration**, **deny-configuration**, **allow-commands**, **deny-commands**, and all user permission bits.
- Users cannot issue the **load override** command when specifying an extended regular expression. Users can only issue the **merge**, **replace**, and **patch** configuration commands.
- If you allow and deny the same commands, the **allow-commands** permissions take precedence over the permissions specified by the **deny-commands**. For example, if you include **allow-commands "request system software add"** and **deny-commands "request system software add"**, the login class user is allowed to install software using the **request system software add** command.
- Regular expressions for **allow-commands** and **deny-commands** can also include the **commit**, **load**, **rollback**, **save**, **status**, and **update** commands.
- If you specify a regular expression for **allow-commands** and **deny-commands** with two different variants of a command, the longest match is always executed.

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

- Related Documentation**
- [Configuring Access Privilege Levels on page 1344](#)
  - [Access Privilege User Permission Flags Overview](#)

## Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

If the **authentication-order** is remote-server then local, Junos OS will retry the local server if the remote-server is unreachable or has timed out. However, if the remote-server rejects the authentication, Junos OS will not retry the authentication.

If none of the configured authentication methods accept the login credentials and if a reject response is received, the login attempt fails. If no response is received from any configured authentication method, the Junos OS consults local password authentication as a last resort.

### Using RADIUS or TACACS+ Authentication

---

You can configure the Junos OS to be both a RADIUS and TACACS+ authentication client.

If an authentication method included in the **[authentication-order]** statement is not available, or if the authentication is available but returns a reject response, the Junos OS tries the next authentication method included in the **authentication-order** statement.

The RADIUS or TACACS+ server authentication might fail because of the following reasons:

- The authentication method is configured, but the corresponding authentication servers are not configured. For instance, the RADIUS and TACACS+ authentication methods are included in the **authentication-order** statement, but the corresponding RADIUS or TACACS+ servers are not configured at the respective **[edit system radius-server]** and **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server does not respond within the timeout period configured at the **[edit system radius-server]** or **[edit system tacplus-server]** hierarchy levels.
- The RADIUS or TACACS+ server is not reachable because of a network problem.

The RADIUS or TACACS+ server authentication might return a reject response because of the following reasons:

- The user profiles of users accessing a router or switch might not be configured on the RADIUS or TACACS+ server.
- The user enters incorrect logon credentials.

### Using Local Password Authentication

You can explicitly configure the password authentication method or use this method as a fallback mechanism when remote authentication servers fail. The password authentication method consults the local user profiles configured at the **[edit system login]** hierarchy level. Users can log in to a router or switch using their local username and password in the following scenarios:

- The password authentication method (password) is explicitly configured as one of the authentication methods in the **[authentication-order authentication-methods]** statement. In this case, the password authentication method is tried if no previous authentication accepts the logon credentials. This is true whether the previous authentication method fails to respond or returns a reject response because of an incorrect username or password.
- The password authentication method is not explicitly configured as one of the authentication methods in the **authentication-order authentication-methods** statement. In this case, the password authentication method is tried only if all configured authentication methods fail to respond. It is not consulted if any configured authentication method returns a reject response because of an incorrect username or password.

### Order of Authentication Attempts

Table 78 on page 1331 describes how the **authentication-order** statement at the **[edit system]** hierarchy level determines the procedure that the Junos OS uses to authenticate users for access to a router or switch.

**Table 78: Order of Authentication Attempts**

Syntax	Order of Authentication Attempts
<b>authentication-order radius;</b>	<ol style="list-style-type: none"> <li>1. Try configured RADIUS authentication servers.</li> <li>2. If RADIUS server is available and authentication is accepted, grant access.</li> <li>3. If RADIUS server is available but authentication is rejected, deny access.</li> <li>4. If RADIUS servers are not available, try password authentication.</li> </ol> <p><b>NOTE:</b> If a RADIUS server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>

Table 78: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
<b>authentication-order [ radius password ];</b>	<ol style="list-style-type: none"> <li>1. Try configured RADIUS authentication servers.</li> <li>2. If RADIUS servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.</li> </ol>
<b>authentication-order [ radius tacplus ];</b>	<ol style="list-style-type: none"> <li>1. Try configured RADIUS authentication servers.</li> <li>2. If RADIUS server is available and authentication is accepted, grant access.</li> <li>3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers.</li> <li>4. If TACACS+ server is available and authentication is accepted, grant access.</li> <li>5. If TACACS+ server is available but authentication is rejected, deny access.</li> <li>6. If both RADIUS and TACACS+ servers are not available, try password authentication.</li> </ol> <p><b>NOTE:</b> If either RADIUS or TACACS+ servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
<b>authentication-order [ radius tacplus password ];</b>	<ol style="list-style-type: none"> <li>1. Try configured RADIUS authentication servers.</li> <li>2. If RADIUS server is available and authentication is accepted, grant access.</li> <li>3. If RADIUS servers fail to respond or return a reject response, try configured TACACS+ servers.</li> <li>4. If TACACS+ server is available and authentication is accepted, grant access.</li> <li>5. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.</li> </ol>
<b>authentication-order tacplus;</b>	<ol style="list-style-type: none"> <li>1. Try configured TACACS+ authentication servers.</li> <li>2. If TACACS+ server is available and authentication is accepted, grant access.</li> <li>3. If TACACS+ server is available but authentication is rejected, deny access.</li> <li>4. If TACACS+ servers are not available, try password authentication.</li> </ol> <p><b>NOTE:</b> If a TACACS+ server is available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>

Table 78: Order of Authentication Attempts (*continued*)

Syntax	Order of Authentication Attempts
<b>authentication-order [ tacplus password ];</b>	<ol style="list-style-type: none"> <li>1. Try configured TACACS+ authentication servers.</li> <li>2. If TACACS+ servers fail to respond or return a reject response, try password authentication, because it is explicitly configured in the authentication order.</li> </ol>
<b>authentication-order [ tacplus radius ];</b>	<ol style="list-style-type: none"> <li>1. Try configured TACACS+ authentication servers.</li> <li>2. If TACACS+ server is available and authentication is accepted, grant access.</li> <li>3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers.</li> <li>4. If RADIUS server is available and authentication is accepted, grant access.</li> <li>5. If RADIUS server is available but authentication is rejected, deny access.</li> <li>6. If both TACACS+ and RADIUS servers are not available, try password authentication.</li> </ol> <p><b>NOTE:</b> If either TACACS+ or RADIUS servers are available, password authentication is not attempted, because it is not explicitly configured in the authentication order.</p>
<b>authentication-order [ tacplus radius password ];</b>	<ol style="list-style-type: none"> <li>1. Try configured TACACS+ authentication servers.</li> <li>2. If TACACS+ server is available and authentication is accepted, grant access.</li> <li>3. If TACACS+ servers fail to respond or return a reject response, try configured RADIUS servers.</li> <li>4. If RADIUS server is available and authentication is accepted, grant access.</li> <li>5. If RADIUS servers fail to respond or return a reject response try password authentication, because it is explicitly configured in the authentication order.</li> </ol>
<b>authentication-order password;</b>	<ol style="list-style-type: none"> <li>1. Try to authenticate the user, using the password configured at the <b>[edit system login]</b> hierarchy level.</li> <li>2. If the authentication is accepted, grant access.</li> <li>3. If the authentication is rejected, deny access.</li> </ol>



**NOTE:** If SSH public keys are configured, SSH user authentication first tries to perform public key authentication before using the authentication methods configured in the authentication-order statement. If you want SSH logins to use the authentication methods configured in the authentication-order statement without first trying to perform public key authentication, do not configure SSH public keys.

In a routing matrix based on a TX Matrix router, the authentication order must be configured only at the configuration groups `re0` and `re1`. The authentication order must not be configured at the `[edit system]` hierarchy. This is because the authentication order for the routing matrix is controlled on the switch-card chassis (or TX Matrix router) or switch-fabric chassis (for TX Matrix Plus router) only.

In Junos OS Release 10.0 and later, the superuser (belonging to the super-user login class) is also authenticated based on the authentication order that is configured for TACACS+, RADIUS, or password authentication using the authentication-order statement. For example, if the only configured authentication order is TACACS+, the superuser can only be authenticated by the TACACS+ server and password authentication cannot be used as an alternative. However, in Junos OS Release 9.6 and earlier, the superuser can use password authentication to login, even if password authentication is not configured explicitly using the authentication-order statement.

#### Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1346](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1369](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1387](#)

## Junos OS User Authentication Methods

The Junos OS supports three methods of user authentication: local password authentication, Remote Authentication Dial-In User Service (RADIUS), and Terminal Access Controller Access Control System Plus (TACACS+).

With local password authentication, you configure a password for each user allowed to log in to the router or switch.

RADIUS and TACACS+ are authentication methods for validating users who attempt to access the router or switch using telnet. They are both distributed client-server systems—the RADIUS and TACACS+ clients run on the router or switch, and the server runs on a remote network system.



You can configure the router or switch to be both a RADIUS and TACACS+ client, and you can also configure authentication passwords in the Junos OS configuration file. You can prioritize the methods to configure the order in which the software tries the different authentication methods when verifying user access.

**Related Documentation**

- *Configuring RADIUS Authentication*
- *Configuring TACACS+ Authentication*
- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1330](#)
- [Configuring RADIUS Authentication \(QFX Series\) on page 1351](#)
- [Configuring TACACS+ Authentication \(QFX Series\) on page 1364](#)

## Junos OS User Accounts Overview

User accounts provide one way for users to access the switch. (Users can access the switch without accounts if you configured RADIUS or TACACS+ servers, as described in [“Junos OS User Authentication Methods” on page 1334](#).) For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

For each user account, you can define the following:

- Username—(Optional) Name that identifies the user. It must be unique within the switch. Do not include spaces, colons, or commas in the username. The username can be up to 64 characters long.
- User’s full name—(Optional) If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
- User identifier (UID)—(Optional) Numeric identifier that is associated with the user account name. The identifier must be in the range from 100 through 64,000 and must be unique within the switch. If you do not assign a UID to a username, the software assigns one when you commit the configuration, preferring the lowest available number.
- You must ensure that the UID is unique. However, it is possible to assign the same UID to different users. If you do this, the CLI displays a warning when you commit the configuration and then assigns the duplicate UID.
- User’s access privilege—(Required) One of the login classes you defined in the **class** statement at the **[edit system login]** hierarchy level, or one of the default classes listed in [“Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies” on page 1338](#).
- Authentication method or methods and passwords that the user can use to access the switch—(Optional) You can use SSH or a Message Digest 5 (MD5) password, or you can enter a plain-text password that Junos OS encrypts using MD5-style encryption before entering it in the password database. For each method, you can specify the user’s password. If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system login user user-name]
user@switch# set authentication plain-text-password
New password: type password here
Retype new password: retry password here
```

The default requirements for plain-text passwords are:

- The password must be between 6 and 128 characters long
- You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- Valid passwords must contain at least one change of case or character class.

For each user account and for root logins, you can configure more than one public RSA or DSA key for user authentication. When a user logs in using a user account or as root, the configured public keys are referenced to determine whether the private key matches any of them.

For SSH authentication, you can also copy the contents of an SSH key file into the configuration.

To load an SSH key file, use the **load-key-file** statement. This statement loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the **load-key-file** statement. To view the SSH key entries, use the configuration mode **show** command. For example:

```
[edit system login user boojum]
user@switch# set authentication load-key-file my-host:.ssh/identity.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@switch# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
207404962528390382038690141584534964170019610608358722961563
475784918273603361276441874265946893207739108344813125957722
625461667999278316123500438660915866283822489746732605661192
181489539813862940327687806538169602027491641637359132693963
44008443 boojum@juniper.net"; # SECRET-DATA
}
```

An account for the user **root** is always present in the configuration. You configure the password for **root** using the **root-authentication** statement, as described in [“Configuring the Root Password” on page 1354](#).

Junos-FIPS and Common Criteria have special password requirements. FIPS and Common Criteria passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the switch, you cannot configure passwords unless they meet this standard.

- Related Documentation**
- [Configuring Junos OS User Accounts on page 1344](#)
  - [Junos OS Login Classes Overview on page 1337](#)

## Junos OS Login Classes Overview

All users who can log in to the router or switch must be in a login class. With login classes, you define the following:

- Access privileges that users have when they are logged in to the router or switch
- Commands and statements that users can and cannot specify
- How long a login session can be idle before it times out and the user is logged out

You can define any number of login classes and then apply one login class to an individual user account.

The Junos operating system (Junos OS) contains a few predefined login classes, which are listed in [Table 79 on page 1337](#). The predefined login classes cannot be modified.

**Table 79: Predefined System Login Classes**

Login Class	Permission Flag Set
operator	clear, network, reset, trace, and view
read-only	view
superuser or super-user	all
unauthorized	None



### NOTE:

- You cannot modify a predefined login class name. If you issue the `set` command on a predefined class name, the Junos OS appends `-local` to the login class name. The following message also appears:  

```
warning: '<class-name>' is a predefined class name; changing to '<class-name>-local'
```
- You cannot issue the `rename` or `copy` command on a predefined login class. Doing so results in the following error message:  

```
error: target '<class-name>' is a predefined class
```

- Related Documentation**
- [Defining Junos OS Login Classes](#)
  - [Defining Junos OS Login Classes on page 1368](#)
  - [Understanding QFabric System Login Classes](#)

## Regular Expressions for Allowing and Denying Junos OS Configuration Mode Hierarchies

Use extended regular expressions to specify which configuration mode hierarchies are denied or allowed. You specify these regular expressions in the **allow/deny-configuration-regexps** and **allow/deny-configuration** statements at the **[edit system login class]** hierarchy level, or by specifying Juniper Networks vendor-specific TACACS+ or RADIUS attributes in your authentication server's configuration. If regular expressions are received during TACACS+ or RADIUS authentication, they merge with any regular expressions configured on the local router or switch.

[Table 80 on page 1338](#) lists common regular expression operators that you can use for allowing or denying configuration mode .

Command regular expressions implement the extended (modern) regular expressions, as defined in POSIX 1003.2.

**Table 80: Configuration Mode Hierarchies—Common Regular Expression Operators**

Operator	Match
	One of two or more terms separated by the pipe. Each term must be a complete standalone expression enclosed in parentheses ( ), with no spaces between the pipe and the adjacent parentheses. For example, <b>(show system alarms) (show system software)</b> .
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, <b>allow-commands "show interfaces\$"</b> means that the user can issue the <b>show interfaces</b> command but cannot issue <b>show interfaces detail</b> or <b>show interfaces extensive</b> .
[ ]	Range of letters or digits. To separate the start and end of a range, use a hyphen ( - ).
( )	A group of commands, indicating a complete, standalone expression to be evaluated; the result is then evaluated as part of the overall expression. Parentheses must be used in conjunction with pipe operators as explained .
*	Zero or more terms.
+	One or more terms.
.	Any character except for a space " " .

### Related Documentation

- *Specifying Access Privileges for Junos OS Configuration Mode Hierarchies*
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1371](#)

## Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands

Use extended regular expressions to specify which operational mode commands are denied or allowed. [Table 81 on page 1339](#) lists common regular expression operators that can be used in the operational mode commands. Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2.

**Table 81: Common Regular Expression Operators to Allow or Deny Operational Mode Commands**

Operator	Match
	One of two or more terms separated by the pipe ( ) symbol. Each term must be a complete standalone expression enclosed in parentheses ( ), with no spaces between the pipe and the adjacent parentheses. For example, <b>(show system alarms) (show system software)</b> .
^	At the beginning of an expression, used to denote where the command begins, and where there might be some ambiguity.
\$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, <b>allow-commands "show interfaces\$"</b> means that the user can issue the <b>show interfaces</b> command but cannot issue the <b>show interfaces detail</b> or <b>show interfaces extensive</b> command.
[ ]	Range of letters or digits. To separate the start and end of a range, use a hyphen ( - ).
( )	A group of commands, indicating a complete, standalone expression to be evaluated; the result is then evaluated as part of the overall expression. Parentheses must always be used in conjunction with pipe operators as explained above.

If a regular expression contains a syntax error, it becomes invalid, and although the user can log in, the permission granted or denied by the regular expression does not take effect. When regular expressions configured on TACACS+ or RADIUS servers merge with regular expressions configured on the router or switch, if the final expression has a syntax error, the overall result is an invalid regular expression. If a regular expression does not contain any operators, all varieties of the command are allowed. For example, if the following statement is included in the configuration, the user can issue the commands **show interfaces detail** and **show interfaces extensive** in addition to showing an individual interface:

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

### Related Documentation

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 192](#)

## Special Requirements for Junos OS Plain-Text Passwords

Junos OS has special requirements when you create plain-text passwords on a router or switch. [Table 82 on page 1340](#) shows the default requirements.

Table 82: Special Requirements for Plain-Text Passwords

Junos OS	Junos-FIPS
The password must be between 6 and 128 characters long.	FIPS passwords must be between 10 and 20 characters long
You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.	You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
Valid passwords must contain at least one change of case or character class.	Passwords must use at least three of the five defined character classes (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters).

You can change the requirements for plain-text passwords.

Junos OS supports the following five character classes for plain-text passwords:

- Lowercase letters
- Uppercase letters
- Numbers
- Punctuation
- Special characters: ! @ # \$ % ^ & \*, + < > ; ;

Control characters are not recommended.

You can include the **plain-text-password** statement at the following hierarchy levels:

- [edit system diag-port-authentication]
- [edit system pic-console-authentication]
- [edit system root-authentication]
- [edit system login user *username* authentication]

The **change-type** statement specifies whether the password is checked for the following:

- The total number of character sets used (**character-set**)
- The total number of character set changes (**set-transitions**)

For example, the following password:

```
MyPassWd@2
```

has four character sets (uppercase letters, lowercase letters, special characters, and numbers) and seven character set changes (**M-y**, **y-P**, **P-a**, **s-W**, **W-d**, **d-@**, and **@-2**).

The **change-type** statement is optional. If you omit the **change-type** option, Junos-FIPS plain-text passwords are checked for character sets, and Junos OS plain-text passwords are checked for character set changes.

The **minimum-changes** statement specifies how many character sets or character set changes are required for the password. This statement is optional. If you do not use the **minimum-changes** statement, character sets are not checked for Junos OS. If the **change-type** statement is configured for the **character-set** option, then the **minimum-changes** value must be 5 or less, because Junos OS only supports five character sets.

The **format** statement specifies the hash algorithm (**md5**, **sha1**, **sha256**, **sha512** or **des**) for authenticating plain-text passwords. This statement is optional. For Junos OS, the default format is **md5**. For Junos-FIPS, only **sha1** is supported.



**NOTE:** Starting with Junos OS Release 13.3, the **sha1** does not enable secure, protected specification of passwords and we recommend that you do not use the **sha1** algorithm to configure passwords. Instead, you can use the **sha256** or **sha512** to specify passwords by using the 256-bit and 512-bit cryptographic hash algorithm respectively for a robust and reliable operation.

The **maximum-length** statement specifies the maximum number of characters allowed in a password. This statement is optional. By default, Junos OS passwords have no maximum; however, only the first 128 characters are significant. Junos-FIPS passwords must be 20 characters or less. The range for Junos OS maximum-length passwords is from 20 to 128 characters.

The **minimum-length** statement specifies the minimum number of characters required for a password. This statement is optional. By default, Junos OS passwords must be at least 6 characters long, and Junos-FIPS passwords must be at least 10 characters long. The range is from 6 to 20 characters.

Changes to password requirements do not take effect until the configuration is committed. When requirements change, only newly created, plain-text passwords are checked; existing passwords are not checked against the new requirements.

The default configuration for Junos OS plain-text passwords is:

```
[edit system login]
passwords {
  change-type character-sets;
  format md5;
  minimum-changes 1;
  minimum-length 6;
}
```

The default configuration for Junos-FIPS plain-text passwords is:

```
[edit system login]
passwords {
  change-type set-transitions;
  format sha1;
```

```
maximum-length 20;  
minimum-changes 3;  
minimum-length 10;  
}
```



**NOTE:** Transitioning to FIPS mode is allowed only when the encrypted password is a FIPS-compliant hash algorithm. Also, configuring passwords that are not FIPS-compliant is not allowed when you are in FIPS mode.

**Related  
Documentation**

- [Changing the Requirements for Junos OS Plain-Text Passwords](#)
- [Configuring the Root Password on page 170](#)
- [Changing the Requirements for Junos OS Plain-Text Passwords](#)
- [Configuring the Root Password on page 1354](#)



## CHAPTER 12

# Configuration

- [Configuration Tasks on page 1343](#)
- [Configuration Examples on page 1378](#)
- [Configuration Statements on page 1392](#)

### Configuration Tasks

---

- [Configuring Access Privilege Levels on page 1344](#)
- [Configuring Login Tips on page 1344](#)
- [Configuring Junos OS User Accounts on page 1344](#)
- [Configuring LLDP on page 1345](#)
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1346](#)
- [Configuring Local User Template Accounts for User Authentication on page 1347](#)
- [Configuring Management Access on page 1349](#)
- [Configuring RADIUS System Accounting on page 1349](#)
- [Configuring RADIUS Authentication \(QFX Series\) on page 1351](#)
- [Configuring Remote Template Accounts for User Authentication on page 1354](#)
- [Configuring the Root Password on page 1354](#)
- [Configuring SNMP on page 1356](#)
- [Configuring SSH Host Keys for Secure Copying of Data on page 1359](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 1361](#)
- [Configuring TACACS+ Authentication \(QFX Series\) on page 1364](#)
- [Configuring TACACS+ System Accounting on page 1366](#)
- [Defining Junos OS Login Classes on page 1368](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1369](#)
- [Recovering the Root Password on page 1370](#)
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1371](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 1372](#)
- [Using Junos OS to Configure Logical System Administrators on page 1374](#)

- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1375](#)
- [VSA Match Conditions and Actions on page 1376](#)

## Configuring Access Privilege Levels

Each top-level command-line interface (CLI) command and each configuration statement have an access privilege level associated with it. Users can execute only those commands and configure and view only those statements for which they have access privileges.

To configure access privilege levels, include the **permissions** statement at the **[edit system login class *class-name*]** hierarchy level:

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

### Related Documentation

- [Example: Configuring Access Privilege Levels on page 1381](#)
- [Understanding Junos OS Access Privilege Levels on page 1325](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 192](#)
- *permissions*

## Configuring Login Tips

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

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

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

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

### Related Documentation

- [CLI User Interface Overview on page 39](#)
- [Defining Junos OS Login Classes](#)
- [login-tip on page 281](#)

## Configuring Junos OS User Accounts

User accounts provide one way for users to access the router or switch. For each account, you define the login name for the user and, optionally, information that identifies the user. After you have created an account, the software creates a home directory for the user.

To create user accounts, include the **user** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
user username {
  class class-name;
  class {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
  }
  full-name complete-name;
  uid uid-value;
  class class-name;
}
```

#### Related Documentation

- [Example: Configuring User Accounts on page 1386](#)
- [Example: Configuring User Login Accounts on page 1389](#)
- [Junos OS User Accounts Overview on page 1335](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1369](#)

## Configuring LLDP

Devices use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the device to identify a variety of devices quickly. The result is a LAN that interoperates smoothly and efficiently.

The LLDP protocol cannot be enabled by issuing the **set protocols lldp** statement at the **[edit]** hierarchy level. Enable the LLDP protocol by configuring it on all interfaces or on specific interfaces.

To configure basic LLDP options using the CLI:

1. Configure the advertisement interval in seconds:

```
[edit protocols lldp]
user@switch# set advertisement-interval 45
```

2. Specify the multiplier used in combination with the **advertisement-interval** value to determine the length of time LLDP information is held before it is discarded:

```
[edit protocols lldp]
user@switch# set hold-multiplier 5
```

3. Configure LLDP on all interfaces or on a specific interface:

```
[edit protocols lldp]
user@switch# set interface (LLDP) all
```

4. Configure tracing operations for the LLDP protocol:

```
[edit protocols lldp]
user@switch# set traceoptions file lldptrace
```

#### Related Documentation

## Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication

Using the **authentication-order** statement, you can prioritize the order in which the Junos OS tries the different authentication methods when verifying user access to a router or switch.

To configure the authentication order, include the **authentication-order** statement at the **[edit system]** hierarchy level:

```
[edit system]
authentication-order [ authentication-methods ];
```

Specify one or more of the following authentication methods in the preferred order, from first tried to last tried:

- **radius**—Verify the user using RADIUS authentication services
- **tacplus**—Verify the user using TACACS+ authentication services.
- **password**—Verify the user using the username and password configured locally by including the authentication statement at the **[edit system login user]** hierarchy level.

The CHAP authentication sequence cannot take more than 30 seconds. If it takes longer to authenticate a client, the authentication is abandoned and a new sequence is initiated.

For example, if you configure three RADIUS servers so that the router or switch attempts to contact each server three times, and with each retry the server times out after 3 seconds, then the maximum time given to the RADIUS authentication method before CHAP considers it a failure is 27 seconds. If you add more RADIUS servers to this configuration, they might not be contacted because the authentication process might be abandoned before these servers are tried.

The Junos OS enforces a limit on the number of standing authentication server requests that the CHAP authentication can have at one time. Thus, an authentication server method—RADIUS, for example—might fail to authenticate a client when this limit is exceeded. If it fails, the authentication sequence is reinitiated by the router or switch until authentication succeeds and the link is brought up. However, if the RADIUS servers are not available and if additional authentication methods such as **tacplus** or **password** are configured along with **radius**, the next authentication method is tried.

The following example shows how to configure **radius** and **password** authentication:

```
[edit system]
user@switch# authentication-order [ radius password ];
```

The following example shows how to delete the **radius** statement from the authentication order:

```
[edit system]
user@switch# delete authentication-order radius
```

The following example shows how to insert the **tacplus** statement after the **radius** statement:

```
[edit system]
user@switch# insert authentication-order tacplus after radius
```

#### Related Documentation

- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1330](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1375](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1387](#)
- *authentication-order*

## Configuring Local User Template Accounts for User Authentication

You use local user template accounts when you need different types of templates for authentication. Each template can define a different set of permissions appropriate for the group of users who use that template. These templates are defined locally on the router or switch and referenced by the TACACS+ and RADIUS authentication servers.

When you configure local user templates and a user logs in, Junos OS issues a request to the authentication server to authenticate the user's login name. If a user is authenticated, the server returns the local username to Junos OS, which then determines whether a local username is specified for that login name (**local-username** for TACACS+, **Juniper-Local-User** for RADIUS). If so, Junos OS selects the appropriate local user template locally configured on the router or switch. If a local user template does not exist for the authenticated user, the router or switch defaults to the **remote** template.

To configure different access privileges for users who share the local user template account, include the **allow-commands** and **deny-commands** commands in the authentication server configuration file.

To configure a local user template, include the **user local-username** statement at the **[edit system login]** hierarchy level and specify the privileges you want to grant to the local users to whom the template applies:

```
[edit system login]
user local-username {
  full-name "Local user account";
  uid uid-value;
  class class-name;
}
```

This example configures the **sales** and **engineering** local user templates:

```
[edit]
system {
  login {
    user sales {
      uid uid-value;
```

```
        class class-name;
    }
    user engineering {
        uid uid-value;
        class class-name;
    }
}

user = simon {
    ...
    service = junos-exec {
        local-user-name = sales
        allow-commands = "configure"
        deny-commands = "shutdown"
    }
}
user = rob {
    ...
    service = junos-exec {
        local-user-name = sales
        allow-commands = "(request system) | (show rip neighbor)"
        deny-commands = "clear"
    }
}
user = harold {
    ...
    service = junos-exec {
        local-user-name = engineering
        allow-commands = "monitor | help | show | ping | traceroute"
        deny-commands = "configure"
    }
}
user = jim {
    ...
    service = junos-exec {
        local-user-name = engineering
        allow-commands = "show bgp neighbor"
        deny-commands = "telnet | ssh"
    }
}
```

When the login users Simon and Rob are authenticated, the router or switch applies the sales local user template. When login users Harold and Jim are authenticated, the router or switch applies the engineering local user template.

**Related  
Documentation**

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)
- [user \(Access\)](#)
- [user \(Access\) on page 332](#)

## Configuring Management Access

To define the management access settings for the routing platform:

1. Next to Allow Telnet Access, select the check box to allow remote Telnet access to the routing platform.
2. Next to Allow SSH Access, selected the check box to allow remote SSH access to the routing platform.
3. Click **Apply** to apply the configuration.

### Related Documentation

- [Configuring Junos OS User Accounts on page 1344](#)
- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 192](#)
- [Example: Configuring Access Privilege Levels on page 1381](#)

## Configuring RADIUS System Accounting

With RADIUS accounting enabled, Juniper Networks routers or switches, acting as RADIUS clients, can notify the RADIUS server about user activities such as software logins, configuration changes, and interactive commands. The framework for RADIUS accounting is described in RFC 2866.

Tasks for configuring RADIUS system accounting are:

1. [Configuring Auditing of User Events on a RADIUS Server on page 1349](#)
2. [Specifying RADIUS Server Accounting and Auditing Events on page 1350](#)
3. [Configuring RADIUS Server Accounting on page 1350](#)

### Configuring Auditing of User Events on a RADIUS Server

To audit user events, include the following statements at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
destination {
  radius {
    server {
      server-address {
        accounting-port port-number;
        secret password;
        source-address address;
        retry number;
        timeout seconds;
      }
    }
  }
}
```

### Specifying RADIUS Server Accounting and Auditing Events

---

To specify the events you want to audit when using a RADIUS server for authentication, include the **events** statement at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
```

**events** is one or more of the following:

- **login**—Audit logins
- **change-log**—Audit configuration changes
- **interactive-commands**—Audit interactive commands (any command-line input)

### Configuring RADIUS Server Accounting

---

To configure RADIUS server accounting, include the **server** statement at the **[edit system accounting destination radius]** hierarchy level:

```
server {
  server-address {
    accounting-port port-number;
    secret password;
    source-address address;
    retry number;
    timeout seconds;
  }
}
```

**server-address** specifies the address of the RADIUS server. To configure multiple RADIUS servers, include multiple **server** statements.



**NOTE:** If no RADIUS servers are configured at the **[edit system accounting destination radius]** statement hierarchy level, the Junos OS uses the RADIUS servers configured at the **[edit system radius-server]** hierarchy level.

**accounting-port *port-number*** specifies the RADIUS server accounting port number.

The default port number is 1813.



**NOTE:** If you enable RADIUS accounting at the **[edit access profile *profile-name* accounting-order]** hierarchy level, accounting is triggered on the default port of 1813 even if you do not specify a value for the **accounting-port** statement.

You must specify a secret (password) that the local router or switch passes to the RADIUS client by including the **secret** statement. If the password contains spaces, enclose the entire password in quotation marks (" ").



In the **source-address** statement, specify a source address for the RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address. The source address is a valid IPv4 address configured on one of the router or switch interfaces.

Optionally, you can specify the number of times that the router or switch attempts to contact a RADIUS authentication server by including the **retry** statement. By default, the router or switch retries three times. You can configure the router or switch to retry from 1 through 10 times.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a RADIUS server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure the timeout to be from 1 through 90 seconds.

If you use the **enhanced-accounting** statement at the **[edit system radius-options]** hierarchy level, the RADIUS attributes such as access method, remote port, and access privileges can be audited. You can limit the number of attribute values to be displayed for auditing by using the **enhanced-avs-max <number>** statement at the **[edit system accounting]** hierarchy level.

```
[edit system radius-options]
enhanced-accounting;

[edit system accounting]
enhanced-avs-max <number>;
```

## Configuring RADIUS Authentication (QFX Series)

RADIUS authentication is a method of authenticating users who attempt to access the router or switch. Tasks to configure RADIUS authentication are:



**NOTE:** The **source-address** statement is not supported at the **[edit system radius-options]** or **[edit system-radius-server name]** hierarchies on the QFabric system.

- [Configuring RADIUS Server Details on page 1351](#)
- [Configuring MS-CHAPv2 for Password-Change Support on page 1352](#)
- [Specifying a Source Address for the Junos OS to Access External RADIUS Servers on page 1353](#)

### Configuring RADIUS Server Details

To use RADIUS authentication on the router or switch, configure information about one or more RADIUS servers on the network by including one **radius-server** statement at the **[edit system]** hierarchy level for each RADIUS server:

```
[edit system]
radius-server server-address {
  accounting-port port-number;
  port number;
  retry number;
```

```
secret password;  
source-address source-address;  
timeout seconds;  
}
```

**server-address** is the address of the RADIUS server.

You can specify a port on which to contact the RADIUS server. By default, port number **1812** is used (as specified in RFC 2865). You can also specify an accounting port to send accounting packets. The default is **1813** (as specified in RFC 2866).

You must specify a password in the **secret password** statement. If the password contains spaces, enclose it in quotation marks. The secret used by the local router or switch must match that used by the server.

Optionally, you can specify the amount of time that the local router or switch waits to receive a response from a RADIUS server (in the **timeout** statement) and the number of times that the router or switch attempts to contact a RADIUS authentication server (in the **retry** statement). By default, the router or switch waits 3 seconds. You can configure this to be a value from 1 through 90 seconds. By default, the router or switch retries connecting to the server three times. You can configure this to be a value from 1 through 10 times.

You can use the **source-address** statement to specify a logical address for individual or multiple RADIUS servers.

To configure multiple RADIUS servers, include multiple **radius-server** statements.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the **user** statement at the **[edit system login]** hierarchy level, as described in [“Overview of Template Accounts for RADIUS and TACACS+ Authentication” on page 1318](#).

You can also configure RADIUS authentication at the **[edit access]** and **[edit access profile]** hierarchy level. Junos OS uses the following search order to determine which set of servers are used for authentication:

1. **[edit access profile profile-name radius-server server-address]**
2. **[edit access radius-server server-address]**
3. **[edit system radius-server server-address]**

---

### Configuring MS-CHAPv2 for Password-Change Support

You can configure the Microsoft implementation of the Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) on the router or switch to support changing of passwords. This feature provides users accessing a router or switch the option of changing the password when the password expires, is reset, or is configured to be changed at the next login.

Before you configure MS-CHAPv2 for password-change support, ensure that you:

- Configure the RADIUS server authentication parameters

- Set the **authentication-order** to use the RADIUS server for the initial password attempt

To configure MS-CHAP-v2, include the following statements at the **[edit system radius-options]** hierarchy level:

```
[edit system radius-options]
password-protocol mschap-v2;
```

The following example shows statements for configuring the MS-CHAPv2 password protocol, password authentication order, and user accounts:

```
[edit]
system {
  authentication-order [ radius password ];
  radius-server {
    192.168.69.149 secret "$9$G-j.5Qz6tpBk.1hrlXxUjiq5Qn/C"; ## SECRET-DATA
  }
  radius-options {
    password-protocol mschap-v2;
  }
  login {
    user bob {
      class operator;
    }
  }
}
```

### Specifying a Source Address for the Junos OS to Access External RADIUS Servers

You can specify which source address Junos OS uses when accessing your network to contact an external RADIUS server for authentication. You can also specify which source address Junos OS uses when contacting a RADIUS server for sending accounting information.

To specify a source address for a RADIUS server, include the **source-address** statement at the **[edit system radius-server server-address]** hierarchy level:

```
[edit system radius-server server-address]
source-address source-address;
```

**source-address** is a valid IP address configured on one of the router or switch interfaces.

#### Related Documentation

- [Example: Configuring RADIUS Authentication on page 1384](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1387](#)
- [Juniper Networks Vendor-Specific RADIUS Attributes on page 1321](#)
- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)
- [Example: Configuring RADIUS Template Accounts on page 1390](#)
- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1375](#)
- [Junos OS User Authentication Methods on page 1334](#)

## Configuring Remote Template Accounts for User Authentication

By default, the Junos OS uses remote template accounts for user authentication when:

- The authenticated user does not exist locally on the router or switch.
- The authenticated user's record in the authentication server specifies local user, or the specified local user does not exist locally on the router or switch.

To configure the remote template account, include the **user remote** statement at the **[edit system login]** hierarchy level and specify the privileges you want to grant to remote users:

```
[edit system login]
user remote {
  full-name "All remote users";
  uid uid-value;
  class class-name;
}
```

To configure different access privileges for users who share the remote template account, include the **allow-commands** and **deny-commands** statements in the authentication server configuration file.

### Related Documentation

- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)
- [user \(Access\)](#)
- [user \(Access\) on page 332](#)

## Configuring the Root Password

Junos OS is preinstalled on the router or switch. When the router or switch is powered on, it is ready to be configured. Initially, you log in as the user "root" with no password.



**NOTE:** If you configure a blank password using the **encrypted-password** statement at the **[edit system root-authentication]** hierarchy level for root authentication, you can commit a configuration, but you are *not* able to log in as superuser and gain root level access to the router or switch.

After you log in, you should configure the root (superuser) password by including the **root-authentication** statement at the **[edit system]** hierarchy level:

```
[edit system]
root-authentication {
  (encrypted-password "password" | load-key-password URL | plain-text-password);
  ssh-dsa "public-key";
  ssh-rsa "public-key";
}
```

If you configure the **plain-text-password** option, you are prompted to enter and confirm the password:

```
[edit system]
user@switch# set root-authentication plain-text-password
New password: type password here
Retype new password: retry password here
```

To load an SSH key file, enter the **load-key-file** statement. This statement loads RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys.

You can also configure SSH RSA keys and SSH DSA keys to authenticate root logins. You can configure more than one public RSA or DSA key for SSH authentication of root logins as well as for user accounts. When a user logs in as root, the public keys are referenced to determine whether the private key matches any of them.

If you load the SSH keys file, the contents of the file are copied into the configuration immediately after you enter the **load-key-file** statement. To view the SSH keys entries, use the configuration mode **show** command. For example:

```
[edit system]
user@switch# set root-authentication load-key-file my-host:.ssh/identity.pub
.file.19692 | 0 KB | 0.3 kB/s | ETA: 00:00:00 | 100%
[edit system]
user@switch# show
root-authentication {
  ssh-rsa "1024 35 9727638204084251055468226757249864241630322
20740496252839038203869014158453496417001961060835872296
15634757491827360336127644187426594689320773910834481012
68312595772262546166799927831612350043866091586628382248
97467326056611921489539813965561563786211940327687806538
16960202749164163735913269396344008443 boojum@juniper.net"; #
  SECRET-DATA
}
```

Junos-FIPS software has special password requirements. FIPS passwords must be between 10 and 20 characters in length. Passwords must use at least three of the five defined character sets (uppercase letters, lowercase letters, digits, punctuation marks, and other special characters). If Junos-FIPS is installed on the router or switch, you cannot configure passwords unless they meet this standard. If you use the **encrypted-password** option, then a null-password (empty) is not permitted.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

#### Related Documentation

- [Recovering the Root Password on page 1233](#)
- [Example: Configuring the Root Password on page 1385](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 1382](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1386](#)

## Configuring SNMP

SNMP is implemented in the Junos OS Software running on the QFX Series products. By default, SNMP is not enabled. To enable SNMP, you must include the SNMP configuration statements at the **[edit]** hierarchy level.

To configure the minimum requirements for SNMP, include the following statements at the **[edit]** hierarchy level of the configuration:

```
[edit]
snmp {
  community public;
}
```

To configure complete SNMP features, include the following statements at the **[edit]** hierarchy level of the configuration:

```
snmp {
  client-list client-list-name {
    ip-addresses;
  }
  community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
      address restrict;
    }
    logical-system logical-system-name {
      routing-instance routing-instance-name {
        clients {
          addresses;
        }
      }
    }
    routing-instance routing-instance-name {
      clients {
        addresses;
      }
    }
    view view-name;
  }
  contact contact;
  description description;
  filter-duplicates;
  filter-interfaces;
  health-monitor {
    falling-threshold integer;
    interval seconds;
    rising-threshold integer;
  }
  interface [ interface-names ];
  location location;
  name name;
  nonvolatile {
    commit-delay seconds;
  }
}
```

```

}
rmon {
  alarm index {
    description description;
    falling-event-index index;
    falling-threshold integer;
    falling-threshold-interval seconds;
    interval seconds;
    request-type;
    rising-event-index index;
    rising-threshold integer;
    sample-type (absolute-value | delta-value);
    startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
    syslog-subtag syslog-subtag;
    variable oid-variable;
  }
  event index {
    community community-name;
    description description;
    type type;
  }
  history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regular-expression>;
  flag flag;
}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance routing-instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
}

```

```
snmp-community community-index {
  community-name community-name;
  security-name security-name;
  tag tag-name;
}
target-address target-address-name {
  address address;
  address-mask address-mask;
  logical-system logical-system;
  port port-number;
  retry-count number;
  routing-instance routing-instance-name;
  tag-list tag-list;
  target-parameters target-parameters-name;
  timeout seconds;
}
target-parameters target-parameters-name {
  notify-filter profile-name;
  parameters {
    message-processing-model (v1 | v2c | V3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {
        authentication-password authentication-password;
      }
      authentication-none;
      privacy-aes128 {
        privacy-password privacy-password;
      }
      privacy-des {
        privacy-password privacy-password;
      }
      privacy-3des {
        privacy-password privacy-password;
      }
      privacy-none;
    }
  }
  remote-engine engine-id {
    user username {
      authentication-sha {
        authentication-password authentication-password;
      }
      authentication-md5 {
        authentication-password authentication-password;
      }
      authentication-none;
    }
  }
}
```



```

    privacy-aes128 {
        privacy-password privacy-password;
    }
    privacy-des {
        privacy-password privacy-password;
    }
    privacy-3des {
        privacy-password privacy-password;
    }
    privacy-none {
        privacy-password privacy-password;
    }
}
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
}
view view-name {
    oid object-identifier (include | exclude);
}
}

```

- Related Documentation**
- [Understanding the Implementation of SNMP on page 6513](#)
  - [snmp on page 1454](#)

## Configuring SSH Host Keys for Secure Copying of Data

Secure Shell (SSH) uses encryption algorithms to generate a host, server, and session key system that ensures secure data transfer. You can configure SSH host keys to support secure copy (SCP) as an alternative to FTP for the background transfer of data such as configuration archives and event logs. To configure SSH support for SCP, you must complete the following tasks:

- Specify SSH known hosts by including hostnames and host key information in the Routing Engine configuration hierarchy.
- Set an SCP URL to specify the host from which to receive data. Setting this attribute automatically retrieves SSH host key information from the SCP server.
- Verify that the host key is authentic.
- Accept the secure connection. Accepting this connection automatically stores host key information in the local host key database. Storing host key information in the configuration hierarchy automates the secure handshake and allows background data transfer using SCP.

Tasks to configure SSH host keys for secure copying of data are:

1. [Configuring SSH Known Hosts on page 1360](#)
2. [Configuring Support for SCP File Transfer on page 1360](#)
3. [Updating SSH Host Key Information on page 1361](#)

---

### Configuring SSH Known Hosts

To configure SSH known hosts, include the **host** statement, and specify hostname and host key options for trusted servers at the **[edit security ssh-known-hosts]** hierarchy level:

```
[edit security ssh-known-hosts]
host corporate-archive-server, ip-address {
    dsa-key key;
}
host archive-server-url {
    rsa-key key;
}
host server-with-ssh-version-1, ip-address {
    rsa1-key key;
}
```

Host keys are one of the following:

- **dsa-key**—Base64 encoded Digital Signature Algorithm (DSA) key.
- **rsa-key**—Base64 encoded RSA public key algorithm, which supports encryption and digital signatures.
- **rsa1-key**—Base64 encoded RSA public key algorithm, which supports encryption and digital signatures for SSH version 1 and SSH version 2.

---

### Configuring Support for SCP File Transfer

To configure a known host to support background SCP file transfers, include the **archive-sites** statement at the **[edit system archival configuration]** hierarchy level.

```
[edit system archival configuration]
archive-sites {
    scp://username<:password>@host<:port>/url-path;
}
```



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([ ]). For example, "scp://username<:password>@[host]<:port>/url-path";

Setting the **archive-sites** statement to point to an SCP URL triggers automatic host key retrieval. At this point, Junos OS connects to the SCP host to fetch the SSH public key, displays the host key message digest or fingerprint as output to the console, and terminates the connection to the server.

```
user@switch# set system archival configuration archive-sites "<scp-url-path>"
The authenticity of host <my-archive-server (<server-ip-address>)> can't be established.
RSA key fingerprint is <ascii-text key>. Are you sure you want to continue connecting
(yes/no)?
```

To verify that the host key is authentic, compare this fingerprint with a fingerprint that you obtain from the same host using a trusted source. If the fingerprints are identical, accept the host key by entering **yes** at the prompt. The host key information is then stored in the Routing Engine configuration and supports background data transfers using SCP.

### Updating SSH Host Key Information

Typically, SSH host key information is automatically retrieved when you set a URL attribute for SCP using the **archival configuration archive-sites** statement at the **[edit system]** hierarchy level. However, if you need to manually update the host key database, use one of the following methods.

1. [Retrieving Host Key Information Manually on page 1361](#)
2. [Importing Host Key Information from a File on page 1361](#)

#### *Retrieving Host Key Information Manually*

To manually retrieve SSH public host key information, use the **fetch-from-server** option with the **set security ssh-known-hosts** command. You must include a hostname attribute with the **set security ssh-known-hosts fetch-from-server** command to specify the host from which to retrieve the SSH public key.

```
user@switch# set security ssh-known-hosts fetch-from-server <hostname>
```

#### *Importing Host Key Information from a File*

To manually import SSH host key information from the known-hosts file located at **/var/tmp/known-hosts** on the server, include the **load-key-file** option with the **set security ssh-known-hosts** command. You must include the path to the **known-hosts** file with the **set security ssh-known-hosts load-key-file** command to specify the location from which to import host key information.

```
user@switch# set security ssh-known-hosts load-key-file /var/tmp/known-hosts
```

## Configuring SSH Service for Remote Access to the Router or Switch

To configure the router or switch to accept SSH as an access service, include the **ssh** statement at the **[edit system services]** hierarchy level:

```
[edit system services]
ssh {
  ciphers [ cipher-1 cipher-2 cipher-3 ...]
  client-alive-count-max number;
  client-alive-interval seconds;
  connection-limit limit;
  hostkey-algorithm <algorithm | no-algorithm>;
  key-exchange algorithm;
  macs algorithm;
  max-sessions-per-connection number;
  no-passwords;
  no-tcp-forwarding;
  protocol-version [v1 v2];
  rate-limit limit;
  root-login <allow | deny | deny-password>;
}
```

By default, the router or switch supports a limited number of simultaneous SSH sessions and connection attempts per minute. Use the following statements to change the defaults:

- **connection-limit *limit***—Maximum number of simultaneous connections per protocol (IPv4 and IPv6). The range is a value from 1 through 250. The default is 75. When you configure a connection limit, the limit is applicable to the number of SSH sessions per protocol (IPv4 and IPv6). For example, a connection limit of 10 allows 10 IPv6 SSH sessions and 10 IPv4 SSH sessions.
- **max-sessions-per-connection *number***—Include this statement to specify the maximum number of SSH sessions allowed per single SSH connection. This allows you to limit the number of cloned sessions tunneled within a single SSH connection. The default value is 10.
- **rate-limit *limit***—Maximum number of connection attempts accepted per minute (a value from 1 through 250). The default is 150. When you configure a rate limit, the limit is applicable to the number of connection attempts per protocol (IPv4 and IPv6). For example, a rate limit of 10 allows 10 IPv6 SSH session connection attempts per minute and 10 IPv4 SSH session connection attempts per minute.

By default, a user can create an SSH tunnel over a CLI session to a router running Junos OS via SSH. This type of tunnel could be used to forward TCP traffic, bypassing any firewall filters or ACLs, allowing access to resources beyond the router. Use the **no-tcp-forwarding** option to prevent a user from creating an SSH tunnel to a router via SSH.

For information about other configuration settings, see the following topics:

- [Configuring the Root Login Through SSH on page 1363](#)
- [Configuring the SSH Protocol Version on page 1363](#)
- [Configuring the Client Alive Mechanism on page 1363](#)

### Configuring the Root Login Through SSH

By default, users are allowed to log in to the router or switch as **root** through SSH. To control user access through SSH, include the **root-login** statement at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
root-login (allow | deny | deny-password);
```

**allow**—Allows users to log in to the router or switch as root through SSH. The default is **allow**.

**deny**—Disables users from logging in to the router or switch as root through SSH.

**deny-password**—Allows users to log in to the router or switch as root through SSH when the authentication method (for example, RSA) does not require a password.

### Configuring the SSH Protocol Version

By default, both version 1 and version 2 of the SSH protocol are enabled. To configure the router or switch to use only version 1 of the SSH protocol, include the **protocol-version** statement and specify **v1** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 ];
```

To configure the router or switch to use only version 2 of the SSH protocol, include the **protocol-version** statement and specify **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v2 ];
```

To explicitly configure the router or switch to use version 1 and 2 of the SSH protocol, include the **protocol-version** statement and specify **v1** and **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 v2 ];
```

For J Series Services Routers, the export license software supports SSH version 1 only.

### Configuring the Client Alive Mechanism

The client alive mechanism is valuable when the client or server depends on knowing when a connection has become inactive. It differs from the standard keepalive mechanism because the client alive messages are sent through the encrypted channel. The client alive mechanism is not enabled at default. To enable it, configure the **client-alive-count-max** and the **client-alive-interval**. This option applies to SSH protocol version 2 only.

In the following example, unresponsive SSH clients will be disconnected after approximately 100 seconds (20 x 5).

```
[edit system services ssh]
```

```
client-alive-count-max 5;  
client-alive-interval 20;
```

## Configuring TACACS+ Authentication (QFX Series)

TACACS+ authentication is a method of authenticating users who attempt to access the router or switch. Tasks to configure TACACS+ configuration are:

- [Configuring TACACS+ Server Details on page 1364](#)
- [Specifying a Source Address for the Junos OS to Access External TACACS+ Servers on page 1365](#)
- [Configuring the Same Authentication Service for Multiple TACACS+ Servers on page 1365](#)
- [Configuring Juniper Networks Vendor-Specific TACACS+ Attributes on page 1366](#)

---

### Configuring TACACS+ Server Details

To use TACACS+ authentication on the router or switch, configure information about one or more TACACS+ servers on the network by including the **tacplus-server** statement at the **[edit system]** hierarchy level:

```
[edit system]  
tacplus-server server-address {  
  port port-number;  
  secret password;  
  single-connection;  
  timeout seconds;  
}
```

**server-address** is the address of the TACACS+ server.

**port-number** is the TACACS+ server port number.

You must specify a secret (password) by using the **secret** statement. The local router or switch passes the **secret** to the TACACS+ client. If the password included spaces, enclose the password in quotation marks. The secret used by the local router or switch must match that used by the server.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a TACACS+ server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds.

Optionally, you can use the **single-connection** statement to have the software maintain one open Transmission Control Protocol (TCP) connection to the server for multiple requests, rather than opening a connection for each connection attempt.



**NOTE:** Early versions of the TACACS+ server do not support the **single-connection** option. If you specify this option and the server does not support it, Junos OS will be unable to communicate with that TACACS+ server.

---

To configure multiple TACACS+ servers, include multiple **tacplus-server** statements.

On a TX Matrix router, TACACS+ accounting should be configured only under the groups **re0** and **re1**.



**NOTE:** Accounting should not be configured at the **[edit system]** hierarchy level; on a TX Matrix router, control is done under the switch-card chassis only.

To configure a set of users that share a single account for authorization purposes, you create a template user. To do this, include the **user** statement at the **[edit system login]** hierarchy level.

### Specifying a Source Address for the Junos OS to Access External TACACS+ Servers

You can specify which source address Junos OS uses when accessing your network to contact an external TACACS+ server for authentication. You can also specify which source address Junos OS uses when contacting a TACACS+ server for sending accounting information.

To specify a source address for a TACACS+ server for authentication, include the **source-address** statement at the **[edit system tacplus-server server-address]** hierarchy level:

```
[edit system tacplus-server server-address]
source-address source-address;
```

**source-address** is a valid IP address configured on one of the router or switch interfaces.

To specify a source address for a TACACS+ server for system accounting, include the **source-address** statement at the **[edit system accounting destination tacplus server server-address]** hierarchy level:

```
[edit system accounting destination tacplus server server-address]
source-address source-address;
```

**source-address** is a valid IP address configured on one of the router or switch interfaces.

### Configuring the Same Authentication Service for Multiple TACACS+ Servers

To configure the same authentication service for multiple TACACS+ servers, include statements at the **[edit system tacplus-server]** and **[edit system tacplus-options]** hierarchy levels. For information about how to configure a TACACS+ server at the **[edit system tacplus-server]** hierarchy level.

To assign the same authentication service to multiple TACACS+ servers, include the **service-name** statement at the **[edit system tacplus-options]** hierarchy level:

```
[edit system tacplus-options]
service-name service-name;
```

**service-name** is the name of the authentication service. By default, the service name is set to **junos-exec**.

The following example shows how to configure the same authentication service for multiple TACACS+ servers:

```
[edit system]
tacplus-server {
  10.2.2.2 secret "$9$2dgoJGDiqP5ZG9A"; ## SECRET-DATA
  10.3.3.3 secret "$9$2dgoJGDiqP5ZG9A"; ## SECRET-DATA
}
tacplus-options {
  service-name bob;
}
```

### Configuring Juniper Networks Vendor-Specific TACACS+ Attributes

---

The Juniper Networks vendor-specific TACACS+ attributes enable you to configure access privileges for users on a TACACS+ server. They are specified in the TACACS+ server configuration file on a per-user basis. Junos OS retrieves these attributes through an authorization request of the TACACS+ server after authenticating a user. You do not need to configure these attributes to run Junos OS with TACACS+.

To specify these attributes, include a **service** statement of the following form in the TACACS+ server configuration file:

```
service = junos-exec {
  local-user-name = <username-local-to-router>
  allow-commands = "<allow-commands-regex>"
  allow-configuration = "<allow-configuration-regex>"
  deny-commands = "<deny-commands-regex>"
  deny-configuration = "<deny-configuration-regex>"
}
```

This **service** statement can appear in a **user** or **group** statement.

#### Related Documentation

- [Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands on page 1375](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1387](#)
- [Juniper Networks Vendor-Specific TACACS+ Attributes on page 1324](#)
- [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)
- [Junos OS User Authentication Methods on page 1334](#)

### Configuring TACACS+ System Accounting

You can use TACACS+ to track and log software logins, configuration changes, and interactive commands. To audit these events, include the following statements at the **[edit system accounting]** hierarchy level:

```
[edit system accounting]
events [ events ];
destination {
  tacplus {
    server {
```



```

server-address {
  port port-number;
  secret password;
  single-connection;
  timeout seconds;
}
}
}
}

```

Tasks for configuring TACACS+ system accounting are:

1. [Specifying TACACS+ Auditing and Accounting Events on page 1367](#)
2. [Configuring TACACS+ Server Accounting on page 1367](#)

### Specifying TACACS+ Auditing and Accounting Events

To specify the events you want to audit when using a TACACS+ server for authentication, include the **events** statement at the **[edit system accounting]** hierarchy level:

```

[edit system accounting]
events [ events ];

```

**events** is one or more of the following:

- **login**—Audit logins
- **change-log**—Audit configuration changes
- **interactive-commands**—Audit interactive commands (any command-line input)

### Configuring TACACS+ Server Accounting

To configure TACACS+ server accounting, include the **server** statement at the **[edit system accounting destination tacplus]** hierarchy level:

```

[edit system accounting destination tacplus]
server {
  server-address {
    port port-number;
    secret password;
    single-connection;
    timeout seconds;
  }
}

```

**server-address** specifies the address of the TACACS+ server. To configure multiple TACACS+ servers, include multiple **server** statements.



**NOTE:** If no TACACS+ servers are configured at the **[edit system accounting destination tacplus]** statement hierarchy level, Junos OS uses the TACACS+ servers configured at the **[edit system tacplus-server]** hierarchy level.

**port-number** specifies the TACACS+ server port number.

You must specify a secret (password) by using the **secret** statement. The local router or switch passes the **secret** to the TACACS+ client. If the password contains spaces, enclose the entire password in quotation marks (" "). The password used by the local router or switch must match that used by the server.

Optionally, you can specify the length of time that the local router or switch waits to receive a response from a TACACS+ server by including the **timeout** statement. By default, the router or switch waits 3 seconds. You can configure this to be a value in the range from 1 through 90 seconds.

Optionally, you can maintain one open TCP connection to the server for multiple requests, rather than opening a connection for each connection attempt, by including the **single-connection** statement.

To ensure that start and stop requests for accounting of login events are correctly logged in the Accounting file instead of the Administration log file on a TACACS+ server, include either the **no-cmd-attribute-value** statement or the **exclude-cmd-attribute** at the **[edit system tacplus-options]** hierarchy level.

If you use the **no-cmd-attribute-value** statement, the value of the **cmd** attribute is set to a null string in the start and stop requests. If you use the **exclude-cmd-attribute** statement, the **cmd** attribute is totally excluded from the start and stop requests. Both statements support the correct logging of accounting requests in the Accounting file, instead of the Administration file.

```
[edit system tacplus-options]
(no-cmd-attribute-value | exclude-cmd-attribute);
```

**Related Documentation** • [Configuring TACACS+ Authentication \(QFX Series\) on page 1364](#)

## Defining Junos OS Login Classes

To define a login class and its access privileges, include the **class** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
class class-name {
  allow-commands "regular-expression";
  allow-configuration "regular-expression";
  deny-commands "regular-expression";
  deny-configuration "regular-expression";
  idle-timeout minutes;
  permissions [ permissions ];
}
```

**Related Documentation** • [Junos OS Login Classes Overview on page 1337](#)  
• [Junos OS User Accounts Overview on page 1335](#)  
• [Example: Creating Login Classes with Specific Privileges on page 1389](#)  
• [Configuring the Junos OS to Display a System Login Announcement on page 163](#)

## Limiting the Number of User Login Attempts for SSH and Telnet Sessions

You can limit the number of times a user can attempt to enter a password while logging in through SSH or Telnet. The connection is terminated if a user fails to log in after the number of attempts specified. You can also specify a delay, in seconds, before a user can try to enter a password after a failed attempt. In addition, you can specify the threshold for the number of failed attempts before the user experiences a delay in being able to enter a password again.

To specify the number of times a user can attempt to enter a password while logging in, include the **retry-options** statement at the **[edit system login]** hierarchy level:

```
[edit system login]
retry-options {
  tries-before-disconnect number;
  backoff-threshold number;
  backoff-factor seconds;
  maximum-time seconds
  minimum-time seconds;
}
```

You can configure the following options:

- **tries-before-disconnect**—Number of times a user can attempt to enter a password when logging in. The connection closes if a user fails to log in after the number specified. The range is from 1 through 10, and the default is 10.
- **backoff-threshold**—Threshold for the number of failed login attempts before the user experiences a delay in being able to enter a password again. Use the **backoff-factor** option to specify the length of the delay in seconds. The range is from 1 through 3, and the default is 2.
- **backoff-factor**—Length of time, in seconds, before a user can attempt to log in after a failed attempt. The delay increases by the value specified for each subsequent attempt after the threshold. The range is from 5 through 10, and the default is 5 seconds.
- **maximum-time *seconds***—Maximum length of time, in seconds, that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the **maximum-time** value, the connection is closed. The range is from 20 through 300 seconds, and the default is 120 seconds.
- **minimum-time**—Minimum length of time, in seconds, that a connection remains open while a user is attempting to enter a correct password. The range is from 20 through 60, and the default is 40.

### Related Documentation

- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions on page 1391](#)
- [Configuring Junos OS User Accounts on page 1344](#)

## Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



**NOTE:** The root password cannot be recovered on a QFabric system.



**NOTE:** You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.

```
ok boot -s
```

12. At the following prompt, enter **recovery** to start the root password recovery procedure.

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN
for /bin/sh: recovery
```

13. Enter configuration mode in the CLI.

14. Set the root password. For example:

```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: juniper1
Retype new password:
```

16. At the second prompt, reenter the new root password.

17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.

19. Exit operational mode in the CLI.

20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

**Related Documentation**

- [Configuring the Root Password on page 1354](#)

## Specifying Access Privileges for Junos OS Configuration Mode Hierarchies

You can specify extended regular expressions with the **allow-configuration** and **deny-configuration** statements to define user access privileges to parts of the configuration hierarchy. Doing so overrides login class permission bits set for a user. You can also use wildcards to restrict access. When you define access privileges to parts of the configuration hierarchy, do the following tasks:

- Specify the full paths in the extended regular expressions with the **allow-configuration** and **deny-configuration** statements.
- Put parentheses around an extended regular expression that connects two or more expressions with the pipe **|** symbol. For example:

```
[edit system login class class-name]
user@switch# set deny-configuration "(system login class) | (system services)"
```



**NOTE:** Each expression separated by a pipe (**|**) symbol must be a complete standalone expression, and must be enclosed in parentheses (**()**). Do not use spaces between regular expressions separated with parentheses and connected with the pipe (**|**) symbol. You cannot define access to keywords such as **set**, **edit**, or **activate**.

When you explicitly provide access to configuration mode hierarchies or regular expressions using the **allow-configuration** statement, you add to the regular permissions set with the **permissions** statement. If you explicitly deny access to configuration mode hierarchies or regular expressions using the **deny-configuration** statement, you remove permissions for the specified configuration mode hierarchy from the default permissions provided by the **permissions** statement.

To explicitly provide access to an individual configuration mode hierarchy that would otherwise be denied, include the **allow-configuration** statement at the **[edit system login class *class-name*]** hierarchy level:

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

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

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

You can include one **deny-configuration** and one **allow-configuration** statement in each login class.

If you allow and deny the same set of configuration hierarchy levels, regular expressions, or commands, the **allow-configuration** statement permissions take precedence over the permissions specified by the **deny-configuration** statement. For example, if you include **allow-configuration "system services"** and **deny-configuration "system services"**, the login class user can continue to edit the configuration or issue commands at the **edit system services** hierarchy level.

#### Related Documentation

- [Defining Access Privileges Using allow/deny-configuration Statements on page 1390](#)
- [Configuring Access Privilege Levels on page 1344](#)

## Specifying Access Privileges for Junos OS Operational Mode Commands

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

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

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

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

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

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

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



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

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

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

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

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

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

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

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

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

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

**Related  
Documentation**

- [Example: Configuring Access Privileges for Operational Mode Commands on page 1381](#)
- [Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands on page 1339](#)
- *allow-commands*
- *deny-commands*

## Using Junos OS to Configure Logical System Administrators

Using Junos OS, you can partition a single router or switch into multiple logical devices that perform independent routing or switching tasks. When creating logical systems, you must configure logical system administrators and interfaces, assign logical interfaces to logical systems, and configure various other logical system statements.

The master administrator can assign one or more logical system administrators to each logical system. Once assigned to a logical system, administrators are restricted to viewing only configurations of the logical system to which they are assigned and accessing only the operational commands that apply to that particular logical system. This restriction means that these administrators cannot access global configuration statements, and all command output is restricted to the logical system to which the administrators are assigned.

To configure logical system administrators, include the **logical-system *logical-system-name*** statement at the **[edit system login class *class-name*]** hierarchy level and apply the class to the user. For example:

```
[edit]
system {
  login {
    class admin1 {
      permissions all;
      logical-system logical-system-LS1;
    }
    class admin2 {
      permissions view; # Gives users assigned to class admin2 the ability to view
                        # but not to change the configuration.
      logical-system logical-system-LS2;
    }
    user user1 {
      class admin1;
    }
    user user2 {
      class admin2;
    }
  }
}
```



```
}
}
```

Fully implementing logical systems requires that you also configure any protocols, routing statements, switching statements, and policy statements for the logical system.

**Related  
Documentation**

- *Defining Junos OS Login Classes*
- [Defining Junos OS Login Classes on page 1368](#)

## Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands

Use regular expressions to specify which operational or configuration mode commands are allowed or denied when you use a RADIUS or TACACS+ server for user authentication. You can specify the regular expressions using the appropriate Juniper Networks vendor-specific RADIUS or TACACS+ attributes in your authentication server configuration.

You can specify **allow-configuration**, **deny-configuration**, **allow-commands**, or **deny-commands** in a single extended regular expression, enclosing multiple commands in parentheses and separating them using the pipe symbol. For example, you can specify multiple **allow-commands** parameters using: **allow-commands= (cmd1 | cmd2 | cmdn)**. You can specify **user-permissions** as a list of comma-separated values, and not as a regular expression.

On a RADIUS or TACACS+ server, you can also use a simplified version for regular expressions where you specify each individual expression on a separate line. The simplified version is valid for **allow-commands**, **deny-commands**, **allow-configuration**, **deny-configuration**, and **permissions** vendor-specific attributes.

For a RADIUS server, specify the individual regular expressions using the following syntax:

```
Juniper-Allow-Commands+= "cmd1"
Juniper-Allow-Commands+= "cmd2"
Juniper-Allow-Commands+= "cmdn"
Juniper-Deny-Commands+= "cmd1"
Juniper-Deny-Commands+= "cmd2"
Juniper-Deny-Commands+= "cmdn"
Juniper-Allow-Configuration+= "regex1"
Juniper-Allow-Configuration+= "regex2"
Juniper-Allow-Configuration+= "regexn"
Juniper-Deny-Configuration+= "regex1"
Juniper-Deny-Configuration+= "regex2"
Juniper-Deny-Configuration+= "regexn"
Juniper-User-Permissions+= "permission-flag1"
Juniper-User-Permissions+= "permission-flag2"
Juniper-User-Permissions+= "permission-flagn"
```

For TACACS+ server, specify the individual regular expressions using the following syntax:

```
allow-commands1= "cmd1"
allow-commands2= "cmd2"
allow-commandsn= "cmdn"
deny-commands1= "cmd1"
```

```
deny-commands2="cmd2"  
deny-commandsn="cmdn"  
allow-configuration1="regex1"  
allow-configuration2="regex2"  
allow-configurationn="regexn"  
deny-configuration1="regex1"  
deny-configuration2="regex2"  
deny-configurationn="regexn"  
user-permissions1="permission-flag1"  
user-permissions2="permission-flag2"  
user-permissionsn="permission-flagn "
```

**NOTE:**

- Numeric values 1 to *n* in the syntax (for TACACS+ server) must be unique but need not be sequential. For example, the following syntax is valid:

```
allow-commands1="cmd1"  
allow-commands3="cmd3"  
allow-commands2="cmd2"  
deny-commands3="cmd3"  
deny-commands2="cmd2"  
deny-commands1="cmd1"
```

- The limit on the number of lines of individual regular expressions is imposed by the TACACS+ or RADIUS server.
- When you issue the `show cli authorization` command, the command output displays the regular expression in a single line, even if you specify each individual expression on a separate line.

---

For more information about Juniper Networks vendor-specific RADIUS and TACACS+ attributes, see [“Juniper Networks Vendor-Specific RADIUS Attributes” on page 1321](#) and [“Juniper Networks Vendor-Specific TACACS+ Attributes” on page 1324](#).



**NOTE:** When RADIUS or TACACS+ authentication is configured for a router, regular expressions configured on the RADIUS or TACACS+ server merge with any regular expressions configured on the local router at the `[edit system login class]` hierarchy level using the `allow-commands`, `deny-commands`, `allow-configuration`, `deny-configuration`, or `permissions` statements. If the final expression has a syntax error, the overall result is an invalid regular expression.

**Related Documentation**

- [Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1330](#)

## VSA Match Conditions and Actions

Devices support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs). They are configured on RADIUS servers and work in combination with 802.1X authentication. Using VSAs, you

can apply port firewall filter attributes as a subset of match conditions and actions sent from the RADIUS server to the switch as a result of successful 802.1X authentication.

Each term in a VSA configured through the RADIUS server consists of *match conditions* and an *action*. Match conditions are the values or fields that the packet must contain. You can define single, multiple, or no match conditions. If no match conditions are specified for the term, the packet is accepted by default. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Allowed actions are to accept a packet or to discard a packet.

The following guidelines apply when you specify match conditions and actions for VSAs:

- Both **match** and **action** statements are mandatory.
- Any or all options (separated by commas) may be included in each **match** and **action** statement.
- Fields separated by commas will be ANDed if they are of a different type. The same types cannot be repeated.
- For OR cases (for example, match 10.1.1.0/24 OR 11.1.1.0/24), apply multiple VSAs to the 802.1X supplicant.
- In order for the **forwarding-class** option to be applied, the forwarding class must be configured on the switch. If it is not configured on the switch, this option is ignored.

[Table 83 on page 1377](#) describes the match conditions you can specify when configuring a VSA using the **match** command on the RADIUS server. The string that defines a match condition is called a *match statement*.

**Table 83: Match Conditions**

Option	Description
<b>destination-mac</b> <i>mac-address</i>	Destination media access control (MAC) address of the packet.
<b>source-vlan</b> <i>source-vlan</i>	Name of the source VLAN.
<b>source-dot1q-tag</b> <i>tag</i>	Tag value in the 802.1Q header, in the range 0 through 4095.
<b>destination-ip</b> <i>ip-address</i>	Address of the final destination node.
<b>ip-protocol</b> <i>protocol-id</i>	IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms:  <b>ah</b> , <b>egp</b> (8), <b>esp</b> (50), <b>gre</b> (47), <b>icmp</b> (1), <b>igmp</b> (2), <b>ipip</b> (4), <b>ipv6</b> (41), <b>ospf</b> (89), <b>pim</b> (103), <b>rsvp</b> (46), <b>tcp</b> (6), or <b>udp</b> (17)
<b>source-port</b> <i>port</i>	TCP or User Datagram Protocol (UDP) source port field. Normally, you specify this match statement in conjunction with the <b>ip-protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text options listed under <b>destination-port</b> .

Table 83: Match Conditions (*continued*)

Option	Description
<b>destination-port</b> <i>port</i>	<p>TCP or UDP destination port field. Normally, you specify this match in conjunction with the <b>ip-protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed):</p> <p>afs (1483), bgp (179), biff (512), bootpc (68), bootps (67), cvspserver (2401), cmd (514), dhcp (67), domain (53), eklogin (2105), ekshell (2106), exec (512), finger (79), ftp (21), ftp-data (20), http (80), https (443), ident (113), imap (143), kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544), ldap (389), login (513), mobileip-agent (434), mobilip-mn (435), msdp (639), netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123), pop3 (110), pptp (1723), printer (515), radacct (1813), radius (1812), rip (520), rkinit (2108), smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514), telnet (23), tacacs-ds (65), talk (517), tftp (69), timed (525), who (513), xdmcp (177), zephyr-clt (2103), zephyr-hm (2104)</p>

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria. [Table 84 on page 1378](#) shows the actions that you can specify in a term.

Table 84: Actions for VSAs

Option	Description
(allow   deny)	Accept a packet or discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.
<b>forwarding-class</b> <i>class-of-service</i>	<p>(Optional) Classify the packet in one of the following forwarding classes:</p> <ul style="list-style-type: none"> <li>assured-forwarding</li> <li>best-effort</li> <li>expedited-forwarding</li> <li>network-control</li> </ul>
<b>loss-priority</b> (low   medium   high)	(Optional) Set the packet loss priority (PLP) to <b>low</b> , <b>medium</b> , or <b>high</b> . Specify both the forwarding class and loss priority.

#### Related Documentation

- [Filtering 802.1X Supplicants Using RADIUS Server Attributes](#)
- [Understanding 802.1X and VSAs on EX Series Switches](#)
- [Understanding VSAs on page 1321](#)

## Configuration Examples

- [Example: Changing the Requirements for Junos OS Plain-Text Passwords on page 1379](#)
- [Example: Configuring Access Privilege Levels on page 1381](#)
- [Example: Configuring Access Privileges for Operational Mode Commands on page 1381](#)

- [Example: Configuring a Plain-Text Password for Root Logins on page 1382](#)
- [Example: Configuring RADIUS Authentication on page 1384](#)
- [Example: Configuring RADIUS System Accounting on page 1385](#)
- [Example: Configuring the Root Password on page 1385](#)
- [Example: Configuring SSH Authentication for Root Logins on page 1386](#)
- [Example: Configuring User Accounts on page 1386](#)
- [Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication on page 1387](#)
- [Example: Creating Login Classes with Specific Privileges on page 1389](#)
- [Example: Configuring User Login Accounts on page 1389](#)
- [Example: Configuring RADIUS Template Accounts on page 1390](#)
- [Defining Access Privileges Using allow/deny-configuration Statements on page 1390](#)
- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions on page 1391](#)

## Example: Changing the Requirements for Junos OS Plain-Text Passwords

This example shows how to set various maximum and minimum requirements for plain-text passwords to increase password strength.

- [Requirements on page 1379](#)
- [Overview on page 1379](#)
- [Configuration on page 1379](#)

### Requirements

This example requires a device running Junos 12.2 or greater. The **minimum-length** and **maximum-length** password requirements statements are available in earlier releases, however, you must have Junos OS Release 12.2 or greater to configure **minimum-lower-cases**, **minimum-numeric**, **minimum-punctuations**, or **minimum-upper-cases**.

### Overview

You can use a variety of requirements to strengthen plain-text passwords for greater security. Junos OS provides a number of possible configurations at the **[edit system login password]** hierarchy level that allow you to require users to create plain-text passwords that conform to a particular set of requirements that may include such things as length, number of changes, type of characters, numbers, or letter case.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system login password minimum-length 12
```

```
set system login password maximum-length 22
```

```
set system login password minimum-numeric 1
```

```
set system login password minimum-upper-cases 1
```

```
set system login password minimum-lower-cases 1
```

```
set system login password minimum-punctuations 1
```

### *Configuring Requirements for Plain-Text Passwords*

**Step-by-Step Procedure** This example configures password requirements that require the user to create a password that has a minimum length of 12 characters, a maximum length of 22 characters, and that includes at least one lower-case letter, at least one upper-case letter, at least one punctuation character, and at least one numeric character.

1. Navigate to configuration mode in the [system login password] hierarchy level.  

```
user@host> edit
[edit]
user@host# edit system login password
```
2. Set a minimum length requirement of 12 characters and a maximum length requirement of 22 characters for user passwords.  

```
[edit system login password]
user@host# set minimum-length 12
[edit system login password]
user@host# set maximum-length 22
```
3. Require users to set a password that has at least one lower-case letter and at least one upper-case letter.  

```
[edit system login password]
user@host# set minimum-lower-cases 1
[edit system login password]
user@host# set minimum-upper-cases 1
```
4. Require users to set a password that has at least one punctuation-class character and at least one number.  

```
[edit system login password]
user@host# set minimum-punctuations 1
[edit system login password]
user@host# set minimum-numeric 1
```

### *Results*

From configuration mode, confirm your configuration by entering the show command at the edit system login password hierarchy level. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit system login password]
user@host# show
  minimum-length 12;
  maximum-length 22;
```

```

minimum-numeric 1;
minimum-upper-cases 1;
minimum-lower-cases 1;

```

- Related Documentation**
- [Special Requirements for Junos OS Plain-Text Passwords on page 1339](#)
  - *password (Login)*

## Example: Configuring Access Privilege Levels

Create two access privilege classes on the router or switch, one for configuring and viewing user accounts only and the second for configuring and viewing SNMP parameters only:

```

[edit]
system {
  login {
    class user-accounts {
      permissions [ configure admin admin-control ];
    }
    class network-mgmt {
      permissions [ configure snmp snmp-control ];
    }
  }
}

```

- Related Documentation**
- [Configuring Access Privilege Levels on page 1344](#)

## Example: Configuring Access Privileges for Operational Mode Commands

The following example shows how to configure access privileges for different login classes for individual operational mode commands:

```

[edit]
system {
  # This login class has operator privileges and the additional ability
  # to reboot the router.
  login {
    # This login class has operator privileges and the additional ability to reboot the
    # router or switch.
    class operator-and-boot {
      permissions [ clear network reset trace view ];
      allow-commands "request system reboot";
    }
    # This login class has operator privileges but can't use any commands beginning
    # with "set".
    # This login class has operator privileges
    # but cannot use any commands beginning with "set"
    class operator-no-set {
      permissions [ clear network reset trace view ];
      deny-commands "^set";
    }
    # This login class has operator privileges and can install software but not view
    # BGP information, and can issue the show route command, without specifying
    # commands or arguments under it.

```

```
class operator-and-install-but-no-bgp {  
  permissions [ clear network reset trace view ];  
  allow-commands "(request system software add)|(show route$)";  
  deny-commands "show bgp";  
}  
}  
}
```

**Related  
Documentation**

- [Specifying Access Privileges for Junos OS Operational Mode Commands on page 192](#)

## Example: Configuring a Plain-Text Password for Root Logins

This example shows how to configure the authentication methods for the root-level user, whose username is “root”.

- [Requirements on page 1382](#)
- [Overview on page 1382](#)
- [Configuration on page 1382](#)
- [Verification on page 1383](#)

### Requirements

---

No special configuration beyond device initialization is required before configuring this example.

Make sure you understand the requirements for a valid plain-text password. For Junos OS, the The default requirements for plain-text passwords are as follow:

- The password must be between 6 and 128 characters long.
- You can include most character classes in a password (uppercase letters, lowercase letters, numbers, punctuation marks, and other special characters). Control characters are not recommended.
- Valid passwords must contain at least one change of case or character class.

### Overview

---

Junos OS is preinstalled on the router. When the router is powered on, it is ready to be configured. Initially, you log in as the user “root” with no password. To set the root password, you have several options. This example shows you how to enter a plain-text password that Junos OS then encrypts for you.

### Configuration

---

**CLI Quick  
Configuration**

```
[edit system]  
set root-authentication plain-text-password  
New password: new-password  
Retype new password: new-password
```



**Configuring [item]**

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

To configure a plain-text password:

1. Type the set command for plain-text password and press Enter.  

```
[edit]
user@host# set system root-authentication plain-text-password
New password:
```
2. Type the new password next to the **New password:** prompt and press Enter.  

```
user@host# new-password
Retype new password:
```
3. Retype the same password next to the next prompt and press Enter.

**Results**

From configuration mode, confirm your configuration by entering the **show** command. It should look something like this:

```
root-authentication {
  encrypted-password "$1$ASwBkGYd$YUcEwgd0IO4QkRzzlQdmT/"; ## SECRET-DATA
}
```

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

After you have confirmed that the interfaces are configured, enter the **commit** command in configuration mode.

**Verification**

- [Verifying the Configuration of a Plain-Text Password for Root Logins on page 1383](#)

**Verifying the Configuration of a Plain-Text Password for Root Logins**

**Purpose** Verify the configuration of a plain-text password.

**Action** From operational mode, confirm your configuration by entering the **show configuration system** command.

```
user@host> show configuration system
root-authentication {
  encrypted-password "$1$ASwBkGYd$YUcEwgd0IO4QkRzzlQdmT/"; ## SECRET-DATA
}
```

**Meaning** If you use a clear-text password, Junos OS displays the password as an encrypted string so that users viewing the configuration cannot see it. As you enter the password in plain text, Junos OS encrypts it immediately. You do not have to configure Junos OS to encrypt

the password as in some other systems. Plain-text passwords are hidden and marked as ## SECRET-DATA in the configuration.

**Related  
Documentation**

- *root-authentication*
- [Special Requirements for Junos OS Plain-Text Passwords on page 1339](#)
- *Configuring Special Requirements for Plain-Text Passwords*
- *Changing the Requirements for Junos OS Plain-Text Passwords*

## Example: Configuring RADIUS Authentication

The Junos OS supports two protocols for central authentication of users on multiple routers: RADIUS and TACACS+. We recommend RADIUS because it is a multivendor IETF standard, and its features are more widely accepted than those of TACACS+ or other proprietary systems. In addition, we recommend using a one-time-password system for increased security, and all vendors of these systems support RADIUS.

The Junos OS uses one or more template accounts to perform user authentication. You create the template account or accounts, and then configure the user access to use that account. If the RADIUS server is unavailable, the fallback is for the login process to use the local account that set up on the router or switch.

The following example shows how to configure RADIUS authentication:

```
[edit]
system {
  authentication-order [ radius password ];
  root-authentication {
    encrypted-password "$9$aH1j8gqQ1gjyjjhgjgiiii"; # SECRET-DATA
  }
  name-server {
    10.1.1.1;
    10.1.1.2;
  }
}
```

The following example shows how to enable RADIUS authentication and define the shared secret between the client and the server. The secret enables the client and server to determine that they are talking to the trusted peer.

Define a timeout value for each server, so that if there is no response within the specified number of seconds, the router can try either the next server or the next authentication mechanism.

```
[edit]
system {
  radius-server {
    10.1.2.1 {
      secret "$9$aH1j8gqQ1sdjerrhser"; # SECRET-DATA
      timeout 5;
    }
    10.1.2.2 {
```

```

        secret "$9$aH1j8gqQ1csdoiuardwefoiud"; # SECRET-DATA
        timeout 5;
    }
}

```

**Related Documentation**

- [Configuring RADIUS Authentication](#)

## Example: Configuring RADIUS System Accounting

The following example shows three servers (10.5.5.5, 10.6.6.6, and 10.7.7.7) configured for RADIUS accounting.

```

system {
  accounting {
    events [ login change-log interactive-commands ];
    destination {
      radius {
        server {
          10.5.5.5 {
            accounting-port 3333;
            secret $9$dkafeqwrew;
            source-address 10.1.1.1;
            retry 3;
            timeout 3;
          }
          10.6.6.6 secret $9$fe3erqwrez;
          10.7.7.7 secret $9$f34929ftby;
        }
      }
    }
  }
}

```

**Related Documentation**

- [Configuring RADIUS System Accounting on page 1349](#)

## Example: Configuring the Root Password

The following example shows how to configure the root password:

```

[edit]
user@switch# set system root-authentication encrypted-password
"$1$14c5.$sBopasddsdfs0"
[edit]
user@switch# show
system {
  root-authentication {
    encrypted-password "$1$14c5.$sBopasddsdfs0";
  }
}

```

**Related Documentation**

- [Configuring the Root Password on page 170](#)
- [Example: Configuring a Plain-Text Password for Root Logins on page 1382](#)

- [Configuring the Root Password on page 1354](#)

## Example: Configuring SSH Authentication for Root Logins

The following example shows how to configure two public DSA keys for SSH authentication of root logins:

```
[edit system]
root-authentication {
  encrypted-password "$1$1wp5tqMX$uy/u5H7OdXTwfWTmeJWXe/";
  ## SECRET-Data;
  ssh-dsa "2354 95 9304@boojum.per";
  ssh-dsa "0483 02 8362@ecbatana.per";
}
```

### Related Documentation

- [Configuring the Root Password on page 170](#)
- [Special Requirements for Junos OS Plain-Text Passwords on page 1339](#)

## Example: Configuring User Accounts

The following example shows how to create accounts for four router or switch users, and create an account for the template user **remote**. All users use one of the default system login classes. User **alexander** also has two digital signal algorithm (DSA) public keys configured for SSH authentication.

```
[edit]
system {
  login {
    user philip {
      full-name "Philip of Macedonia";
      uid 1001;
      class super-user;
      authentication {
        encrypted-password "$1$poPPeY";
      }
    }
    user alexander {
      full-name "Alexander the Great";
      uid 1002;
      class view;
      authentication {
        encrypted-password "$1$14c5.$sBopasdFFdssdfFFdsdfs0";
        ssh-dsa "8924 37 5678 5678@gaugamela.per";
        ssh-dsa "6273 94 9283@boojum.per";
      }
    }
    user darius {
      full-name "Darius King of Persia";
      uid 1003;
      class operator;
      authentication {
        ssh-rsa "1024 37 12341234@ecbatana.per";
      }
    }
  }
}
```

```

    }
    user anonymous {
        class unauthorized;
    }
    user remote {
        full-name "All remote users";
        uid 9999;
        class read-only;
    }
}
}

```

- Related Documentation**
- *Junos OS User Accounts Overview*
  - *Limiting the Number of User Login Attempts for SSH and Telnet Sessions*

### Example: Configuring System Authentication for RADIUS, TACACS+, and Password Authentication

The following example shows how to configure system authentication for RADIUS, TACACS+, and password authentication.

In this example, only the user Philip and users authenticated by a remote RADIUS server can log in. If a user logs in and is not authenticated by the RADIUS server, the user is denied access to the router or switch. If the RADIUS server is not available, the user is authenticated using the **password** authentication method and allowed access to the router or switch. For more information about the password authentication method, see [“Using Local Password Authentication” on page 1331](#).

When Philip tries to log in to the system, if the RADIUS server authenticates him, he is given access and privileges for the **super-user** class. Local accounts are not configured for other users. When they log in to the system and the RADIUS server authenticates them, they are given access using the same user ID (UID) 9999 and the privileges associated with the **operator** class.

```

[edit]
system {
    authentication-order radius;
    login {
        user philip {
            full-name "Philip";
            uid 1001;
            class super-user;
        }
        user remote {
            full-name "All remote users";
            uid 9999;
            class operator;
        }
    }
}
}

```



**NOTE:** For authorization purposes, you can use a template account to create a single account that can be shared by a set of users at the same time. For example, when you create a remote template account, a set of remote users can concurrently share a single UID. For more information about template accounts, see [“Overview of Template Accounts for RADIUS and TACACS+ Authentication”](#) on page 1318.

When a user logs in to a device, the user's login name is used by the RADIUS or TACACS+ server for authentication. If the user is authenticated successfully by the authentication server and the user is not configured at the `[edit system login user]` hierarchy level, the device uses the default remote template user account for the user, provided a remote template account is configured at the `edit system login user remote` hierarchy level. The remote template account serves as a default template user account for all users that are authenticated by the authentication server but not having a locally configured user account on the device. Such users share the same login class and UID.

To configure an alternate template user, specify the **user-name** parameter returned in the RADIUS authentication response packet. Not all RADIUS servers allow you to change this parameter. The following shows a sample Junos OS configuration:

```
[edit]
system {
  authentication-order radius;
  login {
    user philip {
      full-name "Philip";
      uid 1001;
      class super-user;
    }
    user operator {
      full-name "All operators";
      uid 9990;
      class operator;
    }
    user remote {
      full-name "All remote users";
      uid 9999;
      class read-only;
    }
  }
}
```

Assume your RADIUS server is configured with the following information:

- User Philip with password “olympia”
- User Alexander with password “bucephalus” and username “operator”
- User Darius with password “redhead” and username “operator”
- User Roxane with password “athena”

Philip would be given access as a superuser (**super-user**) because he has his own local user account. Alexander and Darius share UID 9990 and have access as operators. Roxane has no template-user override, so she shares access with all the other remote users, getting read-only access.

- Related Documentation**
- [Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication on page 1346](#)

## Example: Creating Login Classes with Specific Privileges

The following example shows how to create several user classes, each with specific privileges. In this example, you configure timeouts to disconnect the class members after a period of inactivity. Users' privilege levels, and therefore the classes of which they are members, should be dependent on their responsibilities within the organization, and the permissions shown here are only examples.

The first class of users (called "observation") can only view statistics and configuration. They are not allowed to modify any configuration. The second class of users (called "operation") can view and modify the configuration. The third class of users (called "engineering") has unlimited access and control.

```
[edit]
system {
  login {
    class observation {
      idle-timeout 5;
      permissions [ view ];
    }
    class operation {
      idle-timeout 5;
      permissions [ admin clear configure interface interface-control network
        reset routing routing-control snmp snmp-control trace-control
        firewall-control rollback ];
    }
    class engineering {
      idle-timeout 5;
      permissions all;
    }
  }
}
```

- Related Documentation**
- [Defining Junos OS Login Classes](#)

## Example: Configuring User Login Accounts

The following example shows how to configure the local administrator account (**user admin**). If RADIUS fails or becomes unreachable, the login process reverts to password authentication on the local accounts on the router or switch.

```
[edit]
system {
  login {
```

```
user admin {  
  uid 1000;  
  class engineering;  
  authentication {  
    encrypted-password "<PASSWORD>"; # SECRET-DATA  
  }  
}  
}
```

**Related Documentation** • [Configuring Junos OS User Accounts](#)

### Example: Configuring RADIUS Template Accounts

The following example shows how to configure RADIUS template accounts for different users or groups of users:

```
[edit]  
system {  
  login {  
    user observation {  
      uid 1001;  
      class observation;  
    }  
    user operation {  
      uid 1002;  
      class operation;  
    }  
    user engineering {  
      uid 1003;  
      class engineering;  
    }  
  }  
}
```

**Related Documentation** • [Overview of Template Accounts for RADIUS and TACACS+ Authentication on page 1318](#)

### Defining Access Privileges Using allow/deny-configuration Statements

The following examples show how to configure access privileges for individual configuration mode hierarchy levels.

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot configure telnet parameters:

```
[edit system login class class-name]  
user@switch# set deny-configuration "system services telnet"
```

If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot issue login class commands within any login class whose name begins with "m":

```
[edit system login class class-name]  
user@switch# set deny-configuration "system login class m.*"
```



If the following statement is included in the configuration and the user's login class permission bit is set to **all**, the user cannot edit a configuration or issue commands (such as **commit**) at the login class or system services hierarchy levels:

```
[edit system login class class-name]
user@switch# set deny-configuration "(system login class) | (system services)"
```

The following example shows how to configure permissions for individual configuration mode hierarchies:

```
[edit]
system {
  login { # This login class has operator privileges and the additional ability to edit
    # configuration at the system services hierarchy level.
    class only-system-services {
      permissions [ configure ];
      allow-configuration "system services";
    }
    # services commands.
    class all-except-system-services { # This login class has operator privileges but
      # cannot edit any system services configuration.
      permissions [ all ];
      deny-configuration "system services";
    }
  }
}
```

**Related  
Documentation**

- *Specifying Access Privileges Using allow/deny-configuration Statements*
- [Specifying Access Privileges for Junos OS Configuration Mode Hierarchies on page 1371](#)

## Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions

The following example shows how to limit the user to four attempts when the user enters a password while logging in through SSH or Telnet. Set the **backoff-threshold** to 2, the **back-off-factor** to 5 seconds, and the **minimum-time** to 40 seconds. The user experiences a delay of 5 seconds after the second attempt to enter a correct password fails. After each subsequent failed attempt, the delay increases by 5 seconds. After the fourth and final failed attempt to enter a correct password, the user experiences an additional 10-second delay, and the connection closes after a total of 40 seconds.

The additional variables **maximum-time** and **lockout-period** are not set in this example.

```
[edit]
system {
  login {
    retry-options {
      backoff-threshold 2;
      backoff-factor 5;
      minimum-time 40;
      tries-before-disconnect 4;
    }
    password {
    }
  }
}
```

}



**NOTE:** This sample only shows the portion off the [edit system login] hierarchy level being modified.

**Related  
Documentation**

- *Limiting the Number of User Login Attempts for SSH and Telnet Sessions*
- *login*
- [login on page 280](#)

---

## Configuration Statements

---

- [access on page 1394](#)
- [accounting \(Access Profile\) on page 1395](#)
- [accounting-options on page 1396](#)
- [accounting-server on page 1398](#)
- [accounting-stop-on-access-deny on page 1399](#)
- [accounting-stop-on-failure on page 1400](#)
- [advertisement-interval on page 1401](#)
- [agent-address on page 1402](#)
- [archival on page 1403](#)
- [archive-sites \(Configuration File\) on page 1404](#)
- [authentication-order on page 1405](#)
- [authentication-server on page 1406](#)
- [authorization on page 1407](#)
- [categories on page 1408](#)
- [client-list on page 1408](#)
- [client-list-name on page 1409](#)
- [clients on page 1409](#)
- [commit-delay on page 1410](#)
- [community \(SNMP\) on page 1411](#)
- [configuration on page 1412](#)
- [connection-limit on page 1413](#)
- [contact on page 1414](#)
- [disable \(LLDP\) on page 1414](#)
- [falling-threshold \(Health Monitor\) on page 1415](#)
- [filter-duplicates on page 1415](#)
- [full-name on page 1416](#)
- [health-monitor on page 1416](#)

- [hold-multiplier](#) on page 1417
- [idle-timeout \(Access\)](#) on page 1418
- [interface \(LLDP\)](#) on page 1419
- [interval \(Health Monitor\)](#) on page 1420
- [lldp](#) on page 1421
- [lldp-configuration-notification-interval](#) on page 1422
- [location](#) on page 1423
- [management-address](#) on page 1424
- [name](#) on page 1425
- [nas-ip-address](#) on page 1425
- [nonvolatile](#) on page 1426
- [oid](#) on page 1426
- [order](#) on page 1427
- [port \(RADIUS Server\)](#) on page 1428
- [profile](#) on page 1429
- [protocols](#) on page 1430
- [protocol-version](#) on page 1443
- [ptopo-configuration-maximum-hold-time](#) on page 1443
- [ptopo-configuration-trap-interval](#) on page 1444
- [radius](#) on page 1445
- [radius-options \(edit system\)](#) on page 1446
- [radius-server](#) on page 1447
- [rate-limit](#) on page 1448
- [remote-debug-permission](#) on page 1449
- [retry](#) on page 1450
- [rising-threshold \(Health Monitor\)](#) on page 1451
- [root-login](#) on page 1452
- [services \(Switches\)](#) on page 1453
- [snmp](#) on page 1454
- [ssh](#) on page 1458
- [system](#) on page 1459
- [tacplus-options](#) on page 1465
- [targets](#) on page 1466
- [traceoptions \(LLDP\)](#) on page 1467
- [transfer-interval \(Configuration\)](#) on page 1469
- [transfer-on-commit](#) on page 1470
- [trap-group](#) on page 1471

- [trap-options](#) on page 1472
- [user \(Access\)](#) on page 1473
- [version](#) on page 1474

## access

---

**Syntax**    access {  
              address-assignment  
              pool *pool-name*  
              address-pool*pool-name*  
              profile *profile-name* {  
                  accounting (Access Profile) {  
                    accounting-stop-on-access-deny;  
                    accounting-stop-on-failure;  
                    (authentication-order (ldap radius | none);  
                    order (radius | none);  
                  }  
                  radius {  
                    accounting-server [*server-addresses*];  
                    authentication-server [*server-addresses*];  
                  }  
              }  
              }

**Hierarchy Level**    [edit]

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure authentication, authorization, and accounting (AAA) services.  
  
                      The statements are explained separately.



**NOTE:** The [edit access] hierarchy is not available on QFabric systems.

---

**Default**    Not enabled

**Required Privilege Level**    admin—To view this statement in the configuration.  
                                  admin-control—To add this statement to the configuration.

**Related Documentation**    • [Configuring 802.1X RADIUS Accounting \(CLI Procedure\)](#)

## accounting (Access Profile)

<b>Syntax</b>	<pre>accounting {   accounting-stop-on-access-deny;   accounting-stop-on-failure;   order (radius   none); }</pre>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the authentication order for authentication, authorization, and accounting (AAA) services.
<b>Default</b>	Not enabled
<b>Options</b>	<p><b>none</b>—Use no authentication for specified subscribers.</p> <p><b>radius</b>—Use RADIUS authentication for specified subscribers.</p> <p>The remaining statements are explained separately.</p>



**NOTE:** The [edit access] hierarchy is not available on QFabric systems.

<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li> <li>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i></li> <li>• <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i></li> <li>• <i>Configuring RADIUS Accounting</i></li> <li>• <a href="#">Understanding RADIUS Accounting on page 1320</a></li> </ul>

## accounting-options

---

```
Syntax  accounting-options {
        class-usage-profile profile-name {
            destination-classes {
                destination-class-name;
            }
            file filename;
            interval minutes;
            source-classes {
                source-class-name;
            }
        }
        file filename {
            archive-sites {
                site-name;
            }
            files number;
            nonpersistent;
            size bytes;
            start-time time;
            transfer-interval minutes;
        }
        filter-profile profile-name {
            counters {
                counter-name;
            }
            file filename;
            interval minutes;
        }
        interface-profile profile-name {
            fields {
                input-bytes;
                input-errors;
                input-multicast;
                input-packets;
                input-unicast;
                output-bytes;
                output-errors;
                output-multicast;
                output-packets;
                output-unicast;
                rpf-check-bytes;
                rpf-check-packets;
                rpf-check6-bytes;
                rpf-check6-packets;
                unsupported-protocol;
            }
            file filename;
            interval minutes;
        }
        mib-profile profile-name {
            file filename;
            interval minutes;
        }
    }
```

```

object-names {
    mib-object-name;
}
operation (get | get-next | walk);
}
policy-decision-statistics-profile profile-name {
    application-aware-access-list-fields {
        address;
        application;
        application-group;
        input-bytes;
        input-interface;
        input-packets;
        mask;
        output-bytes;
        output-packets;
        subscriber-name;
        timestamp;
        vrf-name;
    }
    file filename;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}

```

<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure options for accounting statistics collection.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding RADIUS Accounting on page 1320</a></li> <li>• <a href="#">Understanding VSAs on page 1321</a></li> <li>• <a href="#">Configuring RADIUS System Accounting on page 1349</a></li> <li>• <a href="#">Configuring Remote Template Accounts for User Authentication on page 1354</a></li> <li>• <a href="#">Configuring Local User Template Accounts for User Authentication on page 1347</a></li> </ul>

## accounting-server

---

<b>Syntax</b>	accounting-server[ <i>server-addresses</i> ];
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the Remote Authentication Dial-In User Service (RADIUS) server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
<b>Default</b>	Not enabled
<b>Options</b>	<i>server-addresses</i> —One or more addresses of RADIUS authentication servers.



**NOTE:** The [edit access] hierarchy is not available on QFabric systems.

---

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>show network-access aaa statistics authentication</i></li><li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li><li>• <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i></li><li>• <a href="#">Understanding RADIUS Accounting on page 1320</a></li></ul>



## accounting-stop-on-access-deny

<b>Syntax</b>	accounting-stop-on-access-deny;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> accounting]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if the AAA server denies access to a supplicant.




**NOTE:** The [edit access] hierarchy is not available on QFabric systems.


<b>Default</b>	Not enabled
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li> <li>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i></li> <li>• <i>show network-access aaa statistics authentication</i></li> <li>• <i>Configuring RADIUS Accounting</i></li> </ul>

## accounting-stop-on-failure

---

<b>Syntax</b>	accounting-stop-on-failure;
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> accounting]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if a supplicant fails AAA authorization, but the RADIUS server grants access. For example, a supplicant might fail AAA authentication because of an internal error such as a timeout.
<div> <b>NOTE:</b> The [edit access] hierarchy is not available on QFabric systems.</div>	
<b>Default</b>	Not enabled
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li><li>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i></li><li>• <i>Understanding 802.1X and RADIUS Accounting on EX Series Switches</i></li><li>• <i>Configuring RADIUS Accounting</i></li><li>• <a href="#">Understanding RADIUS Accounting on page 1320</a></li></ul>

## advertisement-interval



<b>Syntax</b>	<code>advertisement-interval seconds;</code>
<b>Hierarchy Level</b>	[edit protocols lldp], [edit routing-instances <i>routing-instance-name</i> protocols lldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for MX Series and T Series routers. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>For MX Series and T Series routers and EX Series switches, configure an interval for LLDP advertisement.</p> <p>For switches configured for Link Layer Discovery Protocol, configure the frequency at which LLDP advertisements are sent.</p> <p>The <b>advertisement-interval</b> value must be greater than or equal to four times the <b>transmit-delay</b> value, or an error will be returned when you attempt to commit the configuration.</p>
	<div>  <p><b>NOTE:</b> The default value of <b>transmit-delay</b> is 2 seconds. If you configure the <b>advertisement-interval</b> as less than 8 seconds and you do not configure a value for <b>transmit-delay</b>, the default value of <b>transmit-delay</b> is automatically changed to 1 second in order to satisfy the requirement that the <b>advertisement-interval</b> value must be greater than or equal to four times the <b>transmit-delay</b> value.</p> </div>
<b>Default</b>	Disabled.
<b>Options</b>	<p><b>seconds</b>—Interval between LLDP advertisement.</p> <p><b>Default:</b> 30</p> <p><b>Range:</b> 5 through 32768</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring LLDP</i></li> <li>• <a href="#">show lldp on page 1486</a></li> <li>• <i>Configuring LLDP (CLI Procedure)</i></li> <li>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i></li> <li>• <i>transmit-delay</i></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>

## agent-address



---

<b>Syntax</b>	agent-address outgoing-interface;
<b>Hierarchy Level</b>	[edit snmp trap-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is <b>outgoing-interface</b> , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
<b>Options</b>	<b>outgoing-interface</b> —Value of the agent address of all SNMPv1 traps generated by this router or switch. The <b>outgoing-interface</b> option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. <b>Default:</b> Disabled (the agent address is not specified in SNMPv1 traps).
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Agent Address for SNMP Traps</i></li></ul>

## archival

<b>Syntax</b>	<pre> archival {   configuration {     archive-sites {       file://&lt;path&gt;/&lt;filename&gt;;       ftp://username@host:&lt;port&gt;url-path password password;       http://username@host:&lt;port&gt;url-path password password;       pasvftp://username@host:&lt;port&gt;url-path password password;       scp://username@host:&lt;port&gt;url-path password password;     }     transfer-interval interval;     transfer-on-commit;   } } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure copying of the currently active configuration to an archive site. An archive site can be a file, or an FTP or SCP location.
<div>  <b>NOTE:</b> The <code>edit system archival</code> hierarchy is not available on QFabric systems. </div>	
<b>Options</b>	The remaining statements are explained separately.
<div>  <b>NOTE:</b> The <code>[edit system archival]</code> hierarchy is not available on QFabric systems. </div>	
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1263</li> </ul>

## archive-sites (Configuration File)

<b>Syntax</b>	<pre>archive-sites {     file://&lt;path&gt;/&lt;filename&gt;;     ftp://username@host:&lt;port&gt;url-path password password;     http://username@host:&lt;port&gt;url-path password password;     pasvftp://username@host:&lt;port&gt;url-path password password;     scp://username@host:&lt;port&gt;url-path password password; }</pre>
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks ( " ") and enclose the IPv6 host address in brackets ( [ ] ). For example,</p> <pre>"scp://username&lt;:password&gt;@[ipv6-host-address]&lt;:port&gt;/url-path"</pre> <p>If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails.</p> <p>The destination filename is saved in the following format, where <i>n</i> corresponds to the number of the compressed configuration rollback file that has been archived:</p> <pre>router-name_juniper.conf.n.gz_YYYYMMDD_HHMMSS.</pre> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> The [edit system archival] hierarchy is not available on QFabric systems.</p> </div>
<b>Options</b>	<p>The prefix used in the configuration statement determines the form of transfer:</p> <p><b>file://</b> —transfer on a path to a named file</p> <p><b>ftp://</b> —transfer using active FTP server</p> <p><b>pasvftp://</b> —transfer to a device that only accepts passive FTP services</p>

**scp://** —transfer to a known host using background SCP file transfers

<b>Required Privilege Level</b>	system—To view this statement in the configuration.
	system-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Configuring Archive Sites for Transfer of Active Configuration Files on page 1264</a>
	• <a href="#">Junos OS Commit Model for Router or Switch Configuration on page 14</a>
	• <a href="#">configuration on page 1269</a>
	• <a href="#">transfer-on-commit on page 1271</a>

## authentication-order

<b>Syntax</b>	authentication-order [none   password   radius];
<b>Hierarchy Level</b>	[edit <a href="#">access profile</a> <i>profile-name</i> ], [edit <a href="#">system</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending authentication messages.
<b>Default</b>	Not enabled
<b>Options</b>	<b>none</b> —No authentication for specified subscribers.
	<b>password</b> —Password authentication.
	<b>radius</b> —RADIUS authentication.



**NOTE:** The [edit access] hierarchy is not available on QFabric systems.

<b>Required Privilege Level</b>	admin—To view this statement in the configuration.
	admin-control—To add this statement to the configuration.

## authentication-server

---

<b>Syntax</b>	<code>authentication-server [server-addresses];</code>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> radius]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the RADIUS server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.
<b>Options</b>	<b>server-addresses</b> —Configure one or more RADIUS server addresses.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li><li>• <i>show network-access aaa statistics authentication</i></li></ul>



## authorization

---

<b>Syntax</b>	<code>authorization <i>authorization</i>;</code>
<b>Hierarchy Level</b>	[edit snmp community <i>community-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the access authorization for SNMP <b>Get</b> , <b>GetBulk</b> , <b>GetNext</b> , and <b>Set</b> requests.
<b>Options</b>	<p><b><i>authorization</i></b>—Access authorization level:</p> <ul style="list-style-type: none"> <li><b>read-only</b>—Enable <b>Get</b>, <b>GetNext</b>, and <b>GetBulk</b> requests.</li> <li><b>read-write</b>—Enable all requests, including <b>Set</b> requests. You must configure a view to enable <b>Set</b> requests.</li> </ul>
	<div>  <p><b>NOTE:</b> The read-write option is not supported on the QFX3000 QFabric system.</p> </div>
	<b>Default:</b> read-only
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring the SNMP Community String on page 6601</a></li> </ul>

## categories

---

<b>Syntax</b>	<pre>categories {     category; }</pre>
<b>Hierarchy Level</b>	[edit snmp trap-group <i>group-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the types of traps that are sent to the targets of the named trap group.
<b>Default</b>	If you omit the <b>categories</b> statement, all trap types are included in trap notifications.
<b>Options</b>	<b>category</b> —Name of a trap type: <b>authentication</b> , <b>chassis</b> , <b>configuration</b> , <b>link</b> , <b>remote-operations</b> , <b>rmon-alarm</b> , or <b>startup</b> .
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li></ul>

## client-list

---

<b>Syntax</b>	<pre>client-list <i>client-list-name</i> {     <i>ip-addresses</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define a list of SNMP clients.
<b>Options</b>	<b>client-list-name</b> —Name of the client list.  <b>ip-addresses</b> —IP addresses of the SNMP clients to be added to the client list,
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Adding a Group of Clients to an SNMP Community on page 6603</a></li></ul>

## client-list-name

---

<b>Syntax</b>	<code>client-list-name <i>client-list-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp community <i>community-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Add a client list or prefix list to an SNMP community.
<b>Options</b>	<i>client-list-name</i> —Name of the client list or prefix list.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Adding a Group of Clients to an SNMP Community on page 6603</a></li> </ul>

## clients

---


<b>Syntax</b>	<pre>clients {     address &lt;restrict&gt;; }</pre>
<b>Hierarchy Level</b>	<code>[edit snmp community <i>community-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
<b>Default</b>	If you omit the <b>clients</b> statement, all SNMP clients using this community string are authorized to access the switch.
<b>Options</b>	<p><b>address</b>—Address of an SNMP client that is authorized to access this switch. You must specify an address, not a hostname. To specify more than one client, include multiple <b>address</b> options.</p> <p><b>restrict</b>—(Optional) Do not allow the specified SNMP client to access the switch.</p>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Communities</a></li> </ul>

## commit-delay

---

<b>Syntax</b>	commit-delay <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit snmp nonvolatile]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the timer for the SNMP <b>Set</b> reply and start of the commit.
<b>Options</b>	<b>seconds</b> —Delay between an affirmative SNMP <b>Set</b> reply and start of the commit operation. <b>Default:</b> 5 seconds
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Commit Delay Timer</i></li></ul>

## community (SNMP)

<b>Syntax</b>	<pre>community <i>community-name</i> {   authorization <i>authorization</i>;   client-list-name <i>client-list-name</i>;   clients {     address restrict;   }   view <i>view-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.</p>
	<p> <b>NOTE:</b> The <b>authorization read-write</b> option is not supported on the QFX3000 QFabric system.</p>
	<p>The SNMP client application specifies an SNMP community name in <b>Get</b>, <b>GetBulk</b>, <b>GetNext</b>, and <b>Set</b> SNMP requests.</p>
<b>Default</b>	If you omit the <b>community</b> statement, all SNMP requests are denied.
<b>Options</b>	<p><b><i>community-name</i></b>—Community string. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the SNMP Community String on page 6601</a></li> </ul>

## configuration

---

**Syntax**    configuration {  
              transfer-interval *interval*;  
              transfer-on-commit;  
              archive-sites {  
                  file://<path>/<filename>;  
                  ftp://username@host:<port>url-path password password;  
                  http://username@host:<port>url-path password password;  
                  pasvftp://username@host:<port>url-path password password;  
                  scp://username@host:<port>url-path password password;  
              }  
          }

**Hierarchy Level**    [edit system archival]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure the router or switch to periodically transfer its currently active configuration (or after each commit).



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.

---

**Options**    The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**

- [Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site on page 1263](#)
- [archive on page 6780](#)
- [archive-sites on page 1267](#)
- [transfer-interval on page 1270](#)
- [transfer-on-commit on page 1271](#)

## connection-limit

<b>Syntax</b>	<code>connection-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).
<b>Options</b>	<p><b>limit</b>—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).</p> <p><b>Range:</b> 1 through 250</p> <p><b>Default:</b> 75</p>



**NOTE:** The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li> <li>• <i>Configuring DTCP-over-SSH Service for the Flow-Tap Application</i></li> <li>• <i>Configuring Finger Service for Remote Access to the Router</i></li> <li>• <i>Configuring FTP Service for Remote Access to the Router or Switch</i></li> <li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1361</a></li> <li>• <i>Configuring Telnet Service for Remote Access to a Router or Switch</i></li> </ul>

## contact

---

<b>Syntax</b>	<code>contact contact;</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the value of the MIB II <b>sysContact</b> object, which is the contact person for the managed system.
<b>Options</b>	<b>contact</b> —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the System Contact on a Device Running Junos OS</i></li></ul>

## disable (LLDP)

---

<b>Syntax</b>	<code>disable;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ], [edit protocols <a href="#">interface lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Disable the LLDP configuration on the switch or on one or more interfaces.
<b>Default</b>	If you do not configure LLDP, it is disabled on the switch and on specific switch interfaces.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show lldp on page 1486</a></li><li>• <i>Configuring LLDP (CLI Procedure)</i></li><li>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i></li><li>• <a href="#">Configuring LLDP on page 1345</a></li><li>• <a href="#">Understanding LLDP on page 1319</a></li></ul>



## falling-threshold (Health Monitor)

---

<b>Syntax</b>	<code>falling-threshold <i>percentage</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp health-monitor]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
<b>Options</b>	<p><b><i>percentage</i></b>—Lower threshold for the alarm entry.</p> <p><b>Range:</b> 1 through 100</p> <p><b>Default:</b> 70 percent of the maximum possible value</p>
<b>Required Privilege Level</b>	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">rising-threshold on page 1451</a></li> <li>• <a href="#">Configuring Health Monitoring on page 6609</a></li> </ul>

## filter-duplicates

---

<b>Syntax</b>	<code>filter-duplicates;</code>
<b>Hierarchy Level</b>	<code>[edit snmp]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Filter duplicate <b>Get</b> , <b>GetNext</b> , or <b>GetBulk</b> SNMP requests.
<b>Required Privilege Level</b>	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 6516</a></li> <li>• <a href="#">Example: Configuring SNMP on page 6575</a></li> </ul>

## full-name

---

<b>Syntax</b>	<code>full-name <i>complete-name</i>;</code>
<b>Hierarchy Level</b>	[edit system login user]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the complete name of a user.
<b>Options</b>	<i>complete-name</i> —Full name of the user. If the name contains spaces, enclose it in quotation marks.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Junos OS User Accounts</i></li><li>• <i>user</i></li><li>• <a href="#">user on page 332</a></li></ul>

## health-monitor

---

<b>Syntax</b>	<pre>health-monitor {     falling-threshold <i>percentage</i>;     interval <i>seconds</i>;     rising-threshold <i>percentage</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure health monitoring.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li><li>• <a href="#">Understanding Health Monitoring on page 6529</a></li></ul>


## hold-multiplier

---

<b>Syntax</b>	hold-multiplier <i>number</i> ;
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series.
<b>Description</b>	Specify the multiplier used in combination with the <a href="#">advertisement-interval</a> value to determine the length of time LLDP information is held before it is discarded. The default value is 4 (or 120 seconds).
<b>Default</b>	Disabled.
<b>Options</b>	<i>number</i> —A number used as a multiplier. <b>Range:</b> 2 through 10 <b>Default:</b> 4 (or 120 seconds)
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show lldp on page 1486</a></li> <li>• <i>Configuring LLDP (CLI Procedure)</i></li> <li>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i></li> <li>• <a href="#">Configuring LLDP on page 1345</a></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>

## idle-timeout (Access)

---

<b>Syntax</b>	<code>idle-timeout seconds;</code>
<b>Hierarchy Level</b>	<code>[edit access group-profile <i>profile-name</i> ppp],</code> <code>[edit access profile <i>profile-name</i> client <i>client-name</i> ppp]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure the idle timeout for a user. The router might consider a PPP session to be idle because of the following reasons:</p> <ul style="list-style-type: none"><li>• There is no ingress traffic on the PPP session.</li><li>• There is no egress traffic.</li><li>• There is neither ingress or egress traffic on the PPP session.</li><li>• There is no ingress or egress PPP control traffic. This is applicable only if keepalives are enabled.</li></ul>
<b>Options</b>	<p><b>seconds</b>—Number of seconds a user can remain idle before the session is terminated.</p> <p><b>Range:</b> 0 through 4,294,967,295 seconds</p> <p><b>Default:</b> 0</p>
<hr/> <div> <b>NOTE:</b> The [edit access] hierarchy is not available on QFabric systems.</div> <hr/>	
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Group Profile for Defining L2TP Attributes</i></li><li>• <i>Configuring PPP Properties for a Client-Specific Profile</i></li><li>• <i>Applying PPP Attributes to L2TP LNS Subscribers with a User Group Profile</i></li></ul>

## interface (LLDP)

<b>Syntax</b>	interface (all   <i>interface-name</i> ) { disable; power-negotiation { disable; } }
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure Link Layer Discovery Protocol (LLDP) on all interfaces or on a specific interface.



**NOTE:** On EX4300 switches, LLDP cannot be configured on the me0 or vme interface. Issuing the command `set protocols lldp interface me0` generates the following error message:

```
error: name: 'me0': Invalid interface
error: statement creation failed: interface
```

Issuing the command `set protocols lldp interface vme` generates the following error message:

```
error: name: 'vme': Invalid interface
error: statement creation failed: interface
```

<b>Default</b>	None
<b>Options</b>	<p><b>all</b>—All interfaces on the switch.</p> <p><b><i>interface-name</i></b>—Name of a specific interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP (CLI Procedure)</a></li> <li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</a></li> <li>• <a href="#">Configuring LLDP on page 1345</a></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>

## interval (Health Monitor)

---

<b>Syntax</b>	interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit snmp health-monitor]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the interval between sampling of the object being monitored by the health monitor.
<b>Options</b>	<p><i>seconds</i>—Time between samples, in seconds.</p> <p><b>Range:</b> 1 through 2147483647 seconds</p> <p><b>Default:</b> 300 seconds</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li></ul>

## lldp

```
Syntax  lldp {
    advertisement-interval seconds;
    disable;
    hold-multiplier number;
    interface (all | [interface-name]) {
        disable;
        power-negotiation {
            disable;
        }
    }
    lldp-configuration-notification-interval seconds;
    management-address ip-management-address;
    netbios-snooping;
    ptopo-configuration-maximum-hold-time seconds;
    ptopo-configuration-trap-interval seconds;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
        <no-stamp> <replace>;
        flag flag <disable>;
    }
    transmit-delay seconds;
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for QFX Series.

**Description** Configure Link Layer Discovery Protocol (LLDP). The switch uses LLDP to advertise its identity and capabilities on a LAN, as well as to receive information about other network devices. LLDP is defined in the IEEE standard 802.1AB-2005.

The remaining statements are explained separately.



**NOTE:** The `transmit-delay` and `netbios-snooping` options are not available on QFabric systems.



**NOTE:** On EX4300 switches, LLDP cannot be configured on the `me0` or `vme` interface. Issuing the command `set protocols lldp interface me0` generates the following error message:

```
error: name: 'me0': Invalid interface
error: statement creation failed: interface
```

Issuing the command `set protocols lldp interface vme` generates the following error message:

```
error: name: 'vme': Invalid interface
error: statement creation failed: interface
.....
```

<b>Default</b>	LLDP is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show lldp on page 1486</a></li><li>• <i>Configuring LLDP (CLI Procedure)</i></li><li>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i></li><li>• <a href="#">Configuring LLDP on page 1345</a></li><li>• <a href="#">Understanding LLDP on page 1319</a></li></ul>

---

## lldp-configuration-notification-interval

---

<b>Syntax</b>	lldp-configuration-notification-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify how often SNMP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, trap notifications of database changes are disabled.
<b>Default</b>	SNMP trap notifications of LLDP database changes are disabled.
<b>Options</b>	<b>seconds</b> —Interval between trap notifications about LLDP database changes. <b>Range:</b> 0 through 3600
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show lldp on page 1486</a></li></ul>



---


## location

---

<b>Syntax</b>	<code>location <i>location</i>;</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the value of the MIB II <b>sysLocation</b> object, which is the physical location of the managed system.
<b>Options</b>	<b><i>location</i></b> —Location of the local system. You must enclose the name within quotation marks (" ").
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the System Location for a Device Running Junos OS</i></li></ul>

## management-address

---

<b>Syntax</b>	<code>management-address <i>ip-management-address</i>;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the management address of the switch to be used in the LLDP Management type, length, and value (TLV). The Management Address TLV typically contains the IPv4 or IPv6 management address of the local system. Only out-of-band management addresses must be used for the management-address. Other remote managers can use this address to obtain information related to the local device.
<div> <b>NOTE:</b> Ensure that the interface with the configured management address has LLDP enabled using the <code>set protocols lldp interface</code> command. If you configure a customized management address for LLDP on an interface that has LLDP disabled, the <code>show lldp local-information</code> command output will not display the correct interface information.</div>	
<b>Default</b>	The LLDP Management Address TLV uses the IP address of the switch's management Ethernet interface ( <b>me0</b> ), or the IP address of the virtual management Ethernet (VME) interface if the switch is a Virtual Chassis member.
<b>Options</b>	<i>ip-management-address</i> —You can specify either an IPv4 or an IPv6 management address for the switch.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show lldp on page 1486</a></li><li>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i></li><li>• <i>EX Series Switches Interfaces Overview</i></li><li>• <a href="#">Understanding LLDP on page 1319</a></li></ul>

## name

---

<b>Syntax</b>	<code>name <i>name</i>;</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the system name from the command-line interface.
<b>Options</b>	<i>name</i> —System name override.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the System Name</a></li> </ul>

## nas-ip-address

---

<b>Syntax</b>	<code>nas-ip-address <i>ip-address</i>;</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the NAS-IP address for outgoing RADIUS packets.
<b>Options</b>	<i>ip-address</i> —IP address of the network access server (NAS) that requests user authentication.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS Authentication</a></li> <li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li> </ul>

## nonvolatile

---

<b>Syntax</b>	<code>nonvolatile {     <a href="#">commit-delay</a> <i>seconds</i>; }</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure options for SNMP <b>Set</b> requests.  The statement is explained separately.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Commit Delay Timer</i></li><li>• <i>commit-delay</i></li></ul>

## oid

---

<b>Syntax</b>	<code>oid <i>object-identifier</i> (exclude include);</code>
<b>Hierarchy Level</b>	[edit snmp view <i>view-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify an object identifier (OID) used to represent a subtree of MIB objects.
<b>Options</b>	<b>exclude</b> —Exclude the subtree of MIB objects represented by the specified OID.  <b>include</b> —Include the subtree of MIB objects represented by the specified OID.  <b><i>object-identifier</i></b> —OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring MIB Views on page 6605</a></li></ul>

## order

---

<b>Syntax</b>	<code>order (radius   [ <i>accounting-order-data-list</i> ] );</code>
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> accounting]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending accounting messages and updates.
<b>Default</b>	No order specified
<b>Options</b>	<p><b>radius</b>—RADIUS accounting for specified subscribers.</p> <p><b>[ <i>accounting-order-data-list</i> ]</b>— Set of data listing the authentication order to be used, enclosed by brackets. This can be any combination of the authentication methods, up to and including a full list of the entire authentication order.</p>



**NOTE:** The [edit access] hierarchy is not available on QFabric systems.

---

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li> <li>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i></li> <li>• <i>Configuring RADIUS Accounting</i></li> </ul>

## port (RADIUS Server)

---

<b>Syntax</b>	<code>port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit system radius-server <i>address</i> ], [edit system accounting destination radius server <i>address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the port number on which to contact the RADIUS server.
<b>Options</b>	<i>number</i> —Port number on which to contact the RADIUS server. <b>Default:</b> 1812 (as specified in RFC 2865)



**NOTE:** The [edit system accounting] hierarchy is not available on QFabric systems.

---

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring RADIUS Authentication</i></li><li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li></ul>
------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## profile

**Syntax**    `profile profile-name {  
                   accounting (Access Profile) {  
                     accounting-stop-on-access-deny;  
                     accounting-stop-on-failure;  
                     order ( radius | [ accounting-order-data-list ] );  
                   }  
                   authentication-order [authentication-method];  
                   radius {  
                     accounting-server [server-addresses];  
                     authentication-server [server-addresses];  
                   }  
                 }`

**Hierarchy Level**    [edit access]

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                                 Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure an access profile. The access profile contains the entire authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, including the authentication method and order, AAA server addresses, and AAA accounting.

**Default**    Not enabled

**Options**    *profile-name*—Profile name of up to 32 characters.  
                   The remaining statements are explained separately.



**NOTE:** The [edit access] hierarchy is not available on QFabric systems.

**Required Privilege Level**    admin—To view this statement in the configuration.  
                                         admin-control—To add this statement to the configuration.

**Related Documentation**    • *Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch*  
                                         • *Configuring 802.1X RADIUS Accounting (CLI Procedure)*  
                                         • *Configuring RADIUS Accounting*

## protocols

```
Syntax protocols {
    bgp {
        disable;
        accept-remote-nexthop;
        advertise-external <conditional>;
        advertise-inactive;
        (advertise-peer-as | no-advertise-peer-as);
        authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
        authentication-key key;
        authentication-key-chain key-chain;
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            hold-down-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family family-name {
            ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
        }
        graceful-restart {
            disable;
            restart-time seconds;
            stale-routes-time seconds;
        }
        group group-name {
            ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
        }
        hold-time seconds;
        import [ policy-names ];
        include-mp-next-hop;
        keep (all | none);
        local-address address;
    }
}
```



```

local-as autonomous-system <loops number> < alias> <private>;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl ttl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
}

```

```

    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (Liveness Detection) (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable;
}
interface interface-name {
    group-limit limit;
    multicast-router-interface;
    static {
        group ip-address;
    }
}

```

```

        robust-count number;
    }
}
isis {
    disable;
    export [ policy-names ];
    ignore-attached-bit;
    interface interface-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
    }
    checksum;
    csnp-interval (seconds | disable);
    disable;
    hello-padding (adaptive | loose | strict);
    level (1 | 2) {
        disable;
        hello-authentication-key key;
        hello-authentication-type authentication;
        hello-interval seconds;
        hold-time seconds;
        ipv4-multicast-metric number;
        metric metric;
        passive;
        priority number;
    }
    lsp-interval milliseconds;
    mesh-group (value | blocked);
    no-ipv4-multicast;
    no-unicast-topology;
    passive;
    point-to-point;
}
level (1 | 2) {
    disable;
    authentication-key key;
    authentication-type authentication;
    external-preference preference;
    no-csnp-authentication;
}

```

```
no-hello-authentication;
no-psnp-authentication;
preference preference;
prefix-export-limit number;
wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
}
topologies {
    ipv4-multicast;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
}
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name (MSTP) name;
    forward-delay seconds;
```

```

hello-time seconds;
interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
        block;
        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
max-hops hops;
msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
        disable;
        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
ospf {
    disable;
    area area-id {
        area-range ip-prefix </prefix-length> <exact> <override-metric metric > <restrict>;
        context-identifier identifier
        interface interface-name {
            disable;
            authentication {
                md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
                simple-password key-string;
            }
            bandwidth-based-metrics {
                bandwidth value metric number;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
}

```

```
    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication {
```

```

        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
    }
    dead-interval seconds;
    demand-circuit;
    flood-reduction;
    hello-interval seconds;
    ipsec-sa sa-name;
    no-neighbor-down-notification;
    retransmit-interval seconds;
    topology (name | default | ipv4-multicast) {
        disable;
        metric metric;
    }
    transit-delay seconds;
}
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
    overload;
    prefix-export-limit number;
    topology-id number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
}

```

```
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
  }
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  family (inet | inet6) {
    disable;
  }
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
      disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance;
  join-prune-timeout;
  nonstop-routing;
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group group-name;
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
  }
  bootstrap-import [ policy-names ];
}
```



```

bootstrap-export [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix</prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}

```

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
tunnel-devices [ mt-fpc/pic/port ];
}
rip {
  authentication-key password;
  authentication-type type;
  (check-zero | no-check-zero);
  group group-name {
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  export [ policy-names ];
  import [ policy-names ];
  metric-out metric;
  neighbor neighbor-name {
    any-sender;
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
      ... same statements as at the [edit protocols rip group group-name
        bfd-liveness-detection] hierarchy level ...
    }
    (check-zero | no-check-zero);
    import [ policy-names ];
    message-size number;
    metric-in metric;
    receive (both | none | version-1 | version-2);
    route-timeout seconds;
    send (broadcast | multicast | none | version-1);
    update-interval seconds;
  }
  preference preference;
  route-timeout seconds;
  update-interval seconds;
}
holddown seconds;
```

```

import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
traceoptions {
    file filename <files number> <size size> <no-stamp> <world-readable |
        no-world-readable>;
    flag flag;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
}

```

```

    }
    max-age seconds;
  }
  traceoptions {
    file filename <files number> <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
  uplink-failure-detection {
    group group-name {
      link-to-monitor interface-name;
      link-to-disable interface-name;
    }
  }
}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number> <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
}
}

```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure protocols.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Junos OS Routing Protocols Configuration Guide](#)

## protocol-version

---

<b>Syntax</b>	<code>protocol-version <i>version</i>;</code>
<b>Hierarchy Level</b>	[edit system services ssh]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the secure shell (SSH) protocol version.
<b>Default</b>	v2—SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.
<b>Options</b>	<i>version</i> —SSH protocol version: v1, v2, or both.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the SSH Protocol Version on page 1363</a></li> </ul>

## ptopo-configuration-maximum-hold-time

---


<b>Syntax</b>	<code>ptopo-configuration-maximum-hold-time <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure how long to maintain the physical topology database entries. The physical topology identifies the devices on the network and their physical interconnections.
<b>Options</b>	<i>seconds</i> —Time to maintain physical topology database entries. <b>Default:</b> 300 <b>Range:</b> 1 through 2147483647
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show lldp on page 1486</a></li> <li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</a></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>

## ptopo-configuration-trap-interval

---


<b>Syntax</b>	<code>ptopo-configuration-trap-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">lldp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify how often SNMP trap notifications are sent regarding changes in physical topology global statistics.
<b>Default</b>	SNMP trap notifications of changes in physical topology global statistics are disabled.
<b>Options</b>	<b><i>seconds</i></b> —Interval between SNMP trap notifications about physical topology global statistics. <b>Range:</b> 0 through 3600
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.

## radius

<b>Syntax</b>	radius { accounting-server [server-addresses]; authentication-server [server-addresses]; }
<b>Hierarchy Level</b>	[edit access profile <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure the RADIUS servers for authentication and for accounting. To configure multiple RADIUS servers, include multiple <b>radius</b> statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p> <p>The statements are explained separately.</p>
<div>  <b>NOTE:</b> The [edit access] hierarchy is not available on QFabric systems. </div>	
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Connecting a RADIUS Server for 802.1X to an EX Series Switch</i></li> <li>• <i>Configuring 802.1X RADIUS Accounting (CLI Procedure)</i></li> <li>• <i>Filtering 802.1X Supplicants Using RADIUS Server Attributes</i></li> <li>• <i>Configuring RADIUS Accounting</i></li> </ul>


## radius-options (edit system)

---

Syntax	<pre>radius-options {   attributes {     nas-ip-address <i>ip-address</i>;   }   enhanced-accounting;   password-protocol <i>mschap-v2</i>; }</pre>
Hierarchy Level	[edit system]
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2.</p> <p>MS-CHAPv2 password protocol configuration option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<hr/>	
<div> <b>NOTE:</b> The <code>radius-options</code> statement is not available on QFabric systems.</div> <hr/>	
<p><b>enhanced-accounting</b> statement introduced in Junos OS Release 14.1.</p>	
Description	Configure RADIUS options for the NAS-IP address for outgoing RADIUS packets and password protocol used in RADIUS packets.
Options	<p><b>enhanced-accounting</b>—View the attribute values of a logged in user.</p> <p><b>nas-ip-address <i>ip-address</i></b>—IP address of the network access server (NAS) that requests user authentication.</p> <p><b>password-protocol <i>mschap-v2</i></b>—Protocol MS-CHAPv2, used for password authentication and password changing.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <i>Configuring MS-CHAPv2 for Password-Change Support</i></li><li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li><li>• <a href="#">Configuring RADIUS System Accounting on page 1349</a></li><li>• <i>enhanced-accounting</i></li></ul>



## radius-server

<b>Syntax</b>	<pre>radius-server server-address {     accounting-port port-number;     port number;     retry number;     secret password;     source-address source-address;     timeout seconds; }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure a RADIUS server for Point-to-Point Protocol (PPP).</p> <p>To configure multiple RADIUS servers, include multiple <b>radius-server</b> statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.</p>
<b>Options</b>	<p><b>server-address</b>—Address of the RADIUS authentication server.</p> <p>The remaining statements are explained separately.</p>
<div>  <p><b>NOTE:</b> The <b>accounting-port</b> and <b>source-address</b> options are not available on QFabric systems.</p> </div>	
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li> <li>• <a href="#">accounting-port on page 249</a></li> <li>• <a href="#">port on page 1428</a></li> <li>• <a href="#">retry on page 301</a></li> <li>• <a href="#">secret on page 305</a></li> <li>• <a href="#">source-address on page 308</a></li> <li>• <a href="#">timeout on page 322</a></li> </ul>

## rate-limit

---

<b>Syntax</b>	<code>rate-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	[edit system services finger], [edit system services ftp], [edit system services netconf ssh], [edit system services ssh], [edit system services telnet], [edit system services xnm-clear-text], [edit system services xnm-ssl]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the maximum number of connections attempts per protocol (either IPv6 or IPv4) on an access service.
<b>Default</b>	150 connections
<b>Options</b>	<b>rate-limit <i>limit</i></b> —(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6). <b>Range:</b> 1 through 250 <b>Default:</b> 150
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li></ul>

## remote-debug-permission

<b>Syntax</b>	remote-debug-permission (qfabric-admin   qfabric-operator   qfabric-user);
<b>Hierarchy Level</b>	[edit system login user <i>username</i> authentication] [edit system root-authentication]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	(QFabric systems only) Configure authentication classes that permit or deny user access to individual components of the QFabric system.
<b>Default</b>	qfabric-user
<b>Options</b>	<p><b>qfabric-admin</b>—Permits a user to log in to individual QFabric system components, view operations, and change component configurations.</p> <p><b>qfabric-operator</b>—Permits a user to log in to individual QFabric system components and view component operations.</p> <p><b>qfabric-user</b>—Prevents a user from logging in to individual QFabric system components.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring QFabric System Login Classes</i></li> <li>• <a href="#">request component login on page 1480</a></li> <li>• <i>Understanding QFabric System Login Classes</i></li> </ul>

## retry

---

<b>Syntax</b>	<code>retry number;</code>
<b>Hierarchy Level</b>	[edit system radius server <i>server-address</i> ], [edit system accounting destination radius server <i>server-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.
<b>Options</b>	<i>number</i> —Number of retries allowed for contacting a RADIUS server. <b>Range:</b> 1 through 10 <b>Default:</b> 3



**NOTE:** The [edit system accounting] hierarchy is not available on QFabric systems.

---

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring RADIUS Authentication (QFX Series) on page 1351</a></li><li>• <a href="#">Configuring RADIUS Accounting</a></li><li>• <a href="#">timeout on page 322</a></li></ul>

---

## rising-threshold (Health Monitor)

---

<b>Syntax</b>	<code>rising-threshold <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit snmp health-monitor]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.
<b>Options</b>	<p><b><i>percentage</i></b>—Upper threshold for the alarm entry.</p> <p><b>Range:</b> 1 through 100</p> <p><b>Default:</b> 80 percent of the maximum possible value</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li><li>• <a href="#">falling-threshold on page 1415</a></li></ul>

## root-login

---

<b>Syntax</b>	root-login (allow   deny   deny-password);
<b>Hierarchy Level</b>	[edit system services ssh]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Control user access through SSH.
<b>Default</b>	Allow user access through SSH.
<b>Options</b>	<b>allow</b> —Allow users to log in to the router or switch as root through SSH. <b>deny</b> —Disable users from logging in to the router or switch as root through SSH. <b>deny-password</b> —Allow users to log in to the router or switch as root through SSH when the authentication method (for example, RSA authentication) does not require a password.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Root Login Through SSH on page 1363</a></li></ul>

---

## services (Switches)

---

**Syntax**

```
services {
  service-deployment {
    servers address {
      port-number port-number;
    }
    source-address address;
  }
  ssh {
    connection-limit limit;
    protocol-version [v1 v2];
    rate-limit limit;
    root-login (allow | deny | deny-password);
  }
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the switch so that users on remote systems can access the local switch through SSH.

The remaining statements are explained separately.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

## snmp

---

```
Syntax  snmp {
    client-list client-list-name {
        ip-addresses;
    }
    community community-name {
        authorization authorization;
        client-list-name client-list-name;
        clients {
            address restrict;
        }
        logical-system logical-system-name {
            routing-instance routing-instance-name {
                clients {
                    addresses;
                }
            }
        }
        routing-instance routing-instance-name {
            clients {
                addresses;
            }
        }
        view view-name;
    }
    contact contact;
    description description;
    filter-duplicates;
    filter-interfaces;
    health-monitor {
        falling-threshold integer;
        interval seconds;
        rising-threshold integer;
    }
    interface [ interface-names ];
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
        }
    }
}
```



```

    variable oid-variable;
}
event index {
    community community-name;
    description description;
    type type;
}
history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
}
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
        regular-expression>;
    flag flag;
}
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance routing-instance-name;
    targets {
        address;
    }
    version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    source-address address;
}
v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system logical-system;
        port port-number;
        retry-count number;
        routing-instance routing-instance-name;
        tag-list tag-list;
        target-parameters target-parameters-name;
    }
}

```

```
    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
  usm {
    local-engine {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none;
      }
    }
    remote-engine engine-id {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none {
          privacy-password privacy-password;
        }
      }
    }
  }
}
```

```

}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
view view-name {
  oid object-identifier (include | exclude);
}
}
}

```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure SNMP.

The remaining statements are explained separately.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding the Implementation of SNMP on page 6513](#)
- [Configuring SNMP on page 1356](#)

## ssh

---

**Syntax**    ssh {  
              ciphers [ *cipher-1 cipher-2 cipher-3 ...*];  
              client-alive-count-max *seconds*;  
              client-alive-interval *seconds*;  
              connection-limit *limit*;  
              hostkey-algorithm <*algorithm*|*no-algorithm*>;  
              key-exchange <*algorithm*>;  
              macs <*algorithm*>;  
              max-sessions-per-connection <*number*>;  
              no-passwords;  
              no-tcp-forwarding;  
              protocol-version [*v1 v2*];  
              rate-limit *limit*;  
              root-login (*allow* | *deny* | *deny-password*);  
              }

**Hierarchy Level**    [edit system services]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.  
                              **client-alive-interval** and **client-alive-max-count** statements introduced in Junos OS Release 12.2.  
                              **no-passwords** statement introduced in Junos OS Release 13.3.

**Description**    Allow SSH requests from remote systems to the local router or switch.

                      The remaining statements are explained separately.

**Required Privilege Level**    system—To view this statement in the configuration.  
                                  system-control—To add this statement to the configuration.

**Related Documentation**    • [Configuring SSH Service for Remote Access to the Router or Switch on page 1361](#)

## system

```
Syntax  system {
    accounting {
        events [ login change-log interactive-commands ];
        destination {
            radius {
                server {
                    server-address {
                        accounting-port port-number;
                        retry number;
                        secret password;
                        source-address address;
                        timeout seconds;
                    }
                }
            }
        }
        tacplus {
            server {
                server-address {
                    port port-number;
                    secret password;
                    single-connection;
                    timeout seconds;
                }
            }
        }
    }
    archival {
        configuration {
            archive-sites {
                ftp://<username>:<password>@<host>:<port>/<url-path>;
                ftp://<username>:<password>@<host>:<port>/<url-path>;
            }
            transfer-interval interval;
            transfer-on-commit;
        }
    }
    arp {
        aging-timer minutes;
        interfaces;
    }
    authentication-order [ authentication-methods ];
    (compress-configuration-files | no-compress-configuration-files);
    default-address-selection;
    domain-name domain-name;
    domain-search [ domain-list ];
    host-name hostname;
    internet-options {
        icmpv4-rate-limit bucket-size packet-rate packet-rate;
        source-port upper-limit <upper-limit>;
    }
    location {
```

```
altitude feet;  
building name;  
country-code code;  
floor number;  
hcoord horizontal-coordinate;  
lata service-area;  
latitude degrees;  
longitude degrees;  
npa-nxx number;  
postal-code postal-code;  
rack number;  
vcoord vertical-coordinate;  
}  
login {  
  announcement text;  
  class class-name {  
    access-end;  
    access-start;  
    allow-configuration "regular-expression";  
    allowed-days "regular-expression";  
    deny-commands "regular-expression";  
    deny-configuration "regular-expression";  
    idle-timeout minutes;  
    login-tip;  
    permissions [ permissions ];  
  }  
  message text;  
  password {  
    change-type (set-transitions | character-set);  
    format (md5 | sha1 | des);  
    maximum-length length;  
    minimum-changes number;  
    minimum-length length;  
  }  
  retry-options {  
    backoff-factor seconds;  
    backoff-threshold number;  
    minimum-time seconds;  
    tries-before-disconnect number;  
  }  
  user username {  
    authentication {  
      (encrypted-password "password" | plain-text-password);  
      load-key-file URL;  
      remote-debug-permission (qfabric-admin | qfabric-operator | qfabric-user);  
      ssh-rsa "public-key";  
      ssh-dsa "public-key";  
    }  
    uid uid-value;  
    class class-name;  
    full-name complete-name;  
  }  
}  
name-server {  
  address;  
}
```

```

no-multicast-echo;
no-redirects;
no-ping-record-route;
no-ping-time-stamp;
ntp {
    authentication-key number type type value password;
    serveraddress <key key-number> <version value> <prefer>;
}
ports {
    auxiliary {
        disable;
        insecure;
        type terminal-type;
    }
    console {
        disable;
        insecure;
        log-out-on-disconnect;
        type terminal-type;
    }
}
radius-server server-address {
    accounting-port port-number;
    port number;
    retry number;
    secret password;
    source-address source-address;
    timeout seconds;
}
radius-options {
    password-protocol mschap-v2;
}
attributes {
    nas-ip-address ip-address;
}
root-authentication {
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
}
(saved-core-context | no-saved-core-context);
saved-core-files saved-core-files;
services {
    finger {
        connection-limit limit;
        rate-limit limit;
    }
    flow-tap-dtcp {
        ssh {
            connection-limit limit;
            rate-limit limit;
        }
    }
}
ftp {
    connection-limit limit;
    rate-limit limit;
}

```

```
}
service-deployment {
  servers server-address {
    port port-number;
  }
  source-address source-address;
}
ssh {
  root-login (allow | deny | deny-password);
  protocol-version [v1 v2];
  connection-limit limit;
  rate-limit limit;
}
telnet {
  connection-limit limit;
  rate-limit limit;
}
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
  session {
    idle-timeout [ minutes ];
    session-limit [ session-limit ];
  }
}
xnm-clear-text {
  connection-limit limit;
  rate-limit limit;
}
xnm-ssl {
  connection-limit limit;
  local-certificate name;
  rate-limit limit;
}
}
static-host-mapping {
  hostname {
    alias [ alias ];
    inet [ address ];
    sysid system-identifier;
  }
}
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
}
```



```

}
console {
    facility severity;
}
file filename {
    archive {
        files number;
        size maximum-file-size;
        start-time "YYYY-MM-DD.hh:mm";
        transfer-interval minutes;
        (world-readable | no-world-readable);
    }
    explicit-priority;
    facility severity;
    match "regular-expression";
    structured-data {
        brief;
    }
}
host (hostname | other-routing-engine | scc-master) {
    explicit-priority;
    facility-override facility;
    facility severity;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
time-format (millisecond | year | year millisecond);
user (username | *) {
    facility severity;
    match "regular-expression";
}
}
tacplus-options {
    service-name service-name;
    (no-cmd-attribute-value | exclude-cmd-attribute);
}
tacplus-server server-address {
    port
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
time-zone (GMThour-offset | time-zone);
}
tracing {
    destination-override {
        syslog host;
    }
}
use-imported-time-zones;
}

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure system management properties.



**NOTE:** The `radius-server source-address` and `radius-options` statements are not available on the QFabric system.

**Required Privilege** system—To view this statement in the configuration.

**Level** system-control—To add this statement to the configuration.

## tacplus-options

<b>Syntax</b>	<pre> tacplus-options {   (exclude-cmd-attribute   no-cmd-attribute-value);   enhanced-accounting;   service-name <i>service-name</i>;   timestamp-and-timezone; } </pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>no-cmd-attribute-value</b> and <b>exclude-cmd-attribute</b> options introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 11.1 for QFX Series.</p> <p><b>timestamp-and-timezone</b> option introduced in Junos OS Release 12.2.</p> <p><b>enhanced-accounting</b> option introduced in Junos OS Release 14.1.</p>
<b>Description</b>	Configure TACACS+ options for authentication and accounting.
<b>Options</b>	<p><b>enhanced-accounting</b>—View the attribute values of a logged in user.</p> <p><b>exclude-cmd-attribute</b>—Exclude the <b>cmd</b> attribute value completely from start and stop accounting records to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p><b>no-cmd-attribute-value</b>—Set the <b>cmd</b> attribute value to an empty string in the TACACS+ accounting start and stop requests to enable logging of accounting records in the correct log file on a TACACS+ server.</p> <p><b>service-name <i>service-name</i></b>—Name of the authentication service used when you configure multiple TACACS+ servers to use the same authentication service.</p> <p><b>Default:</b> junos-exec</p> <p><b>timestamp-and-timezone</b>—Include this statement if you want start time, stop time, and timezone attributes included in start/stop accounting records.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring TACACS+ Authentication</i></li> <li>• <i>Configuring TACACS+ System Accounting</i></li> <li>• <a href="#">Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication on page 1330</a></li> <li>• <i>enhanced-accounting</i></li> </ul>

## targets

---

<b>Syntax</b>	<code>targets {     <i>address</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit snmp trap-group <i>group-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure one or more systems to receive SNMP traps.
<b>Options</b>	<b><i>address</i></b> —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li></ul>

## traceoptions (LLDP)

**Syntax** `traceoptions {  
     file filename <files number> <size size> <world-readable | no-world-readable> <no-stamp>  
     <replace>;  
     flag flag <disable>;  
}`

**Hierarchy Level** [edit protocols [lldp](#)]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Define tracing operations for the Link Layer Discovery Protocol (LLDP). You can trace messages under LLDP for LLDP and PTOPO MIBs.



**NOTE:** The `traceoptions` statement is not supported on the QFX3000 QFabric system.

**Default** Tracing operations are disabled.

**Options** **file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000

**Default:** 3 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **configuration**—Trace configuration operations.
- **interface**—Trace interface update events.
- **netbios**—Trace NetBIOS events.
- **packet**—Trace packet events.
- **rtsock**—Trace routing socket operations.
- **snmp**—Trace SNMP configuration operations.

- **vlan**—Trace VLAN update events.

**no-stamp**—(Optional) Do not timestamp the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**replace**—(Optional) Replace an existing trace file if there is one rather than appending output to it.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**world-readable**—(Optional) Enable unrestricted file access.



**NOTE:** The **traceoptions** statement is not supported on the QFX3000 QFabric system.

---

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LLDP-MED (CLI Procedure)</a></li><li>• <a href="#">Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</a></li><li>• <a href="#">Configuring LLDP on page 1345</a></li><li>• <a href="#">Understanding LLDP on page 1319</a></li></ul>
------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## transfer-interval (Configuration)

<b>Syntax</b>	<code>transfer-interval <i>interval</i>;</code>
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the router or switch to periodically transfer its currently active configuration to an archive site.



**NOTE:** The `edit system archival` hierarchy is not available on QFabric systems.

**Options** *interval*—Interval at which to transfer the current configuration to an archive site.  
**Range:** 15 through 2880 minutes



**NOTE:** The `[edit system archival]` hierarchy is not available on QFabric systems.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Transfer Interval for Periodic Transfer of the Active Configuration to an Archive Site on page 1263](#)
- [archive on page 6780](#)
- [configuration on page 1269](#)
- [transfer-on-commit on page 1271](#)

## transfer-on-commit

---

<b>Syntax</b>	transfer-on-commit;
<b>Hierarchy Level</b>	[edit system archival configuration]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the router or switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.



**NOTE:** When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (") and enclose the IPv6 host address in brackets ([ ]). For example, "ftp://username<:password>@[ipv6-host-address]<:port>/url-path".

---



**NOTE:** The [edit system archival] hierarchy is not available on QFabric systems.

---

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Transfer of the Current Active Configuration When a Configuration Is Committed on page 1264</a></li><li>• <a href="#">archive on page 6780</a></li><li>• <a href="#">configuration on page 1269</a></li><li>• <a href="#">transfer-interval on page 1270</a></li></ul>



## trap-group

<b>Syntax</b>	<pre> trap-group <i>group-name</i> {     categories {         <i>category</i>;     }     destination-port <i>port-number</i>;     routing-instance <i>instance</i>;     targets {         <i>address</i>;     }     version (all   v1   v2); } </pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
<b>Description</b>	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
<b>Options</b>	<p><b><i>group-name</i></b>—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring SNMP Trap Groups</i></li> </ul>

## trap-options

---

<b>Syntax</b>	<pre>trap-options {     agent-address outgoing-interface;     source-address address; }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Options</i></li></ul>

## user (Access)

<b>Syntax</b>	<pre> user username {   authentication {     (encrypted-password "password"   plain-text-password);     load-key-file URL;     remote-debug-permission (qfabric-admin   qfabric-operator   qfabric-user);     ssh-dsa "public-key" &lt;from hostname&gt;;     ssh-rsa "public-key" &lt;from hostname&gt;;   }   class class-name;   full-name "complete-name";   uid uid-value; }</pre>
<b>Hierarchy Level</b>	[edit system login]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure access permission for individual users.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Junos OS User Accounts on page 1344</a></li> <li>• <a href="#">class on page 262</a></li> </ul>

## version

---

<b>Syntax</b>	version (all   v1   v2);
<b>Hierarchy Level</b>	[edit snmp trap-group <i>group-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the version number of SNMP traps.
<b>Default</b>	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.
<b>Options</b>	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.  v1—Send SNMPv1 traps only.  v2—Send SNMPv2 traps only.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li></ul>

# Administration

- [Routine Monitoring on page 1475](#)
- [Monitoring Commands on page 1476](#)

## Routine Monitoring

---

- [Monitoring SNMP on page 1475](#)

### Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.
- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
```

Alarm			
Index	Variable description	Value	State
32768	Health Monitor: root file system utilization jnxHrStoragePercentUsed.1	58	active
32769	Health Monitor: /config file system utilization jnxHrStoragePercentUsed.2	0	active
32770	Health Monitor: RE 0 CPU utilization jnxOperatingCPU.9.1.0.0	0	active
32773	Health Monitor: RE 0 Memory utilization jnxOperatingBuffer.9.1.0.0	35	active
32775	Health Monitor: jkernel daemon CPU utilization Init daemon	0	active

Chassis daemon	50 active
Firewall daemon	0 active
Interface daemon	5 active
SNMP daemon	11 active
MIB2 daemon	42 active
...	

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system
```

```
sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel  
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:  
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx
```

```
Build date: 2010-09-26 06:00:10 U  
sysObjectID.0 = jnxProductQFX3500  
sysUpTime.0   = 24444184  
sysContact.0  = J Smith  
sysName.0     = Lab QFX3500  
sysLocation.0 = Lab  
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics
```

SNMP statistics:

Input:

```
Packets: 0, Bad versions: 0, Bad community names: 0,  
Bad community uses: 0, ASN parse errors: 0,  
Too bigs: 0, No such names: 0, Bad values: 0,  
Read onlys: 0, General errors: 0,  
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,  
Throttle drops: 0, Duplicate request drops: 0
```

Output:

```
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

- Related Documentation
- [health-monitor on page 1416](#)
  - [show snmp mib on page 6874](#)
  - [show snmp statistics on page 1503](#)

---

## Monitoring Commands

- [clear lldp neighbors](#)
- [clear lldp statistics](#)
- [request component login](#)
- [show ethernet-switching interfaces](#)

- `show lldp`
- `show lldp local-information`
- `show lldp neighbors`
- `show lldp statistics`
- `show route instance`
- `show snmp statistics`
- `ssh`

## clear lldp neighbors

---

<b>Syntax</b>	<code>clear lldp neighbors &lt;interface <i>interface</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear the learned remote neighbor information on all or selected interfaces.
<b>Options</b>	<b>none</b> —Clear the remote neighbor information on all interfaces.  <b>interface <i>interface</i></b> —(Optional) Clear the remote neighbor information from the selected interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show lldp</a></li><li>• <a href="#">Configuring LLDP on page 1345</a></li><li>• <a href="#">Understanding LLDP on page 1319</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear lldp neighbors on page 1478</a> <a href="#">clear lldp neighbors interface on page 1478</a>

### Sample Output

#### clear lldp neighbors

```
user@switch> clear lldp neighbors
```

#### clear lldp neighbors interface

```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```



## clear lldp statistics

---

<b>Syntax</b>	<code>clear lldp statistics</code> <code>&lt;interface <i>interface</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear LLDP statistics on one or more interfaces.
<b>Options</b>	<b>none</b> —Clears LLDP statistics on all interfaces.  <b>interface <i>interface-names</i></b> —(Optional) Clear LLDP statistics on an interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 1345</a></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear lldp statistics on page 1479</a> <a href="#">clear lldp statistics interface on page 1479</a>

### Sample Output

#### clear lldp statistics

```
user@switch> clear lldp statistics
```

#### clear lldp statistics interface

```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

## request component login

---

<b>Syntax</b>	<code>request component login <i>component-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	(QFabric systems only) Log in to a QFabric system component. To gain access to individual components by way of the <b>request component login</b> command, you must first provide the <b>qfabric-admin</b> or <b>qfabric-operator</b> class privilege to your user (for more information, see: <a href="#">remote-debug-permission</a> ).
<b>Options</b>	<b><i>component-name</i></b> —Specify the QFabric system component to which you wish to log in.
<b>Required Privilege Level</b>	admin
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring QFabric System Login Classes</i></li><li>• <a href="#">remote-debug-permission on page 1449</a></li><li>• <i>Understanding QFabric System Login Classes</i></li></ul>
<b>List of Sample Output</b>	<a href="#">request component login (with qfabric-admin Privileges) on page 1480</a> <a href="#">request component login (with qfabric-operator Privileges) on page 1481</a> <a href="#">request component login (with qfabric-user Privileges) on page 1481</a>

## Sample Output

The three sample output displays show the results of attempts to log in to Node device EE3093. The results differ depending on the privilege level assigned to the user.

### request component login (with qfabric-admin Privileges)

```
admin@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-admin@node-ee3093> ?
Possible completions:
clear          Clear information in the system
file           Perform file operations
help           Provide help information
load           Load information from file
monitor        Show real-time debugging information
mtrace         Trace multicast path from source to receiver
op             Invoke an operation script
ping           Ping remote target
quit           Exit the management session
request        Make system-level requests
restart        Restart software process
save           Save information to file
set            Set CLI properties, date/time, craft interface message
show           Show system information
ssh            Start secure shell on another host
start          Start shell
```

```

telnet          Telnet to another host
test           Perform diagnostic debugging
traceroute     Trace route to remote host{master}
qfabric-admin@node-ee3093>

```

#### request component login (with qfabric-operator Privileges)

```

operator@qfabric> request component login EE3093
Warning: Permanently added 'qfabric-node-ee3093,169.254.128.41' (RSA) to the list
of known hosts.
--- JUNOS 11.3I built 2011-11-04 12:46:16 UTC
{master}
qfabric-operator@node-ee3093> ?
Possible completions:
file          Perform file operations
help          Provide help information
load          Load information from file
op            Invoke an operation script
quit          Exit the management session
request       Make system-level requests
save          Save information to file
set           Set CLI properties, date/time, craft interface message
show          Show system information
start         Start shell
test          Perform diagnostic debugging
{master}
qfabric-operator@node-ee3093>

```

#### request component login (with qfabric-user Privileges)

```

user0@qfabric> request component login EE3093
error: User user0 does not have sufficient permissions to login to device ee3093

```

## show ethernet-switching interfaces

<b>Syntax</b>	show ethernet-switching interfaces <brief   detail   summary> <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about switched Ethernet interfaces.
<b>Options</b>	<p><b>none</b>—(Optional) Display brief information for Ethernet-switching interfaces.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display Ethernet-switching information for a specific interface.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Troubleshooting Ethernet Switching on page 1895</a><a href="#">Understanding Bridging and VLANs on page 1527</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li> <li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li> <li>• <a href="#">Understanding FCoE on page 5518</a></li> <li>• <a href="#">Interfaces Overview on page 2389</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching interfaces on page 1483</a> <a href="#">show ethernet-switching interfaces summary on page 1484</a> <a href="#">show ethernet-switching interfaces brief on page 1484</a> <a href="#">show ethernet-switching interfaces detail on page 1484</a> <a href="#">show ethernet-switching interfaces interface-name on page 1485</a>
<b>Output Fields</b>	<a href="#">Table 85 on page 1482</a> lists the output fields for the <b>show ethernet-switching interfaces</b> command. Output fields are listed in the approximate order in which they appear.

**Table 85: show ethernet-switching interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of a switching interface.	All levels
<b>State</b>	Interface state. Values are <b>up</b> or <b>down</b> .	none, <b>brief</b> , <b>detail</b> , <b>summary</b>
<b>VLAN members</b>	Name of a VLAN.	none, <b>brief</b> , <b>detail</b> , <b>summary</b>

Table 85: show ethernet-switching interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Blocking</b>	Forwarding state of the interface: <ul style="list-style-type: none"> <li>• <b>blocked</b>—Traffic is not being forwarded on the interface.</li> <li>• <b>unblocked</b>—Traffic is forwarded on the interface.</li> <li>• <b>MAC limit exceeded</b>—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>MAC move limit exceeded</b>—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>Storm control in effect</b> —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>Storm control shutdown in effect</b> —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires.</li> </ul>	none, <b>brief</b> , <b>detail</b> , <b>summary</b>
<b>Index</b>	VLAN index internal to Junos OS software.	<b>detail</b>
<b>untagged   tagged</b>	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	<b>detail</b>

## Sample Output

### show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces
```

```

Interface  State  VLAN members  Blocking
xe-0/0/0.0  up    T1122         unblocked
xe-0/0/1.0  down  default       - MAC limit exceeded
xe-0/0/2.0  down  default       - MAC move limit exceeded
xe-0/0/3.0  down  default       - Storm control in effect
xe-0/0/4.0  down  default       unblocked
xe-0/0/5.0  down  default       unblocked
xe-0/0/6.0  down  default       unblocked
xe-0/0/7.0  down  default       unblocked
xe-0/0/8.0  down  default       unblocked
xe-0/0/9.0  up    T111         unblocked
xe-0/0/10.0 down  default       unblocked
xe-0/0/11.0 down  default       unblocked
xe-0/0/12.0 down  default       unblocked
xe-0/0/13.0 down  default       unblocked
xe-0/0/14.0 down  default       unblocked
xe-0/0/15.0 down  default       unblocked
xe-0/0/16.0 down  default       unblocked
xe-0/0/17.0 down  default       unblocked
xe-0/0/18.0 down  default       unblocked
xe-0/0/19.0 up    T111         unblocked
xe-0/1/0.0  down  default       unblocked
xe-0/1/1.0  down  default       unblocked
xe-0/1/2.0  down  default       unblocked
xe-0/1/3.0  down  default       unblocked

```

### show ethernet-switching interfaces summary

```
user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0
```

### show ethernet-switching interfaces brief

```
user@switch> show ethernet-switching interfaces brief
Interface  State  VLAN members  Blocking
xe-0/0/0.0  down  default       unblocked
xe-0/0/1.0  down  employee-vlan unblocked
xe-0/0/2.0  down  employee-vlan unblocked
xe-0/0/3.0  down  employee-vlan unblocked
xe-0/0/8.0  down  employee-vlan unblocked
xe-0/0/10.0 down  default       unblocked
xe-0/0/11.0 down  employee-vlan unblocked
```

### show ethernet-switching interfaces detail

```
user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/1.0 Index: 66
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/2.0 Index: 67
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/3.0 Index: 68
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/8.0 Index: 69
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/10.0 Index: 70
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/11.0 Index: 71
State: down
VLANs:
    employee-vlan          tagged      unblocked
```

**show ethernet-switching interfaces interface-name**

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
  Interface  State   VLAN members   Blocking
xe-0/0/0.0  down    default         unblocked
```

## show lldp

**Syntax** `show lldp`  
`<detail>`

**Release Information** Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display information about Link Layer Discovery Protocol (LLDP) and Link Level Discovery Protocol—Media Endpoint Discovery (LLDP-MED) configuration and capabilities on the switch. LLDP and LLDP-MED are used to learn about and to distribute device information on network links.



**NOTE:** LLDP-MED is not available on the QFX Series.

**Options** **none**—Display LLDP information for all interfaces.  
**detail**—(Optional) Display detailed LLDP information for all interfaces.

**Required Privilege Level** view

**Related Documentation**

- [Configuring LLDP \(CLI Procedure\)](#)
- [Configuring LLDP-MED \(CLI Procedure\)](#)
- [Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches](#)
- [Configuring LLDP on page 1345](#)
- [Understanding LLDP on page 1319](#)

**List of Sample Output** [show lldp \(EX3200 switches\) on page 1489](#)  
[show lldp \(EX4300 switches\) on page 1489](#)  
[show lldp detail \(EX4300 switches\) on page 1490](#)

**Output Fields** [Table 86 on page 1486](#) lists the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

**Table 86: show lldp Output Fields**

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be <b>enabled</b> or <b>disabled</b> .  <b>NOTE:</b> If a VLAN that has been configured for untagged packets on an interface also has Layer 2 protocol tunneling (L2PT) enabled for LLDP, the LLDP operating state for that interface is displayed as <b>disabled</b> .	All levels



Table 86: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Advertisement interval</b>	Frequency, in seconds, at which LLDP advertisements are sent.  This value is set by the <code>advertisement-interval</code> configuration statement.	All levels
<b>Transmit delay</b>	Seconds of delay before advertisements are sent to neighbors following a change to a TLV (type, length, or value) element in the LLDP protocol or to the state of the local system, such as a change in hostname or management address. You can set this value to reduce the delay in notifying neighbors of a change in the local system.  This value is set by the <code>transmit-delay</code> configuration statement.	All levels
<b>Hold timer</b>	On EX4300 switches, the hold timer shows the length of time LLDP information is held before it is discarded. The hold timer value is equal to the advertisement interval multiplied by the hold multiplier.  On all other switches, the hold timer shows the value of the hold multiplier.  The hold multiplier value is set by the <code>hold-multiplier</code> configuration statement.	All levels
<b>Notification interval</b>	How often LLDP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, LLDP trap notifications on database changes are disabled.  This value is set by the <code>lldp-configuration-notification-interval</code> configuration statement.	All levels
<b>Config Trap Interval</b>	How often LLDP trap notifications are generated as a result of changes in topology—for example, when an endpoint connects or disconnects. If the interval value is 0, LLDP trap notifications on topology changes are disabled.  This value is set by the <code>ptopo-configuration-trap-interval</code> configuration statement.	All levels
<b>Connection Hold timer</b>	Amount of time the system maintains dynamic topology entries.  This value is set by the <code>ptopo-configuration-maximum-hold-time</code> configuration statement.	All levels
<b>LLDP-MED</b>	LLDP-MED operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>MED fast start count</b>	Number of advertisements sent from a switch to a device, such as a VoIP telephone, when the device is first detected by the switch. These increased advertisements are temporary. After a device and a switch exchange information and can communicate, advertisements are reduced to one per second.  This value is set by using the <code>fast-start</code> configuration statement.  <b>NOTE:</b> <code>fast-start</code> is not available on the QFX Series.	All levels
<b>Interface</b>	Name of the interface for which LLDP configuration information is being reported.	All levels
<b>Parent Interface</b>	Name of the aggregated Ethernet interface, if any, to which the interface belongs.	All levels

Table 86: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .	All levels
Power Negotiation	LLDP power negotiation operating state. The state can be <b>Enabled</b> or <b>Disabled</b> .	All levels
Neighbor count	Total number of new LLDP neighbors detected since the last switch reboot.	<b>detail</b>
Interface	Name of the interface that is advertising VLAN information.	All levels
Vlan-id	VLAN tag associated with the interface sending LLDP frames. If the interface is not a member of a VLAN, the VLAN ID is advertised as 0.	<b>detail</b>
Vlan-name	VLAN name associated with the VLAN ID.	<b>detail</b>
LLDP basic TLVs supported	<p>Basic TLVs supported on the switch:</p> <ul style="list-style-type: none"> <li>• <b>Chassis identifier</b>—TLV that advertises the MAC address associated with the local system.</li> <li>• <b>Port identifier</b>—TLV that advertises the port identification for the specified port in the local system.</li> <li>• <b>Port description</b>—Interface name for the port.</li> <li>• <b>System name</b>—TLV that advertises the user-configured name of the local system.</li> <li>• <b>System description</b>—TLV that advertises the system description containing information about the software and current image running on the system. This information is taken from the software and is not configurable.</li> <li>• <b>System capabilities</b>—TLV that advertises the primary functions performed by the system—for example, bridge or router.</li> <li>• <b>Management address</b>—TLV that advertises the IP management address of the local system.</li> </ul>	<b>detail</b>
Supported LLDP 802 TLVs	<p>802.3 TLVs supported on the switch:</p> <ul style="list-style-type: none"> <li>• <b>MAC/PHY configuration status</b>—TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is based on the physical interface structure and is not configurable.</li> <li>• <b>Power via MDI</b>—TLV that advertises MDI power support, PSE power pair, and power class information.</li> <li>• <b>Link aggregation</b>—TLV that advertises if the interface is aggregated and its aggregated interface ID.</li> <li>• <b>Maximum frame size</b>—TLV that advertises the maximum transmission unit (MTU) of the interface sending LLDP frames.</li> <li>• <b>Port VLAN tag</b>—TLV that advertises the VLAN tag configured on the interface.</li> <li>• <b>Port VLAN name</b>—TLV that advertises the VLAN name configured on the interface.</li> </ul>	<b>detail</b>

Table 86: show lldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Supported LLDP MED TLVs	<p>LLDP-MED TLVs supported on the switch:</p> <ul style="list-style-type: none"> <li>• <b>LLDP MED capabilities</b>—TLV that advertises the primary function of the port. The capabilities values range from 0 through 15: <ul style="list-style-type: none"> <li>• 0—Capabilities</li> <li>• 1—Network Policy</li> <li>• 2—Location Identification</li> <li>• 3—Extended Power via MDI-PSE</li> <li>• 4—Inventory</li> <li>• 5–15—Reserved</li> </ul> </li> <li>• <b>Network policy</b>—TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types—such as voice or streaming video—802.1Q VLAN tagging, and 802.1p priority bits and DiffServ code points.</li> <li>• <b>Endpoint location</b>—TLV that advertises the physical location of the endpoint.</li> <li>• <b>Extended power Via MDI</b>—TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.</li> </ul>	detail

## Sample Output

### show lldp (EX3200 switches)

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                        : 4 seconds
Notification interval              : 0 Second(s)
Config Trap Interval               : 0 seconds
Connection Hold timer              : 300 seconds

LLDP MED                           : Disabled
MED fast start count               : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
all	-	Enabled	Enabled	Enabled

### show lldp (EX4300 switches)

```

user@switch> show lldp
LLDP                               : Enabled
Advertisement interval             : 30 seconds
Transmit delay                     : 2 seconds
Hold timer                        : 120 seconds
Notification interval              : 0 Second(s)
Config Trap Interval               : 0 seconds
Connection Hold timer              : 300 seconds

LLDP MED                           : Disabled
MED fast start count               : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
all	-	Enabled	Enabled	Enabled

**show lldp detail (EX4300 switches)**

```

user@switch> show lldp detail
LLDP : Enabled
Advertisement interval : 30 seconds
Transmit delay : 2 seconds
Hold timer : 120 seconds
Notification interval : 0 Second(s)
Config Trap Interval : 0 seconds
Connection Hold timer : 300 seconds

LLDP MED : Disabled
MED fast start count : 3 Packets

```

Interface	Parent Interface	LLDP	LLDP-MED	Power Negotiation
Neighbor count				
all	-	Enabled	Enabled	Enabled
8				

Interface	Parent Interface	Vlan-id	Vlan-name
xe-3/0/0.0	ae31.0	100	v100
xe-3/0/0.0	ae31.0	101	v101
xe-3/0/0.0	ae31.0	4000	v4000
xe-3/0/1.0	ae31.0	100	v100
xe-3/0/1.0	ae31.0	101	v101
xe-3/0/1.0	ae31.0	4000	v4000
xe-3/0/2.0	ae31.0	100	v100
xe-3/0/2.0	ae31.0	101	v101
xe-3/0/2.0	ae31.0	4000	v4000

**LLDP basic TLVs supported:**

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

**Supported LLDP 802 TLVs:**

MAC/PHY configuration/status, Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port VLAN name.

**Supported LLDP MED TLVs:**

LLDP MED capabilities, Network policy, Endpoint location, Extended power Via MDI.

## show lldp local-information

<b>Syntax</b>	show lldp local-information
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the information that the switch provides in Link Layer Discovery Protocol (LLDP) advertisements to its neighbors.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring LLDP (CLI Procedure)</i></li> <li>• <i>Understanding 802.1X and LLDP and LLDP-MED on EX Series Switches</i></li> <li>• <a href="#">management-address on page 1424</a></li> <li>• <a href="#">Configuring LLDP on page 1345</a></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show lldp local-information (EX Series Switch) on page 1492</a>
<b>Output Fields</b>	<a href="#">Table 87 on page 1491</a> lists the output fields for the <b>show lldp local-information</b> command. Output fields are listed in the approximate order in which they appear.

**Table 87: show lldp local-information Output Fields**

Field Name	Field Description
<b>LLDP Local Information details</b>	<p>Information about the local system (the switch):</p> <ul style="list-style-type: none"> <li>• <b>Chassis ID</b>—MAC address associated with the switch.</li> <li>• <b>System name</b>—User-configured name of the switch.</li> <li>• <b>System descr</b>—System description containing information about the switch model and the current software image running on the switch. This information is taken from the software and is not configurable.</li> </ul>
<b>System Capabilities</b>	Capabilities (such as <b>bridge</b> or <b>router</b> ) that are supported or enabled on the system.
<b>Management Information</b>	<p>Details of the management information: <b>Port Name</b>, <b>Port Address</b> (such as 10.204.34.35), <b>Address Type</b> (such as ipv4 or ipv6), <b>Port ID</b> (SNMP interface index), <b>Port ID Subtype</b>, and <b>Port Subtype</b>.</p> <p>The <b>Port Subtype</b> displays:</p> <ul style="list-style-type: none"> <li>• <b>ifindex(2)</b>—IP address of the switch's management Ethernet interface (<b>me0</b>) or virtual management Ethernet (<b>VME</b>) interface address (for a virtual chassis) is used to manage the switch.</li> <li>• <b>unknown(1)</b>—IP management address has been configured with set <b>protocols lldp management-address</b>.</li> </ul>

Table 87: show lldp local-information Output Fields (*continued*)

Field Name	Field Description
<b>Interface name</b>	Name of the local interface which is configured for either LLDP or LLDP-MED.
<b>Parent Interface</b>	Name of the aggregated Ethernet interface, if any, to which the local interface belongs.
<b>SNMP Index</b>	SNMP interface index.
<b>Interface description</b>	User-configured port description.
<b>Status</b>	Administrative status of the interface: either <b>up</b> or <b>down</b> .
<b>Tunneling</b>	Status of tunneling on the interface: either <b>enabled</b> or <b>disabled</b> .

## Sample Output

### show lldp local-information (EX Series Switch)

```
user@switch> show lldp local-information
```

#### LLDP Local Information details

```
Chassis ID   : 00:1d:b5:aa:b9:f0
System name  : switch
System descr : Juniper Networks, Inc. ex8208 , version 10.4I0 [builder] Build
               date: 2010-11-17 12:38:30 UTC
```

#### System Capabilities

```
Supported   : Bridge Router
Enabled     : Bridge Router
```

#### Management Information

```
Port Name    : -
Port Address : 10.93.54.6
Address Type  : IPv4
Port ID       : 34
Port ID Subtype : local(7)
Port Subtype  : ifIndex(2)
```

Interface name	Parent Interface	SNMP Index	Interface description	Status	Tunneling
me0.0	-	34	-	Down	Disabled
xe-3/0/0.0	ae31.0	769	xe-3/0/0.0	Up	Disabled
xe-3/0/1.0	ae31.0	770	xe-3/0/1.0	Up	Disabled
xe-3/0/2.0	ae31.0	771	xe-3/0/2.0	Up	Disabled
xe-3/0/3.0	ae31.0	772	xe-3/0/3.0	Up	Disabled
xe-3/0/4.0	ae31.0	577	xe-3/0/4.0	Up	Disabled
xe-3/0/5.0	ae31.0	578	xe-3/0/5.0	Up	Disabled
xe-3/0/6.0	ae31.0	579	xe-3/0/6.0	Up	Disabled
xe-3/0/7.0	ae31.0	581	xe-3/0/7.0	Up	Disabled

## show lldp neighbors

**Syntax** <show lldp *neighbors*>  
<interface *interface-ids*>

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display learned information about Link Layer Discovery Protocol (LLDP) on all neighboring interfaces or on selected interfaces.

**Options** **none**—Display learned LLDP information on all neighboring interfaces and devices.

**interface *interface-ids***—(Optional) Display learned LLDP information on the selected interfaces or devices.



**NOTE:** When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors in order to interoperate with a wider variety of converged network adapters (CNAs). As a result, information for those ports will not be listed in the output for this command.

**Required Privilege Level** view

**Related Documentation**

- [Configuring LLDP on page 1345](#)
- [Understanding LLDP on page 1319](#)

**List of Sample Output** [show lldp neighbors on page 1495](#)  
[show lldp neighbors interface on page 1496](#)

**Output Fields** [Table 88 on page 1493](#) lists the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

**Table 88: show lldp neighbors Output Fields**

Field Name	Field Description
Local Interface	List of local interfaces for which neighbor information is available.
Parent Interface	List of aggregated Ethernet interfaces, if any, to which the local interfaces belong.
Chassis ID	List of chassis identifiers for neighbors.
Port info	List of port information gathered from neighbors. This could be the port identifier or port description.
System name	List of system names gathered from neighbors.

Table 88: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
<b>LLDP Neighbor Information</b>	Information about both the local system (the switch) and a neighbor system on the interface (appears when the <b>interface</b> option is used).
<b>Local Information</b>	Information about the local system (appears when the <b>interface</b> option is used).
<b>Index</b>	Local interface index (appears when the <b>interface</b> option is used).
<b>Time to live</b>	Number of seconds for which this information is valid (appears when the <b>interface</b> option is used).
<b>Time mark</b>	Date and timestamp of information (appears when the <b>interface</b> option is used).
<b>Local Interface</b>	Name of the local physical interface (appears when the <b>interface</b> option is used).
<b>Parent Interface</b>	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the <b>interface</b> option is used).
<b>Local Port ID</b>	Local interface SNMP index (appears when the <b>interface</b> option is used).
<b>Ageout Count</b>	Number of times the complete set of information advertised by the neighbor has been deleted from LLDP neighbor information maintained by the local system because the information timeliness interval has expired (appears when the <b>interface</b> option is used).
<b>Neighbor Information</b>	Information about a neighbor system on the interface (appears when the <b>interface</b> option is used).
<b>Chassis type</b>	Type of chassis identifier supplied, such as <b>MAC address</b> (appears when the <b>interface</b> option is used).
<b>Chassis ID</b>	Chassis identifier of the chassis type listed (appears when the <b>interface</b> option is used).
<b>Port type</b>	Type of port identifier supplied, such as <b>locally assigned</b> (appears when the <b>interface</b> option is used).
<b>Port ID</b>	Port identifier of the port type listed (appears when the <b>interface</b> option is used).
<b>Port description</b>	Port description (appears when the <b>interface</b> option is used).
<b>System name</b>	Name supplied by the system on the interface (appears when the <b>interface</b> option is used).
<b>System Description</b>	Description supplied by the system on the interface (appears when the <b>interface</b> option is used).



Table 88: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
System capabilities	Capabilities (such as <b>Bridge</b> , <b>Router</b> , and <b>Telephone</b> ) that are supported or enabled by the system on the interface (appears when the <b>interface</b> option is used).
Management Info	<p>Details of management information: <b>Type</b> (such as <b>ipv4</b> or <b>ipv6</b>), <b>Address</b> (such as <b>10.204.34.35</b>), <b>Port ID</b>, <b>Subtype</b>, <b>Interface Subtype</b>, and organization identifier (<b>OID</b>) (appears when the <b>interface</b> option is used).</p> <p>The <b>Interface Subtype</b> displays:</p> <ul style="list-style-type: none"> <li>• <b>ifindex(2)</b>— IP address of the neighbor's management Ethernet interface (<b>me0</b>) or virtual management Ethernet (<b>VME</b>) interface address (for a virtual chassis) is used to manage the switch.</li> <li>• <b>unknown(1)</b>—Neighbor's IP management address has been configured with set <b>protocols lldp management-address</b>.</li> </ul>
Media Info	Additional details about the endpoint device appear when a device that supports LLDP-MED is attached to the interface. The specific details depend upon the capabilities of the device. Details might include <b>Media endpoint class</b> (such as Class 3 for communication devices such as IP phones), <b>MED Hardware revision</b> , <b>MED Firmware revision</b> , <b>MED Software revision</b> , <b>MED Serial number</b> , <b>MED Manufacturer name</b> , or <b>MED Model name</b> .
Organization Info	One or more entries listing remote information by organizationally unique identifier ( <b>OUI</b> ), <b>Subtype</b> , <b>Index</b> , and <b>Info</b> (appears when the <b>interface</b> option is used).
Age	How long the neighbor has been identified (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).
Local Interface	Name of the local physical interface (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).
Chassis ID	Chassis identifier of the chassis type listed (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).
Port description	Port description (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).
System name	NetBIOS name of the host (appears when the <b>interface</b> option is used and NetBIOS snooping is enabled on the switch).

## Sample Output

### show lldp neighbors

```
user@switch> show lldp neighbors
```

Local Interface	Parent Interface	Chassis Id	Port info	System Name
xe-3/0/4.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/0.0	newyork31
xe-3/0/5.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/1.0	newyork31
xe-3/0/6.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/2.0	newyork31
xe-3/0/7.0	ae31.0	b0:c6:9a:63:80:40	xe-0/0/3.0	newyork31
xe-3/0/0.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/0.0	newyork31
xe-3/0/1.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/1.0	newyork31
xe-3/0/2.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/2.0	newyork31
xe-3/0/3.0	ae31.0	b0:c6:9a:63:80:40	xe-0/1/3.0	newyork31

### show lldp neighbors interface

```
user@switch> show lldp neighbors interface ge-0/0/2
```

#### LLDP Neighbor Information:

##### Local Information:

```
Index: 1 Time to live: 240 Time mark: Wed Dec 1 10:23:24 2010 Age: 29 secs
Local Interface   : ge-0/0/2.0
Parent Interface  : -
Local Port ID     : 507
Ageout Count      : 0
```

##### Neighbour Information:

```
Chassis type      : Mac address
Chassis ID        : 00:1f:12:38:7f:c0
Port type         : Locally assigned
Port ID           : 507
Port description  : ge-0/0/2.0
System name       : bng-148p5-dev
```

```
System Description : Juniper Networks, Inc. ex4200-48p , version 10.4I0 Build
date: 2010-11-30 09:32:17 UTC
```

#### System capabilities

```
Supported : Bridge Router
Enabled   : Bridge Router
```

#### Management Info

```
Type           : IPv4
Address         : 10.204.96.235
Port ID        : 34
Subtype        : 1
Interface Subtype : ifIndex(2)
OID            : 1.3.6.1.2.1.31.1.1.1.1.34
```

```
Media endpoint class: Network Connectivity
```

#### Organization Info

```
OUI      : 0.12.f
Subtype  : 1
Index    : 1
Info     : 22A8360000
```

#### Organization Info

```
OUI      : 0.12.f
Subtype  : 2
Index    : 2
Info     : 030100
```

## show lldp statistics

<b>Syntax</b>	<code>show lldp statistics</code> <code>&lt;interface <i>interface-ids</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display LLDP statistics on all or selected interfaces.
<b>Options</b>	<p><b>none</b>—Display LLDP statistics on all interfaces and devices.</p> <p><b>interface <i>interface-ids</i></b>—(Optional) Display LLDP statistics on the selected devices.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LLDP on page 1345</a></li> <li>• <a href="#">Understanding LLDP on page 1319</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show lldp statistics on page 1497</a>
<b>Output Fields</b>	<a href="#">Table 89 on page 1497</a> lists the output fields for the <b>show lldp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 89: show lldp statistics Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of an interface.	All levels
<b>Received</b>	Total number of LLDP frames received on an interface.	All levels
<b>Unknown-TLVs</b>	Number of unrecognized LLDP TLVs received on an interface.	All levels
<b>With Errors</b>	Number of LLDP frames received that contain errors.	All levels
<b>Discarded TLVs</b>	Number of LLDP TLVs received and then discarded on an interface.	All levels
<b>Transmitted</b>	Total number of LLDP frames transmitted on an interface.	All levels
<b>Untransmitted</b>	Total number of LLDP frames not transmitted on an interface.	All levels

## Sample Output

### show lldp statistics

```
user@switch> show lldp statistics
```

```

Interface  Received  Unknown TLVs  With Errors  Discarded TLVs  Transmitted
Untransmitted
me0.0      0         0             0            0               8003         0

```

ge-0/0/0.0 8002	0	0	0	8003	0
ge-0/0/1.0 8002	0	0	0	8003	0

## show route instance

<b>Syntax</b>	show route instance <brief   detail   summary> <instance-name> <operational>
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	(QFabric systems only) Display routing instance information.
<b>Options</b>	<p><b>none</b>—(Same as <b>brief</b>) Display standard information about all routing instances.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>. (These options are not available with the <b>operational</b> keyword.)</p> <p><b>instance-name</b>—(Optional) Display information for a specified routing instance.</p> <p><b>operational</b>—(Optional) Display operational routing instances.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show route instance on page 1500</a> <a href="#">show route instance detail on page 1500</a> <a href="#">show route instance operational on page 1501</a> <a href="#">show route instance summary on page 1501</a>
<b>Output Fields</b>	Table 90 on page 1499 lists the output fields for the <b>show route instance</b> command. Output fields are listed in the approximate order in which they appear.

Table 90: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels
Operational Routing Instances	( <b>operational</b> keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: <b>forwarding</b> or <b>virtual-router</b> .	All levels
State	State of the routing instance: <b>active</b> or <b>inactive</b> .	<b>detail</b>
Interfaces	Name of interfaces belonging to this routing instance.	<b>detail</b>
Tables	Tables (and number of routes) associated with this routing instance.	<b>detail</b>
Router ID	Identifier for the router.	<b>detail</b>

Table 90: show route instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Primary RIB	Primary table for this routing instance.	<b>brief none summary</b>
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

## Sample Output

### show route instance

```

user@switch> show route instance
Instance      Type
Primary RIB
master        forwarding
              inet.0
              4/0/1

__juniper_private1__ forwarding
              __juniper_private1__.inet.0
              1/0/3

__juniper_private2__ forwarding
              __juniper_private2__.inet.0
              0/0/1

__juniper_private3__ forwarding
              __juniper_private3__.inet.0
              1/0/2

__juniper_private4__ forwarding
              __juniper_private4__.inet.0
              4/0/2

__master.anon__ forwarding

r1            virtual-router

r2            virtual-router

```

### show route instance detail

```

user@switch> show route instance detail
master:
  Router ID: 3.3.3.7
  Type: forwarding      State: Active
  Tables:
    inet.0              : 5 routes (4 active, 0 holddown, 1 hidden)

__juniper_private1__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active
  Interfaces:
    lo0.16385
    bme0.0
  Tables:
    __juniper_private1__.inet.0: 6 routes (1 active, 0 holddown, 3 hidden)

__juniper_private2__:
  Router ID: 0.0.0.0
  Type: forwarding      State: Active
  Interfaces:
    lo0.16384

```

```

Tables:
  __juniper_private2__.inet.0: 1 routes (0 active, 0 holddown, 1 hidden)

__juniper_private3__:
Router ID: 0.0.0.0
Type: forwarding      State: Active
Interfaces:
  bme0.1
Tables:
  __juniper_private3__.inet.0: 4 routes (1 active, 0 holddown, 2 hidden)

__juniper_private4__:
Router ID: 0.0.0.0
Type: forwarding      State: Active
Interfaces:
  bme0.2
Tables:
  __juniper_private4__.inet.0: 8 routes (4 active, 0 holddown, 2 hidden)

__master.anon__:
Router ID: 0.0.0.0
Type: forwarding      State: Active

r1:
Router ID: 0.0.0.0
Type: virtual-router  State: Active
Interfaces:
  xe-0/0/0.0

r2:
Router ID: 0.0.0.0
Type: virtual-router  State: Active
Interfaces:
  xe-0/0/3.0

```

### show route instance operational

```

user@switch> show route instance operational
Operational Routing Instances:

__juniper_private1__
__juniper_private2__
__juniper_private3__
__juniper_private4__
r1---qfabric
r2---qfabric
master

```

### show route instance summary

```

user@switch> show route instance summary

```

Instance	Type	Primary RIB	Active/holddown/hidden
master	forwarding	inet.0	4/0/1
__juniper_private1__	forwarding	__juniper_private1__.inet.0	1/0/3
__juniper_private2__	forwarding	__juniper_private2__.inet.0	0/0/1

__juniper_private3__ forwarding	
__juniper_private3__.inet.0	1/0/2
__juniper_private4__ forwarding	
__juniper_private4__.inet.0	4/0/2
__master.anon__ forwarding	
r1	virtual-router
r2	virtual-router



## show snmp statistics

<b>Syntax</b>	show snmp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear snmp statistics on page 6857</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show snmp statistics on page 1506</a>
<b>Output Fields</b>	<a href="#">Table 91 on page 1503</a> describes the output fields for the <b>show snmp statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 91: show snmp statistics Output Fields

Field Name	Field Description
<b>Input</b>	<p>Information about received packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets(snmplnPkts)</b>—Total number of messages delivered to the SNMP entity from the transport service.</li> <li>• <b>Bad versions—(snmplnBadVersions)</b> Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version.</li> <li>• <b>Bad community names—(snmplnBadCommunityNames)</b> Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity.</li> <li>• <b>Bad community uses—(snmplnBadCommunityUses)</b> Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.</li> <li>• <b>ASN parse errors—(snmplnASNParseErrs)</b> Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.</li> <li>• <b>Too big—(snmplnTooBigs)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmplnNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmplnBadValues)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>Read only—(snmplnReadOnlys)</b> Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of <b>readOnly</b>. Only incorrect implementations of SNMP generate this error.</li> </ul>

Table 91: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> <li>• <b>General errors—(snmpInGenErrs)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Total requests varbinds—(snmpInTotalReqVars)</b> Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP <b>GetRequest</b> and <b>GetNext</b> PDUs.</li> <li>• <b>Total set varbinds—(snmpInSetVars)</b> Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP <b>SetRequest</b> PDUs.</li> <li>• <b>Get requests—(snmpInGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get nexts—(snmpInGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Set requests—(snmpInSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get responses—(snmpInGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Traps—(snmpInTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> <li>• <b>Silent drops—(snmpSilentDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.</li> <li>• <b>Proxy drops—(snmpProxyDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.</li> <li>• <b>Commit pending drops</b>—Number of SNMP packets for <b>Set</b> requests dropped because of a previous pending SNMP <b>Set</b> request on the committed configuration.</li> <li>• <b>Throttle drops</b>—Number of SNMP packets for any requests dropped reaching the throttle limit.</li> </ul>

Table 91: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> <li>• <b>Unknown security models—(snmpUnknownSecurityModels)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine.</li> <li>• <b>Invalid messages—(snmpInvalidMsgs)</b> Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message.</li> <li>• <b>Unknown pdu handlers—(snmpUnknownPDUHandlers)</b> Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type.</li> <li>• <b>Unavailable contexts—(snmpUnavailableContexts)</b> Number of requests received for a context that is known to the SNMP engine, but is currently unavailable.</li> <li>• <b>Unknown contexts—(snmpUnknownContexts)</b> Total number of requests received for a context that is unknown to the SNMP engine.</li> <li>• <b>Unsupported security levels—(usmStatsUnsupportedSecLevels)</b> Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable).</li> <li>• <b>Not in time windows—(usmStatsNotInTimeWindows)</b> Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window.</li> <li>• <b>Unknown user names—(usmStatsUnknownUserNames)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine.</li> <li>• <b>Unknown engine ids—(usmStatsUnknownEngineIDs)</b> Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine.</li> <li>• <b>Wrong digests—(usmStatsWrongDigests)</b> Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value.</li> <li>• <b>Decryption errors—(usmStatsDecryptionErrors)</b> Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.</li> </ul>

Table 91: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Output</b>	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets—(snmpOutPkts)</b> Total number of messages passed from the SNMP entity to the transport service.</li> <li>• <b>Too big—(snmpOutTooBig)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmpOutNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmpOutBadValues)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>General errors—(snmpOutGenErrs)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Get requests—(snmpOutGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get nexts—(snmpOutGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs generated by the SNMP entity.</li> <li>• <b>Set requests—(snmpOutSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get responses—(snmpOutGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs generated by the SNMP entity.</li> <li>• <b>Traps—(snmpOutTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> </ul>

## Sample Output

### show snmp statistics

```

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0

```

## ssh

**List of Syntax**    [Syntax on page 1507](#)  
                           [Syntax \(EX Series Switch and the QFX Series\) on page 1507](#)

**Syntax**    `ssh host`  
                   `<bypass-routing>`  
                   `<inet | inet6>`  
                   `<interface interface-name>`  
                   `<logical-system logical-system-name>`  
                   `<routing-instance routing-instance-name>`  
                   `<source address>`  
                   `<v1 | v2>`

**Syntax (EX Series Switch and the QFX Series)**    `ssh host`  
                                                           `<bypass-routing>`  
                                                           `<inet | inet6>`  
                                                           `<interface interface-name>`  
                                                           `<routing-instance routing-instance-name>`  
                                                           `<source address>`  
                                                           `<v1 | v2>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Use the SSH program to open a connection between a local router or switch and a remote system and execute commands on the remote system. You can issue the **ssh** command from the Junos OS CLI to log in to a remote system or from a remote system to log in to the local router or switch. When executing this command, you include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons:

```
ssh address 'cli-command1 ; cli-command2 '
```

**Options**    **host**—Name or address of the remote system.

**bypass-routing**—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

**inet | inet6**—(Optional) Create an IPv4 or IPv6 connection, respectively.

**interface interface-name**—(Optional) Interface name for the SSH session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

**logical-system logical-system-name**—(Optional) Name of a particular logical system for the SSH attempt.

**routing-instance** *routing-instance-name*—(Optional) Name of the routing instance for the SSH attempt.

**source address**—(Optional) Source address of the SSH connection.

**v1 | v2**—(Optional) Use SSH version 1 or 2, respectively, when connecting to a remote host.

**Additional Information** To configure an SSH (version 1) key for your user account, include the **authentication ssh-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. To configure an SSH (version 2) key for your user account, include the **authentication dsa-rsa** statement at the **[edit system login user *user-name*]** hierarchy level.

You can limit the number of times a user can attempt to enter a password while logging in through SSH. To specify the number of times a user can attempt to enter a password to log in through SSH, include the **retry-options** statement at the **[edit system login]** hierarchy level. For details, see the .

**Required Privilege Level** network

**Related Documentation**

- [Configuring SSH Host Keys for Secure Copying of Data on page 1359](#)

**List of Sample Output** [ssh on page 1508](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

ssh

```
user@switch> ssh cree
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes

Host ?cree' added to the list of known hosts.
boojun@cree's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
```

## PART 5

# Ethernet Features

- [Overview on page 1511](#)
- [Configuration on page 1563](#)
- [Administration on page 1827](#)
- [Troubleshooting on page 1895](#)





## CHAPTER 14

# Overview

- [Enhanced Layer 2 Software \(ELS\) CLI on page 1511](#)
- [Bridging and VLANs on page 1525](#)
- [Layer 2 Networking on page 1540](#)
- [Understanding Q-in-Q Tunneling on page 1547](#)
- [Proxy ARP on page 1551](#)
- [Reflective Relay on page 1553](#)
- [Spanning Trees on page 1554](#)

### Enhanced Layer 2 Software (ELS) CLI

---

- [Getting Started with Enhanced Layer 2 Software on page 1511](#)

### Getting Started with Enhanced Layer 2 Software

- [Understanding Enhanced Layer 2 Software Support on page 1511](#)
- [Using the ELS Translator Tool on page 1512](#)
- [Configuring a VLAN on page 1513](#)
- [Configuring the Native VLAN Identifier on page 1514](#)
- [Configuring Layer 2 Interfaces on page 1514](#)
- [Configuring Layer 3 Interfaces on page 1514](#)
- [Configuring an IRB Interface on page 1515](#)
- [Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface on page 1515](#)
- [Enhanced Layer 2 CLI Configuration Statement and Command Changes on page 1516](#)

### Understanding Enhanced Layer 2 Software Support

---

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

ELS is available on the following EX Series switches and QFX Series devices.

Table 92: ELS Support

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

ELS is supported on the EX4300, EX4600, and EX9200 switches for all Junos OS releases, starting with the initial releases shown in [Table 8 on page 44](#).

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

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



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

### Using the ELS Translator Tool

The ELS Translator is a web-based tool that converts Junos OS Layer 2 configurations to Enhanced Layer 2 Software (ELS) configurations. This conversion tool supports all Juniper Networks EX Series, MX Series, and QFX Series platforms with ELS installed. The ELS Translator is hosted on Juniper Networks Customer Support website for EX Series switches, MX Series Universal Edge routers, and QFX Series switches and is available to registered users, internal users, partners, and premium service contract customers. You need to login using your Juniper Networks user name and password to access the ELS Translator tool.

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

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

1. Log onto your device using the console port.



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

2. Copy your entire existing configuration into another file. Save the file to a remote location. See [“Saving a Configuration to a File” on page 1257](#).
3. Retain the portion of your existing configuration related to management network connectivity (such as `[edit system]`). Delete all other top-level configuration hierarchy levels (such as `[edit interfaces]`, `[edit protocols]`, and `[edit vlans]`). Issue a **commit** operation to remove the deleted configuration hierarchy levels.
4. Perform the software upgrade. Reboot your device to complete the upgrade. See [“Software Installation Overview” on page 122](#)



**NOTE:** Maintain your console port connection during the reboot.

5. [Click](#) to access the ELS translator tool in a web browser. Follow the instructions on the page to update your configuration.
6. Return to your console port connection. When the switch has rebooted to complete the software upgrade, copy the configuration from the ELS Translator Tool onto your switch. See [“Uploading a Configuration File” on page 1261](#).
7. Commit the new configuration.



**NOTE:** It is possible a script might not translate correctly, so review translated scripts carefully before loading the converted configuration on your switch or other device.

## Configuring a VLAN

You can configure one or more VLANs to perform Layer 2 bridging. The Layer 2 bridging functions include integrated routing and bridging (IRB) for support for Layer 2 bridging and Layer 3 IP routing on the same interface. EX Series and QFX Series switches can function as Layer 2 switches, each with multiple bridging, or broadcast, domains that participate in the same Layer 2 network. You can also configure Layer 3 routing support for a VLAN.

To configure a VLAN:

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

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

2. Assign at least one interface to the VLAN:

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

---

### Configuring the Native VLAN Identifier

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

To configure the native VLAN ID:

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

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

2. Configure the native VLAN ID:

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

3. Assign the interface to the native VLAN ID:

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

---

### Configuring Layer 2 Interfaces

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

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for trunk interface mode:

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

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for access interface mode:

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

---

### Configuring Layer 3 Interfaces

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

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

To specify an IP address for the logical unit using IPv4:

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

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

To specify an IP address for the logical unit using IPv6:

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

### Configuring an IRB Interface

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

To configure an IRB interface:

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

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

2. Create an IRB logical interface:

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

3. Associate the IRB interface with the VLAN:

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

### Configuring an Aggregated Ethernet Interface and Configuring LACP on That Interface

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

To configure an aggregated Ethernet interface:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

To configure LACP:

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

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

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

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

---

### Enhanced Layer 2 CLI Configuration Statement and Command Changes

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



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

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

**Table 93: Enhanced Layer 2 CLI Changes**

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   analyzer {     name {       ...     }   } } </pre>	<pre> forwarding-options {   analyzer {     name {       ...     }   } } </pre>	Statements moved to different hierarchy.
<pre> ethernet-switching-options {   authentication-whitelist {     ...   } } </pre>	<pre> switch-options {   ...   authentication-whitelist {     ...   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   bpdu-block {     ...   } } </pre>	<pre> protocols {   layer2-control {     bpdu-block {       ...     }   } } </pre>	Statement moved to different hierarchy.
<pre> ethernet-switching-options {   dot1q-tunneling {     ether-type (0x8100   0x88a8   0x9100);     ...   } } </pre>	<pre> interfaces interface-name {   ether-options {     ethernet-switch-profile {       tag-protocol-id [tpids];     }   } }  interfaces interface-name {   aggregated-ether-options {     ethernet-switch-profile {       tag-protocol-id [tpids];     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   interfaces <i>interface-name</i> {     no-mac-learning;     ...   } } </pre>	<pre> switch-options {   interfaces <i>interface-name</i> {     no-mac-learning;     ...   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   mac-notification {     notification-interval <i>seconds</i>;     ...   } } </pre>	—	Statements deleted.
<pre> ethernet-switching-options {   mac-table-aging-time <i>seconds</i>;   ... } </pre>	<pre> protocols {   l2-learning {     global-mac-table-aging-time <i>seconds</i>;     ...   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   nonstop-bridging; } </pre>	<pre> protocols {   layer2-control {     nonstop-bridging {     }   } } </pre>	Statement moved to different hierarchy.
<pre> ethernet-switching-options {   port-error-disable {     disable-timeout <i>timeout</i>;     ...   } } </pre>	<pre> interfaces <i>interface-name</i> family   ethernet-switching {     recovery-timeout <i>seconds</i>;   } </pre>	Statement replaced with a new statement.
<pre> ethernet-switching-options {   redundant-trunk-group {     group <i>name</i> {       description;       interface <i>interface-name</i> {         primary;       }       preempt-cutover-timer <i>seconds</i>;       ...     }   } } </pre>	<pre> switch-options {   redundant-trunk-group {     group <i>name</i> {       description;       interface <i>interface-name</i> {         primary;       }       preempt-cutover-timer <i>seconds</i>;       ...     }   } } </pre>	Hierarchy renamed.



Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   secure-access-port {     interface (all   <i>interface-name</i>) {       (dhcp-trusted   no-dhcp-trusted );       static-ip <i>ip-address</i> {         mac <i>mac-address</i>;         vlan <i>vlan-name</i>;       }     }   }   vlan (all   <i>vlan-name</i>) {     (arp-inspection   no-arp-inspection );     dhcp-option82 {       disable;       circuit-id {         prefix <i>hostname</i>;         use-interface-description;         use-vlan-id;       }       remote-id {         prefix (<i>hostname</i>   mac   none);         use-interface-description;         use-string <i>string</i>;       }       vendor-id [<i>string</i>];     }     (examine-dhcp   no-examine-dhcp);   }   (ip-source-guard   no-ip-source-guard); } </pre>	<pre> vlangs <i>vlan-name</i> forwarding-options{   dhcp-security {     arp-inspection;     group <i>group-name</i> {       interface <i>interface-name</i> {         static-ip <i>ip-address</i> {           mac <i>mac-address</i>;         }       }     }     overrides {       no-option-82;       trusted;     }   }   ip-source-guard;   no-dhcp-snooping;   option-82 {     circuit-id {       prefix {         host-name;         routing-instance-name;       }       use-interface-description (device           logical);       use-vlan-id;     }     remote-id {       host-name;       use-interface-description (device           logical);       use-string <i>string</i>;     }     vendor-id {       use-string <i>string</i>;     }   } } </pre>	<p>Statements moved to different hierarchy.</p> <p><b>NOTE:</b> The statement <b>examine-dhcp</b> does not exist in the changed hierarchy. Instead, DHCP snooping is enabled automatically when other DHCP security features are enabled on a VLAN. See <i>Configuring Port Security (CLI Procedure)</i> for additional information.</p>
<pre> ethernet-switching-options {   secure-access-port {     dhcp-snooping-file {       location <i>local_pathname</i>   <i>remote_URL</i>;       timeout <i>seconds</i>;       write-interval <i>seconds</i>;     }   } } </pre>	<pre> system [   processes [     dhcp-service     dhcp-snooping-file <i>local_pathname</i>         <i>remote_URL</i>;     write-interval <i>interval</i>;   ] } </pre>	<p>Statement moved to different hierarchy.</p>
<pre> ethernet-switching-options {   secure-access-port vlan (all   <i>vlan-name</i>{     mac-move-limit   } } </pre>	<pre> vlangs <i>vlan-name</i> switch-options {   mac-move-limit } </pre>	<p>Statement moved to different hierarchy.</p>

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> ethernet-switching-options {   static {     vlan <i>vlan-id</i> {       mac <i>mac-address</i> next-hop         <i>interface-name</i>;       ...     }   } } </pre>	<pre> vlangs {   <i>vlan-name</i> {     switch-options {       interface <i>interface-name</i> {         static-mac <i>mac-address</i>;         ...       }     }   } } </pre>	Statement replaced with new statement and moved to different hierarchy.
<pre> ethernet-switching-options {   storm-control {     (...)   } } </pre>	<pre> forwarding-options {   storm-control-profiles <i>profile-name</i> {     (...)   } }  interfaces <i>interface-name</i> unit <i>number</i> family   ethernet-switching {     storm-control <i>storm-control-profile</i>;   } </pre>	Storm control configuration is done in two steps. The first step is to create a storm control profile at the [edit forwarding-options] hierarchy, and the second step is to bind the profile to a logical interface at the [edit interfaces] hierarchy. See <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i> for additional information.
<pre> ethernet-switching-options {   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt;       &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable           no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;;     ...   } } </pre>	—	Statements removed.
<pre> ethernet-switching-options {   unknown-unicast-forwarding {     (...)   } } </pre>	<pre> switch-options {   unknown-unicast-forwarding {     (...)   } } </pre>	Hierarchy renamed.
<pre> ethernet-switching-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	<pre> switch-options {   voip {     interface (all   [<i>interface-name</i>         access-ports]) {       forwarding-class (assured-forwarding           best-effort   expedited-forwarding           network-control);       vlan <i>vlan-name</i>;       ...     }   } } </pre>	Hierarchy renamed.

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> interfaces <i>interface-name</i> {   ether-options {     link-mode <i>mode</i>;     speed (auto-negotiation   <i>speed</i>)   } } </pre>	<pre> interfaces <i>interface-name</i> {   link-mode <i>mode</i>;   speed <i>speed</i> } </pre>	Statements moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       native-vlan-id <i>vlan-id</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   native-vlan-id <i>vlan-id</i> } </pre>	Statement moved to different hierarchy.
<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       port-mode <i>mode</i>     }   } } </pre>	<pre> interfaces <i>interface-name</i> {   unit <i>logical-unit-number</i> {     family ethernet-switching {       interface-mode <i>mode</i>     }   } } </pre>	Statement replaced with a new statement.
<pre> interfaces vlan </pre>	<pre> interfaces irb </pre>	Statement replaced with a new statement.

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> protocols {   igmp-snooping {     traceoptions {       file filename &lt;files number&gt;       &lt;no-stamp&gt; &lt;replace&gt;       &lt;size maximum-file-size&gt;       &lt;world-readable         no-world-readable&gt;;       flag flag &lt;flag-modifier&gt; &lt;disable&gt;;     }     vlan (all   vlan-identifier) {       disable;       data-forwarding {         receiver {           install;           source-vlans vlan-name;         }         source {           groups ip-address;         }       }       immediate-leave;       interface (all   interface-name) {         static {           group multicast-ip-address;         }         proxy {           source-address ip-address;         }       }       robust-count number;     }   } } </pre>	<pre> protocols {   igmp-snooping {     vlan vlan-name {       immediate-leave;       interface interface-name {         group-limit &lt;1..65535&gt;         host-only-interface         multicast-router-interface;         immediate-leave;         static {           group multicast-ip-address {             source &lt;&gt;           }         }       }     }     l2-querier {       source-address ip-address;     }     proxy {       source-address ip-address;     }     query-interval number;     query-last-member-interval number;     query-response-interval number;     robust-count number;     traceoptions {       file filename &lt;files number&gt;       &lt;no-stamp&gt; &lt;replace&gt;       &lt;size maximum-file-size&gt;       &lt;world-readable         no-world-readable&gt;;       flag flag &lt;flag-modifier&gt;;     }   } } </pre>	IGMP snooping is configured on a VLAN.
<pre> vlans {   vlan-name {     dot1q-tunneling {       customer-vlans (id   native   range);       layer2-protocol-tunneling all         protocol-name {         drop-threshold number;         shutdown-threshold number;         ...       }     }   } } </pre>	<pre> interface interface-name {   encapsulation extended-vlan-bridge;   flexible-vlan-tagging;   native-vlan-id number;   unit logical-unit-number {     input-vlan-map action;     output-vlan-map action;     vlan-id number;     vlan-id-list [vlan-id vlan-id-vlan-id];   } } </pre>	Statements replaced with new statements and moved to different hierarchy

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> vlsns {   vlan-name {     filter{       input filter-name       output filter-name;       ...     }   } } </pre>	<pre> vlsns {   vlan-name {     forwarding-options {       filter{         input filter-name         output filter-name;         ...       }     }   } } </pre>	Statements moved to different hierarchy.
<pre> vlsns {   vlan-name {     interface interface-name {       egress;       ingress;       mapping (native (push   swap)   policy           tag (push   swap));       pvlan-trunk;       ...     }   } } </pre>	—	Statements removed. You can assign interfaces to a VLAN using the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching vlan members <i>vlan-name</i> ] hierarchy.
<pre> vlsns {   vlan-name {     isolation-id id-number;     ...   } } </pre>	—	Statement removed.
<pre> vlsns {   vlan-name {     l3-interface vlan.logical-interface-number;     ...   } } </pre>	<pre> vlsns {   vlan-name {     l3-interface irb.logical-interface-number;     ...   } } </pre>	Syntax changed.
<pre> vlsns {   vlan-name {     l3-interface-ingress-counting       layer-3-interface-name;     ...   } } </pre>	—	Statement removed. Ingress traffic is automatically tracked.

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> vlands {   vlan-name {     mac-limit limit action action;     ...   } } </pre>	<pre> vlands {   vlan-name {     switch-options {       interface-mac-limit limit {         packet-action action;         ...       }     }   } }  vlands {   vlan-name {     switch-options {       interface interface-name {         interface-mac-limit limit {           packet-action action;           ...         }       }     }   } } </pre>	Statements moved to different hierarchies and renamed.
<pre> vlands {   vlan-name {     mac-table-aging-time seconds;     ...   } } </pre>	<pre> protocols {   l2-learning {     global-mac-table-aging-time seconds;     ...   } } </pre>	Statement moved to different hierarchy and renamed.
<pre> vlands {   vlan-name {     no-local-switching;     ...   } } </pre>	—	Statement removed.
<pre> vlands {   vlan-name {     no-mac-learning;     ...   } } </pre>	<pre> vlands {   vlan-name {     switch-options {       no-mac-learning limit       ...     }   } } </pre>	Statement moved to different hierarchy.
<pre> vlands {   vlan-name {     primary-vlan vlan-name;     ...   } } </pre>	—	Statement removed.

Table 93: Enhanced Layer 2 CLI Changes (*continued*)

Original Hierarchy	Changed Hierarchy	Change Description
<pre> vans {   vlan-name {     vlan-prune;     ...   } } </pre>	—	Statement removed.
<pre> vans {   vlan-name {     vlan-range vlan-id-low-vlan-id-high;     ...   } } </pre>	<pre> vans {   vlan-name {     vlan-id-list [vlan-id-numbers];     ...   } } </pre>	Statement replaced with new statement.

## Bridging and VLANs

- [Ethernet Ring Protection Switching Overview on page 1525](#)
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1526](#)
- [Understanding Bridging and VLANs on page 1527](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 1534](#)
- [Understanding Integrated Routing and Bridging on page 1539](#)
- [Understanding MAC Learning on page 1540](#)

## Ethernet Ring Protection Switching Overview

*Ethernet ring protection switching* (ERPS) helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. The RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An automatic protection switching (APS) protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.

The following standards provide detailed information on Ethernet ring protection switching:

- IEEE 802.1Q - 1998
- IEEE 802.1D - 2004
- IEEE 802.1Q - 2003

- Draft ITU-T Recommendation G.8032/Y.1344, *Ethernet Ring protection switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet-based networks*

For additional information on configuring Ethernet ring protection switching on EX Series switches, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*.

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *Layer 2 Configuration Guide* for a complete example of Ethernet rings and information about STP loop avoidance and prevention.

**Related  
Documentation**

- [Understanding Ethernet Ring Protection Switching Functionality on page 1534](#)
- *Configuring Ethernet Ring Protection Switching*
- *Example: Ethernet Ring Protection Switching Configuration on MX Routers*
- *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*
- *Ethernet Interfaces Feature Guide for Routing Devices*

## Layer 2 Learning and Forwarding for VLANs Overview

When you configure a VLAN, Layer 2 address learning is enabled by default. The VLAN learns unicast media access control (MAC) addresses to avoid flooding the packets to all the ports in the VLAN. Each VLAN creates a source MAC entry in its source and destination MAC tables for each source MAC address learned from packets received on the ports that belong to the VLAN.



**NOTE:** Traffic is not flooded back onto the interface on which it was received. However, because this “split horizon” occurs at a late stage, the packet statistics displayed by commands such as `show interfaces queue` will include flood traffic.

You can optionally disable MAC learning either for the entire device or for a specific VLAN or logical interface. You can also configure the following Layer 2 learning and forwarding properties:

- Static MAC entries for logical interfaces only
- Limit to the number of MAC addresses learned from a specific logical interface or from all the logical interfaces in a VLAN
- Size of the MAC address table for the VLAN
- MAC accounting for a VLAN

**Related  
Documentation**

- *Layer 2 Learning and Forwarding Overview*



## Understanding Bridging and VLANs

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN. This topic explains the following concepts regarding bridging and VLANs:

- [History of VLANs on page 1527](#)
- [How Bridging of VLAN Traffic Works on page 1527](#)
- [Packets Are Either Tagged or Untagged on page 1529](#)
- [Switch Interface Modes—Access, Trunk, or Tagged Access on page 1529](#)
- [Additional Advantages of Using VLANs on page 1531](#)
- [Maximum VLANs and VLAN Members Per Switch on page 1532](#)
- [A Default VLAN Is Configured on Most Switches on page 1532](#)
- [Assigning Traffic to VLANs on page 1533](#)
- [Forwarding VLAN Traffic on page 1533](#)
- [VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces on page 1534](#)

### History of VLANs

Ethernet LANs were originally designed for small, simple networks that primarily carried text. However, over time, the type of data carried by LANs grew to include voice, graphics, and video. This more complex data, when combined with the ever-increasing speed of transmission, eventually became too much of a load for the original Ethernet LAN design. Multiple packet collisions were significantly slowing down the larger LANs.

The IEEE 802.1D-2004 standard helped evolve Ethernet LANs to cope with the higher data and transmission requirements by defining the concept of *transparent bridging* (generally called simply *bridging*). Bridging divides a single physical LAN (now called a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each VLAN is a collection of some of the LAN nodes grouped together to form individual broadcast domains.

When VLANs are grouped logically by function or organization, a significant percentage of data traffic stays within the VLAN. This relieves the load on the LAN because all traffic no longer has to be forwarded to all nodes on the LAN. A VLAN first transmits packets within the VLAN, thereby reducing the number of packets transmitted on the entire LAN. Because packets whose origin and destination are in the same VLAN are forwarded only within the local VLAN, packets that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. This way, bridging and VLANs limit the amount of traffic flowing across the entire LAN by reducing the possible number of collisions and packet retransmissions within VLANs and on the LAN as a whole.

### How Bridging of VLAN Traffic Works

Because the objective of the IEEE 802.1D-2004 standard was to reduce traffic and therefore reduce potential transmission collisions for Ethernet, a system was implemented to reuse information. Instead of having a switch go through a location process every time

a frame is sent to a node, the transparent bridging protocol allows a switch to record the location of known nodes. When packets are sent to nodes, those destination node locations are stored in address-lookup tables called *Ethernet switching tables*. Before sending a packet, a switch using bridging first consults the switching tables to see if that node has already been located. If the location of a node is known, the frame is sent directly to that node.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The key bridging mechanism used by LANs and VLANs is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. As packets are sent, the switch learns the embedded MAC addresses of the sending nodes and stores them in the Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received on the destination node and the time the address was learned.

Learning allows switches to then do *forwarding*. By consulting the Ethernet switching table to see whether the table already contains the frame's destination MAC address, switches save time and resources when forwarding packets to the known MAC addresses. If the Ethernet switching table does not contain an entry for an address, the switch uses flooding to learn that address.

*Flooding* finds a particular destination MAC address without using the Ethernet switching table. When traffic originates on the switch and the Ethernet switching table does not yet contain the destination MAC address, the switch first floods the traffic to all other interfaces within the VLAN. When the destination node receives the flooded traffic, it can send an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and add the address to its Ethernet switching table.

*Filtering*, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—it learns which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

To keep entries in the Ethernet switching table current, the switch uses a fifth bridging mechanism, *aging*. Aging is the reason that the Ethernet switching table entries include timestamps. Each time the switch detects traffic from a MAC address, it updates the timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the

Ethernet switching table. This aging process eventually flushes unavailable network nodes out of the Ethernet switching table.

### Packets Are Either Tagged or Untagged

---

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q ID. The number of available VLANs and VLAN IDs are listed below:

- On a switch running ELS software, you can configure 4093 VLANs.
- On a switch running non-ELS software, you can configure 4091 VLANs.

Ethernet packets include a tag protocol identifier (TPID) EtherType field, which identifies the protocol being transported. When a device within a VLAN generates a packet, this field includes a value of 0x8100, which indicates that the packet is a VLAN-tagged packet. The packet also has a VLAN ID field that includes the unique 802.1Q ID, which identifies the VLAN to which the packet belongs.

In addition to the TPID EtherType value of 0x8100, switches that run Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style also support values of 0x88a8 (Provider Bridging and Shortest Path Bridging) and 0x9100 (Q-inQ).

For a simple network that has only a single VLAN, all packets include a default 802.1Q tag, which is the only VLAN membership that does not mark the packet as tagged. These packets are untagged packets.

### Switch Interface Modes—Access, Trunk, or Tagged Access

---

Ports, or interfaces, on a switch operate in one of three modes:

- Access mode
- Trunk mode
- Tagged-access mode

#### **Access Mode**

An interface in access mode connects a switch to a single network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. Access interfaces accept only untagged packets.

By default, when you boot a switch that runs Junos OS that does not support ELS and use the factory default configuration, or when you boot such a switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that VLAN instead of **default**.

On a switch that runs Junos OS that supports ELS, the VLAN named **default** is not supported. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist. After you assign an interface to a VLAN, the interface functions in access mode.

For switches that run either type of software, you can also configure a trunk port or interface to accept untagged packets from a user-configured VLAN. For details about this concept (native VLAN), see [“Trunk Mode and Native VLAN” on page 1530](#).

### **Trunk Mode**

Trunk mode interfaces are generally used to connect switches to one another. Traffic sent between switches can then consist of packets from multiple VLANs, with those packets multiplexed so that they can be sent over the same physical connection. Trunk interfaces usually accept only tagged packets and use the VLAN ID tag to determine both the packets' VLAN origin and VLAN destination.

On a switch that runs software that does not support ELS, an untagged packet is not recognized on a trunk port unless you configure additional settings on that port.

On a switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets for protocols such as the Link Aggregation Control Protocol (LACP) and the Link Layer Discovery Protocol (LLDP). However, the trunk port does not recognize untagged data packets unless you configure additional settings on that port.

In the rare case where you want untagged packets to be recognized by a trunk port on switches that run either type of software, you must configure the single VLAN on a trunk port as a *native VLAN*. For more information about native VLANs, see [“Trunk Mode and Native VLAN” on page 1530](#).

### **Trunk Mode and Native VLAN**

On a switch that runs Junos OS that does not support ELS, a trunk port does not recognize packets that do not include VLAN tags, which are also known as untagged packets. On a switch that runs Junos OS that supports ELS, a trunk port recognizes untagged control packets, but it does not recognize untagged data packets. With native VLAN configured, untagged packets that a trunk port normally does not recognize are sent over the trunk interface. In a situation where packets pass from a device, such as an IP phone or printer, to a switch in access mode, and you want those packets sent from the switch over a trunk port, use native VLAN mode. Create a native VLAN by configuring a VLAN ID for it, and specify that the trunk port is a member of the native VLAN.

The switch's trunk port will then treat those packets differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, packets on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag).

There is another native VLAN option for switches that do not support ELS. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

### Tagged-Access Mode

Only switches that run Junos OS that does not use the ELS configuration style support tagged-access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:

- Like access mode, tagged-access mode connects the switch to an access layer device. Unlike access mode, tagged-access mode is capable of accepting VLAN tagged packets.
- Like trunk mode, tagged-access mode accepts VLAN tagged packets from multiple VLANs. Unlike trunk port interfaces, which are connected at the core/distribution layer, tagged-access port interfaces connect devices at the access layer.

Like trunk mode, tagged-access mode also supports native VLAN.



**NOTE:** Control packets are never reflected back on the downstream port.

### Additional Advantages of Using VLANs

In addition to reducing traffic and thereby speeding up the network, VLANs have the following advantages:

- VLANs provide segmentation services traditionally provided by routers in LAN configurations, thereby reducing hardware equipment costs.
- Packets coupled to a VLAN can be reliably identified and sorted into different domains. You can contain broadcasts within parts of the network, thereby freeing up network resources. For example, when a DHCP server is plugged into a switch and starts broadcasting its presence, you can prevent some hosts from accessing it by using VLANs to split up the network.
- For security issues, VLANs provide granular control of the network because each VLAN is identified by a single IP subnetwork. All packets passing in and out of a VLAN are consistently tagged with the VLAN ID of that VLAN, thereby providing easy identification, because a VLAN ID on a packet cannot be altered. (For a switch that runs Junos OS that does not support ELS, we recommend that you avoid using 1 as a VLAN ID, because that ID is a default value.)
- VLANs react quickly to host relocation—this is also due to the persistent VLAN tag on packets.
- On an Ethernet LAN, all network nodes must be physically connected to the same network. In VLANs, the physical location of nodes is not important—you can group

network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or physical location.

### Maximum VLANs and VLAN Members Per Switch

---

The number of VLANs supported per switch varies for each switch. Use the configuration-mode command **set vlans *vlan-name* *vlan-id* ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because you have to assign a specific ID number when you create a VLAN—you could overwrite one of the numbers, but you cannot exceed the limit.

You can, however, exceed the recommended VLAN member maximum for a switch.

On a switch that runs Junos OS that does not support the ELS configuration style, the maximum number of VLAN members allowed on the switch is eight times the maximum number of VLANs that the switch supports (vmember limit =  $\text{vlan max} * 8$ ). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears when you commit the configuration. If you commit the configuration despite the warning, the commit succeeds, but there is a risk of the Ethernet switching process (eswd) failing as a result of memory allocation failure.

On a switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs that the switch supports (vmember limit =  $\text{vlan max} * 24$ ). If the configuration of the switch exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

A QFabric system supports up to 131,008 VLAN members (vmembers) on a single network node group, server node group, or redundant server node group. The number of vmembers is calculated by multiplying the maximum number of VLANs by 32.

For example, to calculate how many interfaces are required to support 4,000 VLANs, divide the maximum number of vmembers (128,000) by the number of configured VLANs (4,000). In this case, 32 interfaces are required.

On network Node groups and server Node groups, you can configure link aggregation groups (LAGs) across multiple interfaces. Each LAG and VLAN combination is considered a vmember.

### A Default VLAN Is Configured on Most Switches

---

Some switches that run Junos OS that do not support the ELS configuration style are preconfigured with a VLAN named **default** that does not tag packets and operates only with untagged packets. On these switches, each interface already belongs to the VLAN named **default** and all traffic uses this VLAN until you configure more VLANs and assign traffic to those VLANs.



**NOTE:** When a Juniper Networks QFX3500 or QFX3600 switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the slot element of the interface name. The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the ethernet-switching family. In order to include FPC 1 through FPC 9 in the default VLAN, add the ethernet-switching family to the configurations for those interfaces.

### Assigning Traffic to VLANs

You can assign traffic on any switch to a particular VLAN by referencing either the interface port of the traffic or the MAC addresses of devices sending traffic.

#### ***Assign VLAN Traffic According to the Interface Port Source***

This method is most commonly used to assign traffic to VLANs. In this case, you specify that all traffic received on a particular switch interface is assigned to a specific VLAN. You configure this VLAN assignment when you configure the switch, by using either the VLAN number (called a VLAN ID) or by using the VLAN name, which the switch then translates into a numeric VLAN ID. This method is referred to simply as creating a VLAN because it is the most commonly used method.

#### ***Assign VLAN Traffic According to the Source MAC Address***

In this case, all traffic received from a specific MAC address is forwarded to a specific egress interface (next hop) on the switch. MAC-based VLANs are either static (named MAC addresses configured one at a time) or dynamic (configured using a RADIUS server).

To configure a static MAC-based VLAN on a switch that supports ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. To configure a static MAC-based VLAN on a switch that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*.

### Forwarding VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q spanning-tree protocols.

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. The same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

To pass traffic from a single device on an access port to a switch and then pass those packets on a trunk port, use the native mode configuration previously discussed under [“Trunk Mode” on page 1530](#).

## VLANs Communicate with Integrated Routing and Bridging Interfaces or Routed VLAN Interfaces

---

Traditionally, switches sent traffic to hosts that were part of the same broadcast domain (VLAN) but routers were needed to route traffic from one broadcast domain to another. Also, only routers performed other Layer 3 functions such as traffic engineering.

Switches that run Junos OS that supports the ELS configuration style perform inter-VLAN routing functions using an integrated routing and bridging (IRB) interface named `irb`, while switches that run Junos OS that does not support ELS perform these functions using a routed VLAN interface (RVI) named `vlan`. These interfaces detect both MAC addresses and IP addresses and route data to Layer 3 interfaces, thereby frequently eliminating the need to have both a switch and a router.



### NOTE:

---

#### Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding FCoE on page 5518](#)
- [Interfaces Overview on page 2389](#)

## Understanding Ethernet Ring Protection Switching Functionality

- [Acronyms on page 1535](#)
- [Ring Nodes on page 1535](#)
- [Ring Node States on page 1535](#)
- [Failure Detection on page 1535](#)
- [Logical Ring on page 1536](#)
- [FDB Flush on page 1536](#)
- [Traffic Blocking and Forwarding on page 1536](#)
- [RAPS Message Blocking and Forwarding on page 1536](#)
- [Dedicated Signaling Control Channel on page 1537](#)
- [RAPS Message Termination on page 1538](#)
- [Multiple Rings on page 1538](#)
- [Node ID on page 1538](#)
- [Bridge Domains with the Ring Port \(MX Series Routers Only\) on page 1538](#)



## Acronyms

---

The following acronyms are used in the discussion about Ethernet ring protection switching:

- MA—Maintenance association
- MEP—Maintenance association end point
- OAM—Operations, administration, and management (Ethernet ring protection switching uses connectivity fault management daemon)
- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—Ring automatic protection switching
- WTR—Wait to restore
- RPL—Ring protection link

## Ring Nodes

---

Multiple nodes are used to form a ring. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL. This node also initiates the RAPS message.

## Ring Node States

---

There are three different states for each node of a specific ring:

- init—Not a participant of a specific ring.
- idle—No failure on the ring; the node is performing normally. For a normal node, traffic is unblocked on both ring ports. For the RPL owner, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- protection—A failure occurred on the ring. For a normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.

There can be only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

## Failure Detection

---

Ethernet ring operation depends on quick and accurate failure detection. The failure condition *signal failure (SF)* is supported. For SF detection, an Ethernet continuity check MEP must be configured for each ring link. For fast protection switching, a 10-ms transmission period for this MEP group is supported. OAM monitors the MEP group's MA and reports SF or SF clear events to the Ethernet ring control module. For this MEP group,

the action profile must be configured to update the interface device IFF\_LINKDOWN flag. OAM updates the IFF\_LINKDOWN flag to notify the Ethernet ring control module.

### Logical Ring

This feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring must operate on the physical interface, not the VLAN.

### FDB Flush

When ring protection switching occurs, normally an *FDB flush* is executed. The Ethernet ring control module uses the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.


### Traffic Blocking and Forwarding

Ethernet ring control uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical interface default STP index state to forwarding or discarding in order to control user traffic.


### RAPS Message Blocking and Forwarding

The router or switch treats the ring automatic protection switching (RAPS) message the same as it treats user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS messages between the two ring ports, as shown in [Figure 18 on page 1536](#), the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces' STP index state. The RAPS message is always sent by the router or switch through the ring ports, as shown in [Figure 19 on page 1536](#). A RAPS message received from a discarding ring port is sent to the Ethernet ring control module, but is not sent to the other ring port.

**Figure 18: Protocol Packets from the Network to the Router**



**Figure 19: Protocol Packets from the Router or Switch to the Network**



Juniper Networks switches and Juniper Networks routers use different methods to achieve these routes.

The switches use forwarding database entries to direct the RAPS messages. The forwarding database entry (keyed by the RAPS multicast address and VLAN) has a composite next hop associated with it—the composite next hop associates the two ring interfaces with the forwarding database entry and uses the split horizon feature to prevent sending the packet out on the interface that it is received on. This is an example of the forwarding database entry relating to the RAPS multicast MAC (a result of the **show ethernet-switching table detail** command):

```
VLAN: v1, Tag: 101, MAC: 01:19:a7:00:00:01, Interface: ERP
Interfaces:      ge-0/0/9.0, ge-0/0/3.0
Type: Static
Action: Mirror
Nexthop index: 1333
```

The routers use an implicit filter to achieve ERP routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:
  - term 1: if [Ethernet type is not OAM Ethernet type (0x8902)]  
          { accept packet }
  - term 2: if [source MAC address belongs to this bridge]  
          { drop packet, our packet loop through the ring and come back to home }
  - term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is DISCARDING]  
          { send to CPU }
- Control channel related terms:
  - if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01) AND[ring port STP status is FORWARDING] AND [Incoming interface IFL equal to control channel IFL]  
      { send packet to CPU and send to the other ring port }
  - default term: accept packet.

### Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain for routers (or the same VLAN for switches) in order to forward RAPS protocol data units (PDUs) between the two ring control physical interfaces. If the router control channel logical interface is not a trunk port, only control logical interfaces will be

configured in ring port configuration. If this router control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured for routers. For switches, always specify either a VLAN name or VLAN ID for all links.

### **RAPS Message Termination**

---

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

### **Multiple Rings**

---

The Ethernet ring control module supports multiple rings in each node (two logical interfaces are part of each ring). However, interconnection of multiple rings is not supported in this release. The interconnection of two rings means that two rings may share the same link or share the same node.

### **Node ID**

---

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address.

For routers only, you can configure this node ID when configuring the ring on the node or automatically select an ID such as STP. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID. The node ID on switches is selected automatically and is not configurable.

### **Bridge Domains with the Ring Port (MX Series Routers Only)**

---

On the routers, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain, you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB Layer 3 interface.

#### **Related Documentation**

- [Ethernet Ring Protection Switching Overview on page 1525](#)
- [Configuring Ethernet Ring Protection Switching](#)
- [Example: Ethernet Ring Protection Switching Configuration on MX Routers](#)
- [Ethernet Interfaces Feature Guide for Routing Devices](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 1670](#)

## Understanding Integrated Routing and Bridging

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). VLANs limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN. For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs, you normally you need a router that connects the VLANs. However, you can accomplish this forwarding on a switch without using a router by configuring an integrated routing and bridging (IRB) interface. (These interfaces are also called routed VLAN interfaces, or RVIs). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

An IRB is a special type of Layer 3 virtual interface named **vlan**. Like normal Layer 3 interfaces, the **vlan** interface needs a logical unit number with an IP address. In fact, to be useful an IRB needs at least two logical units and two IP addresses—you must create units with addresses in each of the subnets associated with the VLANs between which you want traffic to be routed. That is, if you have two VLANs (for example, VLAN **red** and VLAN **blue**) with corresponding subnets, your IRB must have a logical unit with an address in the subnet for **red** and a logical unit with an address in the subnet for **blue**. The switch automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs.



**NOTE:** If you are using a version of Junos OS that supports Enhanced Layer 2 Software (ELS), you can also create a Layer 3 virtual interface named **irb** instead of **vlan**—that is, both statements are supported by ELS

Table 94 on page 1539 shows values you might use when configuring an IRB:

**Table 94: Sample IRB Values**

Property	Settings
VLAN names and tags (IDs)	<b>blue</b> , ID 100 <b>red</b> , ID 200
Subnets associated with VLANs	<b>blue</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>red</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
IRB name	interface <b>irb</b>
IRB units and addresses	logical unit 100: 192.0.2.1/25 logical unit 200: 192.0.2.129/25

For the sake of consistency and to avoid confusion, Table 94 on page 1539 shows IRB logical unit numbers that match the IDs of the corresponding VLANs. However, you do not have

to assign logical unit numbers that match the VLAN IDs—you can use any values for the units. To bind the logical units of the IRB to the appropriate VLANs, you use the [l3-interface](#) statement.

Because IRBs operate at Layer 3, you can use Layer 3 services such as firewall filters or CoS rewriting with them.

[Table 95 on page 1540](#) shows the number of IRBs/RVIs that each QFX platform supports.

**Table 95: Number of Supported IRBs/RVIs by Platform**

Platform	Number of Supported IRBs/RVIs
QFX3500	1200
QFX3000-G	1024
QFX3000-M	1024

**Related Documentation**

- [Example: Configuring Routing Between VLANs on One Switch on page 1576](#)

## Understanding MAC Learning

*MAC learning* is the process of obtaining the MAC addresses of all the nodes on a network.

When a node is first connected to an Ethernet LAN or VLAN, it has no information about the other nodes on the network. As data is sent through the network, data packets include a data frame listing their source and destination MAC addresses. The data frame is forwarded to a target port, which is connected to the second device. The MAC address is learned locally at the target port, which facilitates communications for frames that later enter the target port and contain addresses previously learned from a received frame.

MAC learning can also be enabled on a per-VLAN basis. See [Example: Disabling MAC Learning in a VLAN](#) for further information.

By default, MAC learning is enabled on the QFX Series.

**Related Documentation**

- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 1541](#)
- [Overview of Layer 2 Networking on page 1542](#)

## Layer 2 Networking

- [Introduction to the Media Access Control \(MAC\) Layer 2 Sublayer on page 1541](#)
- [Overview of Layer 2 Networking on page 1542](#)
- [Understanding Layer 2 Broadcasting on page 1544](#)
- [Understanding Unicast on page 1545](#)
- [Understanding the Unified Forwarding Table on page 1545](#)

## Introduction to the Media Access Control (MAC) Layer 2 Sublayer

This topic provides an introduction to the MAC sublayer of the data link layer (Layer 2).

In Layer 2 of a network, the Media Access Control (MAC) sublayer provides addressing and channel access control mechanisms that enable several terminals or network nodes to communicate in a network.

The MAC sublayer acts as an interface between the logical link control (LLC) Ethernet sublayer and Layer 1 (the physical layer). The MAC sublayer emulates a full-duplex logical communication channel in a multipoint network. This channel may provide unicast, multicast, or broadcast communication service. The MAC sublayer uses MAC protocols to prevent collisions.

In Layer 2, multiple devices on the same physical link can uniquely identify one another at the data link layer, by using the MAC addresses that are assigned to all ports on a switch. A MAC algorithm accepts as input a secret key and an arbitrary-length message to be authenticated, and outputs a MAC address.

A MAC address is a 12-digit hexadecimal number (48 bits in long). MAC addresses are usually written in one of these formats:

- MM:MM:MM:SS:SS:SS
- MM-MM-MM-SS-SS-SS

The first half of a MAC address contains the ID number of the adapter manufacturer. These IDs are regulated by an Internet standards body. The second half of a MAC address represents the serial number assigned to the adapter by the manufacturer.

Contrast MAC addressing, which works at Layer 2, with IP addressing, which runs at Layer 3 (networking and routing). One way to remember the difference is that the MAC addresses apply to a physical or virtual node, whereas IP addresses apply to the software implementation of that node. MAC addresses are typically fixed on a per-node basis, whereas IP addresses change when the node moves from one part of the network to another.

IP networks maintain a mapping between the IP and MAC addresses of a node using the Address Resolution Protocol (ARP) table. DHCP also typically uses MAC addresses when assigning IP addresses to nodes.

### Related Documentation

- [Overview of Layer 2 Networking on page 1542](#)
- [Understanding MAC Learning on page 1540](#)

## Overview of Layer 2 Networking

Layer 2, also known as the Data Link Layer, is the second level in the seven-layer OSI reference model for network protocol design. Layer 2 is equivalent to the link layer (the lowest layer) in the TCP/IP network model. Layer 2 is the network layer used to transfer data between adjacent network nodes in a wide area network or between nodes on the same local area network.

A *frame* is a protocol data unit, the smallest unit of bits on a Layer 2 network. Frames are transmitted to and received from devices on the same local area network (LAN). Unlike bits, frames have a defined structure and can be used for error detection, control plane activities and so forth. Not all frames carry user data. The network uses some frames to control the data link itself..

At Layer 2, *unicast* refers to sending frames from one node to a single other node, whereas *multicast* denotes sending traffic from one node to multiple nodes, and *broadcasting* refers to the transmission of frames to all nodes in a network. A *broadcast domain* is a logical division of a network in which all nodes of that network can be reached at Layer 2 by a broadcast.

Segments of a LAN can be linked at the frame level using *bridges*. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN.

*Forwarding* is the relaying of packets from one network segment to another by nodes in the network. On a VLAN, a frame whose origin and destination are in the same VLAN are forwarded only within the local VLAN. A network segment is a portion of a computer network wherein every device communicates using the same physical layer.

Layer 2 contains two sublayers:

- Logical link control (LLC) sublayer, which is responsible for managing communications links and handling frame traffic.
- Media access control (MAC) sublayer, which governs protocol access to the physical network medium. By using the MAC addresses that are assigned to all ports on a switch, multiple devices on the same physical link can uniquely identify one another.

The ports, or interfaces, on a switch operate in either access mode, tagged-access, or trunk mode:

- *Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, all ports on a switch are in access mode.
- *Tagged-Access mode* ports connect to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The port itself belongs to a single VLAN. The frames transmitted over an access interface are normal



Ethernet frames. By default, all ports on a switch are in access mode. Tagged-access mode accommodates cloud computing, specifically scenarios including virtual machines or virtual computers. Because several virtual computers can be included on one physical server, the packets generated by one server can contain an aggregation of VLAN packets from different virtual machines on that server. To accommodate this situation, tagged-access mode reflects packets back to the physical server on the same downstream port when the destination address of the packet was learned on that downstream port. Packets are also reflected back to the physical server on the downstream port when the destination has not yet been learned. Therefore, the third interface mode, tagged access, has some characteristics of access mode and some characteristics of trunk mode:

- *Trunk mode* ports handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to other devices or switches.

With native VLAN configured, frames that do not carry VLAN tags are sent over the trunk interface. If you have a situation where packets pass from a device to a switch in access mode, and you want to then send those packets from the switch over a trunk port, use native VLAN mode. Configure the single VLAN on the switch's port (which is in access mode) as a native VLAN. The switch's trunk port will then treat those frames differently than the other tagged packets. For example, if a trunk port has three VLANs, 10, 20, and 30, assigned to it with VLAN 10 being the native VLAN, frames on VLAN 10 that leave the trunk port on the other end have no 802.1Q header (tag). There is another native VLAN option. You can have the switch add and remove tags for untagged packets. To do this, you first configure the single VLAN as a native VLAN on a port attached to a device on the edge. Then, assign a VLAN ID tag to the single native VLAN on the port connected to a device. Last, add the VLAN ID to the trunk port. Now, when the switch receives the untagged packet, it adds the ID you specified and sends and receives the tagged packets on the trunk port configured to accept that VLAN.

Including the sublayers, Layer 2 on the QFX Series supports the following functionality:

- Unicast, multicast, and broadcast traffic.
- Bridging.
- VLAN 802.1Q—Also known as *VLAN tagging*, this protocol allows multiple bridged networks to transparently share the same physical network link by adding VLAN tags to an Ethernet frame.
- Extension of Layer 2 VLANs across multiple switches using Spanning Tree Protocol (STP) prevents looping across the network.
- *MAC learning*, including per-VLAN MAC learning and Layer 2 learning suppression—This process obtains the MAC addresses of all the nodes on a network
- Link aggregation—This process groups of Ethernet interfaces at the physical layer to form a single link layer interface, also known as a *link aggregation group (LAG)* or LAG bundle

- Storm control on the physical port for unicast, multicast, and broadcast
- STP support, including 802.1d, RSTP, MSTP, and Root Guard

**Related  
Documentation**

- [Understanding Bridging and VLANs on page 1527](#)
- *Understanding Bridging and VLANs*

## Understanding Layer 2 Broadcasting

In a Layer 2 network, *broadcasting* refers to sending traffic to all nodes on a network.

Layer 2 broadcast traffic stays within a local area network (LAN) boundary; known as the *broadcast domain*. Layer 2 broadcast traffic is sent to the broadcast domain using a MAC address of FF:FF:FF:FF:FF:FF. Every device in the broadcast domain recognizes this MAC address and passes the broadcast traffic on to other devices in the broadcast domain, if applicable. Broadcasting can be compared to unicasting (sending traffic to a single node) or multicasting (delivering traffic to a group of nodes simultaneously).

Layer 3 broadcast traffic, however, is sent to all devices in a network using a broadcast network address. For example, if your network address is 192.0.0.0, the broadcast network address is 192.255.255.255. In this case, only devices that belong to the 192.0.0.0 network receive the Layer 3 broadcast traffic. Devices that do not belong to this network drop the traffic.

Broadcasting is used in the following situations:

- Address Resolution Protocol (ARP) uses broadcasting to map MAC addresses to IP addresses. ARP dynamically binds the IP address (the logical address) to the correct MAC address. Before IP unicast packets can be sent, ARP discovers the MAC address used by the Ethernet interface where the IP address is configured.
- Dynamic Host Configuration Protocol (DHCP) uses broadcasting to dynamically assign IP addresses to hosts on a network segment or subnet.
- Routing protocols use broadcasting to advertise routes.

Excessive broadcast traffic can sometimes create a broadcast storm. A broadcast storm occurs when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses that create a snowball effect. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service.

**Related  
Documentation**

- [Overview of Layer 2 Networking on page 1542](#)
- [Understanding Storm Control on page 5272](#)
- *Understanding Bridging and VLANs*
- [Understanding Bridging and VLANs on page 1527](#)

## Understanding Unicast

*Unicasting* is the act of sending data from one node of the network to another. In contrast, multicast transmissions send traffic from one data node to multiple other data nodes.

*Unknown unicast* traffic consists of unicast frames with unknown destination MAC addresses. By default, the switch floods these unicast frames that are traveling in a VLAN to all interfaces that are members of the VLAN. Forwarding this type of traffic to interfaces on the switch can trigger a security issue. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service. This is known as a traffic storm.

To prevent a storm, you can disable the flooding of unknown unicast packets to all interfaces by configuring one VLAN or all VLANs to forward any unknown unicast traffic to a specific trunk interface. (This channels the unknown unicast traffic to a single interface.)

### Related Documentation

- [Overview of Layer 2 Networking on page 1542](#)
- [Understanding Bridging and VLANs on page 1527](#)

## Understanding the Unified Forwarding Table

- [Using the Unified Forwarding Table to Optimize Address Storage on page 1545](#)
- [MAC Address and Host Address Memory Allocation on page 1545](#)
- [LPM Table Memory Allocation on page 1546](#)

### Using the Unified Forwarding Table to Optimize Address Storage

On QFX5100 and EX4600 switches, you can control the allocation of forwarding table memory available to store the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match (LPM) table entries.



**NOTE:** Starting with Junos OS 13.2X51-D15, you can allocate more memory to store prefixes in the range /65 to /127 range.

This feature gives you the flexibility to configure your switch to match the needs of your particular network environment.

### MAC Address and Host Address Memory Allocation

There are several profiles that allocate memory differently for MAC addresses and host addresses. You configure the mix that best meets your needs by choosing the appropriate profile. [Table 96 on page 1546](#) lists the profiles you can choose and the associated maximum values for the MAC address and host table entries.

Table 96: Unified Forwarding Table Profiles

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
<b>l2-profile-one</b>	288K	16K	8K	8K	8K	4K	4K
<b>l2-profile-two</b>	224K	80K	40K	40K	40K	20K	20K
<b>l2-profile-three (default)</b>	160K	144K	72K	72K	72K	36K	36K
<b>l3-profile</b>	96K	208K	104K	104K	104K	52K	52K
<b>lpm-profile</b>	32K	16K	8K	8K	8K	4K	4K

Note that all entries in the host table share the same memory space. If the host table stores the maximum number of entries for any given type, the entire shared table is full and is unable to accommodate *any* entries of any other type. As you can see, different entry types occupy different amounts of memory. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address.

[Table 97 on page 1546](#) lists various valid combinations that the host table can store if you use the **l2-profile-one** profile. Each row in the table represents a case in which the host table is full and cannot accommodate any more entries. .

Table 97: Example Host Table Combinations Using l2-profile-one

IPv4 unicast	IPv6 unicast	IPv4 multicast (*, G)	IPv4 multicast (S, G)	IPv6 multicast (*, G)	IPv6 multicast (S, G)
16K	0	0	0	0	0
12K	2K	0	0	0	0
12K	0	2K	2K	0	0
8K	4K	0	0	0	0
4K	2K	2K	2K	0	0
0	4K	0	0	1K	1K

#### LPM Table Memory Allocation

You configure the memory allocation for LPM table entries differently depending on which version of Junos OS you use. To learn how to configure memory allocation for LPM table entries see [“Configuring the Unified Forwarding Table” on page 1697](#).

- Related Documentation**
- [Configuring the Unified Forwarding Table on page 1697](#)

## Understanding Q-in-Q Tunneling

---

- [Understanding Q-in-Q Tunneling on page 1547](#)

### Understanding Q-in-Q Tunneling



**NOTE:** This topic applies to Junos OS switches with support for the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

Q-in-Q tunneling enables service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs because customers' VLAN (C-VLAN) tags are prepended by the service-provider VLAN (S-VLAN) tag, which allows you to preserve each customers' VLAN IDs without conflict. The Juniper Networks Junos operating system (Junos OS) implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

- [How Q-in-Q Tunneling Works on page 1547](#)
- [How VLAN Translation Works on page 1548](#)
- [Sending and Receiving Untagged Packets on page 1548](#)
- [Disabling MAC Address Learning on page 1549](#)
- [Mapping C-VLANs to S-VLANs on page 1549](#)
- [Constraints for Q-in-Q Tunneling and VLAN Translation on page 1550](#)

#### How Q-in-Q Tunneling Works

---

In Q-in-Q tunneling, as a packet travels from a C-VLAN to an S-VLAN, a service-provider-specific 802.1Q tag is added to the packet. This additional tag is used to segregate traffic into S-VLANs. The original customer 802.1Q tag of the packet is retained and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the additional 802.1Q tag is removed.

When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network. Access interfaces are assumed to be customer-facing and accept both tagged and untagged frames. This topic refers to trunk interfaces as S-VLAN interfaces. This type of interface is also sometimes known as a network-to-network interface (NNI). The topic refers to access interfaces as C-VLAN interfaces. This type of interface is also sometimes known as a user-network interface (UNI).

An interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN (1:1) or many C-VLANs to many S-VLANs (N:N). C-VLAN and S-VLAN tags are unique—for instance, you can have both a C-VLAN tag of 101 and an S-VLAN tag of 101. You can limit the set of accepted customer tags to a range of tags or to discrete values. Class-of-service (CoS) values of C-VLANs are unchanged in the downstream direction. You may copy ingress priority and CoS settings to the S-VLAN.

C-VLAN and S-VLAN interfaces accept priority-tagged packets without any configuration.

### How VLAN Translation Works

---

VLAN translation replaces an incoming C-VLAN tag with an S-VLAN tag instead of adding an additional tag. The C-VLAN tag is therefore lost, so a single-tagged packet is normally untagged when it leaves the S-VLAN (at the other end of the link). If an incoming packet has had Q-in-Q tunneling applied in advance, VLAN translation replaces the outer tag and the inner tag is retained when the packet leaves the S-VLAN at the other end of the link.

To configure VLAN translation, use the *mapping swap* statement at the **[edit vlans interface]** hierarchy level.



**NOTE:** You can configure VLAN translation on access ports only. You cannot configure it on trunk ports, and you cannot configure Q-in-Q tunneling on the same access port.

---

### Sending and Receiving Untagged Packets

---

To enable an interface to send and receive untagged packets, you must specify a native VLAN for a physical interface. When the interface receives an untagged packet, it adds the VLAN ID of the native VLAN to the packet and sends the newly tagged packet to the mapped interface.

To specify a native VLAN, use the **native-vlan-id** statement at the **[edit interfaces interface-name]** hierarchy level. The native VLAN ID must match the C-VLAN or S-VLAN ID or be included in the VLAN ID list specified on the logical interface.

For example, on a logical interface for a C-VLAN interface, you might specify a C-VLAN ID list of 100-200. Then, on the C-VLAN physical interface, you could specify a native VLAN ID of 150. This configuration would work because the native VLAN of 150 is included in the C-VLAN ID list of 100-200.

We recommend configuring a native VLAN when using any of the approaches to map C-VLANs to S-VLANs. If you do not configure a native VLAN on an interface, untagged packets received by the interface are discarded. See the Mapping C-VLANs to S-VLANs section in this topic for information about the methods of mapping C-VLANs to S-VLANs.

### Disabling MAC Address Learning

In a Q-in-Q deployment, customer packets interfaces are transported without any changes to source and destination MAC addresses. You can disable MAC address learning at the global, interface, and VLAN levels:

- To disable learning globally, disable MAC address learning for the switch.
- To disable learning for an interface, disable MAC address learning for all VLANs of which the specified interface is a member.
- To disable learning for a VLAN, disable MAC address learning for a specified VLAN.

### Mapping C-VLANs to S-VLANs

There are three ways to map C-VLANs to S-VLANs:

- [All-in-One Bundling on page 1549](#)
- [Many-to-Many Bundling on page 1549](#)
- [Mapping a Specific Interface on page 1550](#)

If you configure multiple mapping methods, the switch gives priority to mapping a specific interface, then to many-to-many bundling, and last to all-in-one bundling. However, for a particular mapping method, setting up overlapping rules for the same C-VLAN is not supported.

#### *All-in-One Bundling*

All-in-one bundling maps all packets from all C-VLAN interfaces to an S-VLAN.

The C-VLAN interface accepts untagged and single-tagged packets. An S-VLAN 802.1Q tag is then added to these packets, and the packets are sent to the S-VLAN interface, which accepts untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

#### *Many-to-Many Bundling*

Many-to-many bundling is used to specify which C-VLANs are mapped to which S-VLANs.

Use many-to-many bundling when you want a subset of the C-VLANs on the access switch to be part of multiple S-VLANs. With many-to-many bundling, the C-VLAN interfaces accept untagged and single-tagged packets. An S-VLAN 802.1Q tag is then added to these packets, and the packets are sent to the S-VLAN interfaces, which accept untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

### **Mapping a Specific Interface**

Use specific interface mapping when you want to assign an S-VLAN to a specific C-VLAN on an interface. The configuration applies only to the specific interface, not to all access interfaces.

Specific interface mapping has two suboptions: **push** and **swap**. When traffic that is mapped to a specific interface is pushed, the packet retains its original tag as it moves from the C-VLAN to the S-VLAN and an additional S-VLAN tag is added to the packet. When traffic that is mapped to a specific interface is swapped, the incoming tag is replaced with a new VLAN tag. This is sometimes known as VLAN rewriting or VLAN translation.

Typically, this method is used to keep data from different customers separate or to provide individualized treatment of the packets on a certain interface. You might also use this method to map VLAN traffic from different customers to a single S-VLAN.

When using specific interface mapping, the C-VLAN interfaces accept untagged and single-tagged packets, while the S-VLAN interfaces accept untagged, single-tagged, and double-tagged packets.



**NOTE:** The C-VLAN and S-VLAN interfaces accept untagged packets provided that the `native-vlan-id` statement is configured on these interfaces.

---

### **Constraints for Q-in-Q Tunneling and VLAN Translation**

---

Be aware of the following constraints when configuring Q-in-Q tunneling and VLAN translation:

- With releases of Junos OS 13.2X51 previous to 13.2X51-D20, you cannot create a regular VLAN on an interface if you have created an S-VLAN or C-VLAN on that interface for Q-in-Q tunneling. This means that you cannot create an integrated routing and bridging (IRB) interface on that interface because regular VLANs are a required part of IRB configuration. With Junos OS 13.2X51-D25, you can create a regular VLAN on a trunk interface that has an S-VLAN, which means that you can also create an IRB interface on the trunk. In this case, the regular VLAN and S-VLAN on the same trunk interface cannot share the same VLAN ID. Junos OS 13.2X51-D25 does not allow you to create a regular VLAN on an access interface that has a C-VLAN.
- Most access port security features are not supported with Q-in-Q tunneling and VLAN translation.
- Configuring Q-in-Q tunneling and VLAN rewriting/VLAN translation on the same port is not supported.
- You can configure at most one VLAN rewrite/VLAN translation for a given VLAN and interface. For example, you can create no more than one translation for VLAN 100 on interface xe-0/0/0.
- The combined total of VLANs and rules for Q-in-Q tunneling and VLAN translation cannot exceed 6000. For example, you can configure and commit 4000 VLANs and



2000 rules for Q-in-Q tunneling and VLAN translation. However, you cannot configure 4000 VLANs and 2500 rules for Q-in-Q tunneling and VLAN translation. If you try to commit a configuration that exceeds the limit, you see CLI and syslog errors that inform you about the problem.

- MAC addresses are learned from S-VLANs, not C-VLANs.
- Broadcast, unknown unicast, and multicast traffic is forwarded to all members in the S-VLAN.
- The following features are not supported with Q-in-Q tunneling:
  - DHCP relay
  - Fibre Channel over Ethernet
  - IP Source Guard
- The following features are not supported with VLAN rewriting/VLAN translation:
  - Fibre Channel over Ethernet
  - Firewall filter applied to a port or VLAN in the output direction
  - Private VLANs
  - VLAN Spanning Tree Protocol
  - Reflective relay

**Related  
Documentation**

- [Configuring Q-in-Q Tunneling on page 1683](#)

---

## Proxy ARP

- [Understanding Proxy ARP on page 1551](#)

### Understanding Proxy ARP

You can configure proxy Address Resolution Protocol (ARP) to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

- [What Is ARP? on page 1552](#)
- [Proxy ARP Overview on page 1552](#)
- [Best Practices for Proxy ARP on page 1552](#)

## What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

## Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.



**NOTE:** For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

You can configure proxy ARP for each interface. You can also configure proxy ARP for a VLAN by using a routed VLAN interface (RVI).

Two modes of proxy ARP are supported: restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- **Restricted**—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.
- **Unrestricted**—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in Juniper Networks Junos operating system (Junos OS) configurations other than those on the switch). We recommend using restricted mode on the switch.

## Best Practices for Proxy ARP

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP to restricted mode.
- Use restricted mode when configuring proxy ARP on RVIs.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

### Related Documentation

- *Configuring Proxy ARP*
- *proxy-arp*

---

## Reflective Relay

---

- [Understanding Reflective Relay for Use with VEPA Technology on page 1553](#)

### Understanding Reflective Relay for Use with VEPA Technology

Virtual Ethernet Port Aggregator (VEPA) technology aggregates packets generated by virtual machines located on the same server and relays them to a physical switch. The physical switch then provides connectivity between the virtual machines located on the server, so the virtual machines do not communicate with one another. Offloading switching activities from a virtual switch to a physical switch reduces the computing overhead on the virtual servers and takes advantage of the security, filtering, and management features of the physical switch. Reflective relay, also known as “hairpin turn,” enables the physical switch to receive aggregated packets from the virtual machines hosted on the server through the VEPA on the downstream port and send those packets out the same downstream port from which the physical switch received them.

- [VEPA on page 1553](#)
- [Reflective Relay on page 1553](#)

#### VEPA

---

Even though virtual machines are capable of sending packets directly to one another, it is more efficient to pass these aggregated packets from the VEPA to a physical switch. The switch can then send any packets destined for a virtual machine located on the same server to the VEPA.

#### Reflective Relay

---

Reflective relay, also known as a “hairpin turn” or “hairpin mode,” returns aggregated packets to the VEPA by using the same downstream port that initially delivered the aggregated packets from the VEPA to the switch. Reflective relay must be configured on the interface located on the physical switch that receives aggregated packets, such as VEPA packets, because some of these packets might need to be sent back to the server if they are destined for another virtual machine on the same server.

Reflective relay only occurs in two situations:

- When the destination address of the packet was learned on that downstream port
- When the destination has not yet been learned

Reflective relay does not otherwise change the operation of the switch. If the interface to which the virtual machine is connected and the MAC address of the virtual machine packet are not yet included in the Ethernet switching table for the virtual machine's associated VLAN, an entry is added. If the source MAC address of an incoming packet under the respective VLAN is not yet present in the Ethernet switching table, the switch floods the packet on all the other ports that are members of the same VLAN, including the port on which the packet arrived.

- Related Documentation**
- [Understanding Bridging and VLANs](#)
  - [Understanding Bridging and VLANs on page 1527](#)
  - [Example: Configuring Reflective Relay for Use with VEPA Technology](#)

## Spanning Trees

---

- [Overview of Spanning-Tree Protocols on page 1554](#)
- [Understanding MSTP on page 1555](#)
- [Understanding RSTP on page 1556](#)
- [Understanding VSTP on page 1557](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558](#)
- [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1559](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1560](#)

### Overview of Spanning-Tree Protocols

QFX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). The default spanning-tree protocol on the QFX Series is RSTP. RSTP provides faster convergence times than STP. However, some legacy networks require the slower convergence times of basic STP.

The STP support provided for the QFX Series includes:

- IEEE 802.1d
- 802.1w RSTP
- 802.1s MSTP

If your network includes IEEE 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. See [“Configuring STP” on page 1704](#). When you explicitly configure STP, the QFX Series products use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP. If you use virtual LANs (VLANs), you should enable VSTP and use it on your network. See [“Understanding VSTP” on page 1557](#).

You can use the same operational commands (**show spanning-tree bridge** and **show spanning-tree interface**) to check the status of your spanning-tree configuration, regardless of which spanning-tree protocol has been configured.

STP uses bridge protocol data unit (BPDU) packets to exchange information with other switches. BPDUs send hello packets out at regular intervals to exchange information across bridges and detect loops in a network topology. There are two types of BPDUs:

- Configuration BPDUs—These BPDUs contain configuration information about the transmitting switch and its ports, including switch and port MAC addresses, switch priority, port priority, and port cost.

- Topology change notification (TCN) BPDUs—When a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN reaches the root bridge.

STP uses the information provided by the BPDUs to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

### Understanding Spanning Tree Protocols on a QFabric System

Although there is no need to run STP in a QFabric system, you can connect a QFabric system to another Layer 2 device and use STP. STP traffic can only be processed on network Node groups. Other Node groups, such as redundant server Node groups and server Node groups, discard the STP bridge protocol data units (BPDUs) traffic and disable the interface automatically. Server Node groups only process host-facing protocols, whereas Network Node groups process all supported protocols.

#### Related Documentation

- [Understanding BPDUs for STP, RSTP, and MSTP on page 1558](#)
- [Understanding MSTP on page 1555](#)
- [Understanding RSTP on page 1556](#)
- [Understanding VSTP on page 1557](#)

## Understanding MSTP

Although RSTP provides faster convergence time than STP does, it still does not solve a problem inherent in STP: all VLANs within a LAN must share the same spanning tree. To solve this problem, the QFX Series products use Multiple Spanning Tree Protocol (MSTP) to create a loop-free topology in networks with multiple spanning-tree regions.

An MSTP region allows a group of bridges to be modeled as a single bridge. An MSTP region contains multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates more efficient load sharing across redundant links.

An MSTP region can support up to 64 MSTIs, and each instance can support from 1 through 4094 VLANs.

#### Related Documentation

- [Overview of Spanning-Tree Protocols on page 1554](#)
- [Understanding RSTP on page 1556](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 1624](#)

## Understanding RSTP

Juniper Networks QFX Series products use Rapid Spanning Tree Protocol (RSTP) on the network side of the QFX Series to provide quicker convergence time than the base Spanning Tree Protocol (STP) does. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state, which speeds up convergence.

Although STP provides basic loop prevention functionality, it does not provide fast network convergence when there are topology changes. The STP process to determine network state transitions is slower than the RSTP process because it is timer-based. A device must reinitialize every time a topology change occurs. The device must start in the listening state and transition to the learning state and eventually to a forwarding or blocking state. When default values are used for the maximum age (20 seconds) and forward delay (15 seconds), it takes 50 seconds for the device to converge. RSTP converges faster because it uses a handshake mechanism based on point-to-point links instead of the timer-based process used by STP.

For networks with virtual LANs (VLANs), you can use VLAN Spanning Tree Protocol (VSTP), which takes the paths of each VLAN into account when calculating routes. VSTP uses RSTP by default.

An RSTP domain running from the edge outward on a QFX Series product has the following components:

- A *root port*, which is the “best path” to the root device.
- A *designated port*, which indicates that the switch is the designated bridge for the other switch connecting to this port.
- An *alternate port*, which provides an alternate root port.
- A *backup port*, which provides an alternate designated port.

Port assignments change through messages exchanged throughout the domain. An RSTP device generates configuration messages once per hello time interval. If an RSTP device does not receive a configuration message from its neighbor after an interval of three hello times, it determines that the connection with the neighbor is lost. When a *root port* or a *designated port* fails on a device, the device generates a configuration message with the proposal bit set. Once its neighbor device receives this message, it verifies that this configuration message is valid for that port and starts a *synchronizing* operation to ensure that all of its ports are in sync with the new information.

Similar sets of messages propagate through the network, restoring the connectivity very quickly after a topology change (in a well-designed network that uses RSTP, network convergence can take as little as 0.5 seconds). If a device does not receive an agreement to a proposal message it has sent, it returns to the original IEEE 802.D convention.

RSTP was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on the QFX Series.

**Related  
Documentation**

- [Overview of Spanning-Tree Protocols on page 1554](#)
- [Understanding MSTP on page 1555](#)
- [Understanding VSTP on page 1557](#)
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)

## Understanding VSTP

VLAN Spanning Tree Protocol (VSTP) enables Juniper Networks switches to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

You can configure VSTP for a maximum of 509 VLANs.

VSTP and RSTP can be configured at the same time. If you configure VSTP and RSTP at the same time and the switch has more than 253 VLANs, VSTP is configured only for the first 253 VLANs. For the remaining VLANs, only RSTP is configured. RSTP and VSTP are the only spanning-tree protocols that can be configured at the same time on a switch.



**NOTE:** We recommend that you enable VSTP on all VLANs that could receive VSTP bridge protocol data units (BPDUs).

**Related  
Documentation**

- [Overview of Spanning-Tree Protocols on page 1554](#)
- [Understanding RSTP on page 1556](#)
- [Configuring VLAN Spanning Tree Protocol](#)
- [Configuring VLAN Spanning-Tree Protocol on page 1705](#)
- [vstp on page 1759](#)

## Understanding BPDU Protection for STP, RSTP, and MSTP



**NOTE:** Using the original CLI, you can disable BPDU protection on interfaces by issuing the **set ethernet-switching-options bpdu-block *interface-name* disable** command.

A Juniper Networks device Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Bridge protocol data unit (BPDU) protection can help prevent STP misconfigurations that can lead to network outages.

A loop-free network is supported through the exchange of a special type of frame called a BPDU. Receipt of BPDUs on certain interfaces in an STP, RSTP, VSTP, or MSTP topology, however, can lead to network outages. Enable BPDU protection on those interfaces to prevent these outages.

Peer STP applications running on the device interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

However, a user bridge application running on a device connected to the device can also generate BPDUs. If these BPDUs are picked up by STP applications running on the device, they can trigger STP miscalculations, and those miscalculations can lead to network outages.

Enable BPDU protection on device interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If BPDUs are received on a protected interface, the interface is disabled and stops forwarding frames.

Not only can you configure BPDU protection on a device with a spanning tree, but also on a device without a spanning tree. This type of topology typically consists of a non-STP device connected to an STP device through a trunk interface.

To configure BPDU protection on a device with a spanning tree, include the **bpdu-block-on-edge** statement at the **[edit protocols (stp | mstp | rstp)]** hierarchy level. To configure BPDU protection on a device without a spanning tree, include the **bpdu-block** statement at the **[edit ethernet-switching-options interface *interface-name*]** hierarchy level.

If BPDUs are sent to an interface (indicating that the misconfiguration has been corrected), the interface can be unblocked in one of two ways:

- If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires.
- Use the operational mode command **clear ethernet-switching bpdu-error**.

Disabling the BPDU protection configuration does not unblock the interface.



- Related Documentation**
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656](#)
  - [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1559](#)
  - [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1560](#)
  - [Understanding MSTP on page 1555](#)
  - [Understanding RSTP on page 1556](#)
  - [Understanding VSTP on page 1557](#)

## Understanding Loop Protection for STP, RSTP, VSTP, and MSTP

A Juniper Networks device provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing ports from entering a forwarding state that would cause a loop to open in the network.

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the device interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

However, a blocking interface can mistakenly transition to the forwarding state if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the device or software configuration error between the device and its neighbor.

When loop protection is enabled, the spanning-tree topology detects root ports and blocked ports and ensures that both keep receiving BPDUs. If a loop-protection-enabled interface stops receiving BPDUs from its designated port, it reacts as it would react to a problem with the physical connection on this interface. It does not transition the interface to a forwarding state, but instead transitions it to a loop-inconsistent state. The interface recovers and it transitions back to the spanning-tree blocking state as soon as it receives a BPDU.

We recommend that you enable loop protection on all device interfaces that have a chance of becoming root or designated ports. Loop protection is most effective when enabled in the entire switched network. When you enable loop protection, you must configure at least one action (**alarm**, **block**, or both).

An interface can be configured for either loop protection or root protection, but not for both.

- Related Documentation**
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660](#)
  - [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1560](#)
  - [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558](#)

- [Understanding MSTP on page 1555](#)
- [Understanding RSTP on page 1556](#)
- [Overview of Spanning-Tree Protocols on page 1554](#)
- [Understanding VSTP on page 1557](#)

## Understanding Root Protection for STP, RSTP, VSTP, and MSTP

A Juniper Networks device provides Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). A loop-free network is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the device interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

You can also see BPDUs generated when you run a bridge application on a device attached to the device. This can interfere with root port election, which may sometimes lead to the wrong root port being elected through the above process. Root protection allows you to manually enforce the root bridge placement in the network.

Enable root protection on interfaces that should not receive higher-priority BPDUs from the root bridge and should not be elected as the root port. These interfaces become designated ports and are typically located on an administrative boundary. If the bridge receives more STP BPDUs on a port that has root protection enabled, that port transitions to a root-prevented STP state (inconsistency state), and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. After the bridge stops receiving more STP BPDUs on the interface with root protection, the interface returns to a listening state, followed by a learning state, and ultimately back to a forwarding state. Recovery back to the forwarding state is automatic.

When root protection is enabled on an interface, it is enabled for all the STP instances on that interface. The interface is blocked only for instances for which it receives more BPDUs. Otherwise, it participates in the spanning-tree topology.

An interface can be configured for either root protection or loop protection, but not for both.

### Related Documentation

- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1664](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656](#)
- [Understanding MSTP on page 1555](#)
- [Understanding RSTP on page 1556](#)
- [Overview of Spanning-Tree Protocols on page 1554](#)

- [Understanding VSTP on page 1557](#)



## CHAPTER 15

# Configuration

- [Bridging and VLAN Configuration Examples on page 1563](#)
- [Reflective Relay Configuration Example on page 1605](#)
- [STP Configuration Examples on page 1609](#)
- [Bridging and VLAN Configuration Tasks on page 1669](#)
- [Q-in-Q Tunneling Configuration Tasks on page 1683](#)
- [Unified Forwarding Table Configuration Task on page 1697](#)
- [Forwarding Mode Configuration Task on page 1702](#)
- [Proxy ARP Configuration Task on page 1702](#)
- [Reflective Relay Configuration Tasks on page 1703](#)
- [STP Configuration Tasks on page 1704](#)
- [Protocols Configuration Statement on page 1709](#)
- [Unified Forwarding Table Configuration Statements on page 1723](#)
- [Reflective Relay Configuration Statements on page 1726](#)
- [STP Configuration Statements on page 1726](#)
- [VLAN Configuration Statements on page 1760](#)
- [Q-in-Q Configuration Statements on page 1817](#)

### Bridging and VLAN Configuration Examples

---

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563](#)
- [Example: Configuring Routing Between VLANs on One Switch on page 1576](#)
- [Example: Disabling MAC Learning on page 1582](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1583](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)

#### Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches

You can configure Ethernet ring protection switching (ERPS) on connected EX Series or QFX Series switches to prevent fatal loops from disrupting a network. ERPS is similar to

the Spanning Tree Protocol, but ERPS is more efficient because it is customized for ring topologies. You must configure at least three switches to form a ring.

This example shows how to configure Ethernet ring protection switching on four switches that are connected to one another on a dedicated link in a ring topology.



**NOTE:** This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Ethernet Ring Protection Switching on EX Series Switches*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#). This example will also work with QFX Series switches.

- [Requirements on page 1564](#)
- [Overview and Topology on page 1564](#)
- [Configuration on page 1565](#)
- [Verification on page 1575](#)

---

## Requirements

This example uses the following hardware and software components:

- Four connected EX Series switches, that support the Enhanced Layer 2 Software (ELS), that will function as nodes in the ring topology.
- Junos OS Release 13.2X50-D10 or later for EX Series switches.
- Junos OS Release 14.1X53-D10 or later for QFX Series.

Before you begin, be sure you have:

- Configured two trunk interfaces on each of the four switches. See [Table 98 on page 1565](#) for a list of the interface names used in this example.
- Configured a VLAN (with name **erp-control-vlan-1** and ID **100**) on all four switches and associated two network interfaces from each of the four switches with the VLAN. See *Configuring VLANs for the QFX Series OR Configuring VLANs for EX Series Switches (CLI Procedure)*. See [Table 98 on page 1565](#) for a list of the interface names used in this example.
- Configured two more VLANs (one with name **erp-data-1** and vlan ID **101** and a second vlan with the name **erp-data-2** and vlan ID **102**) on all four switches and associated both the east and west interfaces on each switch.

---

## Overview and Topology

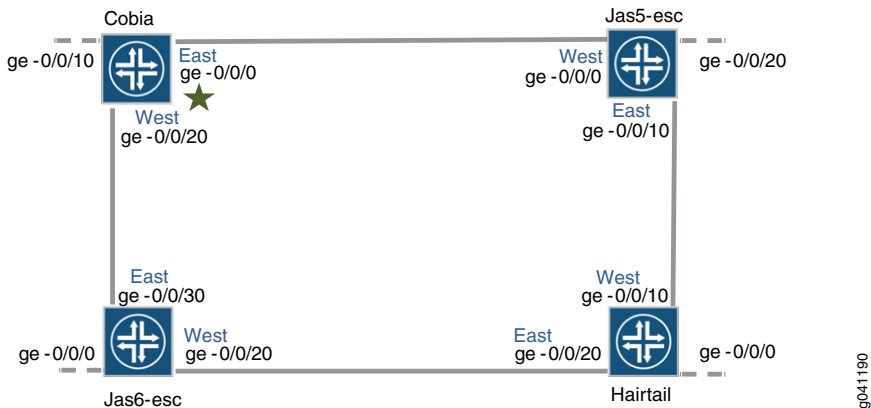
ERPS uses a dedicated physical link, including a control VLAN for trunk ports, between all of the switches to protect the active links. ERPS VLANs are all located on this link and are also blocked by default. When traffic between the switches is flowing with no problems, the active links take care of all traffic. Only if an error occurs on one of the data links would the ERPS control channel take over and start forwarding traffic.



**NOTE:** Trunk ports on switches use a VLAN to create individual control channels for ERPS. When multiple ERPS instances are configured for a ring, there are multiple sets of ring protection links (RPLs) and RPL owners on the ERPS link, and a different channel is blocked for each instance. Nontrunk ports use the physical link as the control channel and protocol data units (PDUs) are untagged, with no VLAN information in the packet.

This example creates one protection ring (called a node ring) named erp1 on four switches connected in a ring by trunk ports as shown in Figure 20 on page 1565. Because the links are trunk ports, VLAN 100 is used for erp1 traffic. The east interface of each switch is connected with the west interface of an adjacent switch. Cobia is the RPL owner, with interface ge-0/0/0 configured as an RPL end interface. The interface ge-0/0/0 of Jas5-esc is configured as the RPL neighbor interface. In the idle state, the RPL end blocks the control VLAN and data channel VLAN for this particular ERP instance—the blocked port on Cobia is marked with a star in Figure 20 on page 1565.

Figure 20: Ethernet Ring Protection Switching Example



In this example, we configure the four switches with the interfaces indicated in both Figure 20 on page 1565 and Table 98 on page 1565.

Table 98: Components to Configure for This Example

Interfaces	Cobia	Jas5-esc	Jas6-esc	Hairtail
East	ge-0/0/0	ge-0/0/10	ge-0/0/30	ge-0/0/20
West	ge-0/0/20	ge-0/0/0	ge-0/0/20	ge-0/0/10
Third	ge-0/0/10	ge-0/0/20	ge-0/0/0	ge-0/0/0

Configuration

- Configuring ERPS on Cobia, the RPL Owner Node on page 1566
- Configuring ERPS on Jas5-esc on page 1568

- [Configuring ERPS on Hairtail on page 1571](#)
- [Configuring ERPS on Jas6-esc on page 1573](#)

### *Configuring ERPS on Cobia, the RPL Owner Node*

#### CLI Quick Configuration

To quickly configure Cobia, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.



**NOTE:** RSTP and ERPS cannot both be configured on a ring port, and RSTP is configured by default. Therefore, in this example, RSTP is disabled on each ring port before configuring ERPS.

```
set protocols rstp interface ge-0/0/0 disable
set protocols rstp interface ge-0/0/20 disable
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 ring-protection-link-owner
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan
  100 ge-0/0/0.0
set protocols protection-group ethernet-ring erp1 east-interface ring-protection-link-end
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
  100 ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
  100 ge-0/0/20.0
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan
  100 ge-0/0/20.0 vlan 100
```

#### Step-by-Step Procedure

To configure ERPS on Cobia:

1. Disable RSTP on the two ports that will use ERPS:
 

```
[edit protocols]
user@switch# set rstp interface ge-0/0/0 disable
user@switch# set rstp interface ge-0/0/20 disable
```
2. Create a node ring named erp1:
 

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Designate Cobia as the RPL owner node:
 

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set ring-protection-link-owner
```
4. Configure the VLANs 101 and 102 as data channels:
 

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel 101
user@switch# set data-channel 102
```
5. Configure the control vlan 100 for this ERP instance on the trunk interface:



```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100
```

6. Configure the east interface of the node ring erp1 with the control channel ge-0/0/0.0 vlan 100 and indicate that this particular ring protection link ends here:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel vlan 100 ge-0/0/0.0
user@switch# set east-interface ring-protection-link-end
```

7. Configure the west interface of the node ring erp1 with the control channel vlan 100 ge-0/0/20.0 :

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel vlan 100 ge-0/0/20.0
```

8. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign 100 as the control VLAN on both interfaces:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel vlan 100 ge-0/0/20.0
user@switch# set east-interface control-channel vlan 100 ge-0/0/20.0
```

**Results** In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
rstp {
  interface ge-0/0/20.0 {
    disable;
  }
  interface ge-0/0/0.0 {
    disable;
  }
}
protection-group {
  ethernet-ring erp1 {
    ring-protection-link-owner;
    east-interface {
      control-channel {
        vlan 100
        ge-0/0/0.0;
      }
      ring-protection-link-end;
    }
    west-interface {
      control-channel {
        vlan 100
        ge-0/0/20.0;
      }
    }
  }
  control-vlan 100;
  data-channel {
    vlan 101-102;
  }
}
```

```
}  
}
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]  
user@switch# show interfaces  
ge-0/0/0 {  
  unit 0 {  
    family ethernet-switching {  
      interface-mode trunk;  
    }  
  }  
}  
ge-0/0/10 {  
  unit 0 {  
    family ethernet-switching {  
      interface-mode trunk;  
    }  
  }  
}  
ge-0/0/20 {  
  unit 0 {  
    family ethernet-switching {  
      interface-mode trunk;  
    }  
  }  
}
```

If you are finished configuring the device, enter **commit** in configuration mode.

### *Configuring ERPS on Jas5-esc*

#### **CLI Quick Configuration**

To quickly configure Jas5-esc, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols rstp interface ge-0/0/10 disable  
set protocols rstp interface ge-0/0/0 disable  
set protocols protection-group ethernet-ring erp1  
set protocols protection-group ethernet-ring erp1 data-channel 101  
set protocols protection-group ethernet-ring erp1 data-channel 102  
set protocols protection-group ethernet-ring erp1 control-vlan 100  
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan  
100 ge-0/0/10.0  
set protocols protection-group ethernet-ring erp1 east-interface control-channel vlan  
100 ge-0/0/0.0  
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan  
100 ge-0/0/0.0
```

```
set protocols protection-group ethernet-ring erp1 west-interface control-channel vlan
100 ge-0/0/10.0
```

### Step-by-Step Procedure

To configure ERPS on Jas5-esc:

1. Disable RSTP on the two ports that will use ERPS:  

```
[edit protocols]
user@switch# set rstp interface ge-0/0/10 disable
user@switch# set rstp interface ge-0/0/0 disable
```
2. Create a node ring named erp1:  

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Configure a control VLAN with ID 100 for the node ring erp1:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100
```
4. Configure two data channels named erp-data-1 and erp-data-2 to define a set of VLAN IDs that belong to a ring instance.  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel vlan 101
user@switch# set data-channel vlan 102
```
5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/10.0:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/10.0 vlan 100
```
6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/0.0 vlan 100:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/0.0 vlan 100
```
7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign vlan # 100 as the control VLAN:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/0.0 vlan 100
user@switch# set east-interface control-channel ge-0/0/10.0 vlan 100
```

**Results** In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
rstp {
  interface ge-0/0/10.0 {
    disable;
  }
  interface ge-0/0/0.0 {
```

```
        disable;
    }
}
protection-group {
    east-interface {
        control-channel {
            ge-0/0/10.0;
            vlan 100;
        }
    }
    west-interface {
        control-channel {
            ge-0/0/0.0;
            vlan 100;
        }
    }
}
control-vlan 100;
data-channel
    vlan 101-102
}
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show interfaces
ge-0/0/0 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
        }
    }
}
ge-0/0/10 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
        }
    }
}
ge-0/0/20 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
        }
    }
}
```

If you are finished configuring the device, enter **commit** in configuration mode.

*Configuring ERPS on Hairtail*

**CLI Quick Configuration** To quickly configure Hairtail, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols rstp interface ge-0/0/10 disable
set protocols rstp interface ge-0/0/20 disable
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/20.0 vlan 100
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/10.0 vlan 100
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0 vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/10.0 vlan 100
```

**Step-by-Step Procedure** To configure ERPS on Hairtail:

1. Disable RSTP on the two ports that will use ERPS:  

```
[edit protocols]
user@switch# set rstp interface ge-0/0/10 disable
user@switch# set rstp interface ge-0/0/20 disable
```
2. Create a node ring named erp1:  

```
[edit protocols]
user@switch# set protection-group ethernet-ring erp1
```
3. Configure the control vlan 100 for the node ring erp1:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100
```
4. Configure two data channels numbered 101 and 102 to define a set of VLAN IDs that belong to a ring instance:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel vlan 101
user@switch# set data-channel vlan 102
```
5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/20.0 vlan 100, and indicate that it connects to a ring protection link:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/20.0 vlan 100
```
6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/10.0 vlan 100 and indicate that it connects to a ring protection link:  

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/10.0 vlan 100
```

7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign 100 as the control VLAN:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/10.0 vlan 100
user@switch# set east-interface control-channel ge-0/0/20.0 vlan 100
```

**Results** In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
rstp {
  interface ge-0/0/10.0 {
    disable;
  }
  interface ge-0/0/20.0 {
    disable;
  }
}
protection-group {
  ethernet-ring erp1 {
    east-interface {
      control-channel {
        ge-0/0/20.0;
        vlan 100;
      }
    }
    west-interface {
      control-channel {
        ge-0/0/10.0;
        vlan 100;
      }
    }
    control-vlan 100;
    data-channel {
      vlan 101-102;
    }
  }
}
```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interface configurations by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show interfaces
ge-0/0/0 {
  unit 0 {
```

```

        family ethernet-switching {
            interface-mode trunk;
        }
    }
}
ge-0/0/10 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
        }
    }
ge-0/0/20 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
        }
    }
}

```

If you are finished configuring the device, enter **commit** in configuration mode.

### Configuring ERPS on Jas6-esc

#### CLI Quick Configuration

To quickly configure Jas6-esc, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set protocols rstp interface ge-0/0/30 disable
set protocols rstp interface ge-0/0/20 disable
set protocols protection-group ethernet-ring erp1
set protocols protection-group ethernet-ring erp1 data-channel 101
set protocols protection-group ethernet-ring erp1 data-channel 102
set protocols protection-group ethernet-ring erp1 control-vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/30.0 vlan 100
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0 vlan 100
set protocols protection-group ethernet-ring erp1 west-interface control-channel
ge-0/0/20.0 vlan 100
set protocols protection-group ethernet-ring erp1 east-interface control-channel
ge-0/0/30.0 vlan 100

```

#### Step-by-Step Procedure

To configure ERPS on Jas6-esc:

1. Disable RSTP on the two ports that will use ERPS:
 

```

[edit protocols]
user@switch# set rstp interface ge-0/0/30 disable
user@switch# set rstp interface ge-0/0/20 disable

```
2. Create a node ring named erp1:
 

```

[edit protocols]
user@switch# set protection-group ethernet-ring erp1

```
3. Configure the control vlan 100 for the node ring erp1:
 

```

[edit protocols protection-group ethernet-ring erp1]
user@switch# set control-vlan 100

```

4. Configure two data channels numbered 101 and 102 to define VLAN IDs that belong to a ring instance.

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set data-channel 101
user@switch# set data-channel 102
```

5. Configure the east interface of the node ring erp1 with the control channel ge-0/0/30.0 :

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set east-interface control-channel ge-0/0/30.0 vlan 100
```

6. Configure the west interface of the node ring erp1 with the control channel ge-0/0/20.0 vlan 100:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0 vlan 100
```

7. Every ring instance on a trunk port has one control VLAN in which ERP packets traverse. The control VLAN also controls data VLANs, if any are configured. Assign vlan number 100 as the control VLAN:

```
[edit protocols protection-group ethernet-ring erp1]
user@switch# set west-interface control-channel ge-0/0/20.0 vlan 100
user@switch# set east-interface control-channel ge-0/0/30.0 vlan 100
```

**Results** In configuration mode, check your ERPS configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@switch# show protocols
  rstp {
    interface ge-0/0/20.0 {
      disable;
    }
    interface ge-0/0/30.0 {
      disable;
    }
  }
  protection-group {
    ethernet-ring erp1 {
      east-interface {
        control-channel {
          vlan 100;
          ge-0/0/30.0;
        }
      }
      west-interface {
        control-channel {
          vlan 100;
          ge-0/0/20.0;
        }
      }
    }
    control-vlan erp-control-vlan-1;
    data-channel {
```



```

        vlan 101-102;
    }
}

```

In configuration mode, check your VLAN configuration by entering the **show vlans** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

In configuration mode, check your interfaces configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@switch# show interfaces
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
    }
  }
}
ge-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
    }
  }
}
ge-0/0/30 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
    }
  }
}

```

## Verification

Verify that ERPS is working correctly.

### Verifying That ERPS Is Working Correctly

**Purpose** Verify that ERPS is working on the four EX switches that function as nodes in the ring topology.

**Action** Check the state of the ring links in the output of the **show protection-group ethernet-ring interface** command. When the ring is configured but not being used (no error exists on the data links), one ERP interface is forwarding traffic and one is discarding traffic. Discarding blocks the ring.

```

user@switch> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group erp1
Interface  Forward State  RPL End  Signal Failure  Admin State
ge-0/0/2.0  discarding     yes      clear           ready
ge-0/0/0.0  forwarding     no       clear           ready

```

To find out what has occurred since the last restart, check the RPS statistics for ring-blocked events. **NR** is a No Request ring block, which means that the switch is not blocking either of the two ERP interfaces. **NR-RB** is a No Request Ring Blocked event, which means that the switch is blocking one of its ERP interfaces and sending a packet out to notify the other switches.

```
user@switch> show protection-group ethernet-ring statistics
Ring Name Local SF Remote SF NR Event NR-RB Event
erp1      2      1      2      3
```

**Meaning** The **show protection-group ethernet-ring interface** command output from the RPL owner node indicates that one interface is forwarding traffic and one is discarding traffic, meaning that the ERP is ready but not active. If at least one interface in the ring is not forwarding, the ring is blocked and therefore ERP is working.

The **show protection-group ethernet-ring statistics** command output indicates that, since the last reboot, both local and remote signal failures have occurred (**Local SF** and **Remote SF**).

The **NR Event** count is 2, indicating that the NR state was entered into twice. **NR** stands for No Request. This means that the switch either originated NR PDUs or received an NR PDU from another switch and stopped blocking the interface to allow ERP to function.

The three **NR-RB** events indicate that on three occasions, this switch either sent out NR-RB PDUs or received NR-RB PDUs from another switch. This occurs when a network problem is resolved and the switch once again blocks the ERP link at one end.

- Related Documentation**
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 1670](#)
  - [Ethernet Ring Protection Switching Overview on page 1525](#)
  - [Understanding Ethernet Ring Protection Switching Functionality on page 1534](#)

## Example: Configuring Routing Between VLANs on One Switch

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs you normally you need a router that connects the VLANs. However, you can accomplish this on a Juniper Networks switch without using a router by configuring an integrated routing and bridging (IRB) interface (also known as a routed VLAN interface—or RVI—in versions of Junos OS that do not support Enhanced Layer 2 Software). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

- [Requirements on page 1577](#)
- [Overview and Topology on page 1577](#)

- [Configure Layer 2 switching for two VLANs on page 1578](#)
- [Verification on page 1580](#)

Requirements


This example uses the following hardware and software components:

- One switch
- Junos OS Release 11.1 or later

Overview and Topology

This example uses an IRB to route traffic between two VLANs on the same switch. The topology is shown in [Figure 21 on page 1577](#).

Figure 21: IRB with One Switch



This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch and configuring an IRB to enable routing between the VLANs. One VLAN, called **blue**, is for the sales and marketing group, and a second, called **red**, is for the customer support team. The sales and support groups each have their own file servers and wireless access points. Each VLAN must have a unique name, tag (VLAN ID), and distinct IP subnet. [Table 99 on page 1577](#) lists the components of the sample topology.

Table 99: Components of the Multiple VLAN Topology

Property	Settings
VLAN names and tag IDs	<b>blue</b> , ID 100 <b>red</b> , ID 200
Subnets associated with VLANs	<b>blue</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>red</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Interfaces in VLAN <b>blue</b>	Sales server port: <b>xe-0/0/4</b> Sales wireless access points: <b>xe-0/0/6</b>
Interfaces in VLAN <b>red</b>	Support server port: <b>xe-0/0/0</b> Support wireless access points: <b>xe-0/0/2</b>
IRB name	interface <b>irb</b>
IRB units and addresses	logical unit 100: 192.0.2.1/25 logical unit 200: 192.0.2.129/25

This configuration example creates two IP subnets, one for the blue VLAN and the second for the red VLAN. The switch bridges traffic within the VLANs. For traffic passing between

two VLANs, the switch routes the traffic using an IRB on which you have configured addresses in each IP subnet.

To keep the example simple, the configuration steps show only a few interfaces and VLANs. Use the same configuration procedure to add more interfaces and VLANs. By default, all interfaces are in access mode, so you do not have to configure the port mode.

### Configure Layer 2 switching for two VLANs

#### CLI Quick Configuration

To quickly configure Layer 2 switching for the two VLANs (**blue** and **red**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:



**NOTE:** The following example uses a version of Junos OS that supports Enhanced Layer 2 Software (ELS). When you use ELS, you create a Layer 3 virtual interface named **irb**. If you are using a version of Junos OS that does not support ELS, you create a Layer 3 virtual interface named **vlan**.

```
[edit]
set interfaces xe-0/0/4 unit 0 description "Sales server port"
set interfaces xe-0/0/4 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/6 unit 0 description "Sales wireless access point port"
set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/0 unit 0 description "Support servers"
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members red
set interfaces xe-0/0/2 unit 0 description "Support wireless access point port"
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members red
set interfaces irb unit 100 family inet address 192.0.2.1/25
set interfaces irb unit 200 family inet address 192.0.2.129/25
set vlans blue l3-interface irb.100
set vlans blue vlan-id 100
set vlans red vlan-id 200
set vlans red l3-interface irb.200
```

#### Step-by-Step Procedure

To configure the switch interfaces and the VLANs to which they belong:

1. Configure the interface for the sales server in the blue VLAN:
 

```
[edit interfaces xe-0/0/4 unit 0]
user@switch# set description "Sales server port"
user@switch# set family ethernet-switching vlan members blue
```
2. Configure the interface for the wireless access point in the blue VLAN:
 

```
[edit interfaces xe-0/0/6 unit 0]
user@switch# set description "Sales wireless access point port"
user@switch# set family ethernet-switching vlan members blue
```
3. Configure the interface for the support server in the red VLAN:
 

```
[edit interfaces xe-0/0/0 unit 0]
user@switch# set description "Support server port"
user@switch# set family ethernet-switching vlan members red
```
4. Configure the interface for the wireless access point in the red VLAN:
 

```
[edit interfaces xe-0/0/2 unit 0]
user@switch# set description "Support wireless access point port"
```

```
user@switch# set family ethernet-switching vlan members red
```

**Step-by-Step Procedure** Now create the VLANs and the IRB. The IRB will have logical units in the broadcast domains of both VLANs.

1. Create the red and blue VLANs by configuring the VLAN IDs for them:  

```
[edit vlans]
user@switch# set blue vlan-id 100
user@switch# set red vlan-id 200
```
2. Create the interface named **irb** with a logical unit in the sales broadcast domain (blue VLAN):  

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.0.2.1/25
```

The unit number is arbitrary and does not have to match the VLAN tag ID. However, configuring the unit number to match the VLAN ID can help avoid confusion.
3. Add a logical unit in the support broadcast domain (red VLAN) to the **irb** interface:  

```
[edit interfaces]
user@switch# set irb unit 200 family inet address 192.0.2.129/25
```
4. Complete the IRB configuration by binding the red and blue VLANs (Layer 2) with the appropriate logical units of the **irb** interface (Layer 3):  

```
[edit vlans]
user@switch# set blue l3-interface irb.100
user@switch# set red l3-interface irb.200
```

Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  xe-0/0/4 {
    unit 0 {
      description "Sales server port";
      family ethernet-switching {
        vlan members blue;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      description "Sales wireless access point port";
      family ethernet-switching {
        vlan members blue;
      }
    }
  }
  xe-0/0/0 {
    unit 0 {
      description "Support server port";
      family ethernet-switching {
        vlan members red;
      }
    }
  }
}
```

```
xe-0/0/2 {
  unit 0 {
    description "Support wireless access point port";
    family ethernet-switching {
      vlan members red;
    }
  }
}
irb {
  unit 100 {
    family inet address 192.0.2.1/25;
  }
  unit 200 {
    family inet address 192.0.2.129/25;
  }
}
}
vllans {
  blue {
    vlan-id 100;
    interface xe-0/0/4.0;
    interface xe-0/0/6.0;
    l3-interface irb 100;
  }
  red {
    vlan-id 200;
    interface xe-0/0/0.0;
    interface xe-0/0/2.0;
    l3-interface irb 200;
  }
}
```



**TIP:** To quickly configure the blue and red VLAN interfaces, issue the `load merge terminal` command, copy the hierarchy, and paste it into the switch terminal window.

---

## Verification

To verify that the **blue** and **red** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 1580](#)
- [Verifying That Traffic Can Be Routed Between the Two VLANs on page 1581](#)

### *Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces*

**Purpose** Verify that the VLANs **blue** and **red** have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
Name      Tag      Interfaces
default   100      xe-0/0/0.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/6.0,
blue      100      xe-0/0/4.0, xe-0/0/6.0,
red       200      xe-0/0/0.0, xe-0/0/2.0, *
mgmt      me0.0*
```

**Meaning** The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **blue** and **red** VLANs have been created. The **blue** VLAN has a tag ID of 100 and is associated with interfaces **xe-0/0/4.0** and **xe-0/0/6.0**. VLAN **red** has a tag ID of 200 and is associated with interfaces **xe-0/0/0.0** and **xe-0/0/2.0**.

### *Verifying That Traffic Can Be Routed Between the Two VLANs*

**Purpose** Verify routing between the two VLANs.

**Action** Verify that the IRB logical units are up:

```
user@switch> show interfaces terse
irb.100          up    up    inet    192.0.2.1/25
irb.200          up    up    inet    192.0.2.129/25
```



**NOTE:** At least one port (access or trunk) with an appropriate VLAN assigned to it must be up for the irb interface to be up.

Verify that switch has created routes that use the IRB logical units:

```
user@switch> show route
192.0.2.0/25      *[Direct/0] 1d 03:26:45
                  > via irb.100
192.0.2.1/32      *[Local/0] 1d 03:26:45
                  Local via irb.100
192.0.2.128/25    *[Direct/0] 1d 03:26:45
                  > via irb.200
192.0.2.129/32    *[Local/0] 1d 03:26:45
                  Local via irb.200
```

List the Layer 3 routes in the switch's Address Resolution Protocol (ARP) table:

```
user@switch> show arp
MAC Address      Address      Name      Flags
00:00:0c:06:2c:0d 192.0.2.7    irb.100    None
00:13:e2:50:62:e0 192.0.2.132  irb.200    None
```

**Meaning** The output of the **show interfaces** and **show route** commands show that the Layer 3 IRB logical units are working and the switch has used them to create direct routes that it will use to forward traffic between the VLAN subnets. The **show arp** command displays

the mappings between the IP addresses and MAC addresses for devices on both **irb.100** (associated with VLAN **blue**) and **irb.200** (associated with VLAN **red**). These two devices can communicate.

- Related Documentation**
- [Understanding Integrated Routing and Bridging on page 1539](#)
  - [irb \(Interfaces\) on page 1787](#)
  - [l3-interface on page 1790](#)

## Example: Disabling MAC Learning

By default, MAC learning is enabled on the QFX Series. This topic provides examples for disabling, enabling, and verifying the operation of MAC learning on the QFX Series. These examples require that you be logged in as the root user to the switch on which you wish to modify MAC learning.



**NOTE:** This task uses Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Disabling MAC Learning*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

- To disable MAC learning in a VLAN:

```
[edit]
user@switch# set vlans vlan10 switch-options interface xe-0/0/0.0 no-mac-learning
```

- To reenab MAC learning:

```
[edit] vlans vlan10 switch-options interface xe-0/0/0.0
user@switch# delete no-mac-learning
```

- To verify the status of MAC learning on the QFX Series:

```
user@switch> show ethernet-switching table
Learning stats: 10 learn msg rcvd, 2 error, 0 forced update
Interface          Local pkts  Transit pkts  Error
xe-0/0/0.0          0           6             1
xe-0/0/22.0          0           0             0
xe-0/0/1.0           0           4             1
xe-0/0/2.0           0           0             0
xe-0/0/3.0           0           0             0
xe-0/0/4.0           0           0             0
xe-0/0/19.0          0           0             0
xe-0/0/18.0          0           0             0
xe-0/0/9.0           0           0             0
```

- Related Documentation**
- [Understanding MAC Learning on page 1540](#)
  - [Disabling MAC Learning on page 1681](#)
  - *no-mac-learning*



## Example: Setting Up Bridging with Multiple VLANs

The QFX Series products use bridging and virtual LANs (VLANs) to connect network devices in a LAN—storage devices, file servers, and other network components—and to segment the LAN into smaller bridging domains.

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure bridging for the QFX Series and how to create two VLANs to segment the LAN:



**NOTE:** This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#). If your switch runs software that does not support ELS, see *Example: Setting Up Bridging with Multiple VLANs*.

- [Requirements on page 1583](#)
- [Overview and Topology on page 1583](#)
- [Configuration on page 1584](#)
- [Verification on page 1586](#)

### Requirements

This example uses the following hardware and software components:

- A configured and provisioned QFX3500 switch
- Junos OS Release 13.2X50-D15 or later for the QFX Series

### Overview and Topology

Switches connect all devices in an office or data center into a single LAN to provide sharing of common resources such as file servers. The default configuration creates a single VLAN, and all traffic on the switch is part of that broadcast domain. Creating separate network segments reduces the span of the broadcast domain and enables you to group related users and network resources without being limited by physical cabling or by the location of a network device in the building or on the LAN.

This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch. One VLAN, called **sales**, is for the sales and marketing group, and a second, called **support**, is for the customer support team. The sales and support groups each have their own dedicated file servers and other resources. For the switch ports to be segmented across the two VLANs, each VLAN must have its own broadcast

domain, identified by a unique name and tag (VLAN ID). In addition, each VLAN must be on its own distinct IP subnet.

The topology used in this example consists of a single QFX3500 switch, with a total of 48 10-Gbps Ethernet ports. (For the purposes of this example, the QSFP+ ports Q0-Q3, which are ports xe-0/1/0 through xe-0/1/15, are excluded.)

**Table 100: Components of the Multiple VLAN Topology**

Property	Settings
Switch hardware	QFX3500 switch configured with 48 10-Gbps Ethernet ports (xe-0/0/0 through xe-0/0/47)
VLAN names and tag IDs	<b>sales</b> , tag 100 <b>support</b> , tag 200
VLAN subnets	<b>sales</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>support</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Interfaces in VLAN <b>sales</b>	File servers: <b>xe-0/0/20</b> and <b>xe-0/0/21</b>
Interfaces in VLAN <b>support</b>	File servers: <b>xe-0/0/46</b> and <b>xe-0/0/47</b>
Unused interfaces	<b>xe-0/0/2</b> and <b>xe-0/0/25</b>

This configuration example creates two IP subnets, one for the sales VLAN and the second for the support VLAN. The switch bridges traffic within a VLAN. For traffic passing between two VLANs, the switch routes the traffic using a Layer 3 routing interface on which you have configured the address of the IP subnet.

To keep the example simple, the configuration steps show only a few devices in each of the VLANs. Use the same configuration procedure to add more LAN devices.

### Configuration

#### CLI Quick Configuration

To quickly configure Layer 2 switching for the two VLANs (**sales** and **support**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/20 unit 0 description "Sales file server port"
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces xe-0/0/46 unit 0 description "Support file server port"
set interfaces xe-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces vlan unit 0 family inet address 192.0.2.0/25
set interfaces vlan unit 1 family inet address 192.0.2.128/25
```

```

set vlans sales l3-interface irb.0
set vlans sales vlan-id 100
set vlans support vlan-id 200
set vlans support l3-interface irb.1

```

**Step-by-Step Procedure** Configure the switch interfaces and the VLANs to which they belong. By default, all interfaces are in access mode, so you do not have to configure the port mode.

1. Configure the interface for the file server in the **sales** VLAN:  

```

[edit interfaces xe-0/0/20 unit 0]
user@switch# set description "Sales file server port"
user@switch# set family ethernet-switching vlan members sales

```
2. Configure the interface for the file server in the **support** VLAN:  

```

[edit interfaces xe-0/0/46 unit 0]
user@switch# set description "Support file server port"
user@switch# set family ethernet-switching vlan members support

```
3. Create the subnet for the **sales** broadcast domain:  

```

[edit interfaces]
user@switch# set vlan unit 0 family inet address 192.0.2.1/25

```
4. Create the subnet for the **support** broadcast domain:  

```

[edit interfaces]
user@switch# set vlan unit 1 family inet address 192.0.2.129/25

```
5. Configure the VLAN tag IDs for the **sales** and **support** VLANs:  

```

[edit vlans]
user@switch# set sales vlan-id 100
user@switch# set support vlan-id 200

```
6. To route traffic between the **sales** and **support** VLANs, define the interfaces that are members of each VLAN and associate a Layer 3 interface:  

```

[edit vlans]
user@switch# set sales l3-interface irb.0
user@switch# set support l3-interface irb.1

```

Display the results of the configuration:

```

user@switch> show configuration
interfaces {
  xe-0/0/20 {
    unit 0 {
      description "Sales file server port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  xe-0/0/46 {
    unit 0 {
      description "Support file server port";
      family ethernet-switching {
        vlan members support;
      }
    }
  }
  vlans {
    unit 0 {

```

```

        family inet address 192.0.2.1/25;
    }
    unit 1 {
        family inet address 192.0.2.129/25;
    }
}
}
}
vlands {
    sales {
        vlan-id 100;
        interface xe-0/0/0.0;
        interface xe-0/0/3.0;
        interface xe-0/0/20.0;
        interface xe-0/0/22.0;
        l3-interface irb0;
    }
    support {
        vlan-id 200;
        interface xe-0/0/24.0;
        interface xe-0/0/26.0;
        interface xe-0/0/44.0;
        interface xe-0/0/46.0;
        l3-interface irb1;
    }
}
}

```



**TIP:** To quickly configure the sales and support VLAN interfaces, issue the **load merge terminal** command. Then copy the hierarchy and paste it into the switch terminal window.

## Verification

Verify that the **sales** and **support** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 1586](#)
- [Verifying That Traffic Is Being Routed Between the Two VLANs on page 1587](#)
- [Verifying That Traffic Is Being Switched Between the Two VLANs on page 1587](#)

### *Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces*

**Purpose** Verify that the **sales** and **support** VLANs have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

**Action** To list all VLANs configured on the switch, use the **show vlans** command:

```

user@switch> show vlans
Name      Tag      Interfaces
default
                                xe-0/0/1.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/5.0,

```

```

xe-0/0/6.0, xe-0/0/7.0, xe-0/0/8.0, xe-0/0/9.0,
xe-0/0/10.0*, xe-0/0/11.0, xe-0/0/12.0, xe-0/0/13.0*,
xe-0/0/14.0, xe-0/0/15.0, xe-0/0/16.0, xe-0/0/17.0,
xe-0/0/18.0, xe-0/0/19.0, xe-0/0/21.0, xe-0/0/23.0*,
xe-0/0/25.0, xe-0/0/27.0, xe-0/0/28.0, xe-0/0/29.0,
xe-0/0/30.0, xe-0/0/31.0, xe-0/0/32.0, xe-0/0/33.0,
xe-0/0/34.0, xe-0/0/35.0, xe-0/0/36.0, xe-0/0/37.0,
xe-0/0/38.0, xe-0/0/39.0, xe-0/0/40.0, xe-0/0/41.0,
xe-0/0/42.0, xe-0/0/43.0, xe-0/0/45.0, xe-0/0/47.0,
xe-0/1/0.0*, xe-0/1/1.0*, xe-0/1/2.0*, xe-0/1/3.0*

sales      100
           xe-0/0/0.0*, xe-0/0/3.0, xe-0/0/20.0, xe-0/0/22.0

support    200
           xe-0/0/0.24, xe-0/0/26.0, xe-0/0/44.0, xe-0/0/46.0*

mgmt
           me0.0*

```

**Meaning** The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **sales** and **support** VLANs have been created. The **sales** VLAN has a tag ID of 100 and is associated with interfaces **xe-0/0/0.0**, **xe-0/0/3.0**, **xe-0/0/20.0**, and **xe-0/0/22.0**. VLAN **support** has a tag ID of 200 and is associated with interfaces **xe-0/0/24.0**, **xe-0/0/26.0**, **xe-0/0/44.0**, and **xe-0/0/46.0**.

#### *Verifying That Traffic Is Being Routed Between the Two VLANs*

**Purpose** Verify routing between the two VLANs.

**Action** List the Layer 3 routes in the switch Address Resolution Protocol (ARP) table:

```

user@switch> show arp
MAC Address      Address      Name      Flags
00:00:0c:06:2c:0d 192.0.2.3    vlan.0    None
00:13:e2:50:62:e0 192.0.2.11   vlan.1    None

```

**Meaning** Sending IP packets on a multiaccess network requires mapping from an IP address to a MAC address (the physical or hardware address). The ARP table displays the mapping between the IP address and MAC address for both **vlan.0** (associated with **sales**) and **vlan.1** (associated with **support**). These VLANs can route traffic to each other.

#### *Verifying That Traffic Is Being Switched Between the Two VLANs*

**Purpose** Verify that learned entries are being added to the Ethernet switching table.

**Action** List the contents of the Ethernet switching table:

```

user@switch> show ethernet-switching table

Ethernet-switching table: 8 entries, 5 learned
VLAN      MAC address  Type  Age  Interfaces
default   *           Flood  -   All-members

```

default	00:00:05:00:00:01	Learn	- xe-0/0/10.0
default	00:00:5e:00:01:09	Learn	- xe-0/0/13.0
default	00:19:e2:50:63:e0	Learn	- xe-0/0/23.0
sales	*	Flood	- All-members
sales	00:00:5e:00:07:09	Learn	- xe-0/0/0.0
support	*	Flood	- All-members
support	00:00:5e:00:01:01	Learn	- xe-0/0/46.0

**Meaning** The output shows that learned entries for the **sales** and **support** VLANs have been added to the Ethernet switching table, and are associated with interfaces **xe-0/0/0.0** and **xe-0/0/46.0**. Even though the VLANs were associated with more than one interface in the configuration, these interfaces are the only ones that are currently operating.

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Understanding Bridging and VLANs on page 1527](#)

## Example: Setting Up Basic Bridging and a VLAN on the QFX Series

The QFX Series products use bridging and virtual LANs (VLANs) to connect network devices—storage devices, file servers, and other LAN components—in a LAN and to segment the LAN into smaller bridging domains.

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure basic bridging and VLANs for the QFX Series:

- [Requirements on page 1588](#)
- [Overview and Topology on page 1588](#)
- [Configuration on page 1589](#)
- [Verification on page 1598](#)

### Requirements

---

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX Series
- A configured and provisioned QFX Series product

### Overview and Topology

---

To use a switch to connect network devices on a LAN, you must at a minimum configure bridging and VLANs. By default, bridging is enabled on all switch interfaces, all interfaces are in access mode, and all interfaces belong to a VLAN called **employee-vlan**, which is

automatically configured. When you plug in access devices—such as desktop computers, file servers, and printers—they are joined immediately into the **employee-vlan** VLAN, and the LAN is up and running.

The topology used in this example consists of a single QFX3500 switch, with a total of 48 10-Gbps Ethernet ports. (For the purposes of this example, the QSFP+ ports Q0-Q3, which are ports xe-0/1/0 through xe-0/1/15, are excluded.) You use the ports to connect devices that have their own power sources. Table 1 details the topology used in this configuration example.

**Table 101: Components of the Basic Bridging Configuration Topology**

Property	Settings
Switch hardware	QFX3500 switch, with 48 10-Gbps Ethernet ports
VLAN name	<b>employee-vlan</b>
VLAN ID	10
Connections to file servers	<b>xe-0/0/17</b> and <b>xe-0/0/18</b>
Direct connections to desktop PCs and laptops	<b>xe-0/0/0</b> through <b>xe-0/0/16</b>
Connections to integrated printer/fax/copier machines	<b>xe-0/0/19</b> through <b>xe-0/0/40</b>
Unused ports	<b>xe-0/0/41</b> through <b>xe-0/0/47</b>

### Configuration

#### CLI Quick Configuration

To quickly configure a VLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans employee-vlan vlan-id 10
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/3 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/4 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/5 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/7 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/8 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/9 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/12 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/15 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/16 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/17 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/18 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/19 unit 0 family ethernet-switching vlan members employee-vlan
```

```

set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/22 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/23 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/25 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/27 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/28 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/29 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/34 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/35 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/36 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/37 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/38 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/39 unit 0 family ethernet-switching vlan members employee-vlan
set interfaces xe-0/0/40 unit 0 family ethernet-switching vlan members employee-vlan

```

### Step-by-Step Procedure

To set up basic bridging and a VLAN:

1. Create a VLAN named employee-vlan and specify the VLAN ID of 10 for it:

```

[edit vlans]
user@switch# set employee-vlan vlan-id 10

```

2. Assign interfaces xe-0/0/0 through xe-0/0/40 to the employee-vlan VLAN:

```

[edit interface]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/1 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/2 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/3 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/4 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/5 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/6 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/7 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/8 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/9 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/10 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/11 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/12 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/13 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/14 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/15 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/16 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/17 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/18 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/19 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/20 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/21 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/22 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/23 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/24 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/25 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/26 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/27 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/28 unit 0 family ethernet-switching vlan members employee-vlan

```



```
user@switch# set xe-0/0/29 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/30 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/31 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/32 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/33 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/34 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/35 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/36 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/37 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/38 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/39 unit 0 family ethernet-switching vlan members employee-vlan
user@switch# set xe-0/0/40 unit 0 family ethernet-switching vlan members employee-vlan
```

3. Connect the two file servers to ports xe-0/0/17 and xe-0/0/18.
4. Connect the desktop PCs and laptops to ports xe-0/0/0 through xe-0/0/16.
5. Connect the integrated printer/fax/copier machines to ports xe-0/0/19 through xe-0/0/40.

**Results** Check the results of the configuration:

```
user@switch> show configuration
xe-0/0/0 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/2 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/3 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/4 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/5 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/6 {
  unit 0 {
    family ethernet-switching {
```

```
        vlan {
            members employee-vlan;
        }
    }
}
xe-0/0/7 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/8 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/9 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/10 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/11 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/12 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
```

```
xe-0/0/13 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/14 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/15 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/16 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/17 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/18 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
xe-0/0/19 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee-vlan;
      }
    }
  }
}
```

```
    }  
  }  
  xe-0/0/20 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/21 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/22 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/23 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/24 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/25 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }  
  xe-0/0/26 {  
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }
```

```
    }  
  }  
}  
xe-0/0/27 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/28 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/29 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/30 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/31 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/32 {  
  unit 0 {  
    family ethernet-switching {  
      vlan {  
        members employee-vlan;  
      }  
    }  
  }  
}  
xe-0/0/33 {  
  unit 0 {  
    family ethernet-switching {
```

```

        vlan {
            members employee-vlan;
        }
    }
}
xe-0/0/34 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/35 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/36 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/37 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/38 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/39 {
    unit 0 {
        family ethernet-switching {
            vlan {
                members employee-vlan;
            }
        }
    }
}
xe-0/0/40 {

```

```
    unit 0 {  
      family ethernet-switching {  
        vlan {  
          members employee-vlan;  
        }  
      }  
    }  
  }
```

---

### Verification

To verify that switching is operational and that **employee-vlan** has been created, perform these tasks:

- [Verifying That the VLAN Has Been Created on page 1598](#)
- [Verifying That Interfaces Are Associated with the Proper VLANs on page 1599](#)

#### ***Verifying That the VLAN Has Been Created***

**Purpose** Verify that the VLAN named **employee-vlan** has been created on the switch.



**Action** List all VLANs configured on the switch:

```

user@switch> show vlans
Routing instance      VLAN name      Tag      Interfaces
default-switch        employee-vlan   10
                      xe-0/0/0.0
                      xe-0/0/1.0
                      xe-0/0/2.0
                      xe-0/0/3.0
                      xe-0/0/4.0
                      xe-0/0/5.0
                      xe-0/0/6.0
                      xe-0/0/7.0
                      xe-0/0/8.0
                      xe-0/0/9.0
                      xe-0/0/10.0
                      xe-0/0/11.0
                      xe-0/0/12.0
                      xe-0/0/13.0
                      xe-0/0/14.0
                      xe-0/0/15.0
                      xe-0/0/16.0
                      xe-0/0/17.0
                      xe-0/0/18.0
                      xe-0/0/19.0
                      xe-0/0/20.0
                      xe-0/0/21.0
                      xe-0/0/22.0
                      xe-0/0/23.0
                      xe-0/0/24.0
                      xe-0/0/25.0
                      xe-0/0/26.0
                      xe-0/0/27.0
                      xe-0/0/28.0
                      xe-0/0/29.0
                      xe-0/0/30.0
                      xe-0/0/31.0
                      xe-0/0/32.0
                      xe-0/0/33.0
                      xe-0/0/34.0
                      xe-0/0/35.0
                      xe-0/0/36.0
                      xe-0/0/37.0
                      xe-0/0/38.0
                      xe-0/0/39.0
                      xe-0/0/40.0
...

```

**Meaning** The `show vlans` command lists the VLANs configured on the switch. This output shows that the VLAN `employee-vlan` has been created.

#### *Verifying That Interfaces Are Associated with the Proper VLANs*

**Purpose** Verify that Ethernet switching is enabled on switch interfaces and that all interfaces are included in the VLAN.

**Action** List all interfaces on which switching is enabled:

```

user@switch> show ethernet-switching interfaces
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/0.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/1.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/2.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/3.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/4.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/5.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/6.0                65535                untagged
                        employee-vlan 10
                        65535 Discarding
Routing Instance Name : default-switch

```

```

Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/7.0                65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/8.0                65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/9.0                65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/10.0               65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/11.0              65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/12.0              65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/13.0              65535                untagged
                        employee-vlan 10
                        65535   Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical   Vlan   TAG   MAC   STP   Logical   Tagging
interface members   limit state   interface flags
xe-0/0/14.0              65535                untagged
                        employee-vlan 10
                        65535   Discarding

```

```

        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/15.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/16.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/17.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/18.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/19.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/20.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/21.0   65535    untagged
        employee-vlan 10
            65535    Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

```

```

Logical interface xe-0/0/22.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/23.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/24.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/25.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/26.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/27.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/28.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )

Logical interface xe-0/0/29.0
Vlan members employee-vlan 10
TAG
MAC limit 65535
STP state Discarding
Logical interface flags Tagging
untagged

```

```

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/30.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/31.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/32.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/33.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/34.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/35.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/36.0  employee-vlan 10 65535
                        Discarding

Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags

```

```

xe-0/0/37.0          65535          untagged
                    employee-vlan 10
                    65535          Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/38.0   65535          untagged
                    employee-vlan 10
                    65535          Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/39.0   65535          untagged
                    employee-vlan 10
                    65535          Discarding
Routing Instance Name : default-switch
Logical Interface flags (DL - disable learning, AD - packet action drop,
                        LH - MAC limit hit, DN - interface down )
Logical      Vlan      TAG  MAC      STP      Logical      Tagging
interface    members   limit state   interface flags
xe-0/0/40.0   65535          untagged
                    employee-vlan 10
                    65535          Discarding
...

```

**Meaning** The `show ethernet-switching interfaces` command lists all interfaces on which switching is enabled (in the **Logical interface** column), along with the VLANs that are active on the interfaces (in the **VLAN members** column). The output in this example shows all the connected interfaces, xe-0/0/0 through xe-0/0/40, are all part of VLAN **employee-vlan**. Notice that the interfaces listed are the logical interfaces, not the physical interfaces. For example, the output shows xe-0/0/0.0 instead of xe-0/0/0. This is because Junos OS creates VLANs on logical interfaces, not directly on physical interfaces.

**Related Documentation**

- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding Bridging and VLANs on page 1527](#)

## Reflective Relay Configuration Example

- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 1605](#)

### Example: Configuring Reflective Relay for Use with VEPA Technology

Reflective relay must be configured on a switch that receives virtual machine aggregated packets, such as Virtual Ethernet Port Aggregator (VEPA) packets, because some of these packets might be sent back to the server destined for another virtual machine on the same server. Reflective relay returns those packets to the original device using the same downstream port that delivered the packets to the switch.



**NOTE:** This example uses Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Reflective Relay for Use with VEPA Technology*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 43.

This example shows how to configure a switch port interface to return packets sent by VEPA on the downstream interface back to the server using the same downstream interface:

- [Requirements on page 1606](#)
- [Overview and Topology on page 1606](#)
- [Configuration on page 1608](#)
- [Verification on page 1608](#)

### Requirements

---

This example uses the following hardware and software components:

- One QFX3500 switch
- One server
- Junos OS Release 12.1 or later for the QFX Series

Before you configure reflective relay on a switch port, be sure you have:

- Configured a server with six virtual machines, VM 1 through VM 6.
- Configured the server with three VLANs named VLAN\_Purple, VLAN\_Orange, and VLAN\_Blue and added two virtual machines to each VLAN.
- Configured the same three VLANs named VLAN\_Purple, VLAN\_Orange, and VLAN\_Blue on one interface.
- Installed and configured VEPA to aggregate the virtual machine packets.

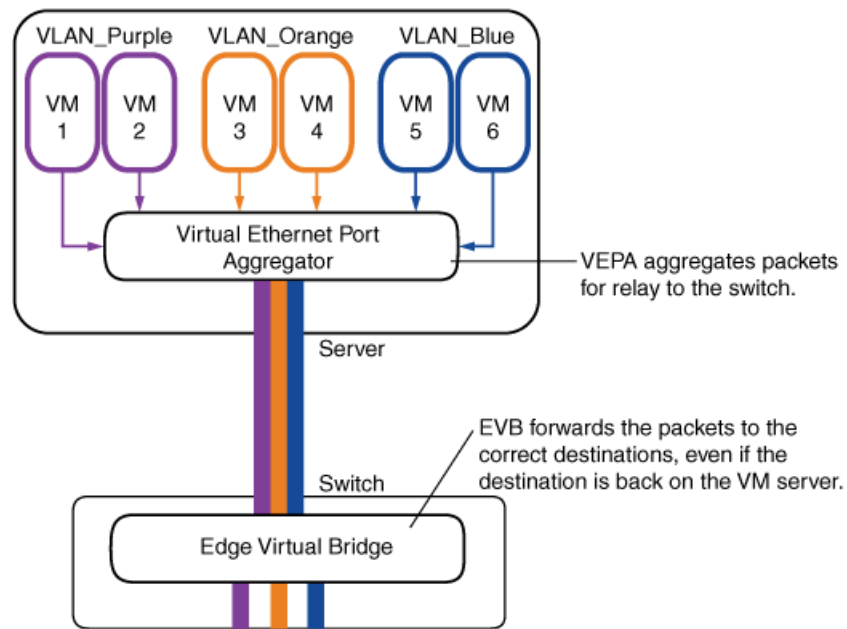
### Overview and Topology

---

In this example, illustrated in [Figure 22 on page 1607](#), a switch is connected to one server that is hosting six virtual machines and is configured with a VEPA for aggregating packets. The server's six virtual machines are VM1 through VM 6, and each virtual machine belongs to one of the three server VLANs, VLAN\_Purple, VLAN\_Orange, or VLAN\_Blue. Instead of the server directly passing packets between virtual machines, packets from any of the three VLANs that are destined for another one of the three VLANs are aggregated using VEPA technology and passed to the switch for processing. You must configure the switch port to accept these aggregated packets on the downstream interface and to return appropriate packets to the server on the same downstream interface after they are processed. [Figure 22 on page 1607](#) shows the topology for this example.



Figure 22: Reflective Relay Topology



g020996

In this example, you configure the physical Ethernet switch port interface for trunk interface mode and reflective relay. Configuring trunk port mode allows the interface to accept VLAN tagged packets. Configuring reflective relay allows the downstream port to return those packets on the same interface. [Table 102 on page 1607](#) shows the components used in this example.

Table 102: Components of the Topology for Configuring Reflective Relay

Component	Description
QFX3500 switch	Switch that supports reflective relay. For a list of switches that support this feature, see <i>QFX Series Software Features Overview</i> .
xe-0/0/2	Switch interface to the server.
Server	Server with virtual machines and VEPA technology.
Virtual machines	Six virtual machines located on the server: V1, V2, V3, V4, V5, and V6.
VLANs	Three VLANs: VLAN_Purple, VLAN_Orange, and VLAN_Blue. Each VLAN has two virtual machine members.
VEPA	Virtual Ethernet port aggregator that aggregates virtual machine packets on the server before the resulting single stream is transmitted to the switch.

## Configuration

---

To configure reflective relay, perform these tasks:

- [Configuring Reflective Relay on the Port on page 1608](#)

### *Configuring Reflective Relay on the Port*

#### CLI Quick Configuration

To quickly configure reflective relay, copy the following commands and paste them into the switch window:

```
[edit]
set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members [VLAN_Blue VLAN_Orange VLAN_Purple]
```

#### Step-by-Step Procedure

To configure reflective relay:

1. Configure the trunk interface mode on the interface:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure reflective relay on the interface to allow it to both accept and send packets:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```

3. Configure the interface for the three VLANs on the server:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members [VLAN_Purple VLAN_Orange VLAN_Blue]
```

#### Results

Check the results of the configuration:

```
[edit interfaces xe-0/0/2]
user@switch# show
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        reflective-relay;
        vlan {
            members [ VLAN_Purple VLAN_Orange VLAN_Blue ];
        }
    }
}
```

## Verification

---

To confirm that reflective relay is enabled and working correctly, perform these tasks:

- [Verifying That Reflective Relay Is Enabled and Working Correctly on page 1608](#)

### *Verifying That Reflective Relay Is Enabled and Working Correctly*

#### Purpose

Verify that reflective relay is enabled and working correctly.

**Action** Use the `show ethernet-switching interfaces detail` command to display the reflective relay status:

```
user@switch> show ethernet-switching interfaces xe-0/0/2 detail
Interface: xe-0/0/2, Index: 66, State: down, Interface mode: Trunk
Reflective Relay Status: Enabled
Ether type for the interface: 0x8100
VLAN membership:
  VLAN_Purple, 802.1Q Tag: 450, tagged, unblocked
  VLAN_Orange, 802.1Q Tag: 460, tagged, unblocked
  VLAN_Blue, 802.1Q Tag: 470, tagged, unblocked
Number of MACs learned on IFL: 0
```

Confirm that reflective relay is working by sending a Layer 2 broadcast message from one virtual machine to another virtual machine located on the same VLAN. Check the switch to verify that the switch sends the packets back on the same interface on which they were received. One way to check this is to set up port mirroring on the switch interface, connect a traffic generator to the mirrored interface, and use the traffic generator to examine packets.

Alternatively, if you do not have a traffic generator available, you can send traffic between two virtual machines with FTP, Telnet, or SSH, while running the `tcpdump` utility on the receiver virtual machine port to capture reflected packets.

**Meaning** The reflective relay status is **Enabled**, meaning that interface `xe-0/0/2` is configured for the trunk interface mode, which accepts VLAN-tagged packets, and for reflective relay, which accepts and returns packets on the same interface.

When the traffic generator shows packets arriving at the switch and returning to the server on the same interface, reflective relay is working.

**Related Documentation**

- [Understanding Reflective Relay for Use with VEPA Technology on page 1553](#)
- [Configuring Port Mirroring](#)
- [interface-mode on page 1785](#)
- [reflective-relay on page 1726](#)

## STP Configuration Examples

- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
- [Example: Configuring Network Regions for VLANs with MSTP on page 1624](#)
- [Example: Connecting an Access Switch to a Distribution Switch on page 1647](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656](#)

- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660](#)
- [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1664](#)

## Example: Configuring Faster Convergence and Improving Network Stability with RSTP

The QFX Series products use Rapid Spanning Tree Protocol (RSTP) to provide a loop-free topology. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state. RSTP provides quicker reconvergence time than original STP because it uses protocol handshake messages rather than fixed timeouts. Eliminating the need to wait for timers to expire makes RSTP more efficient than STP.

This example describes how to configure RSTP on four QFX3500 switches:

- [Requirements on page 1610](#)
- [Overview and Topology on page 1610](#)
- [Configuring RSTP on Switch 1 on page 1612](#)
- [Configuring RSTP on Switch 2 on page 1615](#)
- [Configuring RSTP on Switch 3 on page 1617](#)
- [Configuring RSTP on Switch 4 on page 1620](#)
- [Verification on page 1622](#)

### Requirements

---

This example uses the following hardware and software components:


- Junos OS Release 11.1 for the QFX3500 switches
- Four QFX3500 switches

### Overview and Topology

---

In this example, QFX3500 switches are connected in the topology displayed in [Figure 23 on page 1611](#) to create a loop-free topology.

Figure 23: Network Topology for RSTP



The interfaces shown in [Table 103 on page 1611](#) will be configured for RSTP.



**NOTE:** You can configure RSTP on logical or physical interfaces. This example shows RSTP configured on logical interfaces.

Table 103: Topology for Configuring RSTP on the QFX Series

Components	Settings
Switch 1	<p>The following ports on Switch 1 are connected in this way:</p> <ul style="list-style-type: none"><li>• xe-0/0/9 is connected to Switch 2</li><li>• xe-0/0/13 is connected to Switch 4</li><li>• xe-0/0/11 is connected to Switch 3</li></ul>
Switch 2	<p>The following ports on Switch 2 are connected in this way:</p> <ul style="list-style-type: none"><li>• xe-0/0/14 is connected to Switch 1</li><li>• xe-0/0/18 is connected to Switch 3</li></ul>
Switch 3	<p>The following ports on Switch 3 are connected in this way:</p> <ul style="list-style-type: none"><li>• xe-0/0/26 is connected to Switch 1</li><li>• xe-0/0/28 is connected to Switch 2</li><li>• xe-0/0/24 is connected to Switch 4</li></ul>

Table 103: Topology for Configuring RSTP on the QFX Series (*continued*)

Components	Settings
Switch 4	<p>The following ports on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>xe-0/0/19</b> is connected to Switch 1</li> <li>• <b>xe-0/0/23</b> is connected to Switch 3</li> </ul>
VLAN names and tag IDs	<p><b>sales-vlan</b>, tag 10</p> <p><b>engineering-vlan</b>, tag 20</p> <p><b>publications-vlan</b>, tag 30</p> <p><b>support-vlan</b>, tag 40</p>

This configuration example creates a loop-free topology between four switches using RSTP.

An RSTP topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

### Configuring RSTP on Switch 1

#### CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 1, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the *interface-mode* statement instead of the *port-mode* statement. The *port-mode* statement has been replaced with the *interface-mode* statement.

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
```

```

set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface xe-0/0/13.0 cost 1000
set protocols rstp interface xe-0/0/13.0 mode point-to-point
set protocols rstp interface xe-0/0/9.0 cost 1000
set protocols rstp interface xe-0/0/9.0 mode point-to-point
set protocols rstp interface xe-0/0/11.0 cost 1000
set protocols rstp interface xe-0/0/11.0 mode point-to-point

```

### Step-by-Step Procedure

To configure interfaces and RSTP on Switch 1:

1. Configure the VLANs **sales-vlan**, **engineering-vlan** and **publications-vlan**, and **support-vlan**:  
  

```

[edit vlans]
user@switch1# set sales-vlan description "Sales VLAN"
user@switch1# set sales-vlan vlan-id 10
user@switch1# set engineering-vlan description "Engineering VLAN"
user@switch1# set engineering-vlan vlan-id 20
user@switch1# set publications-vlan description "Publications VLAN"
user@switch1# set publications-vlan vlan-id 30

```
2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:  
  

```

[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]

```
3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the interface-mode statement instead of the port-mode statement. The port-mode statement has been replaced with the interface-mode statement.

- ```

[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching port-mode trunk

```
4. Configure RSTP on the switch:  
  

```

[edit protocols]
user@switch1# rstp bridge-priority 16k
user@switch1# rstp interface xe-0/0/13.0 cost 1000
user@switch1# rstp interface xe-0/0/13.0 mode point-to-point
user@switch1# rstp interface xe-0/0/9.0 cost 1000
user@switch1# rstp interface xe-0/0/9.0 mode point-to-point
user@switch1# rstp interface xe-0/0/11.0 cost 1000
user@switch1# rstp interface xe-0/0/11.0 mode point-to-point

```

**Results** Check the results of the configuration:

```
user@switch1> show configuration
```

```
interfaces {
  xe-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/9 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/11 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 16k;
    interface xe-0/0/13.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/9.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/11.0 {
      cost 1000;
      mode point-to-point;
    }
  }
}
vlands {
  sales-vlan {
    vlan-id 10;
  }
  engineering-vlan {
    vlan-id 20;
  }
}
```



```

    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}

```

### Configuring RSTP on Switch 2

#### CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 2, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 32k
set protocols rstp interface xe-0/0/14.0 cost 1000
set protocols rstp interface xe-0/0/14.0 mode point-to-point
set protocols rstp interface xe-0/0/18.0 cost 1000
set protocols rstp interface xe-0/0/18.0 mode point-to-point

```

#### Step-by-Step Procedure

To configure interfaces and RSTP on Switch 2:

1. Configure the VLANs `sales-vlan`, `engineering-vlan` and `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch2# set sales-vlan description "Sales VLAN"
user@switch2# set sales-vlan vlan-id 10
user@switch2# set engineering-vlan description "Engineering VLAN"
user@switch2# set engineering-vlan vlan-id 20
user@switch2# set publications-vlan description "Publications VLAN"
user@switch2# set publications-vlan vlan-id 30
user@switch2# set support-vlan vlan-description "Support VLAN"
user@switch2# set publications-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the interface-mode statement instead of the port-mode statement. The port-mode statement has been replaced with the interface-mode statement.

```
[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch2# rstp bridge-priority 32k
user@switch2# rstp interface xe-0/0/14.0 cost 1000
user@switch2# rstp interface xe-0/0/14.0 mode point-to-point
user@switch2# rstp interface xe-0/0/18.0 cost 1000
user@switch2# rstp interface xe-0/0/18.0 mode point-to-point
```

**Results** Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
  xe-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/18 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 32k;
    interface xe-0/0/14.0 {
      cost 1000;
    }
  }
}
```

```

        mode point-to-point;
    }
    interface xe-0/0/18.0 {
        cost 1000;
        mode point-to-point;
    }
}
}
vllans {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}
}

```

### Configuring RSTP on Switch 3

#### CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 3, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the interface-mode statement instead of the port-mode statement. The port-mode statement has been replaced with the interface-mode statement.

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 8k
set protocols rstp interface xe-0/0/26.0 cost 1000
set protocols rstp interface xe-0/0/26.0 mode point-to-point
set protocols rstp interface xe-0/0/28.0 cost 1000

```

```

set protocols rstp interface xe-0/0/28.0 mode point-to-point
set protocols rstp interface xe-0/0/24.0 cost 1000
set protocols rstp interface xe-0/0/24.0 mode point-to-point

```

### Step-by-Step Procedure

To configure interfaces and RSTP on Switch 3:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch3# set sales-vlan description "Sales VLAN"
user@switch3# set sales-vlan vlan-id 10
user@switch3# set engineering-vlan description "Engineering VLAN"
user@switch3# set engineering-vlan vlan-id 20
user@switch3# set publications-vlan description "Publications VLAN"
user@switch3# set publications-vlan vlan-id 30
user@switch3# set support-vlan description "Support VLAN"
user@switch3# set publications-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```

[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching port-mode trunk

```

4. Configure RSTP on the switch:

```

[edit protocols]
user@switch3# rstp bridge-priority 8k
user@switch3# rstp interface xe-0/0/26.0 cost 1000
user@switch3# rstp interface xe-0/0/26.0 mode point-to-point
user@switch3# rstp interface xe-0/0/28.0 cost 1000
user@switch3# rstp interface xe-0/0/28.0 mode point-to-point
user@switch3# rstp interface xe-0/0/24.0 cost 1000
user@switch3# rstp interface xe-0/0/24.0 mode point-to-point

```

**Results** Check the results of the configuration:

```

user@switch3> show configuration
interfaces {
  xe-0/0/26 {
    unit 0 {
      family ethernet-switching {

```

```

        port-mode trunk;
        vlan {
            members [10 20 30 40];
        }
    }
}
xe-0/0/28 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members [10 20 30 40];
            }
        }
    }
}
xe-0/0/24 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members [10 20 30 40];
            }
        }
    }
}
}
}
protocols {
    rstp {
        bridge-priority 8k;
        interface xe-0/0/26.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/28.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/24.0 {
            cost 1000;
            mode point-to-point;
        }
    }
    bridge-priority 8k;
}
}
vllans {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
}

```

```

}
publications-vlan {
  vlan-id 30;
}
support-vlan {
  vlan-id 40;
}
}

```

### Configuring RSTP on Switch 4

#### CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 4, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the interface-mode statement instead of the port-mode statement. The port-mode statement has been replaced with the interface-mode statement.

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface xe-0/0/23.0 cost 1000
set protocols rstp interface xe-0/0/23.0 mode point-to-point
set protocols rstp interface xe-0/0/19.0 cost 1000
set protocols rstp interface xe-0/0/19.0 mode point-to-point

```

#### Step-by-Step Procedure

To configure interfaces and RSTP on Switch 4:

1. Configure the VLANs **sales-vlan**, **engineering-vlan**, **publications-vlan**, and **support-vlan**:

```

[edit vlans]
user@switch4# set sales-vlan description "Sales VLAN"
user@switch4# set sales-vlan vlan-id 10
user@switch4# set engineering-vlan description "Engineering VLAN"
user@switch4# set engineering-vlan vlan-id 20
user@switch4# set publications-vlan description "Publications VLAN"
user@switch4# set publications-vlan vlan-id 30
user@switch4# set support-vlan description "Support VLAN"
user@switch4# set support-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]

```

```

user@switch4# set xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring RSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```

[edit interfaces]
user@switch4# set xe-0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set xe-0/0/19 unit 0 family ethernet-switching port-mode trunk

```

4. Configure RSTP on the switch:

```

[edit protocols]
user@switch4# rstp bridge-priority 16k
user@switch4# rstp interface all cost 1000
user@switch4# rstp interface xe-0/0/23.0 cost 1000
user@switch4# rstp interface xe-0/0/23.0 mode point-to-point
user@switch4# rstp interface xe-0/0/19.0 cost 1000
user@switch4# rstp interface xe-0/0/19.0 mode point-to-point

```

**Results** Check the results of the configuration:

```

user@switch4> show configuration
interfaces {
  xe-0/0/23 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
  xe-0/0/19 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members [10 20 30 40];
        }
      }
    }
  }
}
protocols {
  rstp {
    bridge-priority 16k;
    interface xe-0/0/23.0 {
      cost 1000;
    }
  }
}

```

```

        mode point-to-point;
    }
    interface xe-0/0/19.0 {
        cost 1000;
        mode point-to-point;
    }
}
}
vllans {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}
}

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying RSTP Configuration on Switch 1 on page 1622](#)
- [Verifying RSTP Configuration on Switch 2 on page 1623](#)
- [Verifying RSTP Configuration on Switch 3 on page 1623](#)
- [Verifying RSTP Configuration on Switch 4 on page 1623](#)

### Verifying RSTP Configuration on Switch 1

**Purpose** Verify that the RSTP configuration on Switch 1 is correct.

**Action** In operational mode, issue the **show spanning-tree interface** command:

```
user@switch1> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/13.0 | 128:527 | 128:525               | 16384.0019e25040e0      | 1000         | BLK   | ALT  |
| xe-0/0/9.0  | 128:529 | 128:513               | 32768.0019e2503d20      | 1000         | BLK   | ALT  |
| xe-0/0/11.0 | 128:531 | 128:513               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

**Meaning** See the topology in [Figure 23 on page 1611](#). The operational mode command **show spanning-tree interface** shows that **xe-0/0/13.0** is in a forwarding state. The other interfaces on Switch 1 are blocked.



**Verifying RSTP Configuration on Switch 2**

**Purpose** Verify that the RSTP configuration on Switch 2 is correct.

**Action** In operational mode issue the **show spanning-tree interface** command:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/14.0 | 128:513 | 128:513               | 32768.0019e2503d20      | 1000         | BLK   | DESC |
| xe-0/0/18.0 | 128:519 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

**Meaning** See the topology in [Figure 23 on page 1611](#). The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/18.0** is in a forwarding state and the root port. The other interface on Switch 2 is blocked.

**Verifying RSTP Configuration on Switch 3**

**Purpose** Verify that the RSTP configuration on Switch 3 is correct.

**Action** In operational mode, issue the **show spanning-tree interface** command:

```
user@switch3> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/26.0 | 128:513 | 128:513               | 8192.0019e25051e0       | 1000         | FWD   | DESC |
| xe-0/0/28.0 | 128:515 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | DESC |
| xe-0/0/24.0 | 128:517 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | DESC |

**Meaning** See the topology in [Figure 23 on page 1611](#). The operational mode command **show spanning-tree interface** shows that no interface is the root interface.

**Verifying RSTP Configuration on Switch 4**

**Purpose** Verify the RSTP configuration on Switch 4.

**Action** In operational mode, issue the **show spanning-tree interface** command:

```
user@switch4> show spanning-tree interface
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/23.0 | 128:523 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |
| xe-0/0/19.0 | 128:525 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | DESG |

**Meaning** See the topology in [Figure 23 on page 1611](#). The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/23.0** is the root interface and is in the forwarding state.

**Related Documentation**

- [Example: Configuring Network Regions for VLANs with MSTP on page 1624](#)
- [Understanding RSTP on page 1556](#)

## Example: Configuring Network Regions for VLANs with MSTP

Multiple Spanning Tree Protocol (MSTP) is used to create a loop-free topology in networks using multiple spanning-tree regions, each region containing multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates more efficient load sharing across redundant links.

You can create up to 64 MSTI instances for QFX Series products, and each MSTI supports up to 4094 VLANs.

This example describes how to configure MSTP on four QFX3500 switches:

- [Requirements on page 1624](#)
- [Overview and Topology on page 1625](#)
- [Configuring MSTP on Switch 1 on page 1627](#)
- [Configuring MSTP on Switch 2 on page 1630](#)
- [Configuring MSTP on Switch 3 on page 1633](#)
- [Configuring MSTP on Switch 4 on page 1636](#)
- [Verification on page 1639](#)

---

### Requirements

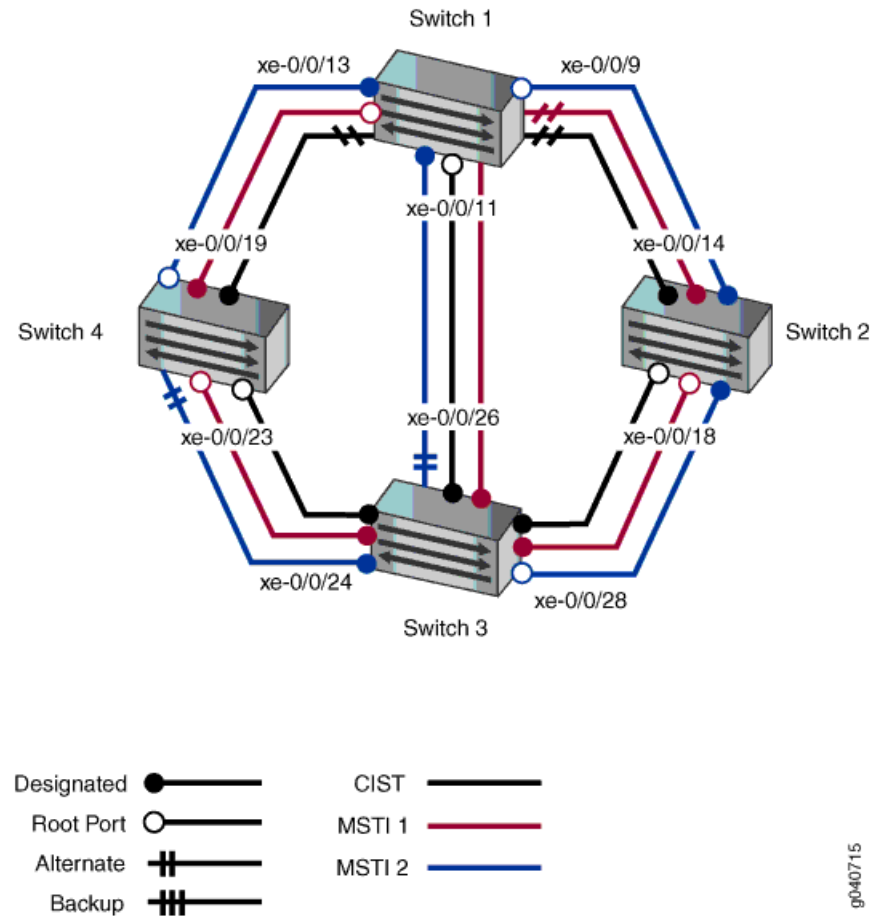
This example uses the following hardware and software components:

- Junos OS Release 11.1 for the QFX3500 switches
- Four QFX3500 switches

## Overview and Topology

When the number of VLANs grows in a network, MSTP provides a more faster way of creating a loop-free topology using MSTIs. Each MSTI in the spanning-tree domain maintains its own tree. Each tree can be mapped to different links, utilizing bandwidth that would be unavailable to a single tree. MSTIs reduce demand on system resources.

Figure 24: Network Topology for MSTP



The interfaces shown in [Table 104 on page 1626](#) will be configured for MSTP.



**NOTE:** You can configure MSTP on logical or physical interfaces. This example shows MSTP configured on logical interfaces.

Table 104: Topology for Configuring MSTP on the QFX Series

| Components             | Settings  |
|------------------------|---|
| Switch 1               | <p>The following ports on Switch 1 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>xe-0/0/9</b> is connected to Switch 2</li> <li>• <b>xe-0/0/13</b> is connected to Switch 4</li> <li>• <b>xe-0/0/11</b> is connected to Switch 3</li> </ul>  |
| Switch 2               | <p>The following ports on Switch 2 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>xe-0/0/14</b> is connected to Switch 1</li> <li>• <b>xe-0/0/18</b> is connected to Switch 3</li> </ul>  |
| Switch 3               | <p>The following ports on Switch 3 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>xe-0/0/26</b> is connected to Switch 1</li> <li>• <b>xe-0/0/28</b> is connected to Switch 2</li> <li>• <b>xe-0/0/24</b> is connected to Switch 4</li> </ul> |
| Switch 4               | <p>The following ports on Switch 4 are connected in this way:</p> <ul style="list-style-type: none"> <li>• <b>xe-0/0/19</b> is connected to Switch 1</li> <li>• <b>xe-0/0/23</b> is connected to Switch 3</li> </ul>  |
| VLAN names and tag IDs | <b>sales-vlan</b> , tag 10<br><b>engineering-vlan</b> , tag 20<br><b>publications-vlan</b> , tag 30<br><b>support-vlan</b> , tag 40   |
| MSTIs                  | 1<br>2  |

The topology in [Figure 24 on page 1625](#) shows a Common Internal Spanning Tree (CIST). The CIST is a single spanning tree connecting all devices in the network. The switch with the highest priority is elected as the root bridge of the CIST.

Also in an MSTP topology are ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.
- The *backup port* is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

In this example, one MSTP region, **region1**, contains Switch 1, Switch 2, Switch 3, and Switch 4. Within the region, four VLANs are created:

- The **sales-vlan** supports sales traffic and has a VLAN tag identifier of 10.
- The **engineering-vlan** supports data traffic and has a VLAN tag identifier of 20.
- The **publications-vlan** supports publications VLAN traffic (for supplicants that fail 802.1X authentication) and has a VLAN tag identifier of 30.
- The **support-vlan** supports video traffic and has a VLAN tag identifier of 40.

The VLANs are associated with specific interfaces on each of the four switches. Two MSTIs, 1 and 2, are then associated with the VLAN tag identifiers, and some MSTP parameters, such as cost, are configured on each switch.

### Configuring MSTP on Switch 1

#### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 1, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface xe-0/0/13.0 cost 1000
set protocols mstp interface xe-0/0/13.0 mode point-to-point
set protocols mstp interface xe-0/0/9.0 cost 1000
set protocols mstp interface xe-0/0/9.0 mode point-to-point
set protocols mstp interface xe-0/0/11.0 cost 1000
set protocols mstp interface xe-0/0/11.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 1 interface xe-0/0/11.0 cost 4000
set protocols mstp msti 2 bridge-priority 8k
set protocols mstp msti 2 vlan [30 40]
```

### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 1:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```
[edit vlans]
user@switch1# set sales-vlan description "Sales VLAN"
user@switch1# set sales-vlan vlan-id 10
user@switch1# set engineering-vlan description "Engineering VLAN"
user@switch1# set engineering-vlan vlan-id 20
user@switch1# set publications-vlan description "Publications VLAN"
user@switch1# set publications-vlan vlan-id 30
user@switch1# set support-vlan description "Support VLAN"
user@switch1# set publications-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```
[edit interfaces]
user@switch1# set xe-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set xe-0/0/11 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch1# mstp configuration-name region1
user@switch1# mstp bridge-priority 16k
user@switch1# mstp interface xe-0/0/13.0 cost 1000
user@switch1# mstp interface xe-0/0/13.0 mode point-to-point
user@switch1# mstp interface xe-0/0/9.0 cost 1000
user@switch1# mstp interface xe-0/0/9.0 mode point-to-point
user@switch1# mstp interface xe-0/0/11.0 cost 4000
user@switch1# mstp interface xe-0/0/11.0 mode point-to-point
user@switch1# mstp msti 1 bridge-priority 16k
user@switch1# mstp msti 1 vlan [10 20]
user@switch1# mstp msti 1 interface xe-0/0/11.0 cost 4000
user@switch1# mstp msti 2 bridge-priority 8k
user@switch1# mstp msti 2 vlan [30 40]
```

### Results

Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
  xe-0/0/13 {
```

```

    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
xe-0/0/9 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
xe-0/0/11 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
}
protocols {
    mstp {
        configuration-name region1;
        bridge-priority 16k;
        interface xe-0/0/13.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/9.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/11.0 {
            cost 4000;
            mode point-to-point;
        }
    }
    msti 1 {

```

```
    bridge-priority 16k;
    vlan [ 10 20 ];
    interface xe-0/0/11.0 {
        cost 4000;
    }
}
msti 2 {
    bridge-priority 8k;
    vlan [ 30 40 ];
}
}
vlangs {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}
```

---

### Configuring MSTP on Switch 2

#### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 2, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the interface-mode statement instead of the port-mode statement. The port-mode statement has been replaced with the interface-mode statement.

```
[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 32k
set protocols mstp interface xe-0/0/14.0 cost 1000
set protocols mstp interface xe-0/0/14.0 mode point-to-point
```



```

set protocols mstp interface xe-0/0/18.0 cost 1000
set protocols mstp interface xe-0/0/18.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 32k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 4k
set protocols mstp msti 2 vlan [30 40]

```

### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 2:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch2# set sales-vlan description "Sales VLAN"
user@switch2# set sales-vlan vlan-id 10
user@switch2# set engineering-vlan description "Engineering VLAN"
user@switch2# set engineering-vlan vlan-id 20
user@switch2# set publications-vlan description "Publications VLAN"
user@switch2# set publications-vlan vlan-id 30
user@switch2# set support-vlan vlan-description "Support VLAN"
user@switch2# set publications-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```

[edit interfaces]
user@switch2# set xe-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch2# set xe-0/0/18 unit 0 family ethernet-switching port-mode trunk

```

4. Configure MSTP on the switch, including the two MSTIs:

```

[edit protocols]
user@switch2# mstp configuration-name region1
user@switch2# mstp bridge-priority 32k
user@switch2# mstp interface xe-0/0/14.0 cost 1000
user@switch2# mstp interface xe-0/0/14.0 mode point-to-point
user@switch2# mstp interface xe-0/0/18.0 cost 1000
user@switch2# mstp interface xe-0/0/18.0 mode point-to-point
user@switch2# mstp interface all cost 1000
user@switch2# mstp msti 1 bridge-priority 32k
user@switch2# mstp msti 1 vlan [10 20]
user@switch2# mstp msti 2 bridge-priority 4k
user@switch2# mstp msti 2 vlan [30 40]

```

**Results** Check the results of the configuration:

```
user@switch2> show configuration
```

```
interfaces {
  xe-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  xe-0/0/18 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 32k;
    interface xe-0/0/14.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/18.0 {
      cost 1000;
      mode point-to-point;
    }
    msti 1 {
      bridge-priority 32k;
      vlan [ 10 20 ];
    }
    msti 2 {
      bridge-priority 4k;
      vlan [ 30 40 ];
    }
  }
}
vlangs {
  sales-vlan {
    vlan-id 10;
  }
  engineering-vlan {
    vlan-id 20;
  }
}
```

```

    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}

```

### Configuring MSTP on Switch 3

#### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 3, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 8k
set protocols mstp interface xe-0/0/26.0 cost 1000
set protocols mstp interface xe-0/0/26.0 mode point-to-point
set protocols mstp interface xe-0/0/28.0 cost 1000
set protocols mstp interface xe-0/0/28.0 mode point-to-point
set protocols mstp interface xe-0/0/24.0 cost 1000
set protocols mstp interface xe-0/0/24.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 4k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 16k
set protocols mstp msti 2 vlan [30 40]

```

### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 3:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```
[edit vlans]
user@switch3# set sales-vlan description "Sales VLAN"
user@switch3# set sales-vlan vlan-id 10
user@switch3# set engineering-vlan description "Engineering VLAN"
user@switch3# set engineering-vlan vlan-id 20
user@switch3# set publications-vlan description "Publications VLAN"
user@switch3# set publications-vlan vlan-id 30
user@switch3# set support-vlan description "Support VLAN"
user@switch3# set publications-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```
[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```
[edit interfaces]
user@switch3# set xe-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set xe-0/0/24 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch3# mstp configuration-name region1
user@switch3# mstp bridge-priority 8k
user@switch3# mstp interface xe-0/0/26.0 cost 1000
user@switch3# mstp interface xe-0/0/26.0 mode point-to-point
user@switch3# mstp interface xe-0/0/28.0 cost 1000
user@switch3# mstp interface xe-0/0/28.0 mode point-to-point
user@switch3# mstp interface xe-0/0/24.0 cost 1000
user@switch3# mstp interface xe-0/0/24.0 mode point-to-point
user@switch3# mstp interface all cost 1000
user@switch3# mstp msti 1 bridge-priority 4k
user@switch3# mstp msti 1 vlan [10 20]
user@switch3# mstp msti 2 bridge-priority 16k
user@switch3# mstp msti 2 vlan [30 40]
```

### Results

Check the results of the configuration:

```
user@switch3> show configuration
interfaces {
  xe-0/0/26 {
```

```

    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
xe-0/0/28 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
xe-0/0/24 {
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members 10;
                members 20;
                members 30;
                members 40;
            }
        }
    }
}
}
protocols {
    mstp {
        configuration-name region1;
        bridge-priority 8k;
        interface xe-0/0/26.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/28.0 {
            cost 1000;
            mode point-to-point;
        }
        interface xe-0/0/24.0 {
            cost 1000;
            mode point-to-point;
        }
    }
}

```

```

msti 1 {
    bridge-priority 4k;
    vlan [ 10 20 ];
}
msti 2 {
    bridge-priority 16k;
    vlan [ 30 40 ];
}
}
}
vlangs {
    sales-vlan {
        vlan-id 10;
    }
    engineering-vlan {
        vlan-id 20;
    }
    publications-vlan {
        vlan-id 30;
    }
    support-vlan {
        vlan-id 40;
    }
}
}

```

#### Configuring MSTP on Switch 4

##### CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 4, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the interface-mode statement instead of the port-mode statement. The port-mode statement has been replaced with the interface-mode statement.

```

[edit]
set vlans sales-vlan description "Sales VLAN"
set vlans sales-vlan vlan-id 10
set vlans engineering-vlan description "Engineering VLAN"
set vlans engineering-vlan vlan-id 20
set vlans publications-vlan description "Publications VLAN"
set vlans publications-vlan vlan-id 30
set vlans support-vlan description "Support VLAN"
set vlans support-vlan vlan-id 40
set interfaces xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces xe-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface xe-0/0/23.0 cost 1000
set protocols mstp interface xe-0/0/23.0 mode point-to-point
set protocols mstp interface xe-0/0/19.0 cost 1000

```

```

set protocols mstp interface xe-0/0/19.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 32k
set protocols mstp msti 2 vlan [30 40]

```

### Step-by-Step Procedure

To configure interfaces and MSTP on Switch 4:

1. Configure the VLANs `sales-vlan`, `engineering-vlan`, `publications-vlan`, and `support-vlan`:

```

[edit vlans]
user@switch4# set sales-vlan description "Sales VLAN"
user@switch4# set sales-vlan vlan-id 10
user@switch4# set engineering-vlan description "Engineering VLAN"
user@switch4# set engineering-vlan vlan-id 20
user@switch4# set publications-vlan description "Publications VLAN"
user@switch4# set publications-vlan vlan-id 30
user@switch4# set support-vlan description "Support VLAN"
user@switch4# set support-vlan vlan-id 40

```

2. Configure the VLANs on the interfaces, including support for the Ethernet switching protocol:

```

[edit interfaces]
user@switch4# set xe-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
user@switch4# set xe-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]

```

3. Configure the port mode for the interfaces:



**NOTE:** If you are configuring MSTP on devices that support the Enhanced Layer 2 Switching (ELS) CLI, use the `interface-mode` statement instead of the `port-mode` statement. The `port-mode` statement has been replaced with the `interface-mode` statement.

```

[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching port-mode trunk

```

4. Configure MSTP on the switch, including the two MSTIs:

```

[edit protocols]
user@switch4# mstp configuration-name region1
user@switch4# mstp bridge-priority 16k
user@switch4# mstp interface all cost 1000
user@switch4# mstp interface xe-0/0/23.0 cost 1000
user@switch4# mstp interface xe-0/0/23.0 mode point-to-point
user@switch4# mstp interface xe-0/0/19.0 cost 1000
user@switch4# mstp interface xe-0/0/19.0 mode point-to-point
user@switch4# mstp msti 1 bridge-priority 16k
user@switch4# mstp msti 1 vlan [10 20]
user@switch4# mstp msti 2 bridge-priority 32k
user@switch4# mstp msti 2 vlan [30 40]

```

**Results** Check the results of the configuration:

```

user@switch4> show configuration
interfaces {

```

```
xe-0/0/23 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members 10;
        members 20;
        members 30;
        members 40;
      }
    }
  }
}
xe-0/0/19 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members 10;
        members 20;
        members 30;
        members 40;
      }
    }
  }
}
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 16k;
    interface xe-0/0/23.0 {
      cost 1000;
      mode point-to-point;
    }
    interface xe-0/0/19.0 {
      cost 1000;
      mode point-to-point;
    }
  }
  msti 1 {
    bridge-priority 16k;
    vlan [ 10 20 ];
  }
  msti 2 {
    bridge-priority 32k;
    vlan [ 30 40 ];
  }
}
vlands {
  sales-vlan {
    vlan-id 10;
  }
  engineering-vlan {
    vlan-id 20;
  }
}
```



```
publications-vlan {  
    vlan-id 30;  
}  
support-vlan {  
    vlan-id 40;  
}  
}
```

---

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying MSTP Configuration on Switch 1 on page 1639](#)
- [Verifying MSTP Configuration on Switch 2 on page 1641](#)
- [Verifying MSTP Configuration on Switch 3 on page 1643](#)
- [Verifying MSTP Configuration on Switch 4 on page 1645](#)

#### *Verifying MSTP Configuration on Switch 1*

**Purpose** Verify the MSTP configuration on Switch 1.

**Action** Use the operational mode commands:

```
user@switch1> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/13.0 | 128:527 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | ROOT |
| xe-0/0/9.0  | 128:529 | 128:513               | 32768.0019e2503d20      | 1000         | BLK   | ALT  |
| xe-0/0/11.0 | 128:531 | 128:513               | 8192.0019e25051e0       | 4000         | BLK   | ALT  |

```
Spanning tree interface parameters for instance 1
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/13.0 | 128:527 | 128:525               | 16385.0019e25040e0      | 1000         | FWD   | ROOT |
| xe-0/0/9.0  | 128:529 | 128:513               | 32769.0019e2503d20      | 1000         | BLK   | ALT  |
| xe-0/0/11.0 | 128:531 | 128:513               | 4097.0019e25051e0       | 4000         | BLK   | ALT  |

```
Spanning tree interface parameters for instance 2
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/13.0 | 128:527 | 128:527               | 8194.0019e25044e0       | 1000         | FWD   | DESG |
| xe-0/0/9.0  | 128:529 | 128:513               | 4098.0019e2503d20       | 1000         | FWD   | ROOT |
| xe-0/0/11.0 | 128:531 | 128:531               | 8194.0019e25044e0       | 1000         | FWD   | DESG |

```
user@switch1> show spanning-tree bridge
```

```
STP bridge parameters
```

```
Context ID : 0
Enabled protocol : MSTP
```

```
STP bridge parameters for CIST
```

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : xe-0/0/13.0
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 2000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Message age : 0
Number of topology changes : 3
Time since last topology change : 921 seconds
Local parameters
  Bridge ID : 16384.00:19:e2:50:44:e0
  Extended system ID : 0
  Internal instance ID : 0
```

```
STP bridge parameters for MSTI 1
```

```
MSTI regional root : 4097.00:19:e2:50:51:e0
Root cost : 2000
Root port : xe-0/0/13.0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 18
Local parameters
  Bridge ID : 16385.00:19:e2:50:44:e0
```

```

Extended system ID          : 0
Internal instance ID        : 1

STP bridge parameters for MSTI 2
MSTI regional root          : 4098.00:19:e2:50:3d:20
Root cost                   : 1000
Root port                   : xe-0/0/9.0
Hello time                  : 2 seconds
Maximum age                 : 20 seconds
Forward delay               : 15 seconds
Hop count                   : 19
Local parameters
  Bridge ID                 : 8194.00:19:e2:50:44:e0
  Extended system ID        : 0
  Internal instance ID      : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

#### *Verifying MSTP Configuration on Switch 2*

**Purpose** Verify the MSTP configuration on Switch 2.

**Action** Use the operational mode commands:

```
user@switch2> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/14.0 | 128:513 | 128:513               | 32768.0019e2503d20      | 1000         | FWD   | DESC |
| xe-0/0/18.0 | 128:519 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |

Spanning tree interface parameters for instance 1

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/14.0 | 128:513 | 128:513               | 32769.0019e2503d20      | 1000         | FWD   | DESC |
| xe-0/0/18.0 | 128:519 | 128:515               | 4097.0019e25051e0       | 1000         | FWD   | ROOT |

Spanning tree interface parameters for instance 2

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/14.0 | 128:513 | 128:513               | 4098.0019e2503d20       | 1000         | FWD   | DESC |
| xe-0/0/18.0 | 128:519 | 128:519               | 4098.0019e2503d20       | 1000         | FWD   | DESC |

```
user@switch2> show spanning-tree bridge
```

STP bridge parameters

```
Context ID : 0
Enabled protocol : MSTP
```

STP bridge parameters for CIST

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : xe-0/0/18.0
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 1000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Message age : 0
Number of topology changes : 1
Time since last topology change : 782 seconds
Local parameters
  Bridge ID : 32768.00:19:e2:50:3d:20
  Extended system ID : 0
  Internal instance ID : 0
```

STP bridge parameters for MSTI 1

```
MSTI regional root : 4097.00:19:e2:50:51:e0
Root cost : 1000
Root port : xe-0/0/18.0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Local parameters
  Bridge ID : 32769.00:19:e2:50:3d:20
```

```

Extended system ID          : 0
Internal instance ID        : 1

STP bridge parameters for MSTI 2
MSTI regional root          : 4098.00:19:e2:50:3d:20
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Local parameters
  Bridge ID                  : 4098.00:19:e2:50:3d:20
  Extended system ID         : 0
  Internal instance ID       : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

#### ***Verifying MSTP Configuration on Switch 3***

**Purpose** Verify the MSTP configuration on Switch 3.

**Action** Use the operational mode commands:

```
user@switch3> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/26.0 | 128:513 | 128:513               | 8192.0019e25051e0       | 1000         | FWD   | DESC |
| xe-0/0/28.0 | 128:515 | 128:515               | 8192.0019e25051e0       | 1000         | FWD   | DESC |
| xe-0/0/24.0 | 128:517 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | DESC |

Spanning tree interface parameters for instance 1

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/26.0 | 128:513 | 128:513               | 4097.0019e25051e0       | 1000         | FWD   | DESC |
| xe-0/0/28.0 | 128:515 | 128:515               | 4097.0019e25051e0       | 1000         | FWD   | DESC |
| xe-0/0/24.0 | 128:517 | 128:517               | 4097.0019e25051e0       | 1000         | FWD   | DESC |

Spanning tree interface parameters for instance 2

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/26.0 | 128:513 | 128:531               | 8194.0019e25044e0       | 1000         | BLK   | ALT  |
| xe-0/0/28.0 | 128:515 | 128:519               | 4098.0019e2503d20       | 1000         | FWD   | ROOT |
| xe-0/0/24.0 | 128:517 | 128:517               | 16386.0019e25051e0      | 1000         | FWD   | DESC |

```
user@switch3> show spanning-tree bridge
```

STP bridge parameters

```
Context ID           : 0
Enabled protocol     : MSTP
```

STP bridge parameters for CIST

```
Root ID              : 8192.00:19:e2:50:51:e0
CIST regional root   : 8192.00:19:e2:50:51:e0
CIST internal root cost : 0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay        : 15 seconds
Number of topology changes : 3
Time since last topology change : 843 seconds
Local parameters
  Bridge ID          : 8192.00:19:e2:50:51:e0
  Extended system ID : 0
  Internal instance ID : 0
```

STP bridge parameters for MSTI 1

```
MSTI regional root   : 4097.00:19:e2:50:51:e0
Hello time           : 2 seconds
Maximum age          : 20 seconds
Forward delay        : 15 seconds
Local parameters
  Bridge ID          : 4097.00:19:e2:50:51:e0
  Extended system ID : 0
  Internal instance ID : 1
```

STP bridge parameters for MSTI 2

```
MSTI regional root   : 4098.00:19:e2:50:3d:20
```

```
Root cost           : 1000
Root port           : xe-0/0/28.0
Hello time          : 2 seconds
Maximum age         : 20 seconds
Forward delay       : 15 seconds
Hop count           : 19
Local parameters
  Bridge ID         : 16386.00:19:e2:50:51:e0
  Extended system ID : 0
  Internal instance ID : 2
```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

#### *Verifying MSTP Configuration on Switch 4*

**Purpose** Verify the MSTP configuration on Switch 4.

**Action** Use the operational mode commands:

```
user@switch4> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/23.0 | 128:523 | 128:517               | 8192.0019e25051e0       | 1000         | FWD   | ROOT |
| xe-0/0/19.0 | 128:525 | 128:525               | 16384.0019e25040e0      | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 1
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/23.0 | 128:523 | 128:517               | 4097.0019e25051e0       | 1000         | FWD   | ROOT |
| xe-0/0/19.0 | 128:525 | 128:525               | 16385.0019e25040e0      | 1000         | FWD   | DESG |

```
Spanning tree interface parameters for instance 2
```

| Interface   | Port ID | Designated<br>port ID | Designated<br>bridge ID | Port<br>Cost | State | Role |
|-------------|---------|-----------------------|-------------------------|--------------|-------|------|
| xe-0/0/23.0 | 128:523 | 128:517               | 16386.0019e25051e0      | 1000         | BLK   | ALT  |
| xe-0/0/19.0 | 128:525 | 128:527               | 8194.0019e25044e0       | 1000         | FWD   | ROOT |

```
user@switch4> show spanning-tree bridge
```

```
STP bridge parameters
```

```
Context ID : 0
Enabled protocol : MSTP
```

```
STP bridge parameters for CIST
```

```
Root ID : 8192.00:19:e2:50:51:e0
Root cost : 0
Root port : xe-0/0/23.0
CIST regional root : 8192.00:19:e2:50:51:e0
CIST internal root cost : 1000
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Message age : 0
Number of topology changes : 4
Time since last topology change : 887 seconds
Local parameters
  Bridge ID : 16384.00:19:e2:50:40:e0
  Extended system ID : 0
  Internal instance ID : 0
```

```
STP bridge parameters for MSTI 1
```

```
MSTI regional root : 4097.00:19:e2:50:51:e0
Root cost : 1000
Root port : xe-0/0/23.0
Hello time : 2 seconds
Maximum age : 20 seconds
Forward delay : 15 seconds
Hop count : 19
Local parameters
  Bridge ID : 16385.00:19:e2:50:40:e0
  Extended system ID : 0
```



```

Internal instance ID          : 1

STP bridge parameters for MSTI 2
MSTI regional root           : 4098.00:19:e2:50:3d:20
Root cost                     : 2000
Root port                    : xe-0/0/19.0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Hop count                     : 18
Local parameters
  Bridge ID                   : 32770.00:19:e2:50:40:e0
  Extended system ID         : 0
  Internal instance ID       : 2

```

**Meaning** The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or the interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

- Related Documentation**
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
  - [Understanding MSTP on page 1555](#)

## Example: Connecting an Access Switch to a Distribution Switch

In large local area networks (LANs), you commonly need to aggregate traffic from a number of access switches into a distribution switch.

This example describes how to connect an access switch to a distribution switch:

- [Requirements on page 1647](#)
- [Overview and Topology on page 1648](#)
- [Configuring the Access Switch on page 1648](#)
- [Configuring the Distribution Switch on page 1652](#)
- [Verification on page 1654](#)

### Requirements

This example uses the following hardware and software components:

- For the distribution switch, one EX 4200-24F switch. This model is designed to be used as a distribution switch for aggregation or collapsed core network topologies and in space-constrained data centers. It has twenty-four 1-Gigabit Ethernet fiber SFP ports and an EX-UM-2XFP uplink module with two 10-Gigabit Ethernet XFP ports.
- For the access switch, one EX 3200-24P, which has twenty-four 1-Gigabit Ethernet ports, all of which support Power over Ethernet (PoE), and an uplink module with four 1-Gigabit Ethernet ports.

- Junos OS Release 11.1 or later for the QFX Series

### Overview and Topology

In a large office that is spread across several floors or buildings, or in a data center, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect a single access switch to a distribution switch.

In the topology, the LAN is segmented into two VLANs, one for the sales department and the second for the support team. One 1-Gigabit Ethernet port on the access switch's uplink module connects to the distribution switch, to one 1-Gigabit Ethernet port on the distribution switch.

[Table 105 on page 1648](#) explains the components of the example topology. The example shows how to configure one of the three access switches. The other access switches could be configured in the same manner.

**Table 105: Components of the Topology for Connecting an Access Switch to a Distribution Switch**

| Property   | Settings  |
|--|---|
| Access switch hardware   | EX 3200-24P, 24 1-Gigabit Ethernet ports, all PoE-enabled ( <b>ge-0/0/0</b> through <b>ge-0/0/23</b> ); one 4-port 1-Gigabit Ethernet uplink module (EX-UM-4SFP)  |
| Distribution switch hardware                                     | EX 4200-24F, 24 1-Gigabit Ethernet fiber SPF ports ( <b>ge-0/0/0</b> through <b>ge-0/0/23</b> ); one 2-port 10-Gigabit Ethernet XFP uplink module (EX-UM-4SFP)  |
| VLAN names and tag IDs   | <b>sales</b> , tag 100 <b>support</b> , tag 200   |
| VLAN subnets   | <b>sales</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>support</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)   |
| Trunk port interfaces  | On the access switch: <b>ge-0/1/0</b> On the distribution switch: <b>ge-0/0/0</b>   |
| Access port interfaces in VLAN <b>sales</b> (on access switch)   | Avaya IP telephones: <b>ge-0/0/3</b> through <b>ge-0/0/19</b> Wireless access points: <b>ge-0/0/0</b> and <b>ge-0/0/1</b> Printers: <b>ge-0/0/22</b> and <b>ge-0/0/23</b> File servers: <b>ge-0/0/20</b> and <b>ge-0/0/21</b> |
| Access port interfaces in VLAN <b>support</b> (on access switch) | Avaya IP telephones: <b>ge-0/0/25</b> through <b>ge-0/0/43</b> Wireless access points: <b>ge-0/0/24</b> Printers: <b>ge-0/0/44</b> and <b>ge-0/0/45</b> File servers: <b>ge-0/0/46</b> and <b>ge-0/0/47</b>                   |
| Unused interfaces on access switch                               | <b>ge-0/0/2</b> and <b>ge-0/0/25</b>  |

### Configuring the Access Switch

To configure the access switch:

#### CLI Quick Configuration

To quickly configure the access switch, copy the following commands and paste them into the switch terminal window:

[edit]

```

set interfaces ge-0/0/0 unit 0 description "Sales Wireless access point port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description "Sales phone port"
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description "Sales printer port"
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description "Sales file server port"
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/24 unit 0 description "Support wireless access point port"
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/26 unit 0 description "Support phone port"
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description "Support printer port"
set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description "Support file server port"
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/1/0 unit 0 description "Uplink module port connection to distribution switch"
set interfaces ge-0/1/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/1/0 unit 0 family ethernet-switching native-vlan-id 1
set interfaces ge-0/1/0 unit 0 family ethernet-switching vlan members [sales support]
set interfaces vlan unit 0 family inet address 192.0.2.1/25
set interfaces vlan unit 1 family inet address 192.0.2.129/25
set vlans sales interface ge-0/0/0.0
set vlans sales interface ge-0/0/3.0
set vlans sales interface ge-0/0/22.0
set vlans sales interface ge-0/0/20.0
set vlans sales l3-interface vlan.0
set vlans sales vlan-id 100
set vlans sales vlan-description "Sales VLAN"
set vlans support interface ge-0/0/24.0
set vlans support interface ge-0/0/26.0
set vlans support interface ge-0/0/44.0
set vlans support interface ge-0/0/46.0
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
set vlans support vlan-description "Support VLAN"

```

#### Step-by-Step Procedure

To configure the access switch:

1. Configure the 1-Gigabit Ethernet interface on the uplink module to be the trunk port that connects to the distribution switch:  

```
[edit interfaces ge-0/1/0 unit 0]user@access-switch# set description "Uplink module port connection to distribution switch"
user@access-switch# set ethernet-switching port-mode trunk
```
2. Specify the VLANs to be aggregated on the trunk port:  

```
[edit interfaces ge-0/1/0 unit 0]user@access-switch# set ethernet-switching vlanmembers [ sales support ]
```
3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):  

```
[edit interfaces ge-0/1/0 unit 0]user@access-switch# set ethernet-switching native-vlan-id 1
```
4. Configure the sales VLAN:  

```
[edit vlans sales]user@access-switch# set vlan-description "Sales VLAN"
user@access-switch# set vlan-id (VLANs) 100
user@access-switch# set l3-interface (VLAN) vlan.0
```
5. Configure the support VLAN:

- ```
[edit vlans support]user@access-switch# set vlan-description "Support
VLAN"user@access-switch# set vlan-id (VLANs) 200user@access-switch# set
l3-interface (VLAN) vlan.1
```
6. Create the subnet for the sales broadcast domain:
 

```
[edit interfaces]user@access-switch# set vlan unit 0 family inet address 192.0.2.1/25
```
  7. Create the subnet for the support broadcast domain:
 

```
[edit interfaces]user@access-switch# set vlan unit 1 family inet address 192.0.2.129/25
```
  8. Configure the interfaces in the sales VLAN:
 

```
[edit interfaces]user@access-switch# set ge-0/0/0 unit 0 description "Sales wireless
access point port"user@access-switch# set ge-0/0/0 unit 0 family ethernet-switching
vlan members salesuser@access-switch# set ge-0/0/3 unit 0 description "Sales phone
port"user@access-switch# set ge-0/0/3 unit 0 family ethernet-switching vlan members
salesuser@access-switch# set ge-0/0/20 unit 0 description "Sales file server
port"user@access-switch# set ge-0/0/20 unit 0 family ethernet-switching vlan members
salesuser@access-switch# set ge-0/0/22 unit 0 description "Sales printer
port"user@access-switch# set ge-0/0/22 unit 0 family ethernet-switching vlan members
sales
```
  9. Configure the interfaces in the support VLAN:
 

```
[edit interfaces]user@access-switch# set ge-0/0/24 unit 0 description "Support
wireless access point port"user@access-switch# set ge-0/0/24 unit 0 family
ethernet-switching vlan members supportuser@access-switch# set ge-0/0/26 unit 0
description "Support phone port"user@access-switch# set ge-0/0/26 unit 0 family
ethernet-switching vlan members supportuser@access-switch# set ge-0/0/44 unit 0
description "Support printer port"user@access-switch# set ge-0/0/44 unit 0 family
ethernet-switching vlan members supportuser@access-switch# set ge-0/0/46 unit 0
description "Support file server port"user@access-switch# set ge-0/0/46 unit 0 family
ethernet-switching vlan members support
```
  10. Configure descriptions and VLAN tag IDs for the sales and support VLANs:
 

```
[edit vlans]user@access-switch# set sales vlan-description "Sales
VLAN"user@access-switch# set sales vlan-id 100user@access-switch# set support
vlan-description "Support VLAN"user@access-switch# set support vlan-id 200
```
  11. To route traffic between the sales and support VLANs and associate a Layer 3 interface with each VLAN:
 

```
[edit vlans]user@access-switch# set sales l3-interface vlan.0user@access-switch#
set support l3-interface vlan.1
```

**Results** Display the results of the configuration:

```
user@access-switch> show
interfaces {
  ge-0/0/0 {
    unit 0 {
      description "Sales wireless access point port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      description "Sales phone port";
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
}
```

```

    }
  }
}
ge-0/0/20 {
  unit 0 {
    description "Sales file server port";
    family ethernet-switching {
      vlan members sales;
    }
  }
}
ge-0/0/22 {
  unit 0 {
    description "Sales printer port";
    family ethernet-switching {
      vlan members sales;
    }
  }
}
ge-0/0/24 {
  unit 0 {
    description "Support wireless access point port";
    family ethernet-switching {
      vlan members support;
    }
  }
}
ge-0/0/26 {
  unit 0 {
    description "Support phone port";
    family ethernet-switching {
      vlan members support;
    }
  }
}
ge-0/0/44 {
  unit 0 {
    description "Support printer port";
    family ethernet-switching {
      vlan members sales;
    }
  }
}
ge-0/0/46 {
  unit 0 {
    description "Support file server port";
    family ethernet-switching {
      vlan members support;
    }
  }
}
ge-0/1/0 {
  unit 0 {
    description "Uplink module port connection to distribution switch";
    family ethernet-switching {
      port-mode trunk;
    }
  }
}

```

```
        vlan members [ sales support ];
        native-vlan-id 1;
    }
}
vlan {
    unit 0 {
        family inet address 192.0.2.1/25;
    }
    unit 1 {
        family inet address 192.0.2.129/25;
    }
}
vlands {
    sales {
        vlan-id 100;
        vlan-description "Sales VLAN";
        l3-interface vlan.0;
    }
    support {
        vlan-id 200;
        vlan-description "Support VLAN";
        l3-interface vlan.1;
    }
}
```



**TIP:** To quickly configure the distribution switch, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

---

### Configuring the Distribution Switch

To configure the distribution switch:

#### CLI Quick Configuration

To quickly configure the distribution switch, copy the following commands and paste them into the switch terminal window:

```
set interfaces ge-0/0/0 description "Connection to access switch"
set interfaces ge-0/0/0 ethernet-switching port-mode trunk
set interfaces ge-0/0/0 ethernet-switching vlan members [ sales support ]
set interfaces vlan unit 0 family inet address 192.0.2.2/25
set interfaces vlan unit 1 family inet address 192.0.2.130/25
set vlans sales vlan-description "Sales VLAN"
set vlans sales vlan-id 100
set vlans sales l3-interface vlan.0
set vlans support vlan-description "Support VLAN"
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
```

**Step-by-Step  
Procedure**

To configure the distribution switch:

1. Configure the interface on the switch to be the trunk port that connects to the access switch:  
  

```
[edit interfaces ge-0/0/0 unit 0]user@distribution-switch# set description "Connection to access switch"
user@distribution-switch# set ethernet-switching port-mode trunk
```
2. Specify the VLANs to be aggregated on the trunk port:  
  

```
[edit interfaces ge-0/0/0 unit 0]user@distribution-switch# set ethernet-switching vlan-members [ sales support ]
```
3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):  
  

```
[edit interfaces]user@distribution-switch# set ge-0/0/0 ethernet-switching native-vlan-id 1
```
4. Configure the sales VLAN:  
  

```
[edit vlans sales]user@distribution-switch# set vlan-description "Sales VLAN"
user@distribution-switch# set vlan-id (VLANs) 100
user@distribution-switch# set l3-interface (VLAN) vlan.0
```
5. Configure the support VLAN:  
  

```
[edit vlans support]user@distribution-switch# set vlan-description "Support VLAN"
user@distribution-switch# set vlan-id (VLANs) 200
user@distribution-switch# set l3-interface (VLAN) vlan.1
```
6. Create the subnet for the sales broadcast domain:  
  

```
[edit interfaces]user@distribution-switch# set vlan unit 0 family inet address 192.0.2.2/25
```
7. Create the subnet for the support broadcast domain:  
  

```
[edit interfaces] user@distribution-switch# set vlan unit 1 family inet address 192.0.2.130/25
```

**Results** Display the results of the configuration:

```
user@distribution-switch> show
interfaces {
  ge-0/0/0 {
    description "Connection to access switch";
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan members [ sales support ];
        native-vlan-id 1;
      }
    }
  }
}
vlan {
  unit 0 {
    family inet address 192.0.2.2/25;
  }
  unit 1 {
    family inet address 192.0.2.130/25;
  }
}
}
```

```
vlan {  
  sales {  
    vlan-id 100;  
    vlan-description "Sales VLAN";  
    l3-interface vlan.0;  
  }  
  support {  
    vlan-id 200;  
    vlan-description "Support VLAN";  
    l3-interface vlan.1;  
  }  
}
```



**TIP:** To quickly configure the distribution switch, issue the **load merge terminal** command, then copy the hierarchy and paste it into the switch terminal window.

---

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the VLAN Members and Interfaces on the Access Switch on page 1654](#)
- [Verifying the VLAN Members and Interfaces on the Distribution Switch on page 1655](#)

### *Verifying the VLAN Members and Interfaces on the Access Switch*

**Purpose** Verify that the **sales** and **support** have been created on the switch.



**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/1.0, ge-0/0/2.0, ge-0/0/4.0, ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0*, ge-0/0/9.0,  ge-0/0/10.0, ge-0/0/11.0*, ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0, ge-0/0/19.0*, ge-0/0/21.0, ge-0/0/23.0, ge-0/0/25.0, ge-0/0/27.0*, ge-0/0/28.0, ge-0/0/29.0, ge-0/0/30.0, ge-0/0/31.0*, ge-0/0/32.0, ge-0/0/33.0, ge-0/0/34.0, ge-0/0/35.0*, ge-0/0/36.0, ge-0/0/37.0, ge-0/0/38.0, ge-0/0/39.0*, ge-0/0/40.0, ge-0/0/41.0, ge-0/0/42.0, ge-0/0/43.0*, ge-0/0/45.0, ge-0/0/47.0, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*
sales	100	ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/20.0, ge-0/0/22.0, ge-0/1/0.0*,
support	200	ge-0/0/24.0*, ge-0/0/26.0, ge-0/0/44.0, ge-0/0/46.0,
mgmt		me0.0*

**Meaning** The output shows the **sales** and **support** VLANs and the interfaces associated with them.

### *Verifying the VLAN Members and Interfaces on the Distribution Switch*

**Purpose** Verify that the **sales** and **support** have been created on the switch.

**Action** List all VLANs configured on the switch:

```
user@switch> show vlans
```

Name	Tag	Interfaces
default		ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0, ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0*, ge-0/0/8.0,  ge-0/0/9.0, ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0*, ge-0/0/19.0, ge-0/0/20.0, ge-0/0/21.0, ge-0/0/22.0*, ge-0/0/23.0, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*
sales	100	ge-0/0/0.0*
support	200	ge-0/0/0.0*
mgmt		me0.0*

**Meaning** The output shows the **sales** and **support** VLANs associated to interface **ge-0/0/0.0**. Interface **ge-0/0/0.0** is the trunk interface connected to the access switch.

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding Bridging and VLANs](#)

## Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations

The QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Configure BPDU protection on interfaces to prevent them from receiving BPDUs that could result in STP misconfigurations, which could lead to network outages.

This example describes how to configure BPDU protection on access interfaces in QFX Series products in an RSTP topology:

- [Requirements on page 1656](#)
- [Overview and Topology on page 1657](#)
- [Configuration on page 1658](#)
- [Verification on page 1658](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series

- Two edged-linked switches in an RSTP topology



**NOTE:** By default, RSTP is enabled on the QFX Series.

### Overview and Topology


A loop-free network is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). However, receipt of BPDUs on certain interfaces in an STP, RSTP, or MSTP topology. It can lead to network outages by triggering an STP misconfiguration. To prevent such outages, enable BPDU protection on those interfaces that should not receive BPDUs.

Enable BPDU protection on switch interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If a BPDU is received on a BPDU-protected interface, the interface is disabled and stops forwarding frames.

Two switches are displayed in [Figure 25 on page 1657](#). In this example, Switch 1 and Switch 2 are configured for RSTP and create a loop-free topology. The interfaces on Switch 2 are access ports.

This example shows you how to configure interface **xe-0/0/5** and interface **xe-0/0/6** as edge ports and how to configure BPDU protection. When BPDU protection is enabled, the interfaces transition to a blocking state when they receive BPDUs.

**Figure 25: BPDU Protection Topology**



g020153

[Table 106 on page 1658](#) shows the components that will be configured for BPDU protection.

**Table 106: Components of the Topology for Configuring BPDU Protection on the QFX Series**

Component	Settings
Switch 1 (Distribution Layer)	Switch 1 is connected to Switch 2 on a trunk interface.
Switch 2 (Access Layer)	Switch 2 has these access ports that require BPDU protection: <ul style="list-style-type: none"> <li>• <code>xe-0/0/5</code></li> <li>• <code>xe-0/0/6</code></li> </ul>

This configuration example uses an RSTP topology. You also can configure BPDU protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.

### Configuration

#### CLI Quick Configuration

To quickly configure BPDU protection on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp interface xe-0/0/5 edge
set protocols rstp interface xe-0/0/6 edge
set protocols rstp bpdu-block-on-edge
```

#### Step-by-Step Procedure

To configure BPDU protection:

1. Configure interface `xe-0/0/5` and interface `xe-0/0/6` on Switch 2 as edge ports:

```
[edit protocols rstp]
user@switch# set interface xe-0/0/5 edge
user@switch# set interface xe-0/0/6 edge
```

2. Configure BPDU protection on all edge ports:

```
[edit protocols rstp]
user@switch# set bpdu-block-on-edge
```

#### Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface xe-0/0/5.0 {
  edge;
}
interface xe-0/0/6.0 {
  edge;
}
bpdu-block-on-edge;
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before BPDU Protection Is Triggered on page 1659](#)
- [Verifying That BPDU Protection Is Working Correctly on page 1659](#)

*Displaying the Interface State Before BPDU Protection Is Triggered*

**Purpose** Before BPDUs are being received from the devices connected to interface **xe-0/0/5** and interface **xe-0/0/6**, confirm the interface state.

**Action** You can verify the interface state using the **show spanning-tree interface** command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/0.0	128:513	128:513	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/1.0	128:514	128:514	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/2.0	128:515	128:515	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/3.0	128:516	128:516	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/4.0	128:517	128:517	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/5.0	128:518	128:518	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/6.0	128:519	128:519	32768.0019e2503f00	20000	FWD	DESG

[output truncated]

**Meaning** The output shows that interface **xe-0/0/5.0** and interface **xe-0/0/6.0** are designated ports in a forwarding state.

*Verifying That BPDU Protection Is Working Correctly*

**Purpose** In this example, the devices connected to Switch 2 start sending BPDUs to interface **xe-0/0/5.0** and interface **xe-0/0/6.0**. Verify that BPDU protection is configured on the interfaces.

**Action** You can verify that BPDU protection is configured on the interfaces by using the **show spanning-tree interface** command:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/0.0	128:513	128:513	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/1.0	128:514	128:514	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/2.0	128:515	128:515	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/3.0	128:516	128:516	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/4.0	128:517	128:517	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/5.0	128:518	128:518	32768.0019e2503f00	20000	BLK	DIS
(Bpdu-Incon)						
xe-0/0/6.0	128:519	128:519	32768.0019e2503f00	20000	BLK	DIS
(Bpdu-Incon)						
xe-0/0/7.0	128:520	128:1	16384.00aabbcc0348	20000	FWD	ROOT
xe-0/0/8.0	128:521	128:521	32768.0019e2503f00	20000	FWD	DESG

[output truncated]

**Meaning** When BPDUs are sent from the devices to interface **xe-0/0/5.0** and interface **xe-0/0/6.0** on Switch 2, the output from the operational mode command **show spanning-tree interface** shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state blocks the interfaces and prevents them from forwarding traffic.

Disabling the BPDU protection configuration on an interface does not unblock the interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. Otherwise, use the operational mode command **clear ethernet-switching bpd-error** to unblock the interface.

If the devices connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state. In such cases, you need to find and repair the misconfiguration on the devices that is triggering the sending of BPDUs to Switch 2.

- Related Documentation**
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
  - [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660](#)
  - [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1664](#)
  - [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558](#)

### Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree

The QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol

(MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing interfaces from moving into a forwarding state that would create a loop in the network.

This example describes how to configure loop protection for an interface for the QFX Series in an RSTP topology:

- [Requirements on page 1661](#)
- [Overview and Topology on page 1661](#)
- [Configuration on page 1663](#)
- [Verification on page 1663](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series
- Three switches in an RSTP topology



**NOTE:** By default, RSTP is enabled for the QFX Series.

### Overview and Topology

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called a bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

A blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor. When this happens, a loop appears in the spanning tree. Loops in a Layer 2 topology cause broadcast, unicast, and multicast frames to continuously circle the looped network. As a switch processes a flood of frames in a looped network, its resources become depleted, and the ultimate result is a network outage.



**NOTE:** An interface can be configured for either loop protection or root protection, but not for both.

Three switches are displayed in [Figure 26 on page 1662](#). In this example, they are configured for RSTP and create a loop-free topology. Interface **xe-0/0/6** is blocking traffic between Switch 3 and Switch 1; thus, traffic is forwarded through interface **xe-0/0/7** on Switch 2. BPDUs are being sent from the root bridge on Switch 1 to both of these interfaces.

This example shows how to configure loop protection on interface **xe-0/0/6** to prevent it from transitioning from a blocking state to a forwarding state and creating a loop in the spanning-tree topology.

Figure 26: Network Topology for Loop Protection




Table 107 on page 1662 shows the components that will be configured for loop protection.

Table 107: Topology for Configuring Loop Protection on the QFX Series

Components	Settings
Switch 1	Switch 1 is the root bridge.
Switch 2	Switch 2 has the root port <b>xe-0/0/7</b> .
Switch 3	Switch 3 is connected to Switch 1 through interface <b>xe-0/0/6</b> .

A spanning-tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you can also configure loop protection for STP or MSTP topologies at the `[edit protocols (mstp | stp)]` hierarchy level.



## Configuration

- CLI Quick Configuration** To quickly configure loop protection on interface **xe-0/0/6**:
- ```
[edit]
set protocols rstp interface xe-0/0/6 bpdu-timeout-action block
```
- Step-by-Step Procedure** To configure loop protection:
1. Configure interface **xe-0/0/6** on Switch 3:
- ```
[edit protocols rstp]
user@switch# set interface xe-0/0/6 bpdu-timeout-action block
```
- Results** Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface xe-0/0/6.0 {
  bpdu-timeout-action {
    block;
  }
}
```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before Loop Protection Is Triggered on page 1663](#)
- [Verifying That Loop Protection Is Working on an Interface on page 1664](#)

### *Displaying the Interface State Before Loop Protection Is Triggered*

- Purpose** Before loop protection is triggered on interface **xe-0/0/6**, confirm that the interface is blocked.

- Action** Display the interface state and role before applying root protection:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/0.0	128:513	128:513	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/1.0	128:514	128:514	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/2.0	128:515	128:515	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/3.0	128:516	128:516	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/4.0	128:517	128:517	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/5.0	128:518	128:518	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/6.0	128:519	128:2	16384.00aabbcc0348	20000	BLK	ALT

[output truncated]

- Meaning** The output from the operational mode command **show spanning-tree interface** shows that **xe-0/0/6.0** is the alternate port and is blocked.

**Verifying That Loop Protection Is Working on an Interface**

**Purpose** Verify that the loop protection configuration on interface **xe-0/0/6**. RSTP has been disabled on interface **xe-0/0/4** on Switch 1. This stops BPDUs from being sent to interface **xe-0/0/6** and triggering loop protection on that interface.

**Action** Display the interface state and role after applying root protection:

```
user@switch> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/0.0	128:513	128:513	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/1.0	128:514	128:514	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/2.0	128:515	128:515	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/3.0	128:516	128:516	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/4.0	128:517	128:517	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/5.0	128:518	128:518	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/6.0	128:519	128:519	32768.0019e2503f00	20000	BLK	DIS

(Loop-Incon)  
[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/6.0** has detected that BPDUs are no longer being forwarded to it and has moved into a loop-inconsistent state. The loop-inconsistent state prevents the interface from transitioning to a forwarding state. The interface recovers and transitions back to its original state as soon as it receives BPDUs.

- Related Documentation**
- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
  - [Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1664](#)
  - [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656](#)
  - [Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1559](#)

**Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees**

QFX Series products provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Root protection increases the efficiency of STP, RSTP, and MSTP by allowing network administrators to enforce the root bridge placement in the network manually.

This example describes how to configure root protection on an interface for the QFX Series.

- [Requirements on page 1665](#)
- [Overview and Topology on page 1665](#)
- [Configuration on page 1667](#)
- [Verification on page 1667](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later for the QFX Series
- Four switches in an RSTP topology

Before you configure the interface for root protection, be sure you have:

- RSTP operating on the switches.



**NOTE:** By default, RSTP is enabled on the QFX Series.

## Overview and Topology

Peer STP applications running on switch interfaces exchange a special type of frame called a bridge protocol data unit (BPDU). Switches communicate interface information using BPDUs to create a loop-free topology that ultimately determines the root bridge and which interfaces block or forward traffic in the spanning tree.

You can also see BPDUs generated when you run a bridge application on a device attached to the switch. This can interfere with root port election, which may sometimes lead to the wrong root port being elected through the above process. Root protection allows you to manually enforce the root bridge placement in the network.

To prevent this from happening, enable root protection on interfaces that should not receive more BPDUs from the root bridge and should not be elected as the root port. These interfaces are typically located on an administrative boundary and are designated ports.

When root protection is enabled on an interface:

- The interface is blocked from becoming the root port.
- Root protection is enabled for all STP instances on that interface.
- The interface is blocked only for instances for which it receives more BPDUs. Otherwise, it participates in the spanning-tree topology.




**NOTE:** An interface can be configured for either root protection or loop protection, but not for both.

Four switches are displayed in [Figure 27 on page 1666](#). In this example, they are configured for RSTP and create a loop-free topology. Interface **xe-0/0/7** on Switch 1 is a designated port on an administrative boundary. It connects to Switch 4. Switch 3 is the root bridge. Interface **xe-0/0/6** on Switch 1 is the root port.

This example shows how to configure root protection on interface **xe-0/0/7** to prevent it from transitioning to become the root port.

Figure 27: Network Topology for Root Protection



[Table 108 on page 1666](#) shows the components that will be configured for root protection.

Table 108: Topology for Configuring Root Protection on the QFX Series

Component	Settings
Switch 1	Switch 1 is connected to Switch 4 through interface <b>xe-0/0/7</b> .
Switch 2	Switch 2 is connected to Switch 1 and Switch 3. Interface <b>xe-0/0/4</b> is the alternate port in the RSTP topology.
Switch 3	Switch 3 is the root bridge and is connected to Switch 1 and Switch 2.

Table 108: Topology for Configuring Root Protection on the QFX Series (*continued*)

Component	Settings
Switch 4	Switch 4 is connected to Switch 1. After loop protection is configured on interface <b>xe-0/0/7</b> , Switch 4 sends more BPDUs that trigger loop protection on interface <b>xe-0/0/7</b> .

A spanning-tree topology contains ports that have specific roles:

- The *root port* is responsible for forwarding data to the root bridge.
- The *alternate port* is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The *designated port* forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you can also configure root protection for STP or MSTP topologies at the **[edit protocols (mstp | stp)]** hierarchy level.

### Configuration

**CLI Quick Configuration** To quickly configure root protection on interface **xe-0/0/7**, copy the following command and paste it into the switch terminal window:

```
[edit]
set protocols rstp interface xe-0/0/7 no-root-port
```

**Step-by-Step Procedure** To configure root protection:

1. Configure interface **xe-0/0/7**:
 

```
[edit protocols rstp]
user@switch#
set interface xe-0/0/7 no-root-port
```

**Results** Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface xe-0/0/7.0 {
  no-root-port;
}
```

### Verification

To confirm that the configuration is working properly, perform these tasks:

- [Displaying the Interface State Before Root Protection Is Triggered on page 1667](#)
- [Verifying That Root Protection Is Working on the Interface on page 1668](#)

#### *Displaying the Interface State Before Root Protection Is Triggered*

**Purpose** Before root protection is triggered on interface **xe-0/0/7**, confirm the interface state.

**Action** Confirm the state of the interfaces before root protection is configured:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/0.0	128:513	128:513	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/1.0	128:514	128:514	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/2.0	128:515	128:515	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/3.0	128:516	128:516	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/4.0	128:517	128:517	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/5.0	128:518	128:2	16384.00aabbcc0348	20000	BLK	ALT
xe-0/0/6.0	128:519	128:1	16384.00aabbcc0348	20000	FWD	ROOT
xe-0/0/7.0	128:520	128:520	32768.0019e2503f00	20000	FWD	DESG

[output truncated]

**Meaning** The output from the operational mode command **show spanning-tree interface** shows that **xe-0/0/7.0** is a designated port in a forwarding state.

#### *Verifying That Root Protection Is Working on the Interface*

**Purpose** A configuration change takes place on Switch 4. A lower bridge priority on Switch 4 causes it to send more BPDUs to interface **xe-0/0/7**. Receipt of more BPDUs on interface **xe-0/0/7** triggers root protection. Verify that root protection is operating on interface **xe-0/0/7**.

**Action** Verify that root protection has been configured and is operating correctly:

```
user@switch> show spanning-tree interface
```

Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-0/0/0.0	128:513	128:513	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/1.0	128:514	128:514	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/2.0	128:515	128:515	32768.0019e2503f00	20000	BLK	DIS
xe-0/0/3.0	128:516	128:516	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/4.0	128:517	128:517	32768.0019e2503f00	20000	FWD	DESG
xe-0/0/5.0	128:518	128:2	16384.00aabbcc0348	20000	BLK	ALT
xe-0/0/6.0	128:519	128:1	16384.00aabbcc0348	20000	FWD	ROOT
xe-0/0/7.0	128:520	128:520	32768.0019e2503f00	20000	BLK	DIS

(Root-Incon)  
[output truncated]

**Meaning** The operational mode command **show spanning-tree interface** shows that interface **xe-0/0/7.0** has transitioned to a loop inconsistent state. The loop inconsistent state blocks the interface and prevents it from becoming a candidate for the root port. When the root bridge no longer receives more STP BPDUs from the interface, the interface recovers and transitions back to a forwarding state. Recovery is automatic.

**Related Documentation**

- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
- [Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660](#)
- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656](#)
- [Understanding Root Protection for STP, RSTP, VSTP, and MSTP on page 1560](#)

## Bridging and VLAN Configuration Tasks

---

- [Adding a Static MAC Address Entry to the Ethernet Switching Table \(CLI Procedure\) on page 1669](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 1670](#)
- [Configuring MAC Limiting \(CLI Procedure\) on page 1672](#)
- [Configuring MAC Table Aging on page 1674](#)
- [Configuring IRB Interfaces on page 1675](#)
- [Configuring Static ARP Entries on page 1676](#)
- [Configuring the Native VLAN Identifier \(CLI Procedure\) on page 1677](#)
- [Configuring VLANs on page 1678](#)
- [Creating a Series of Tagged VLANs on page 1680](#)
- [Disabling MAC Learning on page 1681](#)
- [Configuring MAC Notification \(CLI Procedure\) on page 1682](#)

### Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)

---



**NOTE:** This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

---

The Ethernet switching table, also known as the forwarding table, specifies the known locations of VLAN nodes and the addresses of devices within those nodes. There are two ways to populate the Ethernet switching table on a switch. The easiest method is to let the switch update the table with MAC addresses.

The second way to populate the Ethernet switching table is to manually insert addresses into the table. You can do this to reduce flooding and speed up the switch’s automatic learning process.

Before configuring a static MAC address, be sure that you have:

- Set up the VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*.

To configure an interface to have a static MAC address:

```
[edit vlans vlan-name switch-options interface interface-name]  
user@switch# set static-mac mac-address
```

**Related  
Documentation**

- *Understanding Bridging and VLANs on EX Series Switches*

## Configuring Ethernet Ring Protection Switching (CLI Procedure)

You can configure Ethernet ring protection switching (ERPS) on connected switches to prevent fatal loops from disrupting a network. ERPS is similar to spanning-tree protocols, but ERPS is more efficient than spanning-tree protocols because it is customized for ring topologies. You must configure at least three switches to form a ring. One of the links, called the ring protection link (RPL) end interface, is blocked until another link fails—at this time the RPL link is unblocked, ensuring connectivity.



**NOTE:** Ethernet OAM connectivity fault management (CFM) can be used with ERPS to detect link faults faster in some cases. See *Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)*.

The time needed for switchover to the ERPS link is affected by three settings—link failure detection time, the number of nodes in the ring, and the time it takes to unblock the RPL after a failure is detected.



**NOTE:** Do not configure redundant trunk groups on ERPS interfaces. You can configure VSTP on ERPS interfaces if the VSTP uses a VLAN that is not part of the ERPS control VLAN or data channel VLANs. The total number of ERPS and VSTP or MSTP instances is limited to 253.

Before you begin:

- Optionally, configure two interfaces on each switch as trunk ports. See *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*.
- Configure a VLAN to act as a control VLAN for ERPS if your interfaces are trunk ports. Configure the same VLAN on all switches and associate the two network interfaces from each of the switches with the VLAN. See *Configuring VLANs for EX Series Switches (CLI Procedure)*. If you have multiple ERPS instances, the control VLANs and data channel VLANs must not overlap.
- Data channels are optional on the ERPS link. If you plan to use them, configure a VLAN for each data channel.



To configure ERPS:



**NOTE:** You must configure at least three switches, with only one switch designated as the RPL owner node.

1. RSTP and EPRS cannot both be configured on a ring port, and RSTP is configured by default. Disable RSTP on each switch interface:

```
user@switch# set rstp interface interface-name disable
```

2. Create a node ring on each switch:

```
[edit protocols]
user@switch# set protection-group ethernet-ring ring-name
```

3. Configure a control VLAN for the node ring if the links are trunk ports:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set control-vlan vlan-name-or-vlan-id
```

4. Configure the east interface of the node ring with the control-channel interface. In addition, configure either the east interface or the west interface (but not both) as a link end.

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set east-interface control-channel channel-name
user@switch# set east-interface ring-protection link end
```

5. Configure the west interface of the node ring with the control-channel interface. In addition, configure either the east interface or the west interface (but not both) as a link end.

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set west-interface control-channel control-channel-interface-address
user@switch# set west-interface ring protection link end
```

6. Configure only one switch as the RPL owner node:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set ring-protection-link-owner
```

7. The restore interval configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs). When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL receives notification, restores the link, and waits the length of time indicated by the restore interval before issuing another block on the same link. Optionally, configure the restore interval on each switch:

```
[edit protocols protection-group ethernet-ring ring-name]
user@switch# set restore-interval restore-interval-value
```

8. The guard interval prevents ring nodes from receiving outdated messages (called RAPs). Optionally, configure the guard interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]
user@switch# set guard-interval guard-interval-value
```



**NOTE:** Local settings take priority over global settings.

Global settings are used when no local settings are present. Optionally, you can also configure these global settings on the switch:

- restore interval
- guard interval
- ERP traceoptions: file, page size, file size, flag name

9. Optionally, reconfigure the global guard interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set guard-interval guard-interval-value
```

10. Optionally, reconfigure the global restore interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set restore-interval restore-interval-value
```

11. After detection of a link failure, switching takes place after the hold interval has expired. Optionally, reconfigure the global hold interval on each switch:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set hold-interval hold-interval-value
```

12. Optionally, configure VLANs for data channels on the ERPS link:

```
[edit protocols protection-group ethernet-ring ring name]  
user@switch# set data-channel vlan-name
```

#### Related Documentation

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563](#)
- [Ethernet Ring Protection Switching Overview on page 1525](#)
- [Understanding Ethernet Ring Protection Switching Functionality on page 1534](#)

## Configuring MAC Limiting (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring MAC Limiting \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

This topic describes various ways of configuring a limitation on MAC addresses in packets that are received and forwarded by the switch.



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

The different ways of setting a MAC limit are described in the following sections:

- [Limiting the Number of MAC Addresses Learned by an Interface on page 1673](#)
- [Limiting the Number of MAC Addresses Learned by a VLAN on page 1673](#)

### Limiting the Number of MAC Addresses Learned by an Interface

To secure a port, you can set the maximum number of MAC addresses that can be learned by an interface:

- Set the MAC limit on an interface, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit switch-options]
user@switch# set interface interface-name interface-mac-limit limit packet-action
action
```

After you set a new MAC limit for the interface, the system clears existing entries in the MAC address forwarding table associated with the interface.

### Limiting the Number of MAC Addresses Learned by a VLAN

To limit the number of MAC addresses learned by a VLAN, perform both of the following steps:

1. Set the maximum number of MAC addresses that can be learned by a VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options mac-table-size limit packet-action
action
```

2. Set the maximum number of MAC addresses that can be learned by one or all interfaces in the VLAN, and specify an action that the switch takes after the specified limit is exceeded:

```
[edit vlans]
user@switch# set vlan-name switch-options interface interface-name
interface-mac-limit limit packet-action action
[edit vlans]
user@switch# set vlan-name switch-options interface-mac-limit limit packet-action
action
```



**NOTE:** If you specify a MAC limit and packet action for all interfaces in the VLAN *and* a specific interface in the VLAN, the MAC limit and packet action specified at the specific interface level takes precedence. Also, at the VLAN interface level, only the drop and drop-and-log options are supported.

After you set new MAC limits for a VLAN by using the **mac-table-size** statement or for interfaces associated with a VLAN by using the **interface-mac-limit** statement, the system clears the corresponding existing entries in the MAC address forwarding table.

- Related Documentation**
- [Understanding Bridging and VLANs on EX Series Switches](#)
  - [Configuring Persistent MAC Learning \(CLI Procedure\)](#)

## Configuring MAC Table Aging

MAC table aging ensures that a switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

To manage MAC entries more efficiently, you can configure an entry's aging time, which is the maximum time that an entry can remain in the MAC address table before it is deleted because it has reached its maximum age.



**NOTE:** This task uses Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring MAC Table Aging](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

You can use the **set-mac-table-aging-time** command to configure how long entries remain in the Ethernet switching table before expiring. Here the VLAN is **employee-vlan**:

```
[edit vlans employee-vlan switch-options]  
user@switch# set mac-table-aging-time 200
```

- Related Documentation**
- [Understanding Bridging and VLANs on page 1527](#)
  - [Example: Setting Up Bridging with Multiple VLANs on page 1583](#)
  - [Example: Connecting an Access Switch to a Distribution Switch on page 1647](#)

## Configuring IRB Interfaces

Integrated routing and bridging (IRB) interfaces enable a switch to recognize which packets are being sent to local addresses so that they are bridged whenever possible and are routed only when needed. Whenever packets can be switched instead of routed, several layers of processing are eliminated. Switching also reduces the number of address look-ups.



**NOTE:** In versions of Junos OS that do not support Enhanced Layer 2 Software (ELS), this type of interface is called a routed VLAN interface (RVI).

To configure the routed VLAN interface:

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

```
[edit]
user@switch# set vlans support vlan-id 111
```

2. Assign an interface to the VLAN by specifying the logical interface (with the **unit** statement) and specifying the VLAN name as the member:

```
[edit]
user@switch# set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members
support
```

3. Create the subnet for the VLAN's broadcast domain:

```
[edit]
user@switch# set interfaces irb unit 111 family inet address 111.111.111.1/24
```

4. Bind a Layer 3 interface with the VLAN:

```
[edit]
user@switch# set vlans support l3-interface irb.111
```



**NOTE:** If you are using a version of Junos OS that does not support ELS, you create a Layer 3 virtual interface named **vlan**



**NOTE:** Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

You can display the configuration settings:

```
user@switch> show interfaces irb terse
```

Interface	Admin	Link	Proto	Local	Remote
vlan	up	up			
irb.111	up	up	inet	111.111.111.1/24	

```

user@switch> show vlans
Name      Tag      Interfaces
default
None

```

```
employee-vlan 20
marketing      40
support        111
mgmt
                ge-1/0/0.0, ge-1/0/1.0, ge-1/0/2.0
                ge-1/0/10.0, ge-1/0/20.0, ge-1/0/30.0
                ge-0/0/18.0
                bme0.32769, bme0.32771*

user@switch> show ethernet-switching table
Ethernet-switching table: 1 entries, 0 learned
  VLAN      MAC address      Type      Age Interfaces
  support    00:19:e2:50:95:a0 Static    - Router
```

**Related Documentation**

- [Understanding Integrated Routing and Bridging on page 1539](#)

## Configuring Static ARP Entries

You can create static ARP table entries, which are explicit mappings between IP addresses and MAC addresses.

- To configure a static ARP entry:

```
[edit interfaces interface-name unit logical-unit-number family inet address
address]
user@switch# set arp ip-address (mac | multicast-mac) mac-address
```

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, use the **multicast-mac** statement.

Specify the MAC address as 6 hexadecimal bytes in one of the following formats:  
*nnnnn.nnnnn.nnnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

**Related Documentation**

- [Understanding Static ARP Entries on page 5271](#)
- *arp*

## Configuring the Native VLAN Identifier (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring the Native VLAN Identifier (CLI Procedure)*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 43.

Switches can receive and forward routed or bridged Ethernet frames with 802.1Q VLAN tags. Typically, trunk ports, which connect switches to each other, accept untagged control packets but do not accept untagged data packets. You can enable a trunk port to accept untagged data packets by configuring a native VLAN ID on the interface on which you want the untagged data packets to be received. The logical interface on which untagged packets are to be received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface.

To configure the native VLAN ID by using the command-line interface (CLI):

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

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

2. Configure the native VLAN ID:

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

3. Specify that the logical interface that will receive the untagged data packets is a member of the native VLAN:

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

### Related Documentation

- *Understanding Bridging and VLANs on EX Series Switches*
- *Example: Connecting Access Switches to a Distribution Switch*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588*

## Configuring VLANs

Switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.



**NOTE:** This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#). If your switch runs software that does not support ELS, see [Configuring VLANs](#).

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Specify the description of the VLAN:

```
[edit interfaces interface-name unit 0]
user@switch# set description vlan-description
```

2. Specify the unique name of the VLAN:



**NOTE:** Switches that run Junos OS with the ELS configuration style do not support a default VLAN. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist.

```
[edit interfaces interface-name unit 0]
user@switch# set family ethernet-switching vlan members vlan-name
```

3. Create the subnet for the VLAN:

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

4. Configure the VLAN tag ID or VLAN ID list for the VLAN:

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

or

```
[edit vlans]
user@switch# set vlan-name vlan-id-list [vlan-ids | vlan-id--vlan-id]
```

5. Specify the maximum time that an entry can remain in the forwarding table before it ages out:

```
[edit vlans]
user@switch# set vlan-name switch-options mac-table-aging-time time
```

6. Specify a VLAN firewall filter to be applied to incoming or outgoing packets:

```
[edit vlans]
user@switch# set vlan-name filter (input | output) filter-name
```

### Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Configuring IRB Interfaces on page 1675](#)



- *Creating a Series of Tagged VLANs*
- [Understanding Bridging and VLANs on page 1527](#)

## Creating a Series of Tagged VLANs

When you divide an Ethernet LAN into multiple VLANs, each VLAN is assigned a unique IEEE 802.1Q tag. This tag is associated with each frame in the VLAN, and the network nodes receiving the traffic can use the tag to identify which VLAN a frame is associated with.

Instead of configuring VLANs and 802.1Q tags one at a time for a trunk interface, you can configure a VLAN range to create a series of tagged VLANs.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames can detect which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For example, you could configure the VLAN **employee** and specify a tag range of **10 through 12**. This creates the following VLANs and tags:

- VLAN **employee-10**, tag 10
- VLAN **employee-11**, tag 11
- VLAN **employee-12**, tag 12

Creating tagged VLANs in a series has the following limitations:

- Layer 3 interfaces do not support this feature.
- Because an access interface can only support one VLAN member, access interfaces also do not support this feature.



**NOTE:** This task uses Junos OS for Junos OS for QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Creating a Series of Tagged VLANs*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software”](#) on page 43.

---

To configure a series of tagged VLANs using the CLI (here, the VLAN is **employee**):

1. Configure the series (here, a VLAN series from 120 through 130):

```
[edit]
user@switch# set vlans employee vlan-id-list [ 120-130 ]
```

2. Associate a series of tagged VLANs when you configure an interface in one of two ways:

- Include the name of the series:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/22.0 family ethernet-switching vlanmembers employee
```

- Include the VLAN range:

```
[edit interfaces]
user@switch# set interfaces xe-0/0/22.0 family ethernet-switching vlan members 120-130
```

Associating a series of tagged VLANs to an interface by name or by VLAN range the same result: VLANs **\_\_employee\_120\_\_** through **\_\_employee\_130\_\_** are created.



**NOTE:** When a series of VLANs is created using the `vlan-id-list` command, the VLAN names are preceded and followed by a double underscore.

#### Related Documentation

- [Verifying That a Series of Tagged VLANs Has Been Created on page 1827](#)
- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Example: Setting Up Bridging with Multiple VLANs on page 1583](#)
- [Understanding Bridging and VLANs on page 1527](#)

## Disabling MAC Learning

By default, MAC learning is globally enabled on all node. This topic describes how to disable MAC learning, as well as how to reenable and verify that MAC learning has been enabled or disabled.



**NOTE:** This task supports the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#). If your switch runs software that does not support ELS, see [Disabling MAC Learning](#).

Disabling dynamic MAC learning prevents a node from learning source and destination MAC addresses.

- To disable MAC learning:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# set no-mac-learning
```

- To enable MAC learning:

```
[edit vlans vlan-name switch-options interface interface-name]
user@switch# delete no-mac-learning
user@switch# deactivate no-mac-learning
```

- To verify the status of MAC learning, view the Ethernet MAC learning statistics in operational mode.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 2 entries, 1 learned
  VLAN      MAC address      Type      Age Interfaces
  default    *                Flood     - All-members
  default    00:1f:12:39:90:80 Learn     29 xe-/0/0.0
```

#### Related Documentation

- [Understanding MAC Learning on page 1540](#)
- [Example: Disabling MAC Learning on page 1582](#)
- *no-mac-learning*

## Configuring MAC Notification (CLI Procedure)



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring MAC Notification (CLI Procedure)* or *Configuring MAC Notification*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 43.

When a switch learns or unlearns a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- [Enabling MAC Notification on page 1682](#)
- [Disabling MAC Notification on page 1683](#)
- [Setting the MAC Notification Interval on page 1683](#)

### Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

```
[edit switch-options]
```

```
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, the MAC notification interval is set to 60 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 60
```

### Disabling MAC Notification

MAC notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit switch-options]
user@switch# delete mac-notification
```

To disable MAC notification on a specific interface (here, the interface is ge-0/0/3):

```
[edit switch-options]
user@switch# set interface ge-0/0/3 no-mac-notification
```

### Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit switch-options]
user@switch# set mac-notification notification-interval 5
```

**Related Documentation**

- *Verifying That MAC Notification Is Working Properly*

## Q-in-Q Tunneling Configuration Tasks

- [Configuring Q-in-Q Tunneling on page 1683](#)
- [Configuring All-in-One Bundling on page 1691](#)
- [Configuring Many-to-Many Bundling on page 1692](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 1695](#)

### Configuring Q-in-Q Tunneling

Q-in-Q tunneling and VLAN translation allow service providers to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Data centers can use Q-in-Q tunneling and VLAN translation to isolate customer traffic within a single site or to enable customer traffic flows between cloud data centers in different geographic locations.

Q-in-Q tunneling adds a service VLAN tag before the customer's 802.1Q VLAN tags. The Juniper Networks Junos operating system implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.



**NOTE:** This task uses a Junos OS release that supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Q-in-Q Tunneling*.

With releases of Junos OS 13.2X51 previous to 13.2X51-D20, you cannot create a regular VLAN on an interface if you have created an S-VLAN or C-VLAN on that interface for Q-in-Q tunneling. This means that you cannot create an integrated routing and bridging (IRB) interface on that interface because regular VLANs are a required part of IRB configuration. With Junos OS 13.2X51-D25, you can create a regular VLAN on a trunk interface that has an S-VLAN, which means that you can also create an IRB interface on the trunk. In this case, the regular VLAN and S-VLAN on the same trunk interface cannot share the same VLAN ID. Junos OS 13.2X51-D25 does not allow you to create a regular VLAN on an access interface that has a C-VLAN.

Before setting up Q-in-Q tunneling, make sure you have created and configured the necessary customer VLANs on the neighboring switches. See *Configuring VLANs*.

- [Using the Different Mapping Methods on page 1684](#)
- [Configuring All-in-One Bundling on page 1684](#)
- [Configuring Many-to-Many Bundling on page 1686](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 1689](#)

### Using the Different Mapping Methods

---

Once you have created the required VLANs on the neighboring switches, configure Q-in-Q tunneling using one of the three methods to map customer VLANs (C-VLANs) to service-provider-defined service VLANs (S-VLANs):

- All-in-one bundling maps all packets from all C-VLAN interfaces to an S-VLAN. For information about how to use this method, see [“Configuring All-in-One Bundling” on page 1684](#).
- Use many-to-many bundling when you want a subset of the C-VLANs on the access switch to be part of multiple S-VLANs. For information about how to use this method, see [“Configuring Many-to-Many Bundling” on page 1686](#).
- Use specific interface mapping when you want to assign an S-VLAN to a specific C-VLAN on an interface. For information about how to use this method, see [“Configuring a Specific Interface Mapping with VLAN ID Translation Option” on page 1689](#).

### Configuring All-in-One Bundling

---

You can configure Q-in-Q tunneling using the all-in-one bundling method, which forwards all packets that ingress on a C-VLAN interface to an S-VLAN. (Packets are forwarded to

the S-VLAN regardless of whether they are tagged or untagged prior to ingress.) Using this approach saves you the effort of specifying a specific mapping for each C-VLAN.

First configure the S-VLAN and its interface:

1. Create the S-VLAN and assign a VLAN ID for it.

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

2. Assign a logical interface (unit) to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```



**NOTE:** Do not use logical interface unit 0. You must later bind a VLAN tag ID to the unit you specify in this step, and you cannot bind a VLAN tag ID to unit 0.

3. Enable the interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the S-VLAN interface to send and receive untagged packets:

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

6. Bind the logical interface (unit) of the interface that you specified in step 2 to the VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration creates S-VLAN v10, makes xe-1/1/1.10 a member of that VLAN, enables Q-in-Q tunneling on interface xe-1/1/1, enables xe-1/1/1 to accept untagged packets, and binds the VLAN ID of S-VLAN v10 to a logical interface of xe-1/1/1.

```
set vlans v10 vlan-id 10
set vlans v10 interface xe-1/1/1.10
set interfaces xe-1/1/1 flexible-vlan-tagging
set interfaces xe-1/1/1 set native-vlan-id 10
set interfaces xe-1/1/1 encapsulation extended-vlan-bridge
set interfaces xe-1/1/1 unit 10 vlan-id 10
```

Now configure all-in-one bundling on a C-VLAN interface:

1. Assign a logical interface (unit) of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags :

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

3. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

4. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

5. Configure a logical interface to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id-list vlan-id-numbers
```

6. Configure the system to add an S-VLAN tag (outer tag) as packets travel from a C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# input-vlan-map push
```

7. Configure the system to remove the S-VLAN tag when packets are forwarded (internally) from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, maps packets from C-VLANs 100 through 200 to S-VLAN 10, and enables xe-0/0/1 to accept untagged packets. If a packet originates in C-VLAN 100 and needs to be sent across the S-VLAN, a tag with VLAN ID 10 is added to the packet. When a packet is forwarded (internally) from the S-VLAN interface to interface xe-0/0/1, the tag with VLAN ID 10 is removed.

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 100-200
set interfaces xe-0/0/1 native-vlan-id 150
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop
```

### Configuring Many-to-Many Bundling

---

You can configure Q-in-Q tunneling using the many-to-many bundling method, which maps packets from multiple C-VLANs to multiple S-VLANs. This method is convenient for mapping a range of C-VLANs without having to specify each one individually. (You can also use this method to configure only one C-VLAN to be mapped to an S-VLAN.)

First configure the S-VLANs and assign them to an interface:

1. Create one of the S-VLANs and assign a VLAN ID for it.

```
[edit vlans vlan-name]
```

```
user@switch# vlan-id vlan-id-number
```

2. Repeat step 1 for the other S-VLANs.

3. Assign a logical interface (unit) to be a member of one of the S-VLANs. Do not use logical interface unit 0.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```

4. Repeat step 3 to assign the other logical interfaces on the same physical interface to be a member of other S-VLANs.



5. Enable the physical interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

6. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

7. Enable the S-VLAN interface to send and receive untagged packets:

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

8. Bind one of the logical units of the interface to the VLAN ID for one of the S-VLANs.

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

9. Repeat step 8 to bind the VLAN IDs for the other S-VLANs to the other logical units of the interface:

For example, the following configuration creates S-VLANs v10 and v30 and associates them with interface xe-1/1/1. It also enables Q-in-Q tunneling, enables xe-1/1/1 to accept untagged packets, and maps incoming C-VLAN packets to S-VLANs v10 and v30.

```
set vlans v10 vlan-id 10
set vlans v30 vlan-id 30
set vlans v10 interface xe-1/1/1.10
set vlans v30 interface xe-1/1/1.30
set interfaces xe-1/1/1 flexible-vlan-tagging
set interfaces xe-1/1/1 set native-vlan-id 10
set interfaces xe-1/1/1 encapsulation extended-vlan-bridge
set interfaces xe-1/1/1 unit 10 vlan-id 10
set interfaces xe-1/1/1 unit 30 vlan-id 30
```

To configure the many-to-many bundling method on a C-VLAN interface, perform the following steps for each customer:

1. Assign a logical interface (unit) of one C-VLAN interface to be a member of one S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

2. Repeat step 1 to assign another C-VLAN interface (physical interface) to be a member of another S-VLAN.

3. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the C-VLAN interface to send and receive untagged packets:

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

6. For each physical interface, configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id-list vlan-id-numbers
```

To configure only one C-VLAN to be mapped to an S-VLAN, specify only one VLAN ID after *vlan-id-list*.

7. For each physical interface, configure the system to add an S-VLAN tag (outer tag) as packets travel from the C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# input-vlan-map push
```

8. For each physical interface, configure the system to remove the S-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, and maps packets from C-VLANs 10 through 20 to S-VLAN 10. The configuration for customer 2 makes xe-0/0/2.30 a member of S-VLAN v30, enables Q-in-Q tunneling, and maps packets from C-VLANs 30 through 40, 50 through 60, and 70 through 80 to S-VLAN 30. Both interfaces are configured to accept untagged packets.

If a packet originates in C-VLAN 10 and needs to be sent over the S-VLAN, a tag with a VLAN ID 10 is added to the packet. If a packet is forwarded internally from the S-VLAN interface to xe-0/0/1.10, the tag with VLAN ID 10 is removed. The same principles apply to the C-VLANs configured on interface xe-0/0/2.



**NOTE:** Notice that you can use the same tag value for an S-VLAN and C-VLAN. For example, the configuration for customer 1 maps C-VLAN ID 10 to S-VLAN ID 10. C-VLAN and S-VLAN tags use separate name spaces, so this configuration is allowed.

Configuration for customer 1:

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 10-20
set interfaces xe-0/0/1 native-vlan-id 15
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop
```

Configuration for customer 2:

```
set vlans v30 interface xe-0/0/2.30
set interfaces xe-0/0/2 flexible-vlan-tagging
set interfaces xe-0/0/2 encapsulation extended-vlan-bridge
set interfaces xe-0/0/2 unit 30 vlan-id-list 30-40
set interfaces xe-0/0/2 unit 30 vlan-id-list 50-60
set interfaces xe-0/0/2 unit 30 vlan-id-list 70-80
set interfaces xe-0/0/2 native-vlan-id 75
set interfaces xe-0/0/2 unit 30 input-vlan-map push
set interfaces xe-0/0/2 unit 30 output-vlan-map pop
```

### Configuring a Specific Interface Mapping with VLAN ID Translation Option

You can configure Q-in-Q tunneling by mapping packets from a specified C-VLAN to a specified S-VLAN. In addition, you can configure the system to replace a C-VLAN tag with an S-VLAN tag or replace an S-VLAN tag with a C-VLAN tag (instead of double tagging). This is called VLAN translation or VLAN rewriting. VLAN translation is particularly useful if a service provider's Layer 2 network that connects a customer's sites does not support double tagged packets.

When you use VLAN translation, both ends of the link normally must be able to swap the tags appropriately. That is, both ends of the link must be configured to swap the C-VLAN tag for the S-VLAN tag and swap the S-VLAN tag for the C-VLAN tag so that traffic in both directions is tagged appropriately while in transit and after arrival.

First configure the S-VLAN and its interface:

1. Create the S-VLAN and assign a VLAN ID for it.

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

2. Assign a logical interface to be a member of the S-VLAN. Do not use unit 0.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

3. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

4. Enable the S-VLAN interface to send and receive untagged packets:

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

5. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

6. Bind the logical interface (unit) of the interface that you specified earlier to the VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration creates S-VLAN v200, makes xe-1/1/1.200 a member of that VLAN, enables Q-in-Q tunneling on interface xe-1/1/1, enables xe-1/1/1 to accept untagged packets, and binds a logical interface of xe-1/1/1 to the VLAN ID of VLAN v200.

```
set vlans v200 vlan-id 200
set vlans v200 interface xe-1/1/1.200
set interfaces xe-1/1/1 flexible-vlan-tagging
set interfaces xe-1/1/1 set native-vlan-id 10
set interfaces xe-1/1/1 encapsulation extended-vlan-bridge
set interfaces xe-1/1/1 unit 200 vlan-id 200
```

Now configure a specific interface mapping with optional VLAN ID translation on the C-VLAN interface:

1. Assign a logical interface of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]  
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]  
user@switch# flexible-vlan-tagging
```

3. Enable the C-VLAN interface to send and receive untagged packets:

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

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]  
user@switch# encapsulation extended-vlan-bridge
```

5. Configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]  
user@switch# vlan-id number
```

6. Configure the system to remove the existing C-VLAN tag and replace it with the S-VLAN tag when packets ingress on the C-VLAN interface and are forwarded to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]  
user@switch# input-vlan-map swap
```

7. Configure the system to remove the existing S-VLAN tag and replace it with the C-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]  
user@switch# output-vlan-map swap
```

8. To configure an S-VLAN and associate it with the appropriate C-VLAN interface:

```
[edit vlans vlan-name]  
user@switch# interface interface-name
```

For example, the following configuration on C-VLAN interface xe-0/0/1 enables Q-in-Q tunneling, enables xe-0/0/1 to accept untagged packets, and maps incoming packets from C-VLAN 150 to logical interface 200, which is a member of S-VLAN 200. Also, when packets egress from C-VLAN interface xe-0/0/1 and travel to the S-VLAN interface, the C-VLAN tag of 150 is removed and replaced with the S-VLAN tag of 200. When packets travel from the S-VLAN interface to the C-VLAN interface, the S-VLAN tag of 200 is removed and replaced with the C-VLAN tag of 150.

```
set vlans v200 interface xe-0/0/1.200  
set interfaces xe-0/0/1 flexible-vlan-tagging  
set interfaces xe-0/0/1 set native-vlan-id 10  
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge  
set interfaces xe-0/0/1 unit 200 vlan-id 150  
set interfaces xe-0/0/1 unit 200 output-vlan-map swap  
set interfaces xe-0/0/1 unit 200 input-vlan-map swap
```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 1547](#)

## Configuring All-in-One Bundling

You can configure Q-in-Q tunneling using the all-in-one bundling method, which forwards all packets that ingress on a C-VLAN interface to an S-VLAN. (Packets are forwarded to the S-VLAN regardless of whether they are tagged or untagged prior to ingress.) Using this approach saves you the effort of specifying a specific mapping for each C-VLAN.

First configure the S-VLAN and its interface:

1. Create the S-VLAN and assign a VLAN ID for it.

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

2. Assign a logical interface (unit) to be a member of the S-VLAN.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```



**NOTE:** Do not use logical interface unit 0. You must later bind a VLAN tag ID to the unit you specify in this step, and you cannot bind a VLAN tag ID to unit 0.

3. Enable the interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the S-VLAN interface to send and receive untagged packets:

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

6. Bind the logical interface (unit) of the interface that you specified in step 2 to the VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

For example, the following configuration creates S-VLAN v10, makes xe-1/1/1.10 a member of that VLAN, enables Q-in-Q tunneling on interface xe-1/1/1, enables xe-1/1/1 to accept untagged packets, and binds the VLAN ID of S-VLAN v10 to a logical interface of xe-1/1/1.

```
set vlans v10 vlan-id 10
set vlans v10 interface xe-1/1/1.10
set interfaces xe-1/1/1 flexible-vlan-tagging
set interfaces xe-1/1/1 set native-vlan-id 10
set interfaces xe-1/1/1 encapsulation extended-vlan-bridge
set interfaces xe-1/1/1 unit 10 vlan-id 10
```

Now configure all-in-one bundling on a C-VLAN interface:

1. Assign a logical interface (unit) of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags :

```
[edit interfaces interface-name]
```

```
user@switch# flexible-vlan-tagging
```

3. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

4. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

5. Configure a logical interface to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id-list vlan-id-numbers
```

6. Configure the system to add an S-VLAN tag (outer tag) as packets travel from a C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# input-vlan-map push
```

7. Configure the system to remove the S-VLAN tag when packets are forwarded (internally) from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, maps packets from C-VLANs 100 through 200 to S-VLAN 10, and enables xe-0/0/1 to accept untagged packets. If a packet originates in C-VLAN 100 and needs to be sent across the S-VLAN, a tag with VLAN ID 10 is added to the packet. When a packet is forwarded (internally) from the S-VLAN interface to interface xe-0/0/1, the tag with VLAN ID 10 is removed.

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 100-200
set interfaces xe-0/0/1 native-vlan-id 150
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop
```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 1547](#)
- [Configuring Many-to-Many Bundling on page 1686](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 1689](#)

## Configuring Many-to-Many Bundling

You can configure Q-in-Q tunneling using the many-to-many bundling method, which maps packets from multiple C-VLANs to multiple S-VLANs. This method is convenient

for mapping a range of C-VLANs without having to specify each one individually. (You can also use this method to configure only one C-VLAN to be mapped to an S-VLAN.)

First configure the S-VLANs and assign them to an interface:

1. Create one of the S-VLANs and assign a VLAN ID for it.

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

2. Repeat step 1 for the other S-VLANs.

3. Assign a logical interface (unit) to be a member of one of the S-VLANs. Do not use logical interface unit 0.

```
[edit vlans vlan-name]
user@switch# interface interface-name.unit-number
```

4. Repeat step 3 to assign the other logical interfaces on the same physical interface to be a member of other S-VLANs.

5. Enable the physical interface to transmit packets with two 802.1Q VLAN tags:

```
[edit interfaces interface-name]
user@switch# flexible-vlan-tagging
```

6. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
user@switch# encapsulation extended-vlan-bridge
```

7. Enable the S-VLAN interface to send and receive untagged packets:

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

8. Bind one of the logical units of the interface to the VLAN ID for one of the S-VLANs.

```
[edit interfaces interface-name unit logical-unit-number]
user@switch# vlan-id number
```

9. Repeat step 8 to bind the VLAN IDs for the other S-VLANs to the other logical units of the interface:

For example, the following configuration creates S-VLANs v10 and v30 and associates them with interface xe-1/1/1. It also enables Q-in-Q tunneling, enables xe-1/1/1 to accept untagged packets, and maps incoming C-VLAN packets to S-VLANs v10 and v30.

```
set vlans v10 vlan-id 10
set vlans v30 vlan-id 30
set vlans v10 interface xe-1/1/1.10
set vlans v30 interface xe-1/1/1.30
set interfaces xe-1/1/1 flexible-vlan-tagging
set interfaces xe-1/1/1 set native-vlan-id 10
set interfaces xe-1/1/1 encapsulation extended-vlan-bridge
set interfaces xe-1/1/1 unit 10 vlan-id 10
set interfaces xe-1/1/1 unit 30 vlan-id 30
```

To configure the many-to-many bundling method on a C-VLAN interface, perform the following steps for each customer:

1. Assign a logical interface (unit) of one C-VLAN interface to be a member of one S-VLAN.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```

2. Repeat step 1 to assign another C-VLAN interface (physical interface) to be a member of another S-VLAN.

3. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# flexible-vlan-tagging
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

5. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

6. For each physical interface, configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the list of VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id-list vlan-id-numbers
```

To configure only one C-VLAN to be mapped to an S-VLAN, specify only one VLAN ID after *vlan-id-list*.

7. For each physical interface, configure the system to add an S-VLAN tag (outer tag) as packets travel from the C-VLAN interface to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# input-vlan-map push
```

8. For each physical interface, configure the system to remove the S-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# output-vlan-map pop
```

For example, the following configuration makes xe-0/0/1.10 a member of S-VLAN v10, enables Q-in-Q tunneling, and maps packets from C-VLANs 10 through 20 to S-VLAN 10. The configuration for customer 2 makes xe-0/0/2.30 a member of S-VLAN v30, enables Q-in-Q tunneling, and maps packets from C-VLANs 30 through 40, 50 through 60, and 70 through 80 to S-VLAN 30. Both interfaces are configured to accept untagged packets.

If a packet originates in C-VLAN 10 and needs to be sent over the S-VLAN, a tag with a VLAN ID 10 is added to the packet. If a packet is forwarded internally from the S-VLAN interface to xe-0/0/1.10, the tag with VLAN ID 10 is removed. The same principles apply to the C-VLANs configured on interface xe-0/0/2.





**NOTE:** Notice that you can use the same tag value for an S-VLAN and C-VLAN. For example, the configuration for customer 1 maps C-VLAN ID 10 to S-VLAN ID 10. C-VLAN and S-VLAN tags use separate name spaces, so this configuration is allowed.

Configuration for customer 1:

```
set vlans v10 interface xe-0/0/1.10
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 10 vlan-id-list 10-20
set interfaces xe-0/0/1 native-vlan-id 15
set interfaces xe-0/0/1 unit 10 input-vlan-map push
set interfaces xe-0/0/1 unit 10 output-vlan-map pop
```

Configuration for customer 2:

```
set vlans v30 interface xe-0/0/2.30
set interfaces xe-0/0/2 flexible-vlan-tagging
set interfaces xe-0/0/2 encapsulation extended-vlan-bridge
set interfaces xe-0/0/2 unit 30 vlan-id-list 30-40
set interfaces xe-0/0/2 unit 30 vlan-id-list 50-60
set interfaces xe-0/0/2 unit 30 vlan-id-list 70-80
set interfaces xe-0/0/2 native-vlan-id 75
set interfaces xe-0/0/2 unit 30 input-vlan-map push
set interfaces xe-0/0/2 unit 30 output-vlan-map pop
```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 1547](#)
- [Configuring All-in-One Bundling on page 1684](#)
- [Configuring a Specific Interface Mapping with VLAN ID Translation Option on page 1689](#)

## Configuring a Specific Interface Mapping with VLAN ID Translation Option

You can configure Q-in-Q tunneling by mapping packets from a specified C-VLAN to a specified S-VLAN. In addition, you can configure the system to replace a C-VLAN tag with an S-VLAN tag or replace an S-VLAN tag with a C-VLAN tag (instead of double tagging). This is called VLAN translation or VLAN rewriting. VLAN translation is particularly useful if a service provider's Layer 2 network that connects a customer's sites does not support double tagged packets.

When you use VLAN translation, both ends of the link normally must be able to swap the tags appropriately. That is, both ends of the link must be configured to swap the C-VLAN tag for the S-VLAN tag and swap the S-VLAN tag for the C-VLAN tag so that traffic in both directions is tagged appropriately while in transit and after arrival.

First configure the S-VLAN and its interface:

1. Create the S-VLAN and assign a VLAN ID for it.

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

2. Assign a logical interface to be a member of the S-VLAN. Do not use unit 0.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```

3. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# flexible-vlan-tagging
```

4. Enable the S-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

5. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

6. Bind the logical interface (unit) of the interface that you specified earlier to the VLAN ID for the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id number
```

For example, the following configuration creates S-VLAN v200, makes xe-1/1/1.200 a member of that VLAN, enables Q-in-Q tunneling on interface xe-1/1/1, enables xe-1/1/1 to accept untagged packets, and binds a logical interface of xe-1/1/1 to the VLAN ID of VLAN v200.

```
set vlans v200 vlan-id 200
```

```
set vlans v200 interface xe-1/1/1.200
```

```
set interfaces xe-1/1/1 flexible-vlan-tagging
```

```
set interfaces xe-1/1/1 set native-vlan-id 10
```

```
set interfaces xe-1/1/1 encapsulation extended-vlan-bridge
```

```
set interfaces xe-1/1/1 unit 200 vlan-id 200
```

Now configure a specific interface mapping with optional VLAN ID translation on the C-VLAN interface:

1. Assign a logical interface of the C-VLAN interface to be a member of the S-VLAN.

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name.unit-number
```

2. Enable the interface to transmit packets with 802.1Q VLAN tags:

```
[edit interfaces interface-name]
```

```
user@switch# flexible-vlan-tagging
```

3. Enable the C-VLAN interface to send and receive untagged packets:

```
[edit interfaces interface-name]
```

```
user@switch# set native-vlan-id vlan-id
```

4. Enable extended VLAN bridge encapsulation on the interface:

```
[edit interfaces interface-name]
```

```
user@switch# encapsulation extended-vlan-bridge
```

5. Configure a logical interface (unit) to receive and forward any tagged packet whose VLAN ID tag matches the VLAN IDs you specify:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# vlan-id number
```

6. Configure the system to remove the existing C-VLAN tag and replace it with the S-VLAN tag when packets ingress on the C-VLAN interface and are forwarded to the S-VLAN:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# input-vlan-map swap
```

7. Configure the system to remove the existing S-VLAN tag and replace it with the C-VLAN tag when packets are forwarded from the S-VLAN interface to the C-VLAN interface:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
user@switch# output-vlan-map swap
```

8. To configure an S-VLAN and associate it with the appropriate C-VLAN interface:

```
[edit vlans vlan-name]
```

```
user@switch# interface interface-name
```

For example, the following configuration on C-VLAN interface xe-0/0/1 enables Q-in-Q tunneling, enables xe-0/0/1 to accept untagged packets, and maps incoming packets from C-VLAN 150 to logical interface 200, which is a member of S-VLAN 200. Also, when packets egress from C-VLAN interface xe-0/0/1 and travel to the S-VLAN interface, the C-VLAN tag of 150 is removed and replaced with the S-VLAN tag of 200. When packets travel from the S-VLAN interface to the C-VLAN interface, the S-VLAN tag of 200 is removed and replaced with the C-VLAN tag of 150.

```
set vlans v200 interface xe-0/0/1.200
set interfaces xe-0/0/1 flexible-vlan-tagging
set interfaces xe-0/0/1 set native-vlan-id 10
set interfaces xe-0/0/1 encapsulation extended-vlan-bridge
set interfaces xe-0/0/1 unit 200 vlan-id 150
set interfaces xe-0/0/1 unit 200 output-vlan-map swap
set interfaces xe-0/0/1 unit 200 input-vlan-map swap
```

#### Related Documentation

- [Understanding Q-in-Q Tunneling on page 1547](#)
- [Configuring All-in-One Bundling on page 1684](#)
- [Configuring Many-to-Many Bundling on page 1686](#)

## Unified Forwarding Table Configuration Task

- [Configuring the Unified Forwarding Table on page 1697](#)

### Configuring the Unified Forwarding Table

To optimize the way your switch allocates memory for different types of addresses, you can choose a unified forwarding table profile. In addition to choosing this profile, you can also decide how you want memory allocated for longest prefix match (LPM) entries.

- [Configuring an Address-Storage Profile on page 1698](#)
- [Configuring the LPM Allocation on page 1699](#)

### Configuring an Address-Storage Profile

On QFX5100 and EX4600 switches, you can control the allocation of memory available to store the following:

- MAC addresses
- Layer 3 host entries
- Longest prefix match (LPM) table entries

You configure the mix that best meets your needs by choosing the appropriate profile. [Table 109 on page 1698](#) lists the profiles you can choose and the maximum values for the MAC address and host table entries.

Table 109: Unified Forwarding Table Profiles

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
		IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
<b>l2-profile-one</b>	288K	16K	8K	8K	8K	4K	4K
<b>l2-profile-two</b>	224K	80K	40K	40K	40K	20K	20K
<b>l2-profile-three (default)</b>	160K	144K	72K	72K	72K	36K	36K
<b>l3-profile</b>	96K	208K	104K	104K	104K	52K	52K
<b>lpm-profile*</b>	32K	16K	8K	8K	8K	4K	4K

Note that if the host table stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address. For more information about valid combinations of table entries see *Understanding the Unified Forwarding Table*.

To configure the profile that you want, enter and commit the following statement:

[edit]

```
user@switch# set chassis forwarding-options profile-name
```



**NOTE:** When you configure and commit a profile, the PFE process restarts and all the data interfaces on the switch go down and come back up.

The settings for **l2-profile-three** are configured by default. That is, if you do not enter a **set forwarding-options chassis *profile-name*** statement, these settings are configured.

## Configuring the LPM Allocation

In addition to choosing a profile, you can further optimize memory allocation for LPM table entries by configuring how many IPv6 prefixes in the range /65 through /127 you want the switch to store. The switch uses LPM entries during address lookup to match addresses to the most-specific (longest) applicable prefix. The procedures for configuring the LPM table are different depending on which version of Junos OS you are using.

- [Configuring the LPM Table With Junos OS 13.2X51-D10 and 13.2X52-D10 on page 1699](#)
- [Configuring the LPM Table With Junos OS 13.2x51-D15 on page 1700](#)

### Configuring the LPM Table With Junos OS 13.2X51-D10 and 13.2X52-D10

With Junos OS 13.2x51-D10 and 13.2X52-D10, the switch allocates memory for 16 IPv6 prefixes in the range /65 through /127 by default. If you want to use more than 16 IPv6 prefixes in this range, you must enter and commit the following statement:

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix [1-128]
```

Each increment adds support for 16 IPv6 prefixes between /65 and /127, for a maximum of 2048 such prefixes (16 x 128 = 2048). The system supports 16 of these prefixes by default, so to increase the number of supported prefixes, you must enter a value of 2 or greater. For example, if you enter 2, the system will support 32 IPv6 prefixes in the range /65 through /127.



**NOTE:** When you configure and commit the `num-65-127-prefix` value, all the data interfaces on the switch restart. The management interfaces are unaffected.

The LPM table is shared, and each increment that you add for IPv6 prefixes in the range /65 through /127 reduces the number of table entries that are available for IPv4 prefixes and IPv6 prefixes shorter than /65. Note that IPv6 prefixes /65 and longer consume twice as much memory as shorter IPv6 prefixes and four times as much memory as IPv4 prefixes. So, for example, entering the following statement

```
user@switch# set chassis forwarding-options l2-profile-one num-65-127-prefix 2
```

provides for 16 additional IPv6 prefixes /65 or longer (for a total of 32 such prefixes) and reduces the numbers of other prefixes that can be stored, as indicated:

- 32 fewer IPv6 prefixes shorter than /65 (16 IPv6 prefixes /65 or longer consume the same amount of memory as 32 IPv6 prefixes shorter than /65), or
- 64 fewer IPv4 prefixes (16 IPv6 prefixes /65 or longer consume the same amount of memory as 64 IPv4 prefixes)

[Table 110 on page 1700](#) provides examples of valid combinations that the LPM table can store using the **l2** and **l3** profiles. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

Table 110: Example LPM Table Combinations Using l2-and l3 Profiles With Junos OS 13.2X51-D10 and 13.2X52-D10

num-65-127-prefix Value	IPv4 Entries	IPv6 Entries (Prefix <= 64)	IPv6 Entries (Prefix >= 65)
1 (default)	16K-16	0K	16
1 (default)	0K	8K-16	16
1 (default)	8K-16	4K	16
64	4K	4K	1K
64	2K	5K	1K
64	0K	6K	1K
128	4K	2K	2K
128	2K	3K	2K
128	0K	4K	2K



**NOTE:** With Junos OS 13.2X51-D10 and 13.2X52-D10, the `lpm-profile` does not support IPv6 prefixes. If you use this version of Junos OS and also use the `lpm-profile`, do not configure the `num-65-127-prefix` statement. That is, leave it at its default value of 1, which allows for as many as 128K IPv4 prefixes (the maximum possible).

#### Configuring the LPM Table With Junos OS 13.2x51-D15

With Junos OS 13.2X51-D15, you can configure the memory allocation for the LPM table for the `lpm-profile` profile independently of the other profiles. In addition, Junos OS 13.2x51-D15 offers twice as much storage for IPv6 prefixes /65 through /127 (4K instead of 2K) for the `l2` and `l3` profiles.

- [Configuring the l2 and l3 profiles With Junos OS 13.2x51-D15 on page 1700](#)
- [Configuring The lpm-profile With Junos OS 13.2x51-D15 on page 1701](#)

#### Configuring the l2 and l3 profiles With Junos OS 13.2x51-D15

With Junos OS 13.2x51-D15, you can configure the switch to support as many as 4K IPv6 prefixes /65 through /127 if you are using any profile other than the `lpm-profile` profile. To do so, enter and commit the following statement:

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix [0-4]
```

Each increment adds support for 1K IPv6 prefixes between /65 and /127, for a maximum of 4K such prefixes. The default value is 1, which allocates memory for 1K of IPv6 prefixes in this range. Each increment that you add for IPv6 prefixes in the range /65 through /127 reduces the number of table entries that are available for IPv6 prefixes shorter than /65

and IPv4 prefixes. [Table 111 on page 1701](#) shows the numbers of entries that you can allocate by using the **num-65-127-prefix** statement with Junos OS 13.2X51-D15. Once again, each row represents a case in which the table is full and cannot accommodate any more entries.

**Table 111: LPM Table Combinations for l2 and l3 profiles With Junos OS 13.2X51-D15**

num-65-127-prefix Value	IPv4 Entries	IPv6 Entries (Prefix <= 64)	IPv6 Entries (Prefix >= 65)
0	16K	8K	0K
1 (default)	12K	6K	1K
2	8K	4K	2K
3	4K	2K	3K
4	0K	0K	4K



**NOTE:** When you configure the **num-65-127-prefix** value, the PFE process restarts and all the data interfaces on the switch go down and come back up. The management interfaces are unaffected.

#### **Configuring The *lpm-profile* With Junos OS 13.2x51-D15**

If you use the **lpm-profile** profile with Junos OS 13.2x51-D15, you can control whether the switch allocates any memory for IPv6 prefixes /65 through /127. By default, the switch supports the following with this profile:

- 128K IPv4 prefixes
- 16K IPv6 prefixes (all lengths)

You can disable support for IPv6 prefixes /65 through /127 with the **lpm-profile** profile so that there is more memory for IPv6 prefixes shorter than /65. To do so, enter and commit the following statement:

```
user@switch# set chassis forwarding-options profile-name prefix-65-127-disable
```

If you enter this statement, the switch allocates memory for the following:

- 128K IPv4 and IPv6 prefixes shorter than /65
- 0K IPv6 prefixes /65 through /127

For example, if you use the **prefix-65-127-disable** statement, each of the following combinations are valid:

- 100K IPv4 and 28K IPv6 /64 prefixes
- 64K IPv4 and 64K IPv6 /64 prefixes

- 128K IPv4 and 0K IPv6 /64 prefixes
- 0K IPv4 and 128K IPv6 /64 prefixes

**Related Documentation** • [Understanding the Unified Forwarding Table on page 1545](#)

---

## Forwarding Mode Configuration Task

- [Configuring the Forwarding Mode on page 1702](#)

### Configuring the Forwarding Mode

By default, packets are forwarded using store-and-forward mode. You can configure all the interfaces to use cut-through mode instead.

To enable cut-through switching mode, enter the following statement:

```
[edit forwarding-options]
user@switch# set cut-through
```

**Related Documentation** • [cut-through](#)

---

## Proxy ARP Configuration Task

- [Configuring Proxy ARP \(CLI Procedure\) on page 1702](#)

### Configuring Proxy ARP (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches and QFX3500 and QFX3600 switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see [Configuring Proxy ARP \(CLI Procedure\)](#) or [Configuring Proxy ARP](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

You can configure proxy Address Resolution Protocol (ARP) on your switch to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number proxy-arp (restricted |
unrestricted)
```



**BEST PRACTICE:** We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch does not act as a proxy if the source and target IP addresses are on the same subnet. If you decide to use unrestricted mode, disable gratuitous ARP requests on the interface to avoid



a situation wherein the switch's response to a gratuitous ARP request appears to the host to be an indication of an IP conflict.

To configure proxy ARP on an integrated routing and bridging (IRB) interface:

```
[edit interfaces]
user@switch# set irb.logical-unit-number proxy-arp restricted
```

#### Related Documentation

- [Example: Configuring Proxy ARP on an EX Series Switch](#)
- [Verifying That Proxy ARP Is Working Correctly on page 1834](#)
- [Configuring Integrated Routing and Bridging Interfaces \(CLI Procedure\)](#)

## Reflective Relay Configuration Tasks

- [Configuring Reflective Relay on page 1703](#)

### Configuring Reflective Relay

Configure reflective relay when a switch port must return packets on a downstream port. For example, configure reflective relay when a switch port receives aggregated virtual machine packets from a technology such as virtual Ethernet port aggregator (VEPA). When these packets are passed through the switch, reflective relay allows the switch to send those packets back on the same interface that was used for delivery.



**NOTE:** This task uses Junos OS for QFX3500 and QFX3600 switches that supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Reflective Relay*.

Before you begin configuring reflective relay, ensure that you have:

- Configured packet aggregation on the server connected to the port. See your server documentation.
- Configured the port for all VLANs that could be included in aggregated packets..

To configure reflective relay:

1. Configure an Ethernet interface with an interface mode of **trunk**:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type interface-mode trunk
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface for reflective relay:

```
[edit]
```

```
user@switch# set interfaces interface-name unit number family family-type reflective-relay
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching reflective-relay
```

3. Configure the interface for the VLANs that exist on the VM server:

```
[edit]
user@switch# set interfaces interface-name unit number family family-type vlan members
vlan-names
```

For example:

```
[edit]
user@switch# set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members
[VLAN_Purple VLAN_Orange VLAN_Blue]
```

#### Related Documentation

- [Example: Configuring Reflective Relay for Use with VEPA Technology on page 1605](#)
- [Understanding Reflective Relay for Use with VEPA Technology on page 1553](#)

---

## STP Configuration Tasks

- [Configuring STP on page 1704](#)
- [Configuring VLAN Spanning-Tree Protocol on page 1705](#)
- [Unblocking an Interface That Receives BPDUs in Error on page 1708](#)

### Configuring STP

The default spanning-tree protocol on the device is Rapid Spanning Tree Protocol (RSTP). RSTP provides faster convergence times than Spanning Tree Protocol (STP) does. However, some legacy networks require the slower convergence times of basic STP.

If your network includes 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. When you explicitly configure STP, the device uses the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP.

To configure STP using the CLI:

1. Delete the RSTP configuration on the interface (here, the interface is **xe-0/0/5**):

```
[edit]
user@switch# delete protocols rstp interface xe-0/0/5
```

2. Configure STP on the interface:

```
[edit]
user@switch# set protocols stp interface xe-0/0/5
```

3. Commit the configuration:

```
[edit]
user@switch# commit
```

#### Related Documentation

- [show spanning-tree bridge on page 1870](#)
- [show spanning-tree interface on page 1875](#)

- [Overview of Spanning-Tree Protocols on page 1554](#)

## Configuring VLAN Spanning-Tree Protocol

You can configure the VLAN Spanning-Tree Protocol (VSTP) under the following hierarchy levels:



**NOTE:** This task supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring VLAN Spanning Tree Protocol*.

- [edit *logical-systems logical-system-name protocols*]
- [edit *logical-systems logical-system-name routing-instances routing-instance-name protocols*]
- [edit *protocols*]
- [edit *routing-instances routing-instance-name protocols*]

The routing instance type can be either **virtual-switch** or **layer2-control**.

To configure the VLAN Spanning-Tree Protocol:

1. Enable VSTP as the version of spanning-tree protocol to be configured:

```
[edit]
```

```
user@host@ edit ... protocols (STP Type) vstp
```

2. (Optional) For compatibility with older bridges that do not support VSTP, you can run force VSTP to run as the original IEEE 802.1D Spanning-Tree Protocol (STP) version:

```
[edit ... protocols vstp]
```

```
user@host# set force-version stp
```



**NOTE:** If VSTP has been forced to run as the original STP version, you can revert back to VSTP by first removing the force-version statement from the configuration and then entering the *clear spanning-tree protocol-migration* configuration mode command.

3. Configure the interfaces that participate in the VSTP instance.

- a. Enable configuration of the interface:

```
[edit ... protocols vstp]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols vstp interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols vstp interface interface-name]
user@host# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols vstp interface interface-name]
user@host# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols vstp interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port.

You can also enable BPDU root protection for all spanning-tree protocol instances on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see *Checking the Status of Spanning-Tree Instance Interfaces*.

4. Enable configuration of a VLAN instance:

```
[edit ... protocols vstp]
user@host# edit vlan vlan-id
```

5. Configure the bridge priority

```
[edit ... protocols vstp vlan vlan-id]
user@host# set bridge-priority bridge-priority
```

For more information, see *Bridge Priority for Election of Root Bridge and Designated Bridge*.

## 6. Configure hello BPDU timers.

- a. Configure the maximum expected arrival time of hello BPDUs:

```
[edit ... protocols vstp vlan vlan-id]
user@host# set max-age seconds
```

- b. Configure the time interval at which the root bridge transmits configuration BPDUs:

```
[edit ... protocols vstp vlan vlan-id]
user@host# set hello-time seconds
```

7. (Optional) By default, the bridge port remains in the listening and learning states for 15 seconds before transitioning to the forwarding state. You can specify a delay from 4 through 20 seconds instead:

```
[edit ... protocols vstp vlan vlan-id]
user@host# set forward-delay seconds
```

8. Configure the interfaces that participate in the VSTP instance.

- a. Enable configuration of the interface:

```
[edit ... protocols vstp vlan vlan-id]
user@host# edit interface interface-name
```

- b. Configure the interface priority:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set priority interface-priority
```

- c. (Optional) By default, the interface link cost is determined by the link speed. You can configure the interface link cost to control which bridge is the designated bridge and which port is the designated port:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set cost interface-link-cost
```

- d. Configure the interface link mode to identify point-to-point links:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set mode (p2p | shared)
```

Specify **p2p** if the link is point to point. Specify **shared** if the link is a shared media.

- e. (Optional) Configure the interface as an edge port:

```
[edit ... protocols vstp vlan vlan-id interface interface-name]
user@host# set edge
```

Edge ports do not expect to receive bridge protocol data unit (BPDU) packets. If a BPDU packet is received for an edge port, the port becomes a nonedge port

You can also enable BPDU root protection for all spanning-tree protocol instances on the interface. BPDU root protect ensures the port is the spanning-tree designated port. If the port receives superior BPDU packets, root protect moves this port to a root-prevented spanning-tree state. For configuration details, see *Checking the Status of Spanning-Tree Instance Interfaces*.

9. Verify the VSTP configuration:

```
[edit]
... { # Optional logical system and/or routing instance
  protocols (STP Type) {
    vstp {
```

```

force-version stp; # Optional.
interface interface-name {
    priority interface-priority;
    cost interface-link-cost; # Optional.
    mode (p2p | shared);
    edge; # Optional.
}
vlan vlan-id {
    bridge-priority bridge-priority;
    max-age seconds;
    hello-time seconds;
    forward-delay seconds; # Optional.
    interface interface-name {
        priority interface-priority;
        cost interface-link-cost; # Optional.
        mode (p2p | shared);
        edge; # Optional.
    }
}
}
}
}

```

#### Related Documentation

- *Spanning-Tree Protocols Supported*
- *RSTP or VSTP Forced to Run as IEEE 802.1D STP*
- *Reverting to RSTP or VSTP from Forced IEEE 802.1D STP*
- *VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Overview*
- *VPLS Multihomed Layer 2 Ring and MPLS Infrastructure Topology*

## Unblocking an Interface That Receives BPDUs in Error



**NOTE:** BPDU block protection is disabled on Node devices.

Devices use bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could trigger a spanning-tree misconfiguration. If BPDUs are received on a BPDU-protected interface, the interface transitions to a blocking state and stops forwarding frames.

After you fix the misconfiguration that triggered the sending of BPDUs to an interface, you can unblock the interface and return it to service.



**NOTE:** This task describes how to use both the original CLI and the Enhanced Layer 2 Software (ELS) CLI. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

To unblock an interface after fixing the misconfiguration that triggered the BPDUs and return it to service:

- (Original CLI) Automatically unblock an interface by configuring a timer that expires (here, the interface is **xe-0/0/6**):

```
[edit ethernet-switching-options]
user@switch# set bpd-blockdisable-timeout 30 interface xe-0/0/6
```

- (ELS CLI) Automatically unblock an interface by configuring a timer that expires (here, the interface is **xe-0/0/6**):

```
[edit protocols layer2-control]
user@switch# set bpd-blockdisable-timeout 30 interface xe-0/0/6
```

- Manually unblock an interface using the operational mode command:

```
user@switch> clear ethernet-switching bpd-error interface xe-0/0/6
```

**Related  
Documentation**

- [Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656](#)
- [Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558](#)

---

## Protocols Configuration Statement

- [protocols on page 1710](#)

## protocols

```
Syntax protocols {
    bgp {
        disable;
        accept-remote-nexthop;
        advertise-external <conditional>;
        advertise-inactive;
        (advertise-peer-as | no-advertise-peer-as);
        authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
        authentication-key key;
        authentication-key-chain key-chain;
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            hold-down-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            session-mode (automatic | multihop | single-hop);
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        cluster cluster-identifier;
        damping;
        description text-description;
        export [ policy-names ];
        family family-name {
            ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
        }
        graceful-restart {
            disable;
            restart-time seconds;
            stale-routes-time seconds;
        }
        group group-name {
            ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
        }
        hold-time seconds;
        import [ policy-names ];
        include-mp-next-hop;
        keep (all | none);
        local-address address;
```



```

local-as autonomous-system <loops number> < alias> <private>;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tll-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
dcbx {
    disable;
    interface (interface-name | all) {
        disable;
        application-map application-map-name;
        applications {
            no-auto-negotiation;
        }
        enhanced-transmission-selection {
            no-auto-negotiation;
            no-recommendation-tlv;
            recommendation-tlv {
                no-auto-negotiation;
            }
        }
    }
}

```

```
    }
    dcbx-version (auto-negotiate | ieee-dcbx | dcbx-version-1.01);
    priority-flow-control {
        no-auto-negotiation;
    }
}
}
iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address {
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (Liveness Detection) (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
            <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
igmp-snooping {
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
vlan vlan-name {
    disable;
}
interface interface-name {
    group-limit limit;
    multicast-router-interface;
    static {
        group ip-address;
    }
}
```

```

        robust-count number;
    }
}
isis {
    disable;
    export [ policy-names ];
    ignore-attached-bit;
    interface interface-name {
        bfd-liveness-detection {
            authentication {
                algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        checksum;
        csnp-interval (seconds | disable);
        disable;
        hello-padding (adaptive | loose | strict);
        level (1 | 2) {
            disable;
            hello-authentication-key key;
            hello-authentication-type authentication;
            hello-interval seconds;
            hold-time seconds;
            ipv4-multicast-metric number;
            metric metric;
            passive;
            priority number;
        }
        lsp-interval milliseconds;
        mesh-group (value | blocked);
        no-ipv4-multicast;
        no-unicast-topology;
        passive;
        point-to-point;
    }
    level (1 | 2) {
        disable;
        authentication-key key;
        authentication-type authentication;
        external-preference preference;
        no-csnp-authentication;
    }
}

```

```
no-hello-authentication;
no-psnp-authentication;
preference preference;
prefix-export-limit number;
wide-metrics-only;
}
loose-authentication-check;
lsp-lifetime seconds;
max-areas number;
no-adjacency-holddown;
no-authentication-check;
no-ipv4-routing;
overload {
    advertise-high-metrics;
    timeout seconds;
}
reference-bandwidth reference-bandwidth;
rib-group {
    inet group-name;
}
topologies {
    ipv4-multicast;
}
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    disable;
    family inet {
        shortcuts {
            multicast-rpf-routes;
        }
    }
}
}
lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (LLDP) (all | interface-name) {
        disable;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable> <match
            regex>;
        flag flag (detail | disable | receive | send);
    }
}
}
mstp {
    disable;
    bpdu-timeout-action;
    bridge-priority priority;
    configuration-name (MSTP) name;
    forward-delay seconds;
```

```

hello-time seconds;
interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
        block;
        alarm;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
max-age seconds;
max-hops hops;
msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
        disable;
        cost cost;
        edge;
        mode mode;
        priority priority;
    }
}
revision-level revision-level;
traceoptions {
    file filename <files number > <size size > <no-stamp | world-readable |
    no-world-readable>;
    flag flag;
}
}
ospf {
    disable;
    area area-id {
        area-range ip-prefix </prefix-length> <exact> <override-metric metric > <restrict>;
        context-identifier identifier
        interface interface-name {
            disable;
            authentication {
                md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
                simple-password key-string;
            }
            bandwidth-based-metrics {
                bandwidth value metric number;
            }
            bfd-liveness-detection {
                authentication {
                    algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
                    meticulous-keyed-sha-1 | simple-password);
                    key-chain key-chain-name;
                    loose-check;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
        }
    }
}

```

```

    full-neighbors-only;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
dead-interval seconds;
dynamic-neighbors;
flood-reduction;
hello-interval seconds;
interface-type (nbma | p2mp | p2p);
metric metric;
neighbor address <eligible>;
no-eligible-backup;
no-interface-state-traps;
no-neighbor-down-notification;
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
poll-interval seconds;
priority number;
retransmit-interval seconds;
secondary;
te-metric metric;
topology (name | default | ipv4-multicast) {
    disable;
    bandwidth-based-metrics {
        bandwidth value;
        metric number;
    }
    metric metric;
}
transit-delay seconds;
}
network-summary-export [ policy-names ];
network-summary-import [ policy-names ];
nssa {
    area-range ip-prefix</prefix-length> <exact> <override-metric metric> <restrict>;
    default-lsa {
        default-metric metric;
        metric-type type;
        type-7;
    }
    (summaries | no-summaries);
}
stub <default-metric metric> <summaries | no-summaries>;
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication {

```

```

        md5 key-id key key-string <start-time YYYY-MM-DD.hh:mm>;
        simple-password key-string;
    }
    dead-interval seconds;
    demand-circuit;
    flood-reduction;
    hello-interval seconds;
    ipsec-sa sa-name;
    no-neighbor-down-notification;
    retransmit-interval seconds;
    topology (name | default | ipv4-multicast) {
        disable;
        metric metric;
    }
    transit-delay seconds;
}
}
database-protection {
    ignore-count number;
    ignore-time seconds;
    maximum-lsa number;
    reset-time seconds;
    warning-only;
    warning-threshold percent;
}
export [ policy-names ];
external-preference preference;
graceful-restart {
    disable;
    helper-disable <both | restart-signaling | standard>;
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
}
import [ policy-names ];
no-nssa-abr;
no-rfc-1583;
overload <timeout seconds>;
preference preference;
prefix-export-limit number;
reference-bandwidth reference-bandwidth;
rib-group group-name;
topology (default | ipv4-multicast | name) {
    overload;
    prefix-export-limit number;
    topology-id number;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
traffic-engineering {
    advertise-unnumbered-interfaces;
    credibility-protocol-preference;
    ignore-lsp-metrics;
}

```

```
    multicast-rpf-routes;
    no-topology;
    shortcuts <lsp-metric-into-summary>;
  }
}
pim {
  disable;
  assert-timeout seconds;
  dense-groups {
    addresses;
  }
  dr-election-on-p2p;
  export;
  family (inet | inet6) {
    disable;
  }
  graceful-restart {
    disable;
    restart-duration seconds;
  }
  import [ policy-names ];
  interface interface-name {
    accept-remote-source;
    disable;
    family (inet | inet6) {
      disable;
    }
    hello-interval seconds;
    mode (dense | sparse | sparse-dense);
    neighbor-policy [ policy-names ];
    override-interval milliseconds;
    priority number;
    propagation-delay milliseconds;
    reset-tracking-bit;
    version version;
  }
  join-load-balance;
  join-prune-timeout;
  nonstop-routing;
  override-interval milliseconds;
  propagation-delay milliseconds;
  reset-tracking-bit;
  rib-group group-name;
  rp {
    auto-rp {
      (announce | discovery | mapping);
      (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
      family (inet | inet6) {
        export [ policy-names ];
        import [ policy-names ];
        priority number;
      }
    }
  }
  bootstrap-import [ policy-names ];
}
```



```

bootstrap-export [ policy-names ];
bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix</prefix-length>;
        }
    }
}
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}

```

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
tunnel-devices [ mt-fpc/pic/port ];
}
rip {
  authentication-key password;
  authentication-type type;
  (check-zero | no-check-zero);
  group group-name {
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
  }
  export [ policy-names ];
  import [ policy-names ];
  metric-out metric;
  neighbor neighbor-name {
    any-sender;
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
      ... same statements as at the [edit protocols rip group group-name
        bfd-liveness-detection] hierarchy level ...
    }
    (check-zero | no-check-zero);
    import [ policy-names ];
    message-size number;
    metric-in metric;
    receive (both | none | version-1 | version-2);
    route-timeout seconds;
    send (broadcast | multicast | none | version-1);
    update-interval seconds;
  }
  preference preference;
  route-timeout seconds;
  update-interval seconds;
}
holddown seconds;
```

```

import [ policy-names ];
message-size number;
metric-in metric;
receive (both | none | version-1 | version-2);
rib-group group-name;
route-timeout seconds;
send (broadcast | multicast | none | version-1);
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
update-interval seconds;
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
}
traceoptions {
    file filename <files number> <size size> <no-stamp> <world-readable |
        no-world-readable>;
    flag flag;
}
}
stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            alarm;
            block;
        }
        cost cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
}

```

```

    }
    max-age seconds;
  }
  traceoptions {
    file filename <files number> <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
  uplink-failure-detection {
    group group-name {
      link-to-monitor interface-name;
      link-to-disable interface-name;
    }
  }
}
vstp {
  bpdu-block-on-edge;
  disable (Spanning Trees);
  force-version (Spanning Trees) stp;
  vlan (Spanning Trees) vlan-id {
    bridge-priority (Spanning Trees) priority;
    forward-delay (Spanning Trees) seconds;
    hello-time (Spanning Trees) seconds;
    interface (Spanning Trees) (all | interface-name) {
      bpdu-timeout-action (Spanning Trees) {
        block (Spanning Trees);
        log (Spanning Trees);
      }
      cost (Spanning Trees) cost;
      disable (Spanning Trees);
      edge (Spanning Trees);
      mode (Spanning Trees) mode;
      no-root-port (Spanning Trees);
      priority (Spanning Trees) priority;
    }
    max-age (Spanning Trees) seconds;
    traceoptions (Spanning Trees) {
      file filename <files number> <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
}
}
}

```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure protocols.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • [Junos OS Routing Protocols Configuration Guide](#)

## Unified Forwarding Table Configuration Statements

- [forwarding-options \(chassis\) on page 1724](#)
- [num-65-127-prefix on page 1725](#)
- [prefix-65-127-disable on page 1725](#)

## forwarding-options (chassis)

**Syntax** forwarding options *profile-name* {  
     num-65-127-prefix *value*  
     lpm-profile *prefix-65-127-disable*  
 }

**Hierarchy Level** [edit *chassis*]

**Release Information** Statement introduced in Junos 13.2 for the QFX Series.

**Description** Configure a unified forwarding table profile to allocate the amount a memory available for the following:

- MAC addresses
- Layer 3 host entries
- Longest prefix match table entries

**Options** *profile-name*—name of the profile to use for memory allocation in the unified forwarding table. [Table 112 on page 1724](#) lists the profiles you can choose and the associated values for each type of entry.

**Table 112: Unified Forwarding Table Profiles**

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
l2-profile-one	288K	16K	8K	8K	8K	4K	4K
l2-profile-two	224K	80K	40K	40K	40K	20K	20K
l2-profile-three (default)	160K	144K	72K	72K	72K	36K	36K
l3-profile	96K	208K	104K	104K	104K	52K	52K
lpm-profile*	32K	16K	8K	8K	8K	4K	4K

\* This profile supports only IPv4 in Junos OS 13.2X51-D10. With Junos OS 13.2X51-D15 it supports IPv4 and IPv6.

Note that if the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see [“Understanding the Unified Forwarding Table” on page 1545](#).

You configure the memory allocation for LPM table entries differently depending on whether you use Junos OS 13.2X51-D10 or Junos OS 13.2X51-D15 and later. To learn

how to configure memory allocation for LPM table entries see [“Configuring the Unified Forwarding Table” on page 1697](#).

**Required Privilege  
Level**

- Related  
Documentation**
- [Understanding the Unified Forwarding Table on page 1545](#)
  - [Configuring the Unified Forwarding Table on page 1697](#)

## num-65-127-prefix

---

**Syntax** num-65-127-prefix *value*

**Hierarchy Level** [edit [chassis forwarding-options](#) *profile-name*]

**Release Information** Statement introduced in Junos 13.2 for the QFX Series.

**Description** Configure the number of supported IPv6 prefixes in the range /65 through /127.

- Options**
- value**—With Junos OS 13.2X51D10: Value in the range 1 through 128. Each increment adds support for 16 IPv6 addresses with prefixes between /65 and /127, for a maximum of 2048 such addresses (16 x 128 = 2048).
- value**—With Junos OS 13.2X51D15: Value in the range 0 through 4. Each increment adds support for 1K IPv6 addresses with prefixes between /65 and /127, for a maximum of 4K such addresses.

**Required Privilege  
Level**

- Related  
Documentation**
- [Configuring the Unified Forwarding Table on page 1697](#)

## prefix-65-127-disable

---

**Syntax** prefix-65-127-disable

**Hierarchy Level** [edit [chassis forwarding-options](#) lpm-profile]

**Release Information** Statement introduced in Junos 13.2X51-D15 for the QFX Series.

**Description** Disable support in the longest prefix match (LPM) table for IPv6 prefixes in the range /65 through /127.

**Required Privilege  
Level**

- Related  
Documentation**
- [Configuring the Unified Forwarding Table on page 1697](#)

## Reflective Relay Configuration Statements

---

- [reflective-relay](#) on page 1726

### reflective-relay

---

<b>Syntax</b>	reflective-relay;
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family ethernet-switching]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure a switch interface to return packets back to a device on the same interface that was used to deliver the packets.
<b>Default</b>	Switch interfaces are not configured for reflective relay.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Reflective Relay for Use with VEPA Technology</i></li><li>• <i>Configuring Reflective Relay</i></li></ul>

## STP Configuration Statements

---

- [alarm \(STP\)](#) on page 1727
- [block](#) on page 1728
- [bpdu-block](#) on page 1729
- [bpdu-block-on-edge](#) on page 1730
- [bpdu-timeout-action](#) on page 1731
- [bridge-priority](#) on page 1732
- [configuration-name \(MSTP\)](#) on page 1733
- [cost \(STP\)](#) on page 1734
- [disable \(STP\)](#) on page 1735
- [disable-timeout \(BPDU\)](#) on page 1736
- [edge \(STP\)](#) on page 1737
- [forward-delay](#) on page 1738
- [force-version](#) on page 1739
- [hello-time](#) on page 1740
- [interface \(Spanning Trees\)](#) on page 1741
- [interface \(BPDU\)](#) on page 1742
- [interface \(STP\)](#) on page 1743



- [max-age](#) on page 1744
- [max-hops](#) on page 1745
- [mode \(STP\)](#) on page 1746
- [msti](#) on page 1747
- [mstp](#) on page 1748
- [no-root-port](#) on page 1749
- [priority \(STP\)](#) on page 1750
- [revision-level](#) on page 1751
- [rstp](#) on page 1752
- [stp](#) on page 1753
- [traceoptions \(STP\)](#) on page 1754
- [vlan \(STP\)](#) on page 1758
- [vstp](#) on page 1759

## alarm (STP)


<b>Syntax</b>	<code>alarm;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ], [edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ], [edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ], [edit protocols <a href="#">vstp vlan <i>vlan-id</i> interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For interfaces configured for loop protection, configure the software to generate a message to be sent to the system log file to record the loop-protection event.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP</a> on page 1624</li> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP</a> on page 1610</li> <li>• <a href="#">Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree</a> on page 1660</li> <li>• <a href="#">Understanding Loop Protection for STP, RSTP, VSTP, and MSTP</a> on page 1559</li> <li>• <a href="#">Understanding VSTP</a> on page 1557</li> <li>• <a href="#">show spanning-tree bridge</a> on page 1870</li> <li>• <a href="#">show spanning-tree interface</a></li> </ul>

## block

---

<b>Syntax</b>	block;
<b>Hierarchy Level</b>	[edit protocols mstp (Spanning Trees) <a href="#">interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ], [edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ], [edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ], [edit protocols <a href="#">vstp vlan</a> <i>vlan-id</i> <a href="#">interface</a> (all   <i>interface-name</i> ) <a href="#">bpdu-timeout-action</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure loop protection on a specific interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660</a></li><li>• <a href="#">Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1559</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul>

## bpdu-block

<b>Syntax</b>	<pre>bpdu-block {   interface (all   [<i>interface-name</i>]);   disable-timeout <i>timeout</i>; }</pre>
<b>Hierarchy Level</b>	<ul style="list-style-type: none"> <li>For platforms with ELS CLI: [edit protocols layer2-control]</li> <li>For platforms with Original CLI: [edit ethernet-switching-options]</li> </ul>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure BPDU protection on an interface. If the interface receives BPDUs, it is disabled.
<div>  <b>NOTE:</b> BPDU block protection is disabled on Node devices. </div> <p>The statements are explained separately.</p>	
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li><a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li><a href="#">Unblocking an Interface That Receives BPDUs in Error on page 1708</a></li> <li><a href="#">clear ethernet-switching bpdu-error on page 1837</a></li> <li><a href="#">show spanning-tree bridge on page 1870</a></li> <li><a href="#">show spanning-tree interface</a></li> </ul>

## bpdu-block-on-edge

---

<b>Syntax</b>	bpdu-block-on-edge;
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ], [edit protocols <a href="#">rstp</a> ], [edit protocols <a href="#">vstp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">clear ethernet-switching bpdu-error on page 1837</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul>

## bpdu-timeout-action

<b>Syntax</b>	bpdu-timeout-action { alarm; block; }
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">vstp vlan <i>vlan-id</i> interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the BPDU timeout action on a specific interface. You must configure at least one action ( <b>alarm</b> , <b>block</b> , or both).
	The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li>• <a href="#">Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on page 1660</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on page 1559</a></li> <li>• <a href="#">Understanding VSTP on page 1557</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface</a></li> </ul>

## bridge-priority

---

<b>Syntax</b>	<code>bridge-priority <i>priority</i>;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ], [edit protocols <a href="#">mstp</a> <i>msti-id</i> ], [edit protocols <a href="#">rstp</a> ], [edit protocols <a href="#">stp</a> ], [edit protocols <a href="#">vstp</a> <i>vlan</i> <i>vlan-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.
<b>Options</b>	<b><i>priority</i></b> —Bridge priority. It can be set only in increments of 4096. <b>Range:</b> 0 through 61,440 <b>Default:</b> 32,768
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul>

---

## configuration-name (MSTP)

---

<b>Syntax</b>	configuration-name <i>configuration-name</i> ;
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the configuration name. The configuration name is the MSTP region name carried in the MSTP BPDUs.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <i>show spanning-tree interface</i></li></ul>

## cost (STP)

---

<b>Syntax</b>	<code>cost cost;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">mstp msti msti-id interface interface-name</a> ], [edit protocols rstp (Spanning Trees) <a href="#">interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">vstp vlan vlan-id interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.
<b>Default</b>	Link cost is determined by the link speed.
<b>Options</b>	<b>cost</b> —Link cost associated with the port. <b>Range:</b> 1 through 200,000,000
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul>



## disable (STP)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ], [edit protocols <a href="#">mstp</a> <a href="#">interface</a> <i>interface-name</i> ], [edit protocols <a href="#">mstp</a> <a href="#">msti</a> <i>msti-id</i> <a href="#">vlan</a> ( <i>vlan-id</i>   <i>vlan-name</i> ) <a href="#">interface</a> <i>interface-name</i> ], [edit protocols <a href="#">rstp</a> ], [edit protocols <a href="#">rstp</a> <a href="#">interface</a> <i>interface-name</i> ], [edit protocols <a href="#">stp</a> ], [edit protocols <a href="#">stp</a> <a href="#">interface</a> <i>interface-name</i> ], [edit protocols <a href="#">vstp</a> ], [edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> <a href="#">interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Disable STP, MSTP, RSTP, or VSTP on the switch or on a specific interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li>• <a href="#">Understanding MSTP on page 1555</a></li> <li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li> <li>• <a href="#">Understanding VSTP on page 1557</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface</a></li> </ul>

## disable-timeout (BPDU)

---

<b>Syntax</b>	<code>disable-timeout <i>timeout</i>;</code>
<b>Hierarchy Level</b>	[edit ethernet-switching-options <a href="#">bpdu-block</a> ] [edit protocols layer2-control <a href="#">bpdu-block</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For interfaces configured for BPDU protection, specify the amount of time an interface receiving BPDUs is disabled.
<b>Default</b>	The disable timeout is not enabled.
<b>Options</b>	<b>timeout:</b> Length of time, in seconds, that the interface receiving BPDUs is disabled. Once the timeout expires, the interface is brought back into service. <b>Range:</b> 10 through 3600 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656</a></li><li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface on page 1875</a></li></ul>

## edge (STP)

<b>Syntax</b>	edge;
<b>Hierarchy Level</b>	[edit protocols <b>mstp interface</b> (all   <i>interface-name</i> )], [edit protocols <b>mstp msti</b> <i>msti-id interface interface-name</i> ], [edit protocols <b>rstp interface</b> (all   <i>interface-name</i> )], [edit protocols <b>stp interface</b> (all   <i>interface-name</i> )], [edit protocols <b>vstp vlan</b> <i>vlan-id interface</i> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure interfaces as edge interfaces. Edge interfaces immediately transition to a forwarding state.
<b>Default</b>	Edge interfaces are not enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li>• <a href="#">Understanding MSTP on page 1555</a></li> <li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li> <li>• <a href="#">Understanding VSTP on page 1557</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface</a></li> </ul>

## forward-delay

---

<b>Syntax</b>	<code>forward-delay <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit protocols <i>mstp</i>],</code> <code>[edit protocols <i>rstp</i>],</code> <code>[edit protocols <i>stp</i>],</code> <code>[edit protocols <i>vstp</i> <i>vlan</i> <i>vlan-id</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify how long a bridge interface remains in the listening and learning states before transitioning to the forwarding state.
<b>Options</b>	<b><i>seconds</i></b> —Number of seconds the bridge interface remains in the listening and learning states. <b>Range:</b> 4 through 30 seconds <b>Default:</b> 15 seconds
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul>

---

## force-version

---

<b>Syntax</b>	force-version stp;
<b>Hierarchy Level</b>	[edit protocols <a href="#">vstp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Force VLAN Spanning Tree Protocol (VSTP) to use the STP protocol instead of the default protocol, RSTP.
<b>Options</b>	<b>stp</b> —Spanning Tree Protocol
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <i>show spanning-tree interface</i></li><li>• <a href="#">Understanding VSTP on page 1557</a></li></ul>

## hello-time

---

<b>Syntax</b>	<code>hello-time seconds;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ], [edit protocols <a href="#">rstp</a> ], [edit protocols <a href="#">rstp</a> ], [edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the time interval at which the root bridge transmits configuration BPDUs.
<b>Options</b>	<b>seconds</b> —Number of seconds between transmissions of configuration BPDUs. <b>Range:</b> 1 through 10 seconds <b>Default:</b> 2 seconds
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface</a></li></ul>

## interface (Spanning Trees)

<b>Syntax</b>	<pre> interface <i>interface-name</i> {   arp-on-stp;   bpdu-timeout-action     block;     log;   cost <i>cost</i>;   disable;   edge;   mode <i>mode</i>;   no-root-port;   priority <i>priority</i>; }</pre>
<b>Hierarchy Level</b>	<pre> [edit protocols mstp], [edit protocols mstp msti <i>msti-id</i>], [edit protocols rstp], [edit protocols stp], [edit protocols vstp vlan (all   <i>vlan-id</i>   <i>vlan-name</i>)]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement updated in Junos OS Release 9.4 for EX Series switches to add VSTP support.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure an interface.</p> <p>The <b>edge</b>, <b>mode</b>, and <b>no-root-port</b> options are not available at the <code>[edit protocols mstp msti <i>msti-id</i>]</code> hierarchy level.</p>
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of an interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge</a></li> <li>• <a href="#">show spanning-tree interface</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on EX Series Switches</a></li> <li>• <a href="#">Example: Faster Convergence and Improved Network Stability with RSTP on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li>• <a href="#">Configuring VSTP (CLI Procedure)</a></li> </ul>

- [show spanning-tree bridge on page 1870](#)

## interface (BPDU)

---

<b>Syntax</b>	<code>interface (all   <i>interface-name</i>);</code>
<b>Hierarchy Level</b>	[edit ethernet-switching-options <a href="#">bpdu-block</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Apply BPDU protection to all interfaces or one or more interfaces.
<b>Options</b>	<b>all</b> —All interfaces.  <b><i>interface-name</i></b> —Name of the interface.
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface on page 1875</a></li></ul>



## interface (STP)

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     disable;     cost <i>cost</i>;     edge;     mode <i>mode</i>;     no-root-port;     priority <i>priority</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit protocols <a href="#">mstp</a>], [edit protocols <a href="#">mstp msti</a>], [edit protocols <a href="#">rstp</a>], [edit protocols <a href="#">stp</a>], [edit protocols <a href="#">vstp vlan</a> <i>vlan-id</i>]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure an interface.
<b>Options</b>	<p><i>interface-name</i>—Name of a Gigabit Ethernet interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li>• <a href="#">Understanding RSTP on page 1556</a></li> <li>• <a href="#">Understanding MSTP on page 1555</a></li> <li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li> <li>• <a href="#">Understanding VSTP on page 1557</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface</a></li> </ul>

## max-age

---

<b>Syntax</b>	<code>max-age seconds;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ], [edit protocols <a href="#">rstp</a> ], [edit protocols <a href="#">stp</a> ], [edit protocols <a href="#">vstp</a> <a href="#">vlan</a> <i>vlan-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the maximum age of received protocol BPDUs.
<b>Options</b>	<b>seconds</b> —Maximum age of received protocol BPDUs. <b>Range:</b> 6 through 40 seconds <b>Default:</b> 20 seconds
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface on page 1875</a></li></ul>

---

## max-hops

---

<b>Syntax</b>	<code>max-hops hops;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Multiple Spanning Tree Protocol (MSTP), configure the maximum number of hops that a BPDU can be forwarded in the MSTP region.
<b>Options</b>	<p><i>hops</i> — Number of hops the BPDU can be forwarded.</p> <p><b>Range:</b> 1 through 255 hops</p> <p><b>Default:</b> 20 hops</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface on page 1875</a></li></ul>

## mode (STP)

---

<b>Syntax</b>	<code>mode mode;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">mstp msti msti-id interface interface-name</a> ], [edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">vstp vlan vlan-id interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link mode to identify point-to-point links.
<b>Default</b>	For a full-duplex link, the default link mode is <b>point-to-point</b> . For a half-duplex link, the default link mode is <b>shared</b> .
<b>Options</b>	<i>mode</i> —Link mode: <ul style="list-style-type: none"><li>• <b>point-to-point</b>—Link is point to point.</li><li>• <b>shared</b>—Link is shared media.</li></ul>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface on page 1875</a></li></ul>

## msti

<b>Syntax</b>	<pre> msti <i>msti-id</i> {   vlan (<i>vlan-id</i>   <i>vlan-name</i>);   interface <i>interface-name</i> {     disable;     cost <i>cost</i>;     edge;     mode <i>mode</i>;     priority <i>priority</i>;   } } </pre>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the Multiple Spanning Tree Instance (MSTI) identifier for Multiple Spanning Tree Protocol (MSTP). MSTI IDs are local to each region, so you can reuse the same MSTI ID in different regions.
<b>Default</b>	MSTI is disabled.
<b>Options</b>	<p><i>msti-id</i> —MSTI identifier.</p> <p><b>Range:</b> 1 through 4094. The Common Instance Spanning Tree (CIST) is always MSTI 0.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface on page 1875</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Understanding MSTP on page 1555</a></li> </ul>

## mstp

---

**Syntax**    mstp {  
              disable;  
              bpdu-timeout-action;  
              bridge-priority *priority*;  
              configuration-name (MSTP) *name*;  
              forward-delay *seconds*;  
              hello-time *seconds*;  
              interface (all | *interface-name*) {  
                  bpdu-timeout-action {  
                      block;  
                      alarm;  
                  }  
                  disable;  
                  cost *cost*;  
                  edge;  
                  mode *mode*;  
                  no-root-port;  
                  priority *priority*;  
              }  
              max-age *seconds*;  
              max-hops *hops*;  
              msti *msti-id* {  
                  vlan (*vlan-id* | *vlan-name*);  
                  interface *interface-name* {  
                      disable;  
                      cost *cost*;  
                      edge;  
                      mode *mode*;  
                      priority *priority*;  
                  }  
              }  
              traceoptions {  
                  file *name* <replace> <size *size*> <files *number*> <no-stamp>  
                      <(world-readable | no-world-readable)>;  
                  flag *flag* <*flag-modifier*> <disable>;  
              }  
              revision-level *revision-level*;  
          }

**Hierarchy Level**    [edit protocols]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure Multiple Spanning Tree Protocol (MSTP). MSTP is defined in the IEEE 802.1Q-2003 specification and is used to create a loop-free topology in networks with multiple spanning-tree regions.

The statements are explained separately.

**Default**    MSTP is disabled.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Understanding MSTP on page 1555</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface on page 1875</a></li> </ul>

## no-root-port

<b>Syntax</b>	no-root-port;
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">rstp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">stp interface</a> (all   <i>interface-name</i> )], [edit protocols <a href="#">vstp vlan <i>vlan-id</i> interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure an interface to be a spanning tree designated port. If the bridge receives more STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. When the bridge stops receiving more STP BPDUs on the root-protected interface, interface traffic is no longer blocked.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on page 1664</a></li> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li> <li>• <a href="#">Understanding VSTP on page 1557</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface on page 1875</a></li> </ul>

## priority (STP)

---

<b>Syntax</b>	<code>priority <i>priority</i>;</code>
<b>Hierarchy Level</b>	[edit protocols <b>mstp</b> <b>interface</b> (all   <i>interface-name</i> )], [edit protocols <b>mstp</b> <b>msti</b> <i>msti-id</i> <b>interface</b> <i>interface-name</i> ], [edit protocols <b>rstp</b> <b>interface</b> (all   <i>interface-name</i> )], [edit protocols <b>stp</b> <b>interface</b> (all   <i>interface-name</i> )], [edit protocols <b>vstp</b> <b>vlan</b> <i>vlan-id</i> <b>interface</b> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the interface priority to control which interface is elected as the root port.
<b>Options</b>	<b>priority</b> —Interface priority. The interface priority must be set in increments of 16. <b>Range:</b> 0 through 240 <b>Default:</b> 128
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li><li>• <a href="#">Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610</a></li><li>• <a href="#">Understanding MSTP on page 1555</a></li><li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li><li>• <a href="#">Understanding VSTP on page 1557</a></li><li>• <a href="#">show spanning-tree bridge on page 1870</a></li><li>• <a href="#">show spanning-tree interface on page 1875</a></li></ul>



## revision-level

---

<b>Syntax</b>	<code>revision-level <i>revision-level</i>;</code>
<b>Hierarchy Level</b>	[edit protocols <a href="#">mstp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Multiple Spanning Tree Protocol (MSTP), set the revision number of the MSTP configuration.
<b>Default</b>	The revision number is disabled.
<b>Options</b>	<i>revision-level</i> —Revision number of the MSTP region configuration. <b>Range:</b> 0 through 65535
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Network Regions for VLANs with MSTP on page 1624</a></li> <li>• <a href="#">Understanding MSTP on page 1555</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface on page 1875</a></li> </ul>

## rstp

---

**Syntax**    `rstp {  
    disable;  
    bpdu-block-on-edge;  
    bridge-priority priority;  
    forward-delay seconds;  
    hello-time seconds;  
    interface (all | interface-name) {  
        bpdu-timeout-action {  
            block;  
            alarm;  
        }  
        disable;  
        cost cost;  
        edge;  
        mode mode;  
        no-root-port;  
        priority priority;  
    }  
    max-age seconds;  
    traceoptions {  
        file name <replace> <size size> <files number> <no-stamp>  
        <(world-readable | no-world-readable)>;  
        flag flag <flag-modifier> <disable>;  
    }  
}`

**Hierarchy Level**    [edit protocols]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure Rapid Spanning Tree Protocol (RSTP). RSTP is defined in the IEEE 802.1D-2004 specification and is used to prevent loops in Layer 2 networks, providing shorter convergence times than those provided with basic STP.

The statements are explained separately.

**Default**    RSTP is enabled on all Ethernet switching interfaces.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
- [Understanding RSTP on page 1556](#)
- [show spanning-tree bridge on page 1870](#)
- [show spanning-tree interface on page 1875](#)

## stp

<b>Syntax</b>	<pre> stp {   disable;   bridge-priority <i>priority</i>;   forward-delay <i>seconds</i>;   hello-time <i>seconds</i>;   interface (all   <i>interface-name</i>) {     disable;     bpdu-timeout-action {       block;       alarm;     }     cost <i>cost</i>;     edge;     mode <i>mode</i>;     no-root-port;     priority <i>priority</i>;   }   max-age <i>seconds</i>;   traceoptions {     file <i>name</i> &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;no-stamp&gt;       &lt;(world-readable   no-world-readable)&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;;   } } </pre>
<b>Hierarchy Level</b>	[edit protocols]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>When you explicitly configure STP, a switch uses the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP (defined in the IEEE 802.1D 1998 specification).</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	STP is disabled; by default, RSTP is enabled on all Ethernet switching ports.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on page 1656</a></li> <li>• <a href="#">Configuring STP on page 1704</a></li> <li>• <a href="#">Overview of Spanning-Tree Protocols on page 1554</a></li> <li>• <a href="#">show spanning-tree bridge on page 1870</a></li> <li>• <a href="#">show spanning-tree interface on page 1875</a></li> </ul>

## traceoptions (STP)

---

Syntax	<pre>traceoptions {     file <i>name</i> &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;no-stamp&gt;     &lt;(world-readable   no-world-readable)&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
Hierarchy Level	<pre>[edit protocols <i>mstp</i>], [edit protocols <i>rstp</i>], [edit protocols <i>stp</i>], [edit protocols <i>vstp</i> vlan <i>vlan-id</i>] [edit protocols layer2-control]</pre>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.



**NOTE:** traceoptions is not supported on QFabric systems.

**Description** Set STP protocol-level tracing options.

**Default** Traceoptions is disabled.

**Options** **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file *name***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file `/var/log/stp-log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

**Range:** 2 through 1000 files

**Default:** 1 trace file only

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the STP-specific tracing options:

- **all**—Trace all operations.
- **all-failures**—Trace all failure conditions.
- **bpdu**—Trace BPDU reception and transmission.

- **bridge-detection-state-machine**—Trace the bridge detection state machine.
- **events**—Trace events of the protocol state machine.
- **port-information-state-machine**—Trace the port information state machine.
- **port-migration-state-machine**—Trace the port migration state machine.
- **port-receive-state-machine**—Trace the port receive state machine.
- **port-role-select-state-machine**—Trace the port role selection state machine.
- **port-role-transit-state-machine**—Trace the port role transit state machine.
- **port-transmit-state-machine**—Trace the port transmit state machine.
- **port-state-transit-state-machine**—Trace the port state transit state machine.
- **ppmd**—Trace the state and events for the ppm process.
- **state-machine-variables**—Trace when the state machine variables change.
- **timers**—Trace protocol timers.
- **topology-change-state-machine**—Trace the topology change state machine.

The following are the global tracing options:

- **all**—All tracing operations.
- **config-internal**—Trace configuration internals.
- **general**—Trace general events.
- **normal**—All normal events.

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **parse**—Trace configuration parsing.
- **policy**—Trace policy operations and actions.
- **regex-parse**—Trace regular-expression parsing.
- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.O**. When the **trace-file** again reaches its maximum size, **trace-file.O** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.O**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring Network Regions for VLANs with MSTP on page 1624](#)
  - [Example: Configuring Faster Convergence and Improving Network Stability with RSTP on page 1610](#)
  - [Understanding RSTP on page 1556](#)
  - [Understanding MSTP on page 1555](#)
  - [Overview of Spanning-Tree Protocols on page 1554](#)
  - [Understanding VSTP on page 1557](#)
  - [show spanning-tree bridge on page 1870](#)
  - [show spanning-tree interface on page 1875](#)

## vlan (STP)

```
Syntax  vlan (vlan-id | vlan-name) {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
            bpdu-timeout-action {
                block;
                alarm;
            }
            cost cost;
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
        }
        max-age seconds;
        traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
            flag flag;
        }
    }
```

**Hierarchy Level** [edit protocols **mstp** **msti** *msti-id*],  
[edit protocols **vstp**]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the VLANs for a Multiple Spanning Tree Instance (MSTI).

The remaining statements are explained separately.



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

**Default** Not enabled.

**Options** *vlan-id*—Numeric VLAN identifier.

*vlan-name*—Name of the VLAN.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.



- Related Documentation**
- [Example: Configuring Network Regions for VLANs with MSTP on page 1624](#)
  - [Understanding MSTP on page 1555](#)
  - [Understanding VSTP on page 1557](#)

## vstp

**Syntax**

```
vstp {
  disable;
  bpd-block-on-edge;
  force-version stp;
  vlan (vlan-id | vlan-name) {
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpd-timeout-action {
        alarm;
        block;
      }
      cost cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    traceoptions {
      file name <replace> <size size> <files number> <no-stamp>
        <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure VLAN Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding VSTP on page 1557](#)
  - [show spanning-tree bridge on page 1870](#)
  - [show spanning-tree interface on page 1875](#)

## VLAN Configuration Statements

---

- [\[edit vlans\] Configuration Statement Hierarchy on the QFX Series on page 1761](#)
- [control-channel on page 1764](#)
- [control-vlan on page 1765](#)
- [data-channel on page 1766](#)
- [description \(VLAN\) on page 1767](#)
- [dhcp-relay on page 1768](#)
- [east-interface on page 1773](#)
- [ethernet-ring on page 1774](#)
- [filter \(VLANs\) on page 1775](#)
- [forwarding-options on page 1776](#)
- [guard-interval on page 1781](#)
- [hold-interval \(Protection Group\) on page 1782](#)
- [interface \(VLANs\) on page 1782](#)
- [interface-mac-limit on page 1783](#)
- [interface-mode on page 1785](#)
- [irb \(Interfaces\) on page 1787](#)
- [l3-interface \(VLAN\) on page 1790](#)
- [mac \(Static MAC-Based VLANs\) on page 1791](#)
- [mac-limit on page 1791](#)
- [mac-notification on page 1792](#)
- [mac-statistics on page 1793](#)
- [mac-table-aging-time on page 1794](#)
- [mac-table-size on page 1795](#)
- [members on page 1797](#)
- [native-vlan-id on page 1798](#)
- [notification-interval on page 1799](#)
- [packet-action on page 1800](#)
- [port-mode on page 1803](#)
- [protection-group on page 1804](#)
- [restore-interval on page 1805](#)
- [ring-protection-link-end on page 1806](#)
- [ring-protection-link-owner on page 1806](#)
- [service-id on page 1807](#)
- [switch-options on page 1808](#)
- [static \(Static MAC-Based VLANs\) on page 1809](#)

- [static-mac](#) on page 1810
- [vlan-id](#) (VLANs) on page 1811
- [vlan-id-list](#) on page 1812
- [vlan-rewrite](#) on page 1813
- [vlan-tagging](#) on page 1813
- [vlans](#) on page 1814
- [west-interface](#) on page 1817

## [edit vlans] Configuration Statement Hierarchy on the QFX Series

This topic lists supported and unsupported configuration statements in the **[edit vlans]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *QFX Series Virtual Chassis Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit vlans\] Hierarchy Level](#) on page 1761
- [Unsupported Statements in the \[edit vlans\] Hierarchy Level](#) on page 1763

### Supported Statements in the [edit vlans] Hierarchy Level

The following hierarchy shows the **[edit vlans]** configuration statements supported on one or more of the EX Series switches:

```

vlans {
  vlan-name {
    description text-description;
    domain-type bridge;
    forwarding-options {
      dhcp-security {
        arp-inspection;
        group group-name {
          interface interface-name {
            static-ip ip-address {
              mac mac-address;
            }
          }
          overrides {
            no-option82;
            trusted;
          }
        }
      }
    }
  }
  ip-source-guard;
}

```

```
no-dhcp-snooping;
option-82 {
  circuit-id {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical);
    use-vlan-id;
  }
  remote-id {
    host-name;
    use-interface-description (device | logical);
    use-string string;
  }
  vendor-id {
    use-string string;
  }
}
}
filter {
  input filter-name;
  output filter-name;
}
flood {
  input filter-name;
}
}
l3-interface irb.logical-unit-number;
multicast-snooping-options {
  flood-groups [group-names];
  forwarding-cache {
    threshold {
      reuse threshold;
      suppress threshold;
    }
  }
}
graceful-restart {
  disable;
  restart-duration duration;
}
host-outbound-traffic {
  dot1p bits;
  forwarding-class forwarding-class;
}
multichassis-lag-replicate-state;
nexthop-hold-time time;
options {
  syslog {
    level level;
    mark interval;
    upto level;
  }
}
}
traceoptions {
```

```
file filename {
  files number;
  no-world-readable;
  size file-size;
  world-readable;
}
flag flag {
  disable;
}
}
switch-options {
  interface interface-name {
    interface-mac-limit limit {
      packet-action action;
    }
    static-mac mac-address;
  }
  interface-mac-limit limit {
    packet-action action;
  }
  mac-move-limit limit {
    packet-action action;
  }
  mac-table-size limit {
    packet-action drop;
  }
  no-mac-learning;
}
vlan-id number;
vlan-id-list [vlan-id | vlan-id-vlan-id];
}
```

Unsupported Statements in the [edit vlans] Hierarchy Level

All statements in the [edit vlans] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 113: Unsupported [edit vlans] Configuration Statements on EX Series Switches

Statement	Hierarchy Level
NOTE: Variables, such as <i>filename</i> , are not shown in the statements or hierarchies.	
mcae-mac-synchronize	[edit vlans]
no-irb-layer-2-copy	[edit vlans]

Related Documentation

- [Understanding Bridging and VLANs on page 1527](#)

## control-channel

---

<b>Syntax</b>	<code>control-channel <i>channel-name</i> {     vlan <i>vlan-id</i>;     interface name <i>interface-name</i> }</code>
<b>Hierarchy Level</b>	[edit protocols protection-group <b>ethernet-ring</b> <i>name</i> ( <b>east-interface</b>   <b>west-interface</b> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
<b>Description</b>	Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.
<b>Options</b>	<b>vlan <i>vlan-id</i></b> —If the control channel logical interface is a trunk port, then a dedicated <b>vlan <i>vlan-id</i></b> defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the <b>vlan-id</b> when the control channel logical interface is the trunk port.  <b>interface name <i>interface-name</i></b> —Interface name of the control channel.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li><li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li></ul>

## control-vlan

---

<b>Syntax</b>	control-vlan ( <i>vlan-id</i>   <i>vlan-name</i> )
<b>Hierarchy Level</b>	[edit protocols protection-group <a href="#">ethernet-ring</a> ] [edit protocols protection-group <a href="#">ethernet-ring</a> name (east-interface  west-interface)]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
<b>Description</b>	Specify the VLAN that carries the protocol data units (PDUs) between the nodes in the protected Ethernet ring. This is a control VLAN, meaning that it carries data for one instance of an Ethernet ring protection switching (ERPS) in the control channel. Use a control VLAN on trunk port interfaces. One control channel can contain multiple control VLANs.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</i></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li> </ul>

## data-channel

---

<b>Syntax</b>	<code>data-channel {     vlan <i>number</i>; }</code>
<b>Hierarchy Level</b>	[edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
<b>Description</b>	<p>For Ethernet ring protection, configure a data channel to define a set of VLAN IDs that belong to a ring instance.</p> <p>VLANs specified in the data channel use the same topology used by the ERPS PDU in the control channel. Therefore, if a ring interface is blocked in the control channel, all traffic in the data channel is also blocked on that interface.</p>
<b>Options</b>	<b>vlan <i>number</i></b> —Specify (by VLAN ID) one or more VLANs that belong to a ring instance.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Ethernet Ring Protection Using Ring Instances for Load Balancing</i></li><li>• <i>Example: Configuring Load Balancing Within Ethernet Ring Protection for MX Series Routers</i></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li><li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li></ul>



---

## description (VLAN)

---

<b>Syntax</b>	<code>description <i>text-description</i>;</code>
<b>Hierarchy Level</b>	[edit vlans <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Provide a textual description for the VLAN. The text has no effect on the operation of the VLAN or switch.
<b>Options</b>	<i>text-description</i> —Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. If the text includes spaces, enclose the entire text in quotation marks.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li><li>• <a href="#">Understanding Bridging and VLANs on page 1527</a></li><li>• <a href="#">show vlans on page 1886</a></li></ul>

## dhcp-relay

---

```
Syntax  dhcp-relay {
    active-server-group server-group-name;
    authentication {
        password password-string;
        username-include {
            circuit-type;
            delimiter delimiter-character;
            domain-name domain-name-string;
            interface-name;
            logical-system-name;
            mac-address;
            option-60;
            option-82 <circuit-id> <remote-id>;
            routing-instance-name;
            user-prefix user-prefix-string;
        }
    }
    dhcpv6 {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
    }
    group group-name {
        active-server-group server-group-name;
        authentication {
            ...
        }
        dynamic-profile profile-name {
            ...
        }
        interface interface-name {
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
```

```

bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
}
}
overrides {
    ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
    ...
}
relay-agent-interface-id {
    ...
}
service-profile dynamic-profile-name;
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
}
overrides {

```

```
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
  }
  relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  }
  OBSOLETE – duplicate-clients-on-interface;
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
  group group-name {
    active-server-group server-group-name;
    authentication {
      ...
    }
  }
  dynamic-profile profile-name {
    ...
  }
  interface interface-name {
    exclude;
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          version (0 | 1 | automatic);
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          detection-time {
            threshold milliseconds;
          }
          session-mode (automatic | multihop | singlehop);
          holddown-interval milliseconds;
        }
      }
    }
  }
  overrides {
    ...
  }
```

```

        service-profile dynamic-profile-name;
        trace;
        upto upto-interface-name;
    }
    overrides {
        ...
    }
    relay-option-82 {
        ...
    }
    service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    OBSOLETE - no-arp;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
server-group {

```

```
server-group-name {  
    server-ip-address;  
}  
}  
service-profile dynamic-profile-name;  
}
```

**Hierarchy Level** [edit forwarding-options],  
[edit vlans forwarding-options]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the switch and enable the switch to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcpr-relay** and **dhcprv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router at the same time.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

## east-interface

**Syntax**

```
east-interface {
  node-id mac-address;
  control-channel channel-name {
    vlan number;
    interface name interface-name
  }
  interface-name
  ring-protection-link-end;
}
```

**Hierarchy Level** [edit protocols protection-group **ethernet-ring** *ring-name*]

**Release Information** Statement introduced in Junos OS Release 9.4.  
Statement introduced in Junos OS Release 12.1 for EX Series switches.  
Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

**Description** Define one of the two interface ports for Ethernet ring protection, the other being defined by the **west-interface** statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

EX Series switches do not use the node-id statement--the node ID is automatically configured on the switches using the MAC address.



**NOTE:** Always configure this port first, before configuring the **west-interface** statement.



**NOTE:** The Node ID is not configurable on EX Series switches. The node ID is automatically configured using the MAC address.

The statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Ethernet Ring Protection Switching Overview on page 1525](#)
- [Ethernet Ring Protection Using Ring Instances for Load Balancing](#)
- [west-interface on page 1817](#)
- [ethernet-ring on page 1774](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 1670](#)

## ethernet-ring

---

**Syntax**    ethernet-ring *ring-name* {  
              control-vlan (*vlan-id* | *vlan-name*);  
              data-channel {  
                  vlan *number*  
              }  
              east-interface {  
                  control-channel *channel-name* {  
                      vlan *number*;  
                      interface name *interface-name*  
                  }  
              }  
              guard-interval *number*;  
              node-id *mac-address*;  
              restore-interval *number*;  
              ring-protection-link-owner;  
              west-interface {  
                  control-channel *channel-name* {  
                      vlan *number*;  
                  }  
              }  
          }

**Hierarchy Level**    [edit protocols protection-group]

**Release Information**    Statement introduced in Junos OS Release 9.4.  
                              Statement introduced in Junos OS Release 12.1 for EX Series switches.  
                              Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.

**Description**    For Ethernet PICs on MX Series routers or for EX Series switches, , specify the Ethernet ring in an Ethernet ring protection switching configuration.

**Options**    *ring-name*—Name of the Ethernet protection ring.  
  
              The remaining statements are explained separately.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Ethernet Ring Protection Switching Overview on page 1525](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 1670](#)



## filter (VLANs)

---

<b>Syntax</b>	<code>filter (input   output) <i>filter-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit vlans <i>vlan-name</i>]</code> <code>[edit vlans <i>vlan-name</i> forwarding-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Apply a firewall filter to traffic ingressing or egressing a VLAN.
<b>Default</b>	All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.
<b>Options</b>	<p><b><i>filter-name</i></b>—Name of a firewall filter defined at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to VLAN ingress traffic.</p> <p><b>output</b>—Apply a firewall filter to VLAN egress traffic.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Firewall Filters on page 5290</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5209</a></li> </ul>

## forwarding-options

---

```
Syntax forwarding-options {
  dhcp-relay {
    active-server-group server-group-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        mac-address;
        option-60;
        option-82 <circuit-id> <remote-id>;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
  }
  dhcpv6 {
    active-server-group server-group-name;
    authentication {
      password password-string;
      username-include {
        circuit-type;
        client-id;
        delimiter delimiter-character;
        domain-name domain-name-string;
        interface-name;
        logical-system-name;
        relay-agent-interface-id;
        relay-agent-remote-id;
        relay-agent-subscriber-id;
        routing-instance-name;
        user-prefix user-prefix-string;
      }
    }
  }
  dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
  }
  group group-name {
    active-server-group server-group-name;
    authentication {
      ...
    }
    dynamic-profile profile-name {
      ...
    }
  }
  interface interface-name {
    exclude;
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    }
  }
}
```

```

method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
overrides {
  ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
service-profile dynamic-profile-name;
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode(automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
}

```

```
overrides {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
OBSOLETE – duplicate-clients-on-interface;
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
}
dynamic-profile profile-name {
    ...
}
interface interface-name {
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    ...
}
```

```

    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
  }
  overrides {
    ...
  }
  OBSOLETE - relay-option-60 {
    ...
  }
  relay-option-82 {
    ...
  }
  service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  allow-snooped-clients;
  always-write-giaddr;
  always-write-option-82;
  client-discover-match <option60-and-option82>;
  disable-relay;
  interface-client-limit number;
  layer2-unicast-replies;
  no-allow-snooped-clients;
  no-bind-on-request;
  proxy-mode;
  replace-ip-source-with;
  send-release-on-delete;
  trust-option-82;
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
}

```

```
    }
  }
  server-group {
    server-group-name {
      server-ip-address;
    }
  }
  service-profile dynamic-profile-name;
}
dhcp-security {
  arp-inspection;
  group group-name {
    interface interface-name {
      static-ip ip-address {
        mac mac-address;
      }
    }
    overrides {
      no-option82;
      trusted;
      untrusted;
    }
  }
}
ip-source-guard;
no-dhcp-snooping;
option-82 {
  circuit-id {
    prefix {
      host-name;
      logical-system-name;
      routing-instance-name;
    }
    use-interface-description (device | logical);
    use-vlan-id;
  }
  remote-id {
    host-name hostname;
    use-interface-description (device | logical);
    use-string string;
  }
  vendor-id {
    use-string string;
  }
}
}
fip-security {
  examine-vn2vf;
  examine-vn2vn {
    beacon-period milliseconds;
  }
  fc-map fc-map-value;
  interface interface-name {
    (fcoe-trusted | no-fcoe-trusted;)
  }
}
}
```

<b>Hierarchy Level</b>	<a href="#">[edit]</a> <a href="#">[edit vlans]</a>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Configure traffic forwarding.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

## guard-interval

<b>Syntax</b>	<code>guard-interval <i>number</i>;</code>
<b>Hierarchy Level</b>	<a href="#">[edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i>]</a>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
<b>Description</b>	When a link goes down, the ring protection link (RPL) activates. When the downed link comes back up, the RPL link receives notification, restores the link, and waits for the restore interval before issuing another block on the same link. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
<b>Options</b>	<i>number</i> —Guard timer interval, in milliseconds. <b>Range:</b> 10 through 2000 ms <b>Default:</b> 500 ms
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li> </ul>

## hold-interval (Protection Group)

---

<b>Syntax</b>	<code>hold-interval <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit protocols protection-group <a href="#">ethernet-ring <i>name</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4.
<b>Description</b>	Specify the hold-off timer interval <i>for all rings</i> in 100 millisecond (ms) increments.
<b>Options</b>	<i>number</i> —Hold-timer interval, in milliseconds. <b>Range:</b> 0 through 10,000 ms <b>Default:</b> 100 ms
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li></ul>

## interface (VLANs)

---

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     mapping (native (push   swap)   tag (push   swap)); }</pre>
<b>Hierarchy Level</b>	[edit vlans <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For a specific VLAN, configure an interface.
<b>Options</b>	<i>interface-name</i> —Name of the interface.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li><li>• <a href="#">Configuring VLANs</a></li><li>• <a href="#">Understanding Bridging and VLANs</a></li></ul>



## interface-mac-limit

<b>Syntax</b>	<pre>interface-mac-limit <i>limit</i> {     <b>packet-action</b> drop; }</pre>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> switch-options],          [edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>],          [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],          [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],          [edit routing-instances <i>routing-instance-name</i> switch-options],          [edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],          [edit switch-options],          [edit switch-options interface <i>interface-name</i>],          [edit switch-options interface <i>interface-name</i>],          [edit vlans <i>vlan-name</i> switch-options],          [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
<b>Description</b>	<p>(MX Series routers or EX Series switches only) Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.</p>



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at configuration` statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

---

**Default** For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

**Options** *limit*—Maximum number of MAC addresses learned from an interface.

**Range:** 1 through 131,071 MAC addresses per interface



The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1526](#)
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## interface-mode

<b>Syntax</b>	interface-mode (access   trunk);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	<div>  <p><b>NOTE:</b> This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see <a href="#">port-mode</a>. For ELS details, see “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 43.</p> </div> <p>(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the <b>trunk</b> option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id</b> or <b>vlan-id-list</b> statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <div>  <p><b>NOTE:</b> On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure <b>interface-mode</b> and <b>irb</b> for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> </div>
<b>Options</b>	<p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring a Logical Interface for Access Mode</li> <li>Configuring a Logical Interface for Trunk Mode</li> </ul>

- *Example: Connecting Access Switches to a Distribution Switch*

## irb (Interfaces)

```

Syntax  irb {
        accounting-profile name;
        description text;

        (gratuitous-arp-reply | no-gratuitous-arp-reply);
        hold-time up milliseconds down milliseconds;
        mtu bytes;
        no-gratuitous-arp-request;

        traceoptions {
            flag flag;
        }
        (traps | no-traps);
        unit logical-unit-number {
            accounting-profile name;
            bandwidth rate;
            description text;
            disable;
            encapsulation type;
            family inet {
                accounting {
                    destination-class-usage;
                    source-class-usage {
                        input;
                        output;
                    }
                }
            }
            address ipv4-address {
                arp ip-address (mac | multicast-mac) mac-address <publish>;
                broadcast address;
                preferred;
                primary;
                vrrp-group group-number {
                    (accept-data | no-accept-data);
                    advertise-interval seconds;
                    advertisements-threshold number;
                    authentication-key key;
                    authentication-type authentication;
                    fast-interval milliseconds;
                    (preempt | no-preempt) {
                        hold-time seconds;
                    }
                }
                priority number;
                track {
                    interface interface-name {
                        bandwidth-threshold bandwidth;
                        priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/mask routing-instance instance-name priority-cost cost;
                }
            }
            virtual-address [ addresses ];
        }
    }

```

```
        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
    }
}
```

```

    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

**Hierarchy Level** [edit interfaces *interface-name*

<b>Release Information</b>	Statement introduced in Junos OS Release 12.3R2 for EX Series switches. <b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Configure the properties of a specific integrated bridging and routing (IRB) interface.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>[edit interfaces] Hierarchy Level</i></li><li>• <i>[edit interfaces] Configuration Statement Hierarchy on EX Series Switches</i></li></ul>

---

## l3-interface (VLAN)

---

<b>Syntax</b>	<code>l3-interface (vlan.logical-interface-number   irb.logical-interface-number);</code>
<b>Hierarchy Level</b>	<code>[edit vlans vlan-name]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. <b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between VLANs. Traffic between VLANs must be routed, which requires a common Layer 3 interface.
<b>Default</b>	No Layer 3 (routing) interface is associated with the VLAN.
<b>Options</b>	<code>vlan.logical-interface-number</code> —Number of the logical interface. Use the <b>unit</b> number that you used when you created the <b>vlan</b> interface with a <b>set interfaces vlan unit</b> statement.



**NOTE:** Use this statement with versions of Junos OS that do not support Enhanced Layer 2 Software (ELS).

---

`irb.logical-interface-number`—Logical interface defined with a **set interfaces irb** statement.

---



**NOTE:** Use this statement with versions of Junos OS that support Enhanced Layer 2 Software (ELS).

---

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ethernet-switching interfaces on page 1482</a></li><li>• <a href="#">show vlans on page 1886</a></li></ul>



## mac (Static MAC-Based VLANs)

<b>Syntax</b>	<code>mac mac-address {     next-hop interface-name; }</code>
<b>Hierarchy Level</b>	<code>[edit ethernet-switching-options static vlan vlan-name]</code>
<b>Description</b>	Specify the MAC address to add to the Ethernet switching table.  The remaining statement is explained separately.
<b>Options</b>	<i>mac-address</i> —MAC address
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)</i></li> </ul>

## mac-limit

<b>Syntax</b>	<code>mac-limit number;</code>
<b>Hierarchy Level</b>	<code>[edit vlans vlan-name]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the number of MAC addresses allowed on a VLAN.
<b>Default</b>	MAC limit is disabled.
<b>Options</b>	<i>number</i> —Maximum number of MAC addresses. <b>Range:</b> 1 through 32768



**NOTE:** This statement is not supported on QFabric systems.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show vlans on page 1886</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li> <li>• <a href="#">Configuring MAC Table Aging</a></li> <li>• <a href="#">Understanding Bridging and VLANs</a></li> </ul>

## mac-notification

---

<b>Syntax</b>	<pre>mac-notification {     notification-interval <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit ethernet-switching-options] [edit switch-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Hierarchy level <b>[edit switch-options]</b> added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.
<b>Description</b>	<p>Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.</p> <p>The remaining statement is explained separately.</p>
<b>Default</b>	MAC notification is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring MAC Notification</i></li><li>• <a href="#">Configuring MAC Notification (CLI Procedure) on page 1682</a></li></ul>

## mac-statistics

<b>Syntax</b>	mac-statistics;
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit logical-systems <i>logical-system-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> switch-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols evpn],</p> <p>[edit switch-options],</p> <p>[edit switch-options],</p> <p>[edit vlans <i>vlan-name</i> switch-options]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers.</p> <p>[edit switch-options] and [edit vlans <i>vlan-name</i> switch-options] hierarchy levels introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	(MX Series routers, EX Series switches, and QFX Series only) For bridge domains or VLANs, enable MAC accounting either for a specific bridge domain or VLAN, or for a set of bridge domains or VLANs associated with a Layer 2 trunk port.
<b>Default</b>	disabled
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i></li> <li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Overview on page 1526</a></li> <li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports</i></li> <li>• <i>Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port</i></li> <li>• <i>Configuring EVPN Routing Instances</i></li> </ul>

## mac-table-aging-time

---

<b>Syntax</b>	<code>mac-table-aging-time seconds;</code>
<b>Hierarchy Level</b>	For platforms without ELS:  [edit ethernet-switching-options], [edit vlans <i>vlan-name</i> ]  For platforms with ELS:  [edit vlans <i>vlan-name</i> switch-options]
<b>Release Information</b>	Statement introduced for specific VLANs in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define how long entries remain in the Ethernet switching table before expiring: <ul style="list-style-type: none"><li>• If you specify this statement at the <b>[ethernet-switching-options]</b> hierarchy level, it applies to all VLANs on the switch.</li><li>• If you specify this statement at the <b>[vlans]</b> hierarchy level, it applies to the specified VLAN.</li></ul>
<b>Default</b>	300 seconds
<b>Options</b>	<b>seconds</b> —Time that entries remain in the Ethernet switching table before being removed. <ul style="list-style-type: none"><li>• <b>Range</b>—60 to 1,000,000 seconds.</li><li>• <b>Default</b>—300 seconds.</li></ul>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li><li>• <a href="#">Configuring MAC Table Aging</a></li><li>• <a href="#">Configuring MAC Table Aging on page 1674</a></li><li>• <a href="#">Understanding Bridging and VLANs on page 1527</a></li><li>• <a href="#">show ethernet-switching statistics aging on page 1858</a></li></ul>

## mac-table-size

<b>Syntax</b>	<code>mac-table-size <i>limit</i> {     <code>packet-action</code> drop; }</code>
<b>Hierarchy Level</b>	[edit bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options], [edit logical-systems <i>logical-system-name</i> switch-options], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options], [edit routing-instances <i>routing-instance-name</i> switch-options], [edit switch-options], [edit vlans <i>vlan-name</i> switch-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for the <b>switch-options</b> statement added in Junos OS Release 9.2. Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch. Support for logical systems added in Junos OS Release 9.6. <b>[edit switch-options]</b> and <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches. Support at the <b>[edit vlans <i>vlan-name</i> switch-options]</b> hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Modify the size of the MAC address table for the bridge domain or VLAN, a set of bridge domains or VLANs associated with a trunk port, or a virtual switch. The default is 5120 MAC addresses.



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the **mac-table-size** statement or changing the **mac-table-size** configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the **mac-table-size** statement or use the **commit at** configuration statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the **clear bridge mac-table** command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

**Options** *limit*—Specify the maximum number of addresses in the MAC address table.

**Range:** 16 through 1,048,575 MAC addresses

**Default:** 5120 MAC addresses There is no default MAC address limit for the **mac-table-size** statement at the **[edit switch-options]** hierarchy level. The number of MAC addresses that can be learned is only limited by the platform, 65,535 MAC addresses for EX Series switches and 1,048,575 MAC addresses for other devices.

The remaining statement is explained separately.

**Required Privilege** routing—To view this statement in the configuration.

**Level** routing-control—To add this statement to the configuration.

**Related  
Documentation**

- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1526](#)
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## members

<b>Syntax</b>	<code>members [(all   <i>names</i>   <i>vlan-ids</i>)];</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> unit 0 family ethernet-switching vlan]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For trunk interfaces, configure the VLANs for which the interface can carry traffic.



**TIP:** To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after `vlan` or `vlands` in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

**Options** `all`—Specify that this trunk interface be a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.



**NOTE:** Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, `all` cannot be the name of a VLAN on the switch.

*names*—Names of one or more VLANs.

*vlan-ids*—Numeric identifiers of one or more VLANs.

**Required Privilege Level**  
`routing`—To view this statement in the configuration.  
`routing-control`—To add this statement to the configuration.

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Understanding Bridging and VLANs on page 1527](#)
- [show ethernet-switching interfaces on page 1482](#)
- [show vlans on page 1886](#)

## native-vlan-id

---

<b>Syntax</b>	<code>native-vlan-id <i>vlan-id</i>;</code>
<b>Hierarchy Level</b>	For platforms without ELS:  <code>[edit <a href="#">interfaces</a> <i>interface-name</i> unit 0 family ethernet-switching],</code>  For platforms with ELS:  <code>[edit <a href="#">interfaces</a> <i>interface-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure the VLAN identifier to associate with untagged packets received on the interface. The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface. To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the <code>[edit <a href="#">interfaces</a> <i>interface-name</i> unit <i>logical-unit-number</i>]</code> hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is combined with the <a href="#">interface-mode</a> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p> <p>When the <b>native-vlan-id</b> statement is combined with the <a href="#">flexible-vlan-tagging</a> statement, untagged packets are accepted on the interfaces that are configured for Q-in-Q tunneling.</p> <p>.</p>
<b>Options</b>	<p><b>vlan-id</b>—Numeric identifier of the VLAN.</p> <p><b>Range:</b> 1 through 4094</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Junos OS Network Interfaces Configuration Guide</a></li><li>• .</li><li>• <a href="#">show ethernet-switching interfaces on page 1482</a></li><li>• <a href="#">show vlans on page 1886</a></li></ul>



## notification-interval

---

<b>Syntax</b>	notification-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit ethernet-switching-options mac-notification] [edit switch-options mac-notification]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Hierarchy level <b>[edit switch-options]</b> added in Junos OS Release 14.1X53-D10 for EX Series and QFX Series.
<b>Description</b>	Configure the MAC notification interval for a switch.  The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications are sent to the network management system every 10 seconds.
<b>Options</b>	<b><i>seconds</i></b> —The MAC notification interval, in seconds. <b>Range:</b> 1 through 60 <b>Default:</b> 30
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring MAC Notification</i></li> <li>• <a href="#">Configuring MAC Notification (CLI Procedure) on page 1682</a></li> </ul>

## packet-action

**Syntax** `packet-action action;`

**Hierarchy Level** [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit logical-systems *logical-system-name* switch-options **interface-mac-limit** *limit*],  
 [edit protocols **l2-learning** global-mac-limit *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* bridge-options **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* protocols evpn interface-mac-limit (vpls)],  
 [edit routing-instances *routing-instance-name* protocols evpn interface *interface-name* interface-mac-limit (vpls)],  
 [edit routing-instances *routing-instance-name* protocols evpn mac-table-size *limit*],  
 [edit routing-instances *routing-instance-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit routing-instances *routing-instance-name* switch-options **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit switch-options **interface-mac-limit** *limit*],  
 [edit switch-options **mac-table-size** *limit*],  
 [edit switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **mac-table-size** *limit*],  
 [edit vlans *vlan-name* switch-options **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options interface *interface-name* **interface-mac-limit** *limit*],  
 [edit vlans *vlan-name* switch-options **mac-table-size** *limit*]

**Release Information** Statement introduced in Junos OS Release 8.4.  
 Support for the **switch-options** statement added in Junos OS Release 9.2.  
 Support for top-level configuration for the **virtual-switch** type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.

Support for logical systems added in Junos OS Release 9.6.

[edit switch-options interface *interface-name* interface-mac-limit *limit*], [edit switch-options interface-mac-limit *limit*], [edit switch-options mac-table-size *limit*], [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit *limit*], [edit vlans *vlan-name* switch-options interface-mac-limit *limit*], and [edit vlans *vlan-name* switch-options mac-table-size *limit*] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.

Support for EVPNs introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge Routers.

Support at the [edit switch-options interface *interface-name* interface-mac-limit *limit*] hierarchy level and hierarchy levels under [edit vlans *vlan-name*] introduced in Junos OS Release 13.2X50-D10 for EX Series switches and Junos OS Release 13.2 for the QFX Series.

**Description** Specify the action taken when packets with new source MAC addresses are received after the MAC address limit is reached. If this statement is not configured, packets with new source MAC addresses are forwarded by default.

**Default**



**NOTE:** On a QFX Series Virtual Chassis, if you include the shutdown option at the [edit vlans *vlan-name* switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level and issue the commit operation, the system generates a commit error. The system does not generate an error if you include the shutdown option at the [edit switch-options interface *interface-name* interface-mac-limit packet-action] hierarchy level.

Disabled. The default is for packets for new source MAC addresses to be forwarded after the MAC address limit is reached.

**Options**

- drop**—Drop packets with new source MAC addresses, and do not learn the new source MAC addresses.
- drop-and-log**—(EX Series switches and QFX Series only) Drop packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- log**—(EX Series switches and QFX Series only) Hold packets with new source MAC addresses, and generate an alarm, an SNMP trap, or a system log entry.
- none**—(EX Series switches and QFX Series only) Forward packets with new source MAC addresses, and learn the new source MAC address.
- shutdown**—(EX Series switches and QFX Series only) Disable the specified interface, and generate an alarm, an SNMP trap, or a system log entry.


**Required Privilege Level**

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

**Related  
Documentation**

- *Configuring EVPN Routing Instances*
- [Configuring MAC Limiting \(CLI Procedure\) on page 1672](#)
- *Configuring Persistent MAC Learning (CLI Procedure)*
- *Layer 2 Learning and Forwarding for Bridge Domains Overview*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1526](#)
- *Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports*
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1526](#)
- *Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port*

## port-mode

<b>Syntax</b>	port-mode (access   tagged-access   trunk);
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family ethernet-switching]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<div>  <p><b>NOTE:</b> This statement does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see <a href="#">interface-mode</a>. For ELS details, see “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 43.</p> </div> <p>Configure whether an interface on the switch operates in access, tagged access, or trunk mode.</p>
<b>Default</b>	All switch interfaces are in access mode.
<b>Options</b>	<p><b>access</b>—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.</p> <p><b>tagged-access</b>—Have the interface operate in tagged-access mode. In this mode, the interface can be in multiple VLANs. Tagged access interfaces typically connect to network devices, such as PCs, printers, IP telephones, and IP cameras.</p> <p><b>trunk</b>—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Reflective Relay</i></li> <li><i>Example: Configuring Reflective Relay for Use with VEPA Technology</i></li> </ul>

## protection-group

```
Syntax  protection-group {
        ethernet-ring ring-name {
            control-vlan (vlan-id | vlan-name);
            data-channel {
                vlan number
            }
            east-interface {
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
            }
            guard-interval number;
            hold-interval number;
            node-id mac-address;
            restore-interval number;
            ring-protection-link-owner RPL owner flag;
            west-interface {
                control-channel channel-name {
                    vlan number;
                    interface name interface-name
                }
            }
            guard-interval number;
            hold-interval
            node-id mac-address;
            restore-interval number;
            traceoptions {
                file filename <no-stamp> <world-readable | no-world-readable> <replace> <size size>;
                flag flag;
            }
        }
    }
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 12.1 for EX Series switches.

**Description** Configure Ethernet ring protection switching (ERPS).

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#)
- [Configuring Ethernet Ring Protection Switching \(CLI Procedure\) on page 1670](#)
- [Ethernet Ring Protection Switching Overview on page 1525](#)

## restore-interval

---

<b>Syntax</b>	<code>restore-interval <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit protocols protection-group <a href="#">ethernet-ring <i>ring-name</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
<b>Description</b>	Configures the number of minutes that the node does not process any Ethernet ring protection (ERP) protocol data units (PDUs).. This configuration is a global configuration and applies to all Ethernet rings if the Ethernet ring does not have a more specific configuration for this value. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.
<b>Options</b>	<i>number</i> —Specify the restore interval. <b>Range:</b> 5 through 12 minutes
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li> </ul>

## ring-protection-link-end

---

<b>Syntax</b>	ring-protection-link-end;
<b>Hierarchy Level</b>	[edit protocols protection-group <a href="#">ethernet-ring</a> <i>ring-name</i> ( <a href="#">east-interface</a>   <a href="#">west-interface</a> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Specify that the port is one side of a ring protection link (RPL) by setting the RPL end flag.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li><li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li></ul>

## ring-protection-link-owner

---

<b>Syntax</b>	ring-protection-link-owner;
<b>Hierarchy Level</b>	[edit protocols protection-group <a href="#">ethernet-ring</a> <i>ring-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.
<b>Description</b>	Specify the ring protection link (RPL) owner flag in the Ethernet protection ring. Include this statement only once for each ring (only one node can function as the RPL owner).
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li><li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li></ul>



---

## service-id

---

<b>Syntax</b>	<code>service-id <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit switch-options] [edit vlans <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers. Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).
<b>Options</b>	<b>number</b> —A number that identifies a particular service. <b>Range:</b> 1 through 65535
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system control—To add this statement to the configuration.

## switch-options

---

<b>Syntax</b>	<pre>switch-options {   interface <i>interface-name</i> {     interface-mac-limit <i>limit</i> {       packet-action drop;     }     no-mac-learning;     static-mac <i>static-mac-address</i> {       vlan-id <i>number</i>;     }   }   interface-mac-limit <i>limit</i> {     packet-action drop;   }   mac-statistics;   mac-table-size <i>limit</i> {     packet-action drop;   }   no-mac-learning;   service-id <i>number</i>;   vtep-source-interface }</pre>
<b>Hierarchy Level</b>	<pre>[edit <i>number</i>], [edit vlans <i>vlan--name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vlans   <i>vlan-name</i>], [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	<p>Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

---

## static (Static MAC-Based VLANs)

---

**Syntax**   static {  
              vlan *vlan-name* {  
                  mac *mac-address* {  
                      next-hop *interface-name*;  
                  }  
              }  
          }

**Hierarchy Level**   [edit ethernet-switching-options]

**Release Information**   Statement introduced in Junos OS Release 11.1 for EX Series switches.

**Description**       Specify VLAN and MAC addresses to add to the Ethernet switching table.  
  
                      The remaining statements are explained separately.

**Required Privilege**   system—To view this statement in the configuration.  
                  **Level**   system-control—To add this statement to the configuration.


**Related**           • *Adding a Static MAC Address Entry to the Ethernet Switching Table (CLI Procedure)*  
**Documentation**

## static-mac

---

<b>Syntax</b>	<code>static-mac <i>mac-address</i> {     vlan-id <i>number</i>; }</code>
<b>Hierarchy Level</b>	[edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols evpn interface <i>interface-name</i> ] [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Support for logical systems added in Junos OS Release 9.6. [edit vlans <i>vlan-name</i> switch-options interface <i>interface name</i> ] hierarchy level introduced in Junos OS Release 12.3R2 for EX Series switches. Support for EVPNs added in Junos OS Release 13.2 for MX 3D Series routers. The <b>vlan-id</b> option is not available for EVPNs. [edit vlans <i>vlan-name</i> switch-options interface <i>interface name</i> ] hierarchy level introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Configure a static MAC address for a logical interface in a bridge domain or VLAN.  The <b>vlan-id</b> option can be specified for <b>static-macs</b> only if <b>vlan-id all</b> is configured for the bridging domain or VLAN.
<b>Options</b>	<b><i>mac-address</i></b> —MAC address  <b><i>vlan-id number</i></b> —(Optional) VLAN identifier to associate with static MAC address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring EVPN Routing Instances</i></li><li>• <i>Layer 2 Learning and Forwarding for Bridge Domains Overview</i></li><li>• <a href="#">Layer 2 Learning and Forwarding for VLANs Overview on page 1526</a></li></ul>

## vlan-id (VLANs)

<b>Syntax</b>	<code>vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <pre>[edit vlans <i>vlan-name</i> <i>vlan-range</i>]</pre> <p>For platforms without ELS and with ELS:</p> <pre>[edit vlans <i>vlan-name</i>]</pre> <p>For ELS platforms only:</p> <pre>[edit interfaces <i>interface-name</i> unit <i>number</i>] [edit vlans <i>vlan-name</i> <i>vlan-id-list</i>]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.
<b>Default</b>	<p>On a QFX3500 and QFX3500 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.</p> <p>On a QFX5100 switch, if you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of 1. The number zero is reserved for priority tagging and the number 4093 is also reserved.</p>
	<div>  <p><b>NOTE:</b> You can only create up to 4090 VLANs on a QFX5100 switch. If you create more than 4090 VLANs, the interfaces associated with the extra VLANs are not displayed in the <code>show vlans</code> command output. For example, if you create 4094 VLANs, the extra VLANs will not have interfaces associated with the VLANs. The order in which you configure the extra VLANs determines which interfaces are missing from the <code>show vlans</code> command output.</p> </div>
<b>Options</b>	<p><i>number</i> —VLAN tag identifier.</p> <p><b>Range:</b> 0 through 4093.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Setting Up Bridging with Multiple VLANs</i></li> <li>• <i>Understanding Bridging and VLANs</i></li> </ul>

## vlan-id-list

<b>Syntax</b>	<code>vlan-id-list [ <i>vlan-id-numbers</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</p> <p>[edit interfaces <i>interface-name</i> unit 0],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit vlans <i>vlan-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	<p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the <b>trunk</b> option in the <b>interface-mode</b> statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id-list</b> statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p>This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.</p>
<b>Options</b>	<p><b><i>vlan-id-numbers</i></b>—Valid VLAN identifiers. You can combine individual numbers with range lists including a hyphen.</p> <p><b>Range:</b> 0 through 4095</p>



**NOTE:** On EX Series switches and the QFX Series, the range is 0 through 4094.

<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
---------------------------------	--------------------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Bridge Domain</i></li> <li>• <i>Configuring a VLAN</i></li> <li>• <i>Configuring VLANs for EX Series Switches (CLI Procedure)</i></li> <li>• <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i></li> </ul>
------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

## vlan-rewrite

<b>Syntax</b>	vlan-rewrite translate (200 500   201 501)
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge interface-mode trunk] [edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching interface-mode trunk]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid VLANs pass through without translation.
<b>Options</b>	<b>translate 200 500</b> —Translates incoming packets with VLAN 200 to 500.  <b>translate 201 501</b> —Translates incoming packets with VLAN 201 to 501.  <b>translate 202 502</b> —Translates incoming packets with VLAN 202 to 502.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Rewriting a VLAN Tag and Adding a New Tag</i></li> </ul>

## vlan-tagging

<b>Syntax</b>	vlan-tagging;
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> ] [edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.
<b>Default</b>	VLAN tagging is disabled by default.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">vlan-id on page 2746</a></li> <li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2593</a></li> </ul>

## vlan

---

```
Syntax  vlans {
        vlan-name {
            description text-description;
            domain-type bridge;
            forwarding-options {
                dhcp-security {
                    arp-inspection;
                    group group-name {
                        interface interface-name {
                            static-ip ip-address {
                                mac mac-address;
                            }
                        }
                    }
                    overrides {
                        no-option82;
                        trusted;
                        untrusted;
                    }
                }
            }
            ip-source-guard;
            no-dhcp-snooping;
            option-82 {
                circuit-id {
                    prefix {
                        host-name;
                        logical-system-name;
                        routing-instance-name;
                    }
                    use-interface-description (device | logical);
                    use-vlan-id;
                }
                remote-id {
                    host-name hostname;
                    use-interface-description (device | logical);
                    use-string string;
                }
                vendor-id {
                    use-string string;
                }
            }
        }
    }
    fip-security {
        examine-vn2vf;
        examine-vn2vn {
            beacon-period milliseconds;
        }
        fc-map fc-map-value;
        interface interface-name {
            (fcoe-trusted | no-fcoe-trusted;)
        }
    }
}
```



```

l3-interface irb.logical-unit-number;
multicast-snooping-options {
  flood-groups [group-names];
  forwarding-cache {
    threshold {
      reuse threshold;
      suppress threshold;
    }
  }
  graceful-restart {
    disable;
    restart-duration duration;
  }
  host-outbound-traffic {
    dot1p bits;
    forwarding-class forwarding-class;
  }
  multichassis-lag-replicate-state;
  nexthop-hold-time time;
  options {
    syslog {
      level level;
      mark interval;
      upto level;
    }
  }
  traceoptions {
    file filename {
      files number;
      no-world-readable;
      size file-size;
      world-readable;
    }
    flag flag {
      disable;
    }
  }
}
switch-options {
  interface interface-name {
    interface-mac-limit limit {
      packet-action action;
    }
    static-mac mac-address;
  }
  interface-mac-limit limit {
    packet-action action;
  }
  mac-move-limit limit {
    packet-action action;
  }
  mac-table-size limit {
    packet-action drop;
  }
  no-mac-learning;
}

```

```
    vlan-id number;  
    vlan-id-list [vlan-id | vlan-id-vlan-id];  
    vxlan  
  }  
}
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statements for private VLANs and Q-in-Q tunneling introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure VLAN properties on the QFX Series.

**Default** If you use the default factory configuration, all switch interfaces become part of the VLAN **default**.

**Options** *vlan-name*—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.


The remaining statements are described separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing—control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Bridging and VLANs on page 1527](#)
- [Configuring VLANs on page 1678](#)

## west-interface

<b>Syntax</b>	<pre> west-interface {   node-id <i>mac-address</i>;   control-channel <i>channel-name</i> {     vlan <i>number</i>;     interface name <i>interface-name</i>   }   interface-name   ring-protection-link-end; } </pre>
<b>Hierarchy Level</b>	[edit protocols protection-group <b>ethernet-ring</b> <i>ring-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.153-D10 for QFX Series switches.</p>
<b>Description</b>	<p>Define one of the two interface ports for Ethernet ring protection, the other being defined by the <b>east-interface</b> statement at the same hierarchy level. The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.</p>
<div>  <b>NOTE:</b> Always configure this port second, after configuring the <b>east-interface</b> statement. </div>	
<p>The statements are explained separately.</p>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Ethernet Ring Protection Switching Overview on page 1525</a></li> <li>• <a href="#">Ethernet Ring Protection Using Ring Instances for Load Balancing</a></li> <li>• <a href="#">east-interface on page 1773</a></li> <li>• <a href="#">ethernet-ring on page 1774</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Ethernet Ring Protection Switching on EX Series Switches and QFX Switches on page 1563</a></li> <li>• <a href="#">Configuring Ethernet Ring Protection Switching (CLI Procedure) on page 1670</a></li> </ul>

## Q-in-Q Configuration Statements

- [flexible-vlan-tagging on page 1818](#)
- [input-vlan-map on page 1819](#)

- [native-vlan-id on page 1820](#)
- [output-vlan-map on page 1821](#)
- [pop on page 1822](#)
- [push on page 1823](#)
- [swap on page 1824](#)
- [vlan-id-list on page 1825](#)

---

## flexible-vlan-tagging

---

<b>Syntax</b>	flexible-vlan-tagging;
<b>Hierarchy Level</b>	[edit interfaces aex], [edit interfaces ge- <i>fpc/pic/port</i> ], [edit interfaces et- <i>fpc/pic/port</i> ], [edit interfaces ps0], [edit interfaces xe- <i>fpc/pic/port</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Support for aggregated Ethernet added in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers. Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	Support simultaneous transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port, and on pseudowire logical interfaces.  This statement is supported on M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2, IQ2-E, and IQ PICs or in MX Series DPCs, or on Ethernet interfaces for PTX Series Packet Transport Routers or 100-Gigabit Ethernet Type 5 PIC with CFP. This statement is supported on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Mixed Tagging</i></li><li>• <i>Configuring Flexible VLAN Tagging on PTX Series Packet Transport Routers</i></li></ul>

## input-vlan-map

<b>Syntax</b>	<pre>input-vlan-map {   (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);   inner-tag-protocol-id <i>tpid</i>;   inner-vlan-id <i>number</i>;   tag-protocol-id <i>tpid</i>;   vlan-id <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>pop-pop</b>, <b>pop-swap</b>, <b>push-push</b>, <b>swap-push</b>, and <b>swap-swap</b> statements introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only as well as Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite profile to be applied to incoming frames on this logical interface.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <a href="#">output-vlan-map on page 1821</a></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>

## native-vlan-id

---

<b>Syntax</b>	<code>native-vlan-id <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>ge-fpc/pic/port</i> ], [edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	<p>Configure mixed tagging support for untagged packets on a port for the following:</p> <ul style="list-style-type: none"><li>• M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging</li><li>• MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging</li><li>• T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP</li><li>• EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces</li></ul> <p>When the <b>native-vlan-id</b> statement is included with the <a href="#">flexible-vlan-tagging</a> statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <p>The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface. To configure the logical interface, include the <b>vlan-id</b> statement (matching the <b>native-vlan-id</b> statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the <b>native-vlan-id</b> statement is included with the <a href="#">interface-mode</a> statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.</p>
<b>Options</b>	<p><b><i>number</i></b>—VLAN ID number.</p> <p><b>Range:</b> (ACX Series routers and EX Series switches) 0 through 4094.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Mixed Tagging Support for Untagged Packets</i></li><li>• <i>Configuring a Logical Interface for Access Mode</i></li><li>• <a href="#">Configuring the Native VLAN Identifier (CLI Procedure) on page 1677</a></li><li>• <i>Understanding Bridging and VLANs on EX Series Switches</i></li></ul>


- [flexible-vlan-tagging on page 1818](#)
- *Understanding Q-in-Q Tunneling on EX Series Switches*

## output-vlan-map

<b>Syntax</b>	<pre>output-vlan-map {   (pop   pop-pop   pop-swap   push   push-push   swap   swap-push   swap-swap);   inner-tag-protocol-id <i>tpid</i>;   inner-vlan-id <i>number</i>;   tag-protocol-id <i>tpid</i>;   vlan-id <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>pop-pop</b>, <b>pop-swap</b>, <b>push-push</b>, <b>swap-push</b>, and <b>swap-swap</b> statements added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
<b>Description</b>	<p>For Gigabit Ethernet IQ, 10-Port 10-Gigabit Ethernet SFPP interfaces, 100-Gigabit Ethernet Type 5 PIC with CFP only, Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, define the rewrite operation to be applied to outgoing frames on this logical interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags</i></li> <li>• <a href="#">input-vlan-map on page 1819</a></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>


## pop

---

<b>Syntax</b>	pop;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	<div> <b>NOTE:</b> On EX4300 switches, pop is not supported at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b>] hierarchy level.</div> <p>For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2, and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Removing a VLAN Tag</i></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li></ul>



## push

<b>Syntax</b>	<code>push;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">input-vlan-map</a>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <a href="#">output-vlan-map</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>  <a href="#">input-vlan-map</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>  output-vlan-map]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.</p>
<b>Description</b>	<p> <b>NOTE:</b> On EX4300 switches, <code>push</code> is not supported at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p> <p>Specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.</p> <p>You can use this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces; 10-Gigabit Ethernet LAN/WAN PIC; aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces; 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p> <p>If you include the <b>push</b> statement in the configuration, you must also include the <a href="#">pop</a> statement at the <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]</code> hierarchy level.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Stacking a VLAN Tag</i></li> <li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li> </ul>

## swap

---

<b>Syntax</b>	swap;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>input-vlan-map</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <b>output-vlan-map</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3R2 for EX Series switches. Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	<p>Specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.</p> <p>On MX Series routers, you can enter this statement on Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, 10-Gigabit Ethernet LAN/WAN PIC, aggregated Ethernet using Gigabit Ethernet IQ interfaces, and 100-Gigabit Ethernet Type 5 PIC with CFP. On EX Series switches, you can enter this statement on Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Rewriting the VLAN Tag on Tagged Frames</i></li><li>• <i>Configuring Q-in-Q Tunneling (CLI Procedure)</i></li></ul>

## vlan-id-list

<b>Syntax</b>	<code>vlan-id-list [ <i>vlan-id-numbers</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i>]</p> <p>[edit interfaces <i>interface-name</i> unit 0],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit vlans <i>vlan-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.4.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>
<b>Description</b>	<p>Specify a VLAN identifier list to use for a bridge domain or VLAN in trunk mode.</p> <p>Specify the <b>trunk</b> option in the <b>interface-mode</b> statement to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id-list</b> statement to forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID to forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p>This statement also enables you to bind a logical interface to a list of VLAN IDs, thereby configuring the logical interface to receive and forward a frame with a tag that matches the specified VLAN ID list.</p>
<b>Options</b>	<p><b><i>vlan-id-numbers</i></b>—Valid VLAN identifiers. You can combine individual numbers with range lists including a hyphen.</p> <p><b>Range:</b> 0 through 4095</p>



**NOTE:** On EX Series switches and the QFX Series, the range is 0 through 4094.

<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
---------------------------------	--------------------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Bridge Domain</i></li> <li>• <i>Configuring a VLAN</i></li> <li>• <i>Configuring VLANs for EX Series Switches (CLI Procedure)</i></li> <li>• <i>Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances</i></li> </ul>
------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- *Configuring VLAN Identifiers for VLANs and VPLS Routing Instances*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*

## CHAPTER 16

# Administration

- [Routine Monitoring on page 1827](#)
- [Monitoring Commands on page 1835](#)

## Routine Monitoring

---

- [Verifying That MAC Notification Is Working Properly on page 1827](#)
- [Verifying That a Series of Tagged VLANs Has Been Created on page 1827](#)
- [Verifying That a Private VLAN Is Working on page 1829](#)
- [Verifying That Proxy ARP Is Working Correctly on page 1834](#)

### Verifying That MAC Notification Is Working Properly

**Purpose** Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.

**Action** To verify that MAC notification is enabled or disabled and also to verify the MAC notification interval setting.

```
user@switch> show ethernet-switching mac-notification
Notification Status: Enabled
Notification Interval: 60
Notifications Sent      : 0
Notifications Table Maxsize : 256
```

**Meaning** The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 60 seconds.

**Related Documentation**

- [Configuring MAC Notification](#)
- [Configuring MAC Notification \(CLI Procedure\) on page 1682](#)

### Verifying That a Series of Tagged VLANs Has Been Created

**Purpose** Verify that a series of tagged VLANs has been created on the switch.

- Action** 1. Display the VLANs in the ascending order of their VLAN ID:

```
user@switch> show vlans sort-by tag
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

2. Display the VLANs by the alphabetical order of the VLAN name:

```
user@switch> show vlans sort-by name
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

3. Display the VLANs by specifying the VLAN range name (here, the VLAN range name is **employee**):

```
user@switch> show vlans employee
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*

**Meaning** The sample output shows the VLANs configured on the switch. The series of tagged VLANs is displayed: **\_\_employee\_120\_\_** through **\_\_employee\_130\_\_**. Each of the tagged VLANs is configured on the trunk interface **xe-0/0/22.0**. The asterisk (\*) next to the interface name indicates that the interface is **UP**.

When a series of VLANs is created using the **vlan-range** statement, the VLAN names are preceded and followed by a double underscore.

- Related Documentation**
- *Creating a Series of Tagged VLANs*
  - [Creating a Series of Tagged VLANs on page 1680](#)

## Verifying That a Private VLAN Is Working

**Purpose** After creating and configuring private VLANs (PVLANS), verify that they are set up properly.

**Action** 1. To determine whether you successfully created the primary and secondary VLAN configurations:

- For a PVLAN on a single switch, use the **show configuration vlans** command:

```
user@switch> show configuration vlans
community1 {
    interface {
        interface a;
        interface b;
    }
    primary-vlan pvlan;
}
community2 {
    interface {
```

```

        interface d;
        interface e;
    }
    primary-vlan pvlan;
}
pvlan {
    vlan-id 1000;
    interface {
        isolated1;
        isolated2;
        trunk1;
        trunk2;
    }
    no-local-switching;
}

```

- For a PVLAN spanning multiple switches, use the **show vlans extensive** command:

```

user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/7.0*, untagged, access

```

```

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access

```

```

VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/2.0, untagged, access

```

```

VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk

```

```

VLAN: community2, Created at: Tue May 11 18:16:05 2010

```



```

802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, untagged, access
    ge-1/0/6.0*, untagged, access

```

```

VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access
    ge-0/0/1.0*, untagged, access
    ge-0/0/2.0, untagged, access
    ge-0/0/7.0*, untagged, access
    ge-1/0/6.0*, untagged, access

```

```

Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
    __pvlan_primary_ge-0/0/0.0__
    __pvlan_primary_ge-0/0/2.0__
Community VLANs :
    COM1
    community2
Inter-switch-isolated VLAN :
    __pvlan_primary_isiv__

```

2. Use the **show vlans extensive** command to view VLAN information and link status for a PVLAN on a single switch or for a PVLAN spanning multiple switches.

- For a PVLAN on a single switch:

```

user@switch> show vlans pvlan extensive
VLAN: pvlan, Created at: time
802.1Q Tag: vlan-id, Internal index: index-number, Admin State: Enabled,
Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged 6 (Active = 0)
    trunk1, tagged, trunk
    interface a, untagged, access
    interface b, untagged, access
    interface c, untagged, access
    interface d, untagged, access
    interface e, untagged, access
    interface f, untagged, access
    trunk2, tagged, trunk
Secondary VLANs: Isolated 2, Community 2
Isolated VLANs :
    __pvlan_pvlan_isolated1__
    __pvlan_pvlan_isolated2__
Community VLANs :

```

```
community1
community2
```

- For a PVLAN spanning multiple switches:

```
user@switch> show vlans extensive
VLAN: COM1, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 100, Internal index: 3, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/7.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/0.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 5, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 1)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access

VLAN: __pvlan_primary_ge-0/0/2.0__, Created at: Tue May 11 18:16:05 2010
Internal index: 6, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 1 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/2.0, untagged, access

VLAN: __pvlan_primary_isiv__, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 50, Internal index: 7, Admin State: Enabled, Origin: Static
Private VLAN Mode: Inter-switch-isolated, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 0 (Active = 0)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk

VLAN: community2, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 20, Internal index: 8, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 2 (Active = 2)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/1.0*, untagged, access
    ge-1/0/6.0*, untagged, access

VLAN: primary, Created at: Tue May 11 18:16:05 2010
802.1Q Tag: 10, Internal index: 2, Admin State: Enabled, Origin: Static
```

```

Private VLAN Mode: Primary
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 3 (Active = 3), Untagged 5 (Active = 4)
    ge-0/0/20.0*, tagged, trunk
    ge-0/0/22.0*, tagged, trunk, pvlan-trunk
    ge-0/0/23.0*, tagged, trunk, pvlan-trunk
    ge-0/0/0.0*, untagged, access
    ge-0/0/1.0*, untagged, access
    ge-0/0/2.0, untagged, access
    ge-0/0/7.0*, untagged, access
    ge-1/0/6.0*, untagged, access

Secondary VLANs: Isolated 2, Community 2, Inter-switch-isolated 1
Isolated VLANs :
    __pvlan_primary_ge-0/0/0.0__
    __pvlan_primary_ge-0/0/2.0__
Community VLANs :
    COM1
    community2
Inter-switch-isolated VLAN :
    __pvlan_primary_isiv__

```

3. Use the **show ethernet-switching table** command to view logs for MAC learning on the VLANs:

```

user@switch> show ethernet-switching table
Ethernet-switching table: 8 entries, 1 learned

```

VLAN	MAC address	Type	Age	Interfaces
default	*	Flood	-	All-members
pvlan	*	Flood	-	All-members
pvlan	MAC1	Replicated	-	interface a
pvlan	MAC2	Replicated	-	interface c
pvlan	MAC3	Replicated	-	isolated2
pvlan	MAC4	Learn	0	trunk1
__pvlan_pvlan_isolated1__ *		Flood	-	All-members
__pvlan_pvlan_isolated1__ MAC4		Replicated	-	trunk1
__pvlan_pvlan_isolated2__ *		Flood	-	All-members
__pvlan_pvlan_isolated2__ MAC3		Learn	0	isolated2
__pvlan_pvlan_isolated2__ MAC4		Replicated	-	trunk1
community1	*	Flood	-	All-members
community1	MAC1	Learn	0	interface a
community1	MAC4	Replicated	-	trunk1
community2	*	Flood	-	All-members

community2	MAC2	Learn	0 interface c
community2	MAC4	Replicated	- trunk1



**NOTE:** If you have configured a PVLAN spanning multiple switches, you can use the same command on all the switches to check the logs for MAC learning on those switches.

**Meaning** In the output displays for a PVLAN on a single switch, you can see that the primary VLAN contains two community domains (**community1** and **community2**), two isolated ports, and two trunk ports. The PVLAN on a single switch has only one tag (1000), which is for the primary VLAN.

The PVLAN that spans multiple switches contains multiple tags:

- The community domain **COM1** is identified with tag **100**.
- The community domain **community2** is identified with tag **20**.
- The interswitch isolated domain is identified with tag **50**.
- The primary VLAN **primary** is identified with tag **10**.

Also, for the PVLAN that spans multiple switches, the trunk interfaces are identified as **pvlan-trunk**.

- Related Documentation**
- *Creating a Private VLAN on a Single EX Series Switch (CLI Procedure)*
  - *Creating a Private VLAN Spanning Multiple EX Series Switches (CLI Procedure)*
  - *Creating a Private VLAN on a Single Switch*
  - *Creating a Private VLAN Spanning Multiple Switches*

## Verifying That Proxy ARP Is Working Correctly

**Purpose** Verify that the switch is sending proxy ARP messages.

**Action** List the system statistics for ARP:

```
user@switch> show system statistics arp
arp:
  90060 datagrams received
  34 ARP requests received
  610 ARP replies received
  2 resolution request received
  0 unrestricted proxy requests
  0 restricted proxy requests
  0 received proxy requests
  0 unrestricted proxy requests not proxied
  0 restricted proxy requests not proxied
```

```

0 datagrams with bogus interface
0 datagrams with incorrect length
0 datagrams for non-IP protocol
0 datagrams with unsupported op code
0 datagrams with bad protocol address length
0 datagrams with bad hardware address length
0 datagrams with multicast source address
0 datagrams with multicast target address
0 datagrams with my own hardware address
0 datagrams for an address not on the interface
0 datagrams with a broadcast source address
294 datagrams with source address duplicate to mine
89113 datagrams which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
309 ARP requests sent
35 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

**Meaning** The statistics show that two proxy ARP requests were received. The **unrestricted proxy requests not proxied** and **restricted proxy requests not proxied** fields indicate that all the unproxied ARP requests received have been proxied by the switch.

**Related Documentation**

- *Configuring Proxy ARP*
- *Configuring Proxy ARP (CLI Procedure) on page 1702*

## Monitoring Commands

- `clear ethernet-switching bpdu-error`
- `clear ethernet-switching layer2-protocol-tunneling error`
- `clear ethernet-switching layer2-protocol-tunneling statistics`
- `clear ethernet-switching table`
- `clear spanning-tree statistics`
- `show ethernet-switching interfaces`
- `show ethernet-switching layer2-protocol-tunneling interface`
- `show ethernet-switching layer2-protocol-tunneling statistics`
- `show ethernet-switching layer2-protocol-tunneling vlan`
- `show ethernet-switching mac-learning-log`
- `show ethernet-switching mac-notification`
- `show ethernet-switching statistics aging`

- `show ethernet-switching statistics mac-learning`
- `show ethernet-switching table`
- `show spanning-tree bridge`
- `show spanning-tree interface`
- `show spanning-tree mstp configuration`
- `show spanning-tree statistics`
- `show system statistics arp`
- `show vlans`

## clear ethernet-switching bpdu-error

<b>Syntax</b>	<code>clear ethernet-switching bpdu-error interface <i>interface-name</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.1 for EX Series switches. Command updated in Junos OS Release 11.1 for EX Series switches—a BPDU error shuts down the interface and this command brings the interface back up. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear bridge protocol data unit (BPDU) errors from an interface and bring up the interface.
<b>Options</b>	<i>interface-name</i> —Clear BPDU errors on the specified interface.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show spanning-tree interface on page 1875</a></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches</a></li> <li>• <a href="#">Understanding BPDU Protection for STP, RSTP, and MSTP on page 1558</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ethernet-switching bpdu-error interface on page 1837</a>

### Sample Output

#### clear ethernet-switching bpdu-error interface

```
user@switch> clear ethernet-switching bpdu-error interface xe-0/0/1.0
```

## clear ethernet-switching layer2-protocol-tunneling error

---

<b>Syntax</b>	clear ethernet-switching layer2-protocol-tunneling error <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Clear Layer 2 protocol tunneling (L2PT) errors on one or more interfaces. If an interface has been disabled because the amount of Layer 2 protocol traffic exceeded the shutdown threshold or because the switch has detected an error in the network topology or configuration, use this command to reenable the interface.
<b>Options</b>	<b>none</b> —Clears L2PT errors on all interfaces.  <b>interface <i>interface-name</i></b> —(Optional) Clear L2PT errors on the specified interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</i></li><li>• <i>Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</i></li><li>• <i>Configuring Layer 2 Protocol Tunneling</i></li></ul>
<b>List of Sample Output</b>	<a href="#">clear ethernet-switching layer2-protocol-tunneling error on page 1838</a> <a href="#">clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0 on page 1838</a>

### Sample Output

#### clear ethernet-switching layer2-protocol-tunneling error

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error
```

#### clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error interface xe-0/0/1.0
```



## clear ethernet-switching layer2-protocol-tunneling statistics

<b>Syntax</b>	clear ethernet-switching layer2-protocol-tunneling statistics <interface <i>interface-name</i> > <vlan <i>vlan-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Clear Layer 2 protocol tunneling (L2PT) statistics on one or more interfaces or VLANs.
<b>Options</b>	<p><b>none</b>—Clear L2PT statistics on all interfaces and VLANs.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear L2PT statistics on the specified interface.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Clear L2PT statistics on the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling statistics on page 1849</a></li> <li>• <i>Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</i></li> <li>• <i>Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</i></li> <li>• <i>Configuring Layer 2 Protocol Tunneling</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ethernet-switching layer2-protocol-tunneling statistics on page 1839</a> <a href="#">clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0 on page 1839</a> <a href="#">clear ethernet-switching layer2-protocol-tunneling error vlan v2 on page 1839</a>

### Sample Output

#### clear ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics
```

#### clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0


```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics interface xe-0/1/1.0
```

#### clear ethernet-switching layer2-protocol-tunneling error vlan v2

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics vlan v2
```

## clear ethernet-switching table

---

<b>Syntax</b>	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <management-vlan> <persistent-mac < <i>interface</i>   <i>mac-address</i> >> <vlan <i>vlan-name</i> >
<b>Syntax (QFX Series)</b>	clear ethernet-switching table <interface <i>interface-name</i> > <mac <i>mac-address</i> > <persistent-mac < <i>interface</i>   <i>mac-address</i> >> <vlan <i>vlan-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<div> <b>NOTE:</b> On a QFabric system, using this command on an FCoE-enabled VLAN when FCoE sessions are active can cause traffic flooding and FCoE traffic drop. The FCoE sessions are not terminated and the traffic reconverges after a short period of time.</div> <div>Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).</div>
<b>Options</b>	<p><b>none</b>—Clear learned entries in the Ethernet switching table, except for persistent MAC addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.</p> <p><b>mac <i>mac-address</i></b>—(Optional) Clear the specified learned MAC address from the Ethernet switching table.</p> <p><b>management-vlan</b>—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.</p> <p><b>persistent-mac &lt;<i>interface</i>   <i>mac-address</i>&gt;</b>—(Optional) Clear all MAC addresses, including persistent MAC addresses. Use the <b>interface</b> option to clear all MAC addresses on an interface, or use the <b>mac-address</b> option to clear all entries for a specific MAC address.</p> <p>Use this command whenever you move a device in your network that has a persistent MAC address on the switch. If you move the device to another port on the switch and do not clear the persistent MAC address from the original port it was learned on, then the new port will not learn the MAC address and the device will not be able to connect. If the original port is down when you move the device, then the new port</p>

will learn the MAC address and the device can connect—however, unless you cleared the MAC address on the original port, when the port comes back up, the system reinstalls the persistent MAC address in the forwarding table for that port. If this occurs, the address is removed from the new port and the device loses connectivity.

**vlan *vlan-name***—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

**Required Privilege Level**

view

**Related Documentation**

- *show ethernet-switching table*
- [show ethernet-switching table on page 1864](#)
- *Verifying That Persistent MAC Learning Is Working Correctly*

**List of Sample Output** [clear ethernet-switching table on page 1841](#)

**Output Fields** This command produces no output.


## Sample Output

[clear ethernet-switching table](#)

```
user@switch> clear ethernet-switching table
```

## clear spanning-tree statistics

---

<b>List of Syntax</b>	<a href="#">Syntax on page 1842</a> <a href="#">Syntax (EX Series Switches and the QFX Series) on page 1842</a>
<b>Syntax</b>	clear spanning-tree statistics <interface <i>interface-name</i> > <logical-system <i>logical-system-name</i> >
<b>Syntax (EX Series Switches and the QFX Series)</b>	clear spanning-tree statistics <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear Spanning Tree Protocol statistics.
<b>Options</b>	<b>none</b> —Reset STP counters for all interfaces for all routing instances.  <b>interface <i>interface-name</i></b> —(Optional) Clear STP statistics for the specified interface only.  <b>logical-system <i>logical-system-name</i></b> —(Optional) Clear STP statistics on a particular logical system.
<div> <b>NOTE:</b> The <b>logical-system</b> option is not available on QFabric systems.</div>	
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show spanning-tree statistics on page 1883</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear stp statistics on page 1842</a>

### Sample Output

#### clear stp statistics

```
user@host> clear stp statistics
```

## show ethernet-switching interfaces

<b>Syntax</b>	show ethernet-switching interfaces <brief   detail   summary> <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about switched Ethernet interfaces.
<b>Options</b>	<p><b>none</b>—(Optional) Display brief information for Ethernet-switching interfaces.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display Ethernet-switching information for a specific interface.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Troubleshooting Ethernet Switching on page 1895</a><a href="#">Understanding Bridging and VLANs on page 1527</a></li> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li> <li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li> <li>• <a href="#">Understanding FCoE on page 5518</a></li> <li>• <a href="#">Interfaces Overview on page 2389</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching interfaces on page 1844</a> <a href="#">show ethernet-switching interfaces summary on page 1845</a> <a href="#">show ethernet-switching interfaces brief on page 1845</a> <a href="#">show ethernet-switching interfaces detail on page 1845</a> <a href="#">show ethernet-switching interfaces interface-name on page 1846</a>
<b>Output Fields</b>	<a href="#">Table 85 on page 1482</a> lists the output fields for the <b>show ethernet-switching interfaces</b> command. Output fields are listed in the approximate order in which they appear.

**Table 114: show ethernet-switching interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of a switching interface.	All levels
<b>State</b>	Interface state. Values are <b>up</b> or <b>down</b> .	none, <b>brief</b> , <b>detail</b> , <b>summary</b>
<b>VLAN members</b>	Name of a VLAN.	none, <b>brief</b> , <b>detail</b> , <b>summary</b>

Table 114: show ethernet-switching interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Blocking</b>	Forwarding state of the interface: <ul style="list-style-type: none"> <li>• <b>blocked</b>—Traffic is not being forwarded on the interface.</li> <li>• <b>unblocked</b>—Traffic is forwarded on the interface.</li> <li>• <b>MAC limit exceeded</b>—The interface is temporarily disabled because of a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>MAC move limit exceeded</b>—The interface is temporarily disabled because of a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>Storm control in effect</b> —The interface is temporarily disabled because of a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.</li> <li>• <b>Storm control shutdown in effect</b> —The interface is temporarily disabled because of a storm control shutdown error. The disabled interface is automatically restored to service when the disable timeout expires.</li> </ul>	none, <b>brief</b> , <b>detail</b> , <b>summary</b>
<b>Index</b>	VLAN index internal to Junos OS software.	<b>detail</b>
<b>untagged   tagged</b>	Specifies whether the interface forwards IEEE802.1Q-tagged or untagged traffic.	<b>detail</b>

## Sample Output

### show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces
```

Interface	State	VLAN members	Blocking
xe-0/0/0.0	up	T1122	unblocked
xe-0/0/1.0	down	default	– MAC limit exceeded
xe-0/0/2.0	down	default	– MAC move limit exceeded
xe-0/0/3.0	down	default	– Storm control in effect
xe-0/0/4.0	down	default	unblocked
xe-0/0/5.0	down	default	unblocked
xe-0/0/6.0	down	default	unblocked
xe-0/0/7.0	down	default	unblocked
xe-0/0/8.0	down	default	unblocked
xe-0/0/9.0	up	T111	unblocked
xe-0/0/10.0	down	default	unblocked
xe-0/0/11.0	down	default	unblocked
xe-0/0/12.0	down	default	unblocked
xe-0/0/13.0	down	default	unblocked
xe-0/0/14.0	down	default	unblocked
xe-0/0/15.0	down	default	unblocked
xe-0/0/16.0	down	default	unblocked
xe-0/0/17.0	down	default	unblocked
xe-0/0/18.0	down	default	unblocked
xe-0/0/19.0	up	T111	unblocked
xe-0/1/0.0	down	default	unblocked
xe-0/1/1.0	down	default	unblocked
xe-0/1/2.0	down	default	unblocked
xe-0/1/3.0	down	default	unblocked

**show ethernet-switching interfaces summary**

```

user@switch> show ethernet-switching interfaces summary
xe-0/0/0.0
xe-0/0/1.0
xe-0/0/2.0
xe-0/0/3.0
xe-0/0/8.0
xe-0/0/10.0
xe-0/0/11.0

```

**show ethernet-switching interfaces brief**

```

user@switch> show ethernet-switching interfaces brief
Interface  State  VLAN members  Blocking
xe-0/0/0.0  down  default       unblocked
xe-0/0/1.0  down  employee-vlan unblocked
xe-0/0/2.0  down  employee-vlan unblocked
xe-0/0/3.0  down  employee-vlan unblocked
xe-0/0/8.0  down  employee-vlan unblocked
xe-0/0/10.0 down  default       unblocked
xe-0/0/11.0 down  employee-vlan unblocked

```

**show ethernet-switching interfaces detail**

```

user@switch> show ethernet-switching interfaces detail
Interface: xe-0/0/0.0 Index: 65
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/1.0 Index: 66
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/2.0 Index: 67
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/3.0 Index: 68
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/8.0 Index: 69
State: down
VLANs:
    employee-vlan          untagged    unblocked

Interface: xe-0/0/10.0 Index: 70
State: down
VLANs:
    default                untagged    unblocked

Interface: xe-0/0/11.0 Index: 71
State: down
VLANs:
    employee-vlan          tagged      unblocked

```

### show ethernet-switching interfaces interface-name

```
user@switch> show ethernet-switching interfaces xe-0/0/0.0
Interface    State    VLAN members    Blocking
xe-0/0/0.0  down    default          unblocked
```



## show ethernet-switching layer2-protocol-tunneling interface

<b>Syntax</b>	<code>show ethernet-switching-layer2-protocol-tunneling interface</code> <code>&lt;interface-name&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Display information about Layer 2 protocol tunneling (L2PT) on interfaces that have been configured for L2PT.
<b>Options</b>	<b>none</b> —Display L2PT information about all interfaces on which L2PT is enabled. <b>interface-name</b> —(Optional) Display L2PT information for the specified interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling statistics on page 1849</a></li> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling vlan on page 1852</a></li> <li>• <i>Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</i></li> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling statistics on page 1849</a></li> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling vlan on page 1852</a></li> <li>• <i>Configuring Layer 2 Protocol Tunneling</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching layer2-protocol-tunneling interface on page 1848</a> <a href="#">show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0 on page 1848</a>
<b>Output Fields</b>	Table 115 on page 1847 lists the output fields for the <b>show ethernet-switching layer2-protocol-tunneling interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 115: show ethernet-switching layer2-protocol-tunneling interface Output Fields

Field Name	Field Description
<b>Interface</b>	Name of an interface on the switch.
<b>Operation</b>	Type of operation being performed on the interface. Values are <b>Encapsulation</b> and <b>Decapsulation</b> .
<b>State</b>	State of the interface. Values are <b>active</b> and <b>shutdown</b> .
<b>Description</b>	If the interface state is <b>shutdown</b> , displays why the interface is shut down. If the description says <b>Loop detected</b> , it means that the interface is an access interface that has received L2PT-enabled PDUs. Access interfaces should not receive L2PT-enabled PDUs. This scenario might mean that there is a loop in the network.

## Sample Output

### show ethernet-switching layer2-protocol-tunneling interface

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface
```

Layer2 Protocol Tunneling information:

Interface	Operation	State	Description
xe-0/0/0.0	Encapsulation	Shutdown	Shutdown threshold exceeded
xe-0/0/1.0	Decapsulation	Shutdown	Loop detected
xe-0/0/2.0	Decapsulation	Active	

### show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface xe-0/0/0.0
```

Layer2 Protocol Tunneling information:

Interface	Operation	State	Description
xe-0/0/0.0	Encapsulation	Shutdown	Shutdown threshold exceeded

## show ethernet-switching layer2-protocol-tunneling statistics


<b>Syntax</b>	show ethernet-switching-layer2-protocol-tunneling statistics <interface <i>interface-name</i> > <vlan <i>vlan-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Display Layer 2 protocol tunneling (L2PT) statistics for Layer 2 PDU packets received by the switch.
<div>  <b>NOTE:</b> The show ethernet-switching-layer2-protocol-tunneling statistics command does not display L2PT statistics for Layer 2 PDU packets transmitted from the switch. </div>	
<b>Options</b>	<p><b>none</b>—Display L2PT statistics for all interfaces on which you enabled L2PT.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display L2PT statistics for the specified interface.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Display L2PT statistics for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ethernet-switching layer2-protocol-tunneling statistics on page 1839</a></li> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling interface on page 1847</a></li> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling vlan on page 1852</a></li> <li>• <a href="#">show vlans</a></li> <li>• <a href="#">Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</a></li> <li>• <a href="#">Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</a></li> <li>• <a href="#">show vlans on page 1886</a></li> <li>• <a href="#">Configuring Layer 2 Protocol Tunneling</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching layer2-protocol-tunneling statistics on page 1850</a> <a href="#">show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0 on page 1850</a> <a href="#">show ethernet-switching layer2-protocol-tunneling statistics vlan v2 on page 1850</a>
<b>Output Fields</b>	Table 116 on page 1850 lists the output fields for the <b>show ethernet-switching layer2-protocol-tunneling statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 116: show ethernet-switching layer2-protocol-tunneling statistics Output Fields

VLAN	Field Description
<b>VLAN</b>	Name of a VLAN on which L2PT has been configured.
<b>Interface</b>	Name of an interface on which L2PT has been configured.
<b>Protocol</b>	Name of a protocol for which L2PT has been enabled. Values are <b>all</b> , <b>802.1x</b> , <b>802.3ah</b> , <b>cdp</b> , <b>e-lmi</b> , <b>gvrp</b> , <b>lACP</b> , <b>lldp</b> , <b>mmrp</b> , <b>mvrp</b> , <b>stp</b> , <b>udld</b> , <b>vstp</b> , and <b>vtp</b> .
<b>Operation</b>	Type of operation being performed on the interface. Values are <b>Encapsulation</b> and <b>Decapsulation</b> .
<b>Packets</b>	Number of packets that have been encapsulated or de-encapsulated.
<b>Drops</b>	Number of packets that have exceeded the drop threshold and have been dropped.
<b>Shutdowns</b>	Number of times that packets have exceeded the shutdown threshold and the interface has been shut down.

## Sample Output

### show ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics
```

```
Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v1    xe-0/0/0.0  mvrp     Encapsulation  0        0      0
v1    xe-0/0/1.0  mvrp     Decapsulation  0        0      0
v1    xe-0/0/2.0  mvrp     Decapsulation  60634    0      0
v2    xe-0/0/0.0  cdp      Encapsulation  0        0      0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  lldp     Encapsulation  0        0      0
```

### show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics interface xe-0/0/0.0
```

```
Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v1    xe-0/0/0.0  mvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  cdp      Encapsulation  0        0      0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  lldp     Encapsulation  0        0      0
v2    xe-0/0/0.0  mvrp     Encapsulation  0        0      0
v2    xe-0/0/0.0  stp      Encapsulation  0        0      0
v2    xe-0/0/0.0  vtp      Encapsulation  0        0      0
v2    xe-0/0/0.0  vstp     Encapsulation  0        0      0
```

### show ethernet-switching layer2-protocol-tunneling statistics vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics vlan v2
```

```
Layer2 Protocol Tunneling Statistics:
VLAN  Interface  Protocol  Operation  Packets  Drops  Shutdowns
v2    xe-0/0/0.0  cdp      Encapsulation  0        0      0
v2    xe-0/0/0.0  gvrp     Encapsulation  0        0      0
```

v2	xe-0/0/0.0	lldp	Encapsulation	0	0	0
v2	xe-0/0/0.0	mvrp	Encapsulation	0	0	0
v2	xe-0/0/0.0	stp	Encapsulation	0	0	0
v2	xe-0/0/0.0	vtp	Encapsulation	0	0	0
v2	xe-0/0/0.0	vstp	Encapsulation	0	0	0
v2	xe-0/0/1.0	cdp	Decapsulation	0	0	0
v2	xe-0/0/1.0	gvrp	Decapsulation	0	0	0
v2	xe-0/0/1.0	lldp	Decapsulation	0	0	0
v2	xe-0/0/1.0	mvrp	Decapsulation	0	0	0
v2	xe-0/0/1.0	stp	Decapsulation	0	0	0
v2	xe-0/0/1.0	vtp	Decapsulation	0	0	0

## show ethernet-switching layer2-protocol-tunneling vlan

<b>Syntax</b>	show ethernet-switching-layer2-protocol-tunneling vlan <vlan-name>
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Display information about Layer 2 protocol tunneling (L2PT) on VLANs that have been configured for L2PT.
<b>Options</b>	<b>none</b> —Display information about L2PT for the VLANs on which you have configured L2PT. <b>vlan-name</b> —(Optional) Display information about L2PT for the specified VLAN.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling interface on page 1847</a></li> <li>• <a href="#">show ethernet-switching layer2-protocol-tunneling statistics on page 1849</a></li> <li>• <a href="#">show vlans</a></li> <li>• <a href="#">Example: Configuring Layer 2 Protocol Tunneling on EX Series Switches</a></li> <li>• <a href="#">Configuring Layer 2 Protocol Tunneling on EX Series Switches (CLI Procedure)</a></li> <li>• <a href="#">show vlans on page 1886</a></li> <li>• <a href="#">Configuring Layer 2 Protocol Tunneling</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching layer2-protocol-tunneling vlan on page 1853</a> <a href="#">show ethernet-switching layer2-protocol-tunneling vlan v2 on page 1853</a>
<b>Output Fields</b>	Table 117 on page 1852 lists the output fields for the <b>show ethernet-switching layer2-protocol-tunneling vlan</b> command. Output fields are listed in the approximate order in which they appear.

Table 117: show ethernet-switching layer2-protocol-tunneling vlan Output Fields

Field Name	Field Description
VLAN	Name of the VLAN on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are <b>all</b> , <b>802.1x</b> , <b>802.3ah</b> , <b>cdp</b> , <b>e-lmi</b> , <b>gvrp</b> , <b>lACP</b> , <b>lldp</b> , <b>mmrp</b> , <b>mvrp</b> , <b>stp</b> , <b>vstp</b> , and <b>vtp</b> .
Drop Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the switch begins dropping the Layer 2 PDUs.
Shutdown Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the interface is disabled.

## Sample Output

### show ethernet-switching layer2-protocol-tunneling vlan

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan
```

```
Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
                Threshold Threshold
v1             mvrp          100           200
v2             cdp            0             0
v2             cdp            0             0
v2             gvrp           0             0
```

### show ethernet-switching layer2-protocol-tunneling vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan v2
```

```
Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop          Shutdown
                Threshold Threshold
v2             cdp            0             0
v2             cdp            0             0
v2             gvrp           0             0
```

## show ethernet-switching mac-learning-log

<b>Syntax</b>	show ethernet-switching mac-learning-log
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Displays the event log of learned MAC addresses.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching table on page 1864</a></li> <li>• <a href="#">show ethernet-switching interfaces on page 1482</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching mac-learning-log on page 1854</a>
<b>Output Fields</b>	Table 118 on page 1854 lists the output fields for the <b>show ethernet-switching mac-learning-log</b> command. Output fields are listed in the approximate order in which they appear.

Table 118: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
<b>Date and Time</b>	Timestamp in UTC when the MAC operation occurred.
<b>vlan_name</b>	VLAN name. A value defined by the user for all user-configured VLANs. The name of the VLAN on which the MAC is learned.
<b>MAC</b>	Learned MAC address.
<b>Event op</b>	MAC address that are added, learned, deleted, changed or moved from one interface to another interface.
<b>Interface Name</b>	The name of the interface on which the MAC address is learned. When a MAC address is moved, there is another field with the name of the interface. The log displays the name of the interface from where the MAC address moved, and the name of the interface to where the MAC address moved.
<b>Flags</b>	Displays the MAC address flags in which the MAC event occurred. This option is for debugging purposes.

## Sample Output

### show ethernet-switching mac-learning-log

```

user@switch> show ethernet-switching mac-learning-log
Mon Jun 30 13:49:49 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on
ge-1/0/22.0 with flags: 0x2001f    << MAC address that as dynamically learned
Mon Jun 30 13:50:29 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was deleted from
ge-1/0/22.0 with flags: 0x1080    << MAC address that was deleted
Mon Jun 30 13:51:28 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was added to
ge-1/0/22.0 with flags: 0x2013f    << Static MAC address that was added
Mon Jun 30 13:51:46 2014 vlan_name v11+11 mac 00:00:00:01:01:01 was deleted from

```



```
ge-1/0/22.0 with flags: 0x1120 << delete of Static MAC address that was deleted
Mon Jun 30 13:52:03 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was learned on
ge-1/0/22.0 with flags: 0x2001f << MAC address that was dynamically learned
Mon Jun 30 13:52:11 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was moved from
ge-1/0/22.0 to ge-1/0/21.0 with flags: 0x2101f << MAC address that was moved
Mon Jun 30 13:54:24 2014 vlan_name v11+11 mac 00:10:94:00:00:02 was changed on
ge-1/0/21.0 with flags: 0x2113f << MAC address that changed from a dynamic
address to a static address
```

## show ethernet-switching mac-notification

<b>Syntax</b>	show ethernet-switching mac-notification
<b>Release Information</b>	Command introduced in Junos OS Release 9.6 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about MAC notification.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Verifying That MAC Notification Is Working Properly</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching mac-notification (MAC Notification Enabled) on page 1856</a> <a href="#">show ethernet-switching mac-notification (MAC Notification Disabled) on page 1856</a>
<b>Output Fields</b>	Table 119 on page 1856 lists the output fields for the <b>show ethernet-switching mac-notification</b> command. Output fields are listed in the order in which they appear.

Table 119: show ethernet-switching mac-notification Output Fields

Field Name	Field Description
<b>Notification Status</b>	MAC notification status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—MAC notification is enabled.</li> <li>• <b>Disabled</b>—MAC notification is disabled.</li> </ul>
<b>Notification Interval</b>	MAC notification interval in seconds.
<b>Notifications Sent</b>	Number of notifications sent to SNMP when MACs are learned or when MACs age out.
<b>Notifications Table Maxsize</b>	Maximum size of the notification table, which is populated when notifications are sent to the SNMP server.

### Sample Output

#### show ethernet-switching mac-notification (MAC Notification Enabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status           : Enabled
Notification Interval         : 30
Notifications Sent            : 0
Notifications Table Maxsize   : 256

```

### Sample Output

#### show ethernet-switching mac-notification (MAC Notification Disabled)

```

user@switch> show ethernet-switching mac-notification
Notification Status           : Disabled
Notification Interval         : 0

```

Notifications Sent : 0  
Notifications Table Maxsize : 256

## show ethernet-switching statistics aging

<b>Syntax</b>	show ethernet-switching statistics aging <brief   detail>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display media access control (MAC) aging statistics.
<b>Options</b>	<b>none</b> —(Optional) Display MAC aging statistics. <b>brief   detail</b> —(Optional) Display the specified level of output.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ethernet-switching statistics mac-learning on page 1860</a></li> <li>• <a href="#">mac-table-aging-time on page 1794</a></li> <li>• <i>Configuring MAC Table Aging</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching statistics aging on page 1859</a>
<b>Output Fields</b>	Table 120 on page 1858 lists the output fields for the <b>show ethernet-switching statistics aging</b> command. Output fields are listed in the approximate order in which they appear.

Table 120: show ethernet-switching statistics aging Output Fields

Field Name	Field Description	Level of Output
<b>Total age messages received</b>	Total number of aging messages received from the hardware.	All levels
<b>Immediate aging</b>	Aging message indicating that the entry should be removed immediately.	All levels
<b>MAC address seen</b>	Aging message indicating that the MAC address has been detected by hardware and that the aging timer should be stopped.	All levels
<b>MAC address not seen</b>	Aging message indicating that the MAC address has not been detected by the hardware and that the aging timer should be started.	All levels
<b>Error age messages</b>	The received aging message contains the following errors: <ul style="list-style-type: none"> <li>• <b>Invalid VLAN</b>—The VLAN of the packet does not exist.</li> <li>• <b>No such entry</b>—The MAC address and VLAN pair provided by the aging message does not exist.</li> <li>• <b>Static entry</b>—An unsuccessful attempt was made to age out a static MAC entry.</li> </ul>	All levels

## Sample Output

### show ethernet-switching statistics aging

```
user@switch> show ethernet-switching statistics aging
```

```
Total age messages received: 0
```

```
Immediate aging: 0, MAC address seen: 0, MAC address not seen: 0
```

```
Error age messages: 0
```

```
Invalid VLAN: 0, No such entry: 0, Static entry: 0
```

## show ethernet-switching statistics mac-learning

---

<b>Syntax</b>	<code>show ethernet-switching statistics mac-learning</code> <code>&lt;brief   detail&gt;</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.4 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display media access control (MAC) learning statistics.
<b>Options</b>	<b>none</b> —(Optional) Display MAC learning statistics for all interfaces.  <b>brief   detail</b> —(Optional) Display the specified level of output. The default is <b>brief</b> .  <b>interface <i>interface-name</i></b> —(Optional) Display MAC learning statistics for the specified interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ethernet-switching statistics aging</a></li><li>• <a href="#">show ethernet-switching mac-learning-log</a></li><li>• <a href="#">show ethernet-switching table</a></li><li>• <a href="#">show ethernet-switching interfaces</a></li><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch</a></li><li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs for EX Series Switches</a></li><li>• <a href="#">show ethernet-switching statistics aging on page 1858</a></li><li>• <a href="#">show ethernet-switching mac-learning-log on page 1854</a></li><li>• <a href="#">show ethernet-switching table on page 1864</a></li><li>• <a href="#">show ethernet-switching interfaces on page 1482</a></li><li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li><li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching statistics mac-learning on page 1861</a> <a href="#">show ethernet-switching statistics mac-learning detail on page 1862</a> <a href="#">show ethernet-switching statistics mac-learning interface ge-0/0/28 detail on page 1862</a> <a href="#">show ethernet-switching statistics mac-learning interface on page 1862</a> <a href="#">show ethernet-switching statistics mac-learning detail (QFX Series) on page 1862</a>
<b>Output Fields</b>	<a href="#">Table 121 on page 1861</a> lists the output fields for the <b>show ethernet-switching statistics mac-learning</b> command. Output fields are listed in the approximate order in which they appear.

Table 121: show ethernet-switching statistics mac-learning Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the interface for which statistics are being reported. (Displayed in the output under the heading <b>Interface</b> .)	All levels
<b>Learning message from local packets</b>	MAC learning message generated due to packets coming in on the management interface. (Displayed in the output under the heading <b>Local pkts</b> .)	All levels
<b>Learning message from transit packets</b>	MAC learning message generated due to packets coming in on network interfaces. (Displayed in the output under the heading <b>Transit pkts</b> .)	All levels
<b>Learning message with error</b>	<p>MAC learning messages received with errors (Displayed under the heading <b>Error</b>):</p> <ul style="list-style-type: none"> <li>• <b>Invalid VLAN</b>—The VLAN of the packet does not exist.</li> <li>• <b>Invalid MAC</b>—The MAC address is either NULL or a multicast MAC address.</li> <li>• <b>Security violation</b>—The MAC address is not an allowed MAC address.</li> <li>• <b>Interface down</b>—The MAC address is learned on an interface that is down.</li> <li>• <b>Incorrect membership</b>—The MAC address is learned on an interface that is not a member of the VLAN.</li> <li>• <b>Interface limit</b>—The number of MAC addresses learned on the interface has exceeded the limit.</li> <li>• <b>MAC move limit</b>—This MAC address has moved among multiple interfaces too many times in a given interval.</li> <li>• <b>VLAN limit</b>—The number of MAC addresses learned on the VLAN has exceeded the limit.</li> <li>• <b>VLAN membership limit</b>—The number of MAC addresses learned on the interface as a member of the specified VLAN (VLAN membership MAC limit) has exceeded the limit.</li> <li>• <b>Invalid VLAN index</b>—The VLAN of the packet, although configured, does not yet exist in the kernel.</li> <li>• <b>Interface not learning</b>—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked.</li> <li>• <b>No nexthop</b>—The MAC address is learned on an interface that does not have a unicast next hop.</li> <li>• <b>MAC learning disabled</b>—The MAC address is learned on an interface on which MAC learning has been disabled.</li> <li>• <b>Others</b>—The message contains some other error.</li> </ul>	All levels

## Sample Output

### show ethernet-switching statistics mac-learning

```
user@switch> show ethernet-switching statistics mac-learning
```

```
Learning stats: 0 learn msg rcvd, 0 error
Interface      Local pkts      Transit pkts      Error
ge-0/0/0.0     0                0                  0
ge-0/0/1.0     0                0                  0
ge-0/0/2.0     0                0                  0
ge-0/0/3.0     0                0                  0
```

**show ethernet-switching statistics mac-learning detail**

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error
```

```
Interface: ge-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0    Interface down: 0
  Incorrect membership: 0  Interface limit: 0
  MAC move limit: 0      VLAN limit: 0
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0
```

```
Interface: ge-0/0/1.0
Learning message from local packets: 0
Learning message from transit packets: 2
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0    Interface down: 0
  Incorrect membership: 0  Interface limit: 0
  MAC move limit: 0      VLAN limit: 0
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0
```

**show ethernet-switching statistics mac-learning interface ge-0/0/28 detail**

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/28 detail
```

```
Interface: ge-0/0/28.0
Learning message from local packets: 0
Learning message from transit packets: 5
Learning message with error: 0
  Invalid VLAN: 0      Invalid MAC: 0
  Security violation: 0    Interface down: 0
  Incorrect membership: 0  Interface limit: 0
  MAC move limit: 0      VLAN limit: 0
                          VLAN membership limit: 20
  Invalid VLAN index: 0   Interface not learning: 0
  No nexthop: 0          MAC learning disabled: 0
  Others: 0
```

**show ethernet-switching statistics mac-learning interface**

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1
Interface      Local pkts  Transit pkts  Error
ge-0/0/1.0    0           1             1
```

**show ethernet-switching statistics mac-learning detail (QFX Series)**

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error
```

```
Interface: xe-0/0/0.0
Learning message from local packets: 0
Learning message from transit packets: 1
Learning message with error: 0
```



Invalid VLAN:	0	Invalid MAC:	0
Security violation:	0	Interface down:	0
Incorrect membership:	0	Interface limit:	0
MAC move limit:	0	VLAN limit:	0
Invalid VLAN index:	0	Interface not learning:	0
No nexthop:	0	MAC learning disabled:	0
Others:	0		

Interface: xe-0/0/1.0

Learning message from local packets: 0

Learning message from transit packets: 2

Learning message with error: 0

Invalid VLAN:	0	Invalid MAC:	0
Security violation:	0	Interface down:	0
Incorrect membership:	0	Interface limit:	0
MAC move limit:	0	VLAN limit:	0
Invalid VLAN index:	0	Interface not learning:	0
No nexthop:	0	MAC learning disabled:	0
Others:	0		

## show ethernet-switching table

<b>Syntax</b>	show ethernet-switching table <brief   detail   extensive   summary> <interface <i>interface-name</i> > <management-vlan> <sort-by ( <i>name</i>   <i>tag</i> )> <vlan <i>vlan-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series. Output for private VLANs introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Displays the Ethernet switching table.
<b>Options</b>	<p><b>none</b>—(Optional) Display brief information about the Ethernet switching table.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Ethernet switching table for a specific interface.</p> <p><b>management-vlan</b>—(Optional) Display the Ethernet switching table for a management VLAN.</p> <p><b>sort-by (<i>name</i>   <i>tag</i>)</b>—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.</p> <p><b>vlan <i>vlan-name</i></b>—(Optional) Display the Ethernet switching table for a specific VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588</a></li> <li>• <a href="#">Example: Setting Up Bridging with Multiple VLANs</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ethernet-switching table on page 1865</a> <a href="#">show ethernet-switching table (Private VLANs) on page 1866</a> <a href="#">show ethernet-switching table brief on page 1866</a> <a href="#">show ethernet-switching table detail on page 1866</a> <a href="#">show ethernet-switching table extensive on page 1868</a> <a href="#">show ethernet-switching table interface on page 1869</a>
<b>Output Fields</b>	<a href="#">Table 122 on page 1864</a> lists the output fields for the <b>show ethernet-switching table</b> command. Output fields are listed in the approximate order in which they appear.

**Table 122: show ethernet-switching table Output Fields**

Field Name	Field Description	Level of Output
VLAN	Name of a VLAN.	All levels

Table 122: show ethernet-switching table Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>MAC address</b>	MAC address associated with the VLAN.	All levels
<b>Type</b>	Type of MAC address: <ul style="list-style-type: none"> <li>• <b>static</b>—The MAC address is manually created.</li> <li>• <b>learn</b>—The MAC address is learned dynamically from a packet's source MAC address.</li> <li>• <b>flood</b>—The MAC address is unknown and flooded to all members.</li> </ul>	All levels
<b>Age</b>	Time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
<b>Interfaces</b>	Interface associated with learned MAC addresses or with the <b>All-members</b> option (flood entry).	All levels
<b>Learned</b>	For learned entries, the time at which the entry was added to the Ethernet switching table.	<b>detail, extensive</b>

## Sample Output

### show ethernet-switching table

```

user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
VLAN      MAC address      Type      Age  Interfaces
F2         *                Flood     -    All-members
F2         00:00:05:00:00:03 Learn     0    xe-0/0/44.0
F2         00:19:e2:50:7d:e0 Static    -    Router
Linux      *                Flood     -    All-members
Linux      00:19:e2:50:7d:e0 Static    -    Router
Linux      00:30:48:90:54:89 Learn     0    xe-0/0/47.0
T1         *                Flood     -    All-members
T1         00:00:05:00:00:01 Learn     0    xe-0/0/46.0
T1         00:00:5e:00:01:00 Static    -    Router
T1         00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T1         00:19:e2:50:7d:e0 Static    -    Router
T10        *                Flood     -    All-members
T10        00:00:5e:00:01:09 Static    -    Router
T10        00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T10        00:19:e2:50:7d:e0 Static    -    Router
T111       *                Flood     -    All-members
T111       00:19:e2:50:63:e0 Learn     0    xe-0/0/15.0
T111       00:19:e2:50:7d:e0 Static    -    Router
T111       00:19:e2:50:ac:00 Learn     0    xe-0/0/15.0
T2         *                Flood     -    All-members
T2         00:00:5e:00:01:01 Static    -    Router
T2         00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T2         00:19:e2:50:7d:e0 Static    -    Router
T3         *                Flood     -    All-members
T3         00:00:5e:00:01:02 Static    -    Router
T3         00:19:e2:50:63:e0 Learn     0    xe-0/0/46.0
T3         00:19:e2:50:7d:e0 Static    -    Router
T4         *                Flood     -    All-members

```

```

T4          00:00:5e:00:01:03 Static      - Router
T4          00:19:e2:50:63:e0 Learn       0 xe-0/0/46.0
[output truncated]

```

### show ethernet-switching table (Private VLANs)

```

user@switch> show ethernet-switching table
Ethernet-switching table: 10 entries, 3 learned
VLAN      MAC address      Type      Age Interfaces
pvlan     *                Flood     - All-members
pvlan     00:10:94:00:00:02 Replicated - xe-0/0/28.0
pvlan     00:10:94:00:00:35 Replicated - xe-0/0/46.0
pvlan     00:10:94:00:00:46 Replicated - xe-0/0/4.0
c2        *                Flood     - All-members
c2        00:10:94:00:00:02 Learn       0 xe-0/0/28.0
c1        *                Flood     - All-members
c1        00:10:94:00:00:46 Learn       0 xe-0/0/4.0
__pvlan_pvlan_xe-0/0/46.0__ *          Flood     - All-members
__pvlan_pvlan_xe-0/0/46.0__ 00:10:94:00:00:35 Learn 0 xe-0/0/46.0

```

### show ethernet-switching table brief

```

user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
VLAN      MAC address      Type      Age Interfaces
F2        *                Flood     - All-members
F2        00:00:05:00:00:03 Learn       0 xe-0/0/44.0
F2        00:19:e2:50:7d:e0 Static      - Router
Linux     *                Flood     - All-members
Linux     00:19:e2:50:7d:e0 Static      - Router
Linux     00:30:48:90:54:89 Learn       0 xe-0/0/47.0
T1        *                Flood     - All-members
T1        00:00:05:00:00:01 Learn       0 xe-0/0/46.0
T1        00:00:5e:00:01:00 Static      - Router
T1        00:19:e2:50:63:e0 Learn       0 xe-0/0/46.0
T1        00:19:e2:50:7d:e0 Static      - Router
T10       *                Flood     - All-members
T10       00:00:5e:00:01:09 Static      - Router
T10       00:19:e2:50:63:e0 Learn       0 xe-0/0/46.0
T10       00:19:e2:50:7d:e0 Static      - Router
T111     *                Flood     - All-members
T111     00:19:e2:50:63:e0 Learn       0 xe-0/0/15.0
T111     00:19:e2:50:7d:e0 Static      - Router
T111     00:19:e2:50:ac:00 Learn       0 xe-0/0/15.0
T2        *                Flood     - All-members
T2        00:00:5e:00:01:01 Static      - Router
T2        00:19:e2:50:63:e0 Learn       0 xe-0/0/46.0
T2        00:19:e2:50:7d:e0 Static      - Router
T3        *                Flood     - All-members
T3        00:00:5e:00:01:02 Static      - Router
T3        00:19:e2:50:63:e0 Learn       0 xe-0/0/46.0
T3        00:19:e2:50:7d:e0 Static      - Router
T4        *                Flood     - All-members
T4        00:00:5e:00:01:03 Static      - Router
T4        00:19:e2:50:63:e0 Learn       0 xe-0/0/46.0
[output truncated]

```

### show ethernet-switching table detail

```

user@switch> show ethernet-switching table detail
Ethernet-switching table: 57 entries, 17 learned
F2, *

```

```
Interface(s): xe-0/0/44.0
Type: Flood
Nexthop index: 0

F2, 00:00:05:00:00:03
Interface(s): xe-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
Nexthop index: 0

F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, *
Interface(s): xe-0/0/47.0
Type: Flood
Nexthop index: 0

Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, 00:30:48:90:54:89
Interface(s): xe-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T1, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T1, 00:00:05:00:00:01
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

T1, 00:00:5e:00:01:00
Interface(s): Router
Type: Static
Nexthop index: 0

T1, 00:19:e2:50:63:e0
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
Nexthop index: 0

T1, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T10, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T10, 00:00:5e:00:01:09
Interface(s): Router
```

```
Type: Static
Nexthop index: 0

T10, 00:19:e2:50:63:e0
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T10, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

T111, *
Interface(s): xe-0/0/15.0
Type: Flood
Nexthop index: 0
[output truncated]
```

### show ethernet-switching table extensive

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 57 entries, 17 learned
F2, *
Interface(s): xe-0/0/44.0
Type: Flood
Nexthop index: 0

F2, 00:00:05:00:00:03
Interface(s): xe-0/0/44.0
Type: Learn, Age: 0, Learned: 2:03:09
Nexthop index: 0

F2, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, *
Interface(s): xe-0/0/47.0
Type: Flood
Nexthop index: 0

Linux, 00:19:e2:50:7d:e0
Interface(s): Router
Type: Static
Nexthop index: 0

Linux, 00:30:48:90:54:89
Interface(s): xe-0/0/47.0
Type: Learn, Age: 0, Learned: 2:03:08
Nexthop index: 0

T1, *
Interface(s): xe-0/0/46.0
Type: Flood
Nexthop index: 0

T1, 00:00:05:00:00:01
Interface(s): xe-0/0/46.0
Type: Learn, Age: 0, Learned: 2:03:07
```

```

    Nexthop index: 0

T1, 00:00:5e:00:01:00
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T1, 00:19:e2:50:63:e0
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:07
  Nexthop index: 0

T1, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, *
  Interface(s): xe-0/0/46.0
  Type: Flood
  Nexthop index: 0

T10, 00:00:5e:00:01:09
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T10, 00:19:e2:50:63:e0
  Interface(s): xe-0/0/46.0
  Type: Learn, Age: 0, Learned: 2:03:08
  Nexthop index: 0

T10, 00:19:e2:50:7d:e0
  Interface(s): Router
  Type: Static
  Nexthop index: 0

T111, *
  Interface(s): xe-0/0/15.0
  Type: Flood
  Nexthop index: 0
[output truncated]

```

### show ethernet-switching table interface

```

user@switch> show ethernet-switching table interface xe-0/0/1
Ethernet-switching table: 1 unicast entries

```

VLAN	MAC address	Type	Age	Interfaces
V1	*	Flood	-	All-members
V1	00:00:05:00:00:05	Learn	0	xe-0/0/1.0

## show spanning-tree bridge

**List of Syntax** [Syntax on page 1870](#)  
[Syntax \(QFX Series\) on page 1870](#)

**Syntax** show spanning-tree bridge  
 <brief | detail>  
 <msti *msti-id*>  
 <routing-instance *routing-instance-name*>  
 <vlan-id *vlan-id*>

**Syntax (QFX Series)** show spanning-tree bridge  
 <brief | detail>  
 <msti *msti-id*>  
 <vlan-id *vlan-id*>

**Release Information** Command introduced in Junos OS Release 8.4.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display the configured or calculated Spanning Tree Protocol (STP) parameters.

**Options** **none**—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).  
**brief | detail**—(Optional) Display the specified level of output.  
**msti *msti-id***—(Optional) Display STP bridge information for the specified MSTI.  
**routing-instance *routing-instance-name***—(Optional) Display STP bridge information for the specified routing instance.  
**vlan-id *vlan-id***—(Optional) Display STP bridge information for the specified VLAN.

**Required Privilege Level** view

**List of Sample Output** [show spanning-tree bridge routing-instance on page 1871](#)  
[show spanning-tree bridge msti on page 1872](#)  
[show spanning-tree bridge vlan-id \(MSTP\) on page 1873](#)  
[show spanning-tree bridge \(RSTP\) on page 1873](#)  
[show spanning-tree bridge vlan-id \(RSTP\) on page 1874](#)

**Output Fields** [Table 123 on page 1870](#) lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

**Table 123: show spanning-tree bridge Output Fields**

Field Name	Field Description
Routing instance name	Name of the routing instance under which the bridge is configured.
Enabled protocol	Spanning Tree Protocol type enabled.



Table 123: show spanning-tree bridge Output Fields (*continued*)

Field Name	Field Description
Root ID	Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Root cost	Calculated cost to reach the root bridge from the bridge where the command is entered.
Root port	Interface that is the current elected root port for this bridge.
CIST regional root	Bridge ID of the elected MSTP regional root bridge.
CIST internal root cost	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
Hello time	Configured number of seconds between transmissions of configuration bridge protocol data units (BPDUs).
Maximum age	Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).
Forward delay	How long an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Hop count	Configured maximum number of hops a BPDU can be forwarded in the MSTP region.
Message age	Number of elapsed seconds since the most recent BPDU was received.
Number of topology changes	Total number of STP topology changes detected since the routing device last booted.
Time since last topology change	Number of elapsed seconds since the most recent topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	System identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.

## Sample Output

### show spanning-tree bridge routing-instance

```

user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name       : vs1
Enabled protocol           : MSTP

```

```
STP bridge parameters for CIST
  Root ID                : 32768.00:13:c3:9e:c8:80
  Root cost               : 0
  Root port              : ge-10/2/0
  CIST regional root      : 32768.00:13:c3:9e:c8:80
  CIST internal root cost : 22000
  Hello time              : 2 seconds
  Maximum age             : 20 seconds
  Forward delay           : 15 seconds
  Hop count               : 18
  Message age             : 0
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID             : 32768.00:90:69:0b:7f:d1
    Extended system ID    : 1

STP bridge parameters for MSTI 1
  MSTI regional root      : 32769.00:13:c3:9e:c8:80
  Root cost               : 22000
  Root port              : ge-10/2/0
  Hello time              : 2 seconds
  Maximum age             : 20 seconds
  Forward delay           : 15 seconds
  Hop count               : 18
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID             : 32769.00:90:69:0b:7f:d1
    Extended system ID    : 1

STP bridge parameters for MSTI 2
  MSTI regional root      : 32770.00:13:c3:9e:c8:80
  Root cost               : 22000
  Root port              : ge-10/2/0
  Hello time              : 2 seconds
  Maximum age             : 20 seconds
  Forward delay           : 15 seconds
  Hop count               : 18
  Number of topology changes : 1
  Time since last topology change : 1191 seconds
  Local parameters
    Bridge ID             : 32770.00:90:69:0b:7f:d1
    Extended system ID    : 1
```

### show spanning-tree bridge msti

```
user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name      : vs1
Enabled protocol          : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root      : 32769.00:13:c3:9e:c8:80
  Root cost               : 22000
  Root port              : xe-10/2/0
  Hello time              : 2 seconds
  Maximum age             : 20 seconds
  Forward delay           : 15 seconds
  Hop count               : 18
```

```

Number of topology changes      : 1
Time since last topology change : 1191 seconds
Local parameters
  Bridge ID                     : 32769.00:90:69:0b:7f:d1
  Extended system ID            : 1

```

### show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1101 routing-instance vs1 detail
```

```

STP bridge parameters
Routing instance name          : vs1
Enabled protocol               : MSTP

STP bridge parameters for CIST
Root ID                       : 32768.00:13:c3:9e:c8:80
Root cost                     : 0
Root port                     : xe-10/2/0
CIST regional root            : 32768.00:13:c3:9e:c8:80
CIST internal root cost       : 22000
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Hop count                     : 18
Message age                   : 0
Number of topology changes    : 0
Local parameters
  Bridge ID                   : 32768.00:90:69:0b:7f:d1
  Extended system ID          : 1
  Hello time                  : 2 seconds
  Maximum age                 : 20 seconds
  Forward delay               : 15 seconds
  Path cost method            : 32 bit
  Maximum hop count           : 20

```

### show spanning-tree bridge (RSTP)

```
user@host> show spanning-tree bridge
```

```

STP bridge parameters
Routing instance name          : GLOBAL
Enabled protocol               : RSTP
Root ID                       : 28672.00:90:69:0b:3f:d0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Message age                   : 0
Number of topology changes    : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                   : 28672.00:90:69:0b:3f:d0
  Extended system ID          : 0

STP bridge parameters for bridge VLAN 10
Root ID                       : 28672.00:90:69:0b:3f:d0
Hello time                    : 2 seconds
Maximum age                   : 20 seconds
Forward delay                 : 15 seconds
Message age                   : 0
Number of topology changes    : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                   : 28672.00:90:69:0b:3f:d0

```

```
Extended system ID          : 0

STP bridge parameters for bridge VLAN 20
Root ID                     : 28672.00:90:69:0b:3f:d0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Message age                  : 0
Number of topology changes   : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                  : 28672.00:90:69:0b:3f:d0
  Extended system ID         : 0
```

#### show spanning-tree bridge vlan-id (RSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name        : GLOBAL
Enabled protocol             : RSTP

STP bridge parameters for VLAN 10
Root ID                     : 28672.00:90:69:0b:3f:d0
Hello time                   : 2 seconds
Maximum age                  : 20 seconds
Forward delay                : 15 seconds
Message age                  : 0
Number of topology changes   : 58
Time since last topology change : 14127 seconds
Local parameters
  Bridge ID                  : 28672.00:90:69:0b:3f:d0
  Extended system ID         : 0
```

## show spanning-tree interface

<b>List of Syntax</b>	<a href="#">Syntax on page 1875</a> <a href="#">Syntax (EX Series Switches and the QFX Series) on page 1875</a>
<b>Syntax</b>	<pre>show spanning-tree interface &lt;brief   detail&gt; &lt;msti <i>msti-id</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt; &lt;vlan-id <i>vlan-id</i>&gt;</pre>
<b>Syntax (EX Series Switches and the QFX Series)</b>	<pre>show spanning-tree interface &lt;brief   detail&gt; &lt;msti <i>msti-id</i>&gt; &lt;vlan-id <i>vlan-id</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display the configured or calculated interface-level STP parameters.
<b>Options</b>	<p><b>none</b>—Display brief STP interface information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>msti <i>msti-id</i></b>—(Optional) Display STP interface information for the specified MST instance.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display STP interface information for the specified routing instance.</p> <p><b>vlan-id <i>vlan-id</i></b>—(Optional) Display STP interface information for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show spanning-tree interface on page 1876</a> <a href="#">show spanning-tree interface (QFX Series) on page 1877</a> <a href="#">show spanning-tree interface detail on page 1877</a> <a href="#">show spanning-tree interface msti on page 1879</a> <a href="#">show spanning-tree interface vlan-id on page 1879</a> <a href="#">show spanning-tree interface (VSTP) on page 1880</a> <a href="#">show spanning-tree interface vlan-id (VSTP) on page 1880</a>
<b>Output Fields</b>	<p><a href="#">Table 124 on page 1875</a> lists the output fields for the <b>show spanning-tree interface</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 124: show spanning-tree Interface Output Fields**

Field Name	Field Description
<b>Interface name</b>	Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance.

Table 124: show spanning-tree Interface Output Fields (*continued*)

Field Name	Field Description
<b>Port ID</b>	Logical interface identifier configured to participate in the MSTP or VSTP instance.
<b>Designated port ID</b>	Port ID of the designated port for the LAN segment to which this interface is attached.
<b>Designated bridge ID</b>	Bridge ID of the designated bridge for the LAN segment to which this interface is attached.
<b>Port Cost</b>	Configured cost for the interface.
<b>Port State</b>	STP port state: forwarding ( <b>FWD</b> ), blocking ( <b>BLK</b> ), listening, learning, or disabled.
<b>Port Role</b>	MSTP, VSTP, or RSTP port role: designated ( <b>DESG</b> ), backup ( <b>BKUP</b> ), alternate ( <b>ALT</b> ), ( <b>ROOT</b> ), or Root Prevented ( <b>Root-Prev</b> ).
<b>Link type</b>	MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.
<b>Alternate</b>	Identifies the interface as an MSTP, VSTP, or RSTP alternate root port ( <b>Yes</b> ) or nonalternate root port ( <b>No</b> ).
<b>Boundary Port</b>	Identifies the interface as an MSTP regional boundary port ( <b>Yes</b> ) or nonboundary port ( <b>No</b> ).

## Sample Output

### show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32768.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32768.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32768.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32768.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32768.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32768.0090690b47d1	2000	FWD	DESG

```
Spanning tree interface parameters for instance 1
```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32769.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32769.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32769.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32769.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32769.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32769.0090690b47d1	2000	FWD	DESG

## Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32770.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32770.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32770.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32770.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32770.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32770.0090690b47d1	2000	FWD	DESG

## show spanning-tree interface (QFX Series)

```
user@1f0> show spanning-tree interface routing-instance vs1 detail
```

## Spanning tree interface parameters for instance 0

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32768.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32768.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32768.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32768.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32768.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32768.0090690b47d1	2000	FWD	DESG

## Spanning tree interface parameters for instance 1

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32769.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32769.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32769.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32769.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32769.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32769.0090690b47d1	2000	FWD	DESG

## Spanning tree interface parameters for instance 2

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ae1	128:1	128:1	32770.0090690b47d1	1000	FWD	DESG
ge-2/1/2	128:2	128:2	32770.0090690b47d1	20000	FWD	DESG
ge-2/1/5	128:3	128:3	32770.0090690b47d1	29999	FWD	DESG
ge-2/2/1	128:4	128:26	32770.0013c39ec880	20000	FWD	ROOT
xe-9/2/0	128:5	128:5	32770.0090690b47d1	2000	FWD	DESG
xe-9/3/0	128:6	128:6	32770.0090690b47d1	2000	FWD	DESG

## show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
```

## Spanning tree interface parameters for instance 0

```
Interface name           : ae1
Port identifier          : 128.1
Designated port ID      : 128.1
Port cost                : 1000
Port state               : Forwarding
Designated bridge ID     : 32768.00:90:69:0b:47:d1
Port role                : Designated
Link type                : Pt-Pt/NONEDGE
```

```
Boundary port                : No

Interface name                : ge-2/1/2
Port identifier               : 128.2
Designated port ID           : 128.2
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

Interface name                : ge-2/1/5
Port identifier               : 128.3
Designated port ID           : 128.3
Port cost                     : 29999
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

Interface name                : ge-2/2/1
Port identifier               : 128.4
Designated port ID           : 128.26
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:13:c3:9e:c8:80
Port role                     : Root
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

Interface name                : xe-9/2/0
Port identifier               : 128.5
Designated port ID           : 128.5
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No

Interface name                : xe-9/3/0
Port identifier               : 128.6
Designated port ID           : 128.6
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                 : No
```

#### Spanning tree interface parameters for instance 1

```
Interface name                : ae1
Port identifier               : 128.1
Designated port ID           : 128.1
Port cost                     : 1000
Port state                    : Forwarding
Designated bridge ID         : 32768.00:90:69:0b:47:d1
```



```

Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/2
Port identifier      : 128.2
Designated port ID   : 128.2
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/1/5
Port identifier      : 128.3
Designated port ID   : 128.3
Port cost            : 29999
Port state           : Forwarding
Designated bridge ID : 32768.00:90:69:0b:47:d1
Port role           : Designated
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

Interface name       : ge-2/2/1
Port identifier      : 128.4
Designated port ID   : 128.26
Port cost            : 20000
Port state           : Forwarding
Designated bridge ID : 32768.00:13:c3:9e:c8:80
Port role           : Root
Link type           : Pt-Pt/NONEDGE
Boundary port       : No

...

```

### show spanning-tree interface msti

```

user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
Spanning tree interface parameters for instance 1

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
xe-7/0/0	128:1	128:1	32769.0090690b4fd1	2000	FWD	DESG
ge-5/1/0	128:2	128:2	32769.0090690b4fd1	20000	FWD	DESG
ge-5/1/1	128:3	128:3	32769.0090690b4fd1	20000	FWD	DESG
ae1	128:4	128:1	32769.0090690b47d1	10000	BLK	ALT
ge-5/1/4	128:5	128:3	32769.0090690b47d1	20000	BLK	ALT
xe-7/2/0	128:6	128:6	32769.0090690b47d1	2000	FWD	ROOT

### show spanning-tree interface vlan-id

```

user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

```

Interface	Port ID	Designated port ID	Designated bridge ID	Port Cost	State	Role
ge-11/0/5	128:1	128:1	32768.0090690b7fd1	20000	FWD	DESG
ge-11/0/6	128:2	128:1	32768.0090690b7fd1	20000	BLK	BKUP
ge-11/1/0	128:3	128:2	32768.0090690b4fd1	20000	BLK	ALT
ge-11/1/1	128:4	128:3	32768.0090690b4fd1	20000	BLK	ALT

ge-11/1/4	128:5	128:1	32768.0090690b47d1	20000	BLK	ALT
xe-10/0/0	128:6	128:5	32768.0090690b4fd1	2000	BLK	ALT
xe-10/2/0	128:7	128:4	32768.0090690b47d1	2000	FWD	ROOT

**show spanning-tree interface (VSTP)**

```
user@host> show spanning-tree interface
```

```
Spanning tree interface parameters for instance 0
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

```
Spanning tree interface parameters for VLAN 10
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

```
Spanning tree interface parameters for VLAN 20
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

**show spanning-tree interface vlan-id (VSTP)**

```
user@host> show spanning-tree interface vlan-id 10
```

```
Spanning tree interface parameters for VLAN 10
```

Interface	Port ID	Designated port ID	Designated bridge ID	Cost	State	Role
ge-1/0/1	128:1	128:1	28672.0090690b3fe0	20000	FWD	DESG
ge-1/0/2	128:2	128:2	28672.0090690b3fe0	20000	FWD	DESG

## show spanning-tree mstp configuration

<b>List of Syntax</b>	<a href="#">Syntax on page 1881</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 1881</a>
<b>Syntax</b>	show spanning-tree mstp configuration <brief   detail> <routing-instance <i>routing-instance-name</i> >
<b>Syntax (EX Series Switch and the QFX Series)</b>	show spanning-tree mstp configuration <brief   detail>
<b>Release Information</b>	Command introduced in Junos OS Release 8.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the MSTP configuration.
<b>Options</b>	<b>none</b> —Display MSTP configuration information.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>routing-instance <i>routing-instance-name</i></b> —(Optional) Display MSTP configuration information for the specified routing instance.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show spanning-tree mstp configuration detail on page 1882</a> <a href="#">show spanning-tree mstp configuration detail (QFX Series) on page 1882</a>
<b>Output Fields</b>	<a href="#">Table 125 on page 1881</a> lists the output fields for the <b>show spanning-tree mstp configuration</b> command. Output fields are listed in the approximate order in which they appear.

**Table 125: show spanning-tree mstp configuration Output Fields**

Field Name	Field Description
<b>Context id</b>	Internally generated identifier.
<b>Region name</b>	MSTP region name carried in the MSTP BPDUs.
<b>Revision</b>	Revision number of the MSTP configuration.
<b>Configuration digest</b>	Numerical value derived from the VLAN-to-instance mapping table.
<b>MSTI</b>	MST instance identifier.
<b>Member VLANs</b>	VLAN identifiers associated with the MSTI.

## Sample Output

### show spanning-tree mstp configuration detail

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

### show spanning-tree mstp configuration detail (QFX Series)

```
user@1f0> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier      : 1
Region name            : henry
Revision               : 3
Configuration digest    : 0x6da4b5c4fd587757eef35675365e1

MSTI      Member VLANs
  0 0-99,101-199,201-4094
  1 100
  2 200
```

## show spanning-tree statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 1883</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 1883</a>
<b>Syntax</b>	<pre>show spanning-tree statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show spanning-tree statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>   vlan <i>vlan-id</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for QFX Series switches.</p>
<b>Description</b>	Display STP statistics.
<b>Options</b>	<p><b>none</b>—Display brief STP statistics.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display STP statistics for the specified interface.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display STP statistics for the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show spanning-tree statistics routing-instance on page 1884</a> <a href="#">show spanning-tree statistics interface routing-instance detail on page 1884</a>
<b>Output Fields</b>	<p><a href="#">Table 126 on page 1883</a> lists the output fields for the <b>show spanning-tree statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 126: show spanning-tree statistics Output Fields**

Field Name	Field Description
Message type	Type of message being counted.
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
BPDUs sent in last interval	Number of BPDUs sent within a specified interval.
BPDUs received in last interval	Number of BPDUs received within a specified interval.

Table 126: show spanning-tree statistics Output Fields (*continued*)

Field Name	Field Description
<b>Interface</b>	Interface for which the statistics are being displayed.
<b>Next BPDU transmission</b>	Number of seconds until the next BPDU is scheduled to be sent.

## Sample Output

### show spanning-tree statistics routing-instance

```
user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type           : bpdus
BPDUs sent             : 1396
BPDUs received         : 1027
BPDUs sent in last interval : 5      (duration: 4 sec)
BPDUs received in last interval: 4    (duration: 4 sec)
```

### show spanning-tree statistics interface routing-instance detail

```
user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface  BPDUs sent  BPDUs received  Next BPDU
                                     transmission
ge-11/1/4      7           190           0
```

## show system statistics arp

<b>Syntax</b>	show system statistics arp
<b>Release Information</b>	Command introduced in Junos OS Release 9.6 for EX Series switches.
<b>Description</b>	Display system-wide Address Resolution Protocol (ARP) statistics.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Proxy ARP on an EX Series Switch</i></li> <li>• <a href="#">Verifying That Proxy ARP Is Working Correctly on page 1834</a></li> </ul>

## show system statistics arp

```

user@switch> show system statistics arp
arp:
    90060 datagrams received
    34 ARP requests received
    610 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 unrestricted proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast target address
    0 datagrams with my own hardware address
    0 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    294 datagrams with source address duplicate to mine
    89113 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    309 ARP requests sent
    35 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

```

## show vlans

**Syntax** `show vlans`  
`<brief | detail | extensive>`  
`<dot1q-tunneling>`  
`<sort-by (tag | name)>`  
`<vlan-range-name>`

**Release Information** Command introduced in Junos OS Release 11.1 for the QFX Series.  
Option **dot1q-tunneling** added in Junos OS Release 12.1 for the QFX Series.

**Description** Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the **show vlans** command displays both tagged and untagged membership for those VLANs.



**NOTE:** When a series of VLANs is created using the `vlan-range` statement, such VLAN names are preceded and followed by a double underscore. For example, a series of VLANs using the VLAN range 1 through 3 and the base VLAN name `marketing` would be displayed as `__marketing_1__`, `__marketing_2__`, and `__marketing_3__`.



**NOTE:** To display an 802.1X supplicant successfully authenticated in multiple-suppliant mode with dynamic VLAN movement, use the `show vlans vlan-name extensive` operational mode command, where `vlan-name` is the dynamic VLAN.

**Options** **none**—Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**sort-by (tag | name)**—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

**vlan-range-name**—(Optional) Display VLANs in ascending order of VLAN range names.

**Required Privilege Level** `view`

**Related Documentation**

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series on page 1588](#)
- [Example: Setting Up Bridging with Multiple VLANs](#)
- [Understanding Bridging and VLANs](#)
- [show ethernet-switching interfaces on page 1482](#)



**List of Sample Output**

- [show vlans on page 1889](#)
- [show vlans \(Private VLANs\) on page 1889](#)
- [show vlans brief on page 1890](#)
- [show vlans detail on page 1890](#)
- [show vlans extensive \(Port-Based\) on page 1891](#)
- [show vlans \(Q-in-Q Tunneling\) on page 1892](#)
- [show vlans extensive \(Q-in-Q Tunneling\) on page 1892](#)
- [show vlans extensive \(Q-in-Q Tunneling and L2TP\) on page 1892](#)
- [show vlans sort-by tag on page 1892](#)
- [show vlans sort-by name on page 1893](#)
- [show vlans tag on page 1894](#)

**Output Fields** Table 89 on page 1497 lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

**Table 127: show vlans Output Fields**

Field Name	Field Description	Level of Output
<b>Name</b>	Name of a VLAN.	none, <b>brief</b>
<b>Tag</b>	802.1Q tag applied to this VLAN. If <b>none</b> is displayed, no tag is applied.	All levels
<b>Interfaces</b>	Interface associated with learned MAC addresses or <b>All-members</b> option (flood entry). An asterisk (*) beside the interface indicates that the interface is <b>UP</b> .	All levels
<b>Address</b>	IP address.	none, <b>brief</b>
<b>Ports Active /Total</b>	Number of interfaces associated with a VLAN: <b>Active</b> indicates interfaces that are <b>UP</b> , and <b>Total</b> indicates interfaces that are active and inactive.	<b>brief</b>
<b>VLAN</b>	Name of a VLAN.	<b>detail, extensive</b>
<b>Admin state</b>	State of the interface. Values are:  <b>enabled</b> —The interface is turned on, and the physical link is operational and can pass packets.	<b>detail,extensive</b>
<b>MAC learning Status</b>	Indicates if MAC learning is disabled.	<b>detail, extensive</b>
<b>Description</b>	Description for the VLAN.	<b>detail,extensive</b>
<b>Primary IP</b>	Primary IP address associated with a VLAN.	<b>detail</b>
<b>Number of interfaces</b>	Number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.	<b>detail, extensive</b>
<b>STP</b>	Spanning tree associated with a VLAN.	<b>detail,extensive</b>
<b>Tagged interfaces</b>	Tagged interfaces with which a VLAN is associated.	<b>detail,extensive</b>

Table 127: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
Untagged interfaces	Untagged interfaces with which a VLAN is associated.	detail. extensive
Dot1q Tunneling Status	Indicates if Q-in-Q tunneling is enabled.	extensive
Customer VLAN ranges	List of customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).	extensive
Private VLAN Mode	The private VLAN mode for this VLAN. Values include <b>Primary</b> , <b>Isolated</b> , and <b>Community</b> .	extensive
Primary VLAN	Primary VLAN tag for this secondary VLAN.	extensive
Internal Index	VLAN index internal to Junos OS software.	extensive
Origin	Manner in which the VLAN was created: <b>static</b> or <b>learn</b> .	extensive
Protocol	Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X,	extensive
IP addresses	IP address associated with a VLAN.	extensive
Number of MAC entries	For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.	extensive
Number of mapping rules	Number of mapping rules for Q-in-Q tunneling ( <b>Push</b> ) and VLAN translation ( <b>Swap</b> ).	
Secondary VLANs	Secondary VLANs associated with a primary VLAN.	extensive
Isolated VLANs	Isolated VLANs associated with a primary VLAN.	extensive
Community VLANs	Community VLANs associated with a primary VLAN.	extensive
VLANs summary	VLAN counts: <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of VLANs on the switch.</li> <li>• <b>Configured VLANs</b>—Number of VLANs that are based on user-configured settings.</li> <li>• <b>Internal VLANs</b>—Number of VLANs created by the system with no explicit configuration or protocol—for example, the <b>default</b> VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership.</li> <li>• <b>Temporary VLANs</b>—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them.</li> </ul>	All levels

Table 127: show vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Dot1q VLANs summary</b>	802.1Q VLAN counts: <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of 802.1Q-tagged and untagged VLANs on the switch.</li> <li>• <b>Tagged VLANs</b>—Number of 802.1Q-tagged VLANs.</li> <li>• <b>Untagged VLANs</b>—Number of untagged 802.1Q VLANs.</li> <li>• <b>Private VLAN</b>—Counts of the following kinds of 802.1Q private VLANs (PVLANS):           <ul style="list-style-type: none"> <li>• <b>Primary VLANs</b>—Number of primary forwarding private VLANs.</li> <li>• <b>Community VLANs</b>—Number of community transporting and forwarding private VLANs.</li> <li>• <b>Isolated VLANs</b>—Number of isolated receiving and forwarding private VLANs.</li> <li>• <b>Inter-switch-isolated VLANs</b>—Number of inter-switch isolated receiving and forwarding private VLANs.</li> </ul> </li> </ul>	All levels
<b>Dot1q Tunneled VLANs summary</b>	Q-in-Q-tunneled VLAN counts: <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of Q-in-Q-tunneled VLANs on the switch.</li> <li>• <b>Private VLAN</b>—Counts of primary, community, and isolated Q-in-Q-tunneled private VLANs (PVLANS).</li> </ul>	All levels

## Sample Output

### show vlans

```
user@switch> show vlans
```

Name	Tag	Interfaces
default	None	xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0, xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0, xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0, xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0, xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0
v0001	1	xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0
v0002	2	None
v0003	3	None
v0004	4	None
v0005	5	None

### show vlans (Private VLANs)

```
user@switch> show vlans
```

Name	Tag	Interfaces
__pvlan_pvlan_xe-0/0/46.0__		

```

c1                xe-0/0/44.0*, xe-0/0/46.0*
c2                xe-0/0/4.0*, xe-0/0/44.0*
default           xe-0/0/28.0*, xe-0/0/44.0*
pvlan             500
                  None
                  xe-0/0/4.0*, xe-0/0/28.0*, xe-0/0/44.0*, xe-0/0/46.0*

```

## show vlans brief

```
user@switch> show vlans brief
```

Name	Tag	Address	Ports Active/Total
default	None		0/23
v0001	1		0/4
v0002	2		0/0
v0003	3		0/0
v0004	4		0/0
v0005	5		0/0
v0006	6		0/0
v0007	7		0/0
v0008	8		0/0
v0009	9		0/0
v0010	10		0/2
v0011	11		0/0
v0012	12		0/0
v0013	13		0/0
v0014	14		0/0
v0015	15		0/0
v0016	16		0/0

## show vlans detail

```
user@switch> show vlans detail
```

```
VLAN: default, Tag: Untagged, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 23 (Active = 0)
```

```
STP: None, RTG: None
```

```
Untagged interfaces: xe-0/0/34.0, xe-0/0/33.0, xe-0/0/32.0, xe-0/0/31.0,
xe-0/0/30.0, xe-0/0/29.0, xe-0/0/28.0, xe-0/0/27.0, xe-0/0/26.0,
xe-0/0/25.0, xe-0/0/19.0, xe-0/0/18.0, xe-0/0/17.0, xe-0/0/16.0,
xe-0/0/15.0, xe-0/0/14.0, xe-0/0/13.0, xe-0/0/11.0, xe-0/0/9.0, xe-0/0/8.0,
xe-0/0/3.0, xe-0/0/2.0, xe-0/0/1.0,
```

```
Tagged interfaces: None
```

```
VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 4 (Active = 0)
```

```
Dot1q Tunneling Status: Enabled
```

```
STP: None, RTG: None
```

```
Untagged interfaces: None
```

```
Tagged interfaces: xe-0/0/24.0, xe-0/0/23.0, xe-0/0/22.0, xe-0/0/21.0,
```

```
VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
```

```
Description: None
```

```
Primary IP: None, Number of interfaces: 0 (Active = 0)
```

```
STP: None, RTG: None
```

```
Untagged interfaces: None
```

```
Tagged interfaces: None
```

```

VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
Description: None
Primary IP: None, Number of interfaces: 0 (Active = 0)
STP: None, RTG: None
Untagged interfaces: None
Tagged interfaces: None

VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)

```

### show vlans extensive (Port-Based)

```

user@switch> show vlans extensive
VLAN: default, created at Mon Feb  4 12:13:47 2008
Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
Description: None
Customer VLAN ranges:
    1-4100
Protocol: Port based
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
    xe-0/0/34.0 (untagged, access)
    xe-0/0/33.0 (untagged, access)
    xe-0/0/32.0 (untagged, access)
    xe-0/0/31.0 (untagged, access)
    xe-0/0/30.0 (untagged, access)
    xe-0/0/29.0 (untagged, access)
    xe-0/0/28.0 (untagged, access)
    xe-0/0/27.0 (untagged, access)
    xe-0/0/26.0 (untagged, access)
    xe-0/0/25.0 (untagged, access)
    xe-0/0/19.0 (untagged, access)
    xe-0/0/18.0 (untagged, access)
    xe-0/0/17.0 (untagged, access)
    xe-0/0/16.0 (untagged, access)
    xe-0/0/15.0 (untagged, access)
    xe-0/0/14.0 (untagged, access)
    xe-0/0/13.0 (untagged, access)
    xe-0/0/11.0 (untagged, access)
    xe-0/0/9.0 (untagged, access)
    xe-0/0/8.0 (untagged, access)
    xe-0/0/3.0 (untagged, access)
    xe-0/0/2.0 (untagged, access)
    xe-0/0/1.0 (untagged, access)

Secondary VLANs: Isolated 1, Community 1
Isolated VLANs :
    __pvlan_pvlan_xe-0/0/3.0__
Community VLANs :
    comm1

VLAN: v0001, created at Mon Feb  4 12:13:47 2008
Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)

```

```
xe-0/0/24.0 (tagged, trunk)
xe-0/0/23.0 (tagged, trunk)
xe-0/0/22.0 (tagged, trunk)
xe-0/0/21.0 (tagged, trunk)
```

```
VLAN: v0002, created at Mon Feb  4 12:13:47 2008
Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None
```

```
VLAN: v0003, created at Mon Feb  4 12:13:47 2008
Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
Description: None
Protocol: Port based, Layer 3 interface: None
IP addresses: None
STP: None, RTG: None.
Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
None
```

#### show vlans (Q-in-Q Tunneling)

```
user@switch> show vlans dot1q-tunneling
Name      Tag      Interfaces
sv100     100      xe-0/0/4.0*, xe-0/0/15.0*
```

#### show vlans extensive (Q-in-Q Tunneling)

```
user@switch> show vlans sv100 extensive
VLAN: sv100, Created at: Sat Sep 10 12:53:52 2011
802.1Q Tag: 100, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
    10-20
    40-50
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 1), Untagged 0 (Active = 0)
    ge-0/0/0.0, tagged, trunk

Number of mapping rules:
    Push 1 (Active = 0), Policy 0 (Active = 0), Swap 0 (Active = 0)

    xe-0/0/3.0*, 300, push
```

#### show vlans extensive (Q-in-Q Tunneling and L2TP)

```
user@switch> show vlans v1 extensive
VLAN: v1, Created at: Fri Mar 2 05:07:38 2012
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled
```

#### show vlans sort-by tag

```
user@switch> show vlans sort-by tag
Name      Tag      Interfaces
default   None
__vlan-x_1__  1
```

__vlan-x_2__	2	None
__vlan-x_3__	3	None
__vlan-x_4__	4	None
__vlan-x_5__	5	None
__vlan-x_6__	6	None
__vlan-x_7__	7	None
__vlan-x_8__	8	None
__vlan-x_9__	9	None
__vlan-x_10__	10	None
__vlan-x_11__	11	None
__vlan-x_12__	12	None
__vlan-x_13__	13	None
__vlan-x_14__	14	None
__vlan-x_15__	15	None
__vlan-x_16__	16	None
__vlan-x_17__	17	None
__vlan-x_18__	18	None
__vlan-x_19__	19	None
__vlan-x_20__	20	None

### show vlans sort-by name

```
user@switch> show vlans sort-by employee
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	

```
__employee_128__ 128    xe-0/0/22.0*
__employee_129__ 129    xe-0/0/22.0*
__employee_130__ 130    xe-0/0/22.0*
__employee_130__ 130    xe-0/0/22.0*
```

### show vlans tag

```
user@switch> show vlans employee
```

Name	Tag	Interfaces
__employee_120__	120	xe-0/0/22.0*
__employee_121__	121	xe-0/0/22.0*
__employee_122__	122	xe-0/0/22.0*
__employee_123__	123	xe-0/0/22.0*
__employee_124__	124	xe-0/0/22.0*
__employee_125__	125	xe-0/0/22.0*
__employee_126__	126	xe-0/0/22.0*
__employee_127__	127	xe-0/0/22.0*
__employee_128__	128	xe-0/0/22.0*
__employee_129__	129	xe-0/0/22.0*
__employee_130__	130	xe-0/0/22.0*



# Troubleshooting

- [Troubleshooting Procedures on page 1895](#)

## Troubleshooting Procedures

---

- [Troubleshooting Ethernet Switching on page 1895](#)

### Troubleshooting Ethernet Switching

**Problem**    **Description:** Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

**Solution**    Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 15 to 1,000,000 seconds.)

```
[edit vlans]
user@switch# set vlans sales mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message,

thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

- Related Documentation**
- *arp*
  - [mac-table-aging-time on page 1794](#)

## PART 6

# OVSDB and VXLAN

- [Overview on page 1899](#)
- [Configuration on page 1917](#)
- [Administration on page 2009](#)
- [Troubleshooting on page 2045](#)



# Overview

- [OVSDB Overview on page 1899](#)
- [VXLAN Overview on page 1912](#)

## OVSDB Overview

- [Open vSwitch Database Support on Juniper Networks Devices on page 1899](#)
- [Understanding the Junos OS Implementation of VXLAN and OVSDB in a VMware NSX for Multi-Hypervisor Environment for the Data Center on page 1900](#)
- [Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices on page 1902](#)
- [Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903](#)
- [Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDB on page 1904](#)
- [Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment on page 1905](#)
- [Open vSwitch Database Schema For Physical Devices on page 1910](#)

## Open vSwitch Database Support on Juniper Networks Devices

[Table 128 on page 1899](#) lists the Juniper Networks devices that support the Open vSwitch Database (OVSDB) management protocol. For each device, the table also includes the OVSDB software package and the initial Junos OS release that must be installed for OVSDB support. The OVSDB software package release must be the same as the Junos OS release running on the device.

Table 128: OVSDB Support on Junos OS Devices

Junos OS Device	OVSDB Software Package	Junos OS Release
MX80 3D Universal Edge Router	<code>jsdn-powerpc-release</code>	14.1R2
MX240, MX480, MX960 3D Universal Edge Routers	<code>jsdn-i386-release</code>	14.1R2
QFX5100 Ethernet Switch	<code>jsdn-i386-release</code>	14.1X53-D10

**Related Documentation** • [Installing Open vSwitch Database Components on Juniper Networks Devices on page 1983](#)

## Understanding the Junos OS Implementation of VXLAN and OVSDDB in a VMware NSX for Multi-Hypervisor Environment for the Data Center

Some Juniper Networks devices support Virtual Extensible LAN (VXLAN) and the Open vSwitch Database (OVSDDB) management protocol. (For information about the Juniper Networks devices on which OVSDDB is supported and the Junos OS release in which support is introduced, see [“Open vSwitch Database Support on Juniper Networks Devices” on page 1899](#).) Support for VXLAN and OVSDDB enables the Juniper Networks devices in a physical network to be integrated into a virtual network.

The implementation of VXLAN and OVSDDB on Juniper Networks devices is supported in a VMware NSX for Multi-Hypervisor environment for the data center. [Table 129 on page 1900](#) outlines the components that compose this environment and products that are typically deployed for each component.

**Table 129: NSX Multi-Hypervisor Components and Products That Can Be Implemented**

Component	Products
Cloud management platform (CMP)	CloudStack OpenStack Custom CMP
Network virtualization platform	NSX for Multi-Hypervisor
Hypervisor	Kernel-based Virtual Machine (KVM) Red Hat VMware ESXi Xen <b>NOTE:</b> Juniper Networks supports KVM and ESXi only.
Virtual switch	Open vSwitch (OVS) NSX vSwitch
SDN controller	NSX multi-hypervisor controller <b>NOTE:</b> Juniper Networks supports NSX multi-hypervisor controller version 4.0.3.
Overlay protocol	VXLAN
MAC learning protocol	OVSDDB

Figure 28 on page 1901 shows a high-level view of the architecture into which the NSX for Multi-Hypervisor platform fits, while Figure 29 on page 1901 provides a more detailed representation of the components in the virtual and physical networks.

Figure 28: High-Level NSX for Multi-Hypervisor Architecture

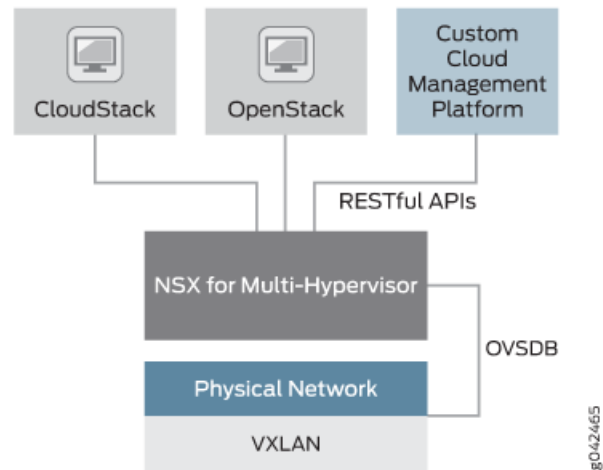



Figure 29: Integration of Juniper Networks Device That Implements VXLAN and OVSDB into NSX for Multi-Hypervisor Environment



In the data center topology shown in Figure 29 on page 1901, the physical and virtual servers need to communicate. To facilitate this communication, a Juniper Networks device that supports VXLAN is strategically deployed so that it serves as a *gateway*, which is also known as a hardware virtual tunnel endpoint (VTEP), at the edge of the physical network. Working in conjunction with the software VTEP, which is deployed at the edge of the virtual network, the hardware VTEP encapsulates packets from resources on physical

server 1 with a VXLAN header, and after the packets traverse the Layer 3 transport network, the software VTEP removes the VXLAN header from the packets and forwards the packets to the appropriate virtual machines (VMs). In essence, the encapsulation and de-encapsulation of packets by the hardware and software VTEPs enables components in the physical and virtual networks to coexist without one needing to understand the workings of the other.

The same Juniper Networks device that acts as hardware VTEP in [Figure 29 on page 1901](#) implements OVSDb, which enables this device to learn the MAC addresses of physical server 1 and other physical servers, and publish the addresses in the OVSDb schema, which was defined for physical devices. In the virtual network, one or more NSX controllers collect the MAC addresses of HOST 1 and other virtual servers, and publish the addresses in the OVSDb schema. Using the OVSDb schema, components in the physical and virtual networks can exchange MAC addresses, as well as statistical information, enabling the components to learn about and reach each other in their respective networks.

**Related  
Documentation**

- [Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices on page 1902](#)
- [Open vSwitch Database Schema For Physical Devices on page 1910](#)

## Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices

The Juniper Networks Junos operating system (Junos OS) implementation of the Open vSwitch Database (OVSDb) management protocol provides a means through which VMware NSX controllers and Juniper Networks devices that support OVSDb can communicate. In an NSX multi-hypervisor environment, NSX controllers and Juniper Networks devices exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to entities in a physical network and vice versa.

To enable communication between NSX controllers and Juniper Networks devices, the Junos OS implementation of OVSDb includes an OVSDb server and an OVSDb client, both of which run on each Juniper Networks device that supports OVSDb.

The OVSDb server on a Juniper Networks device can communicate with an OVSDb client on one or more NSX controllers. To establish a connection between a Juniper Networks device and an NSX controller, you must specify information about the controller (IP address) and the connection (port over which the connection occurs and the communication protocol to be used) on each Juniper Networks device. After the configuration is successfully committed, the connection is established between the management port (fxp0) of the Juniper Networks device and the NSX controller port that you specify in the Junos OS configuration.

The OVSDb server stores and maintains an OVSDb database schema, which is defined for physical devices. This schema contains control and statistical information provided by the OVSDb client on the Juniper Networks devices and NSX controllers. This information is stored in various tables in the schema. The OVSDb client on the Juniper Networks devices and NSX controllers monitors the schema for additions, deletions, and



modifications to this information, and the information is used for various purposes such as learning the MAC addresses of virtual hosts and physical servers.

The schema provides a means through which the Juniper Networks devices and the NSX controllers can exchange information. For example, the Juniper Networks devices capture MAC routes to entities in the physical network and push this information to a table in the schema so that NSX controllers with connections to these Juniper Networks devices can access the MAC routes. Conversely, NSX controllers capture MAC routes to entities in the virtual network and push this information to a table in the schema so that Juniper Networks devices with connections to the NSX controllers can access the MAC routes.

Some of the OVSDb table names include the words *local* or *remote*, for example, *unicast MACs local table* and *unicast MACs remote table*. Information in *local* tables is learned by a Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP), while information in *remote* tables is learned from other software or hardware VTEPs.

#### Related Documentation

- [Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903](#)
- [Open vSwitch Database Schema For Physical Devices on page 1910](#)

## Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers

The Juniper Networks Junos operating system (Junos OS) implementation of the Open vSwitch Database (OVSDb) management protocol provides a means through which VMware NSX controllers and Juniper Networks devices that support OVSDb can communicate. This implementation of OVSDb supports one cluster of NSX controllers, which includes three or five controllers as recommended by VMware.

To implement the OVSDb management protocol on a Juniper Networks device, you must explicitly configure a connection to one NSX controller, using the Junos OS CLI. If the NSX controller to which you explicitly configure a connection is in a cluster, the controller pushes information about other controllers in the same cluster to the device, and the device establishes connections with the other controllers. However, you can also explicitly configure connections with the other controllers in the cluster, using the Junos OS CLI.

A Juniper Networks device exchanges control and statistical data with each NSX controller to which it is connected. Therefore, the benefits of connecting a Juniper Networks device to multiple NSX controllers include redundancy and load-balancing of the controller workload.

Connections to all NSX controllers are made on the management interface of the Juniper Networks device. (The management interface on MX Series routers is fxp0 and on QFX5100 switches is em0 or em1.) To set up a connection between a Juniper Networks device and an NSX controller, you need to configure the following parameters on the Juniper Networks device:

- IP address of the NSX controller.
- The protocol that secures the connection. Secure Sockets Layer (SSL) is the supported protocol.



**NOTE:** The SSL connection requires a private key and certificates, which must be stored in the `/var/db/certs` directory of the Juniper Networks device. For more information about the files, including actions you must take to create and install some of the files, see [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers”](#) on page 1984.

- Number of the port over which the connection is made. The port number of the default port is 6632.

Optionally, you can configure the following connection timers on the Juniper Networks device:

- Inactivity probe duration—The maximum amount of time, in milliseconds, that the connection can be inactive before an inactivity probe is sent. The default value is 0 milliseconds, which means that an inactivity probe is never sent.
- Maximum backoff duration—If an attempt to connect to an NSX controller fails, the maximum amount of time, in milliseconds, before the device can make the next attempt. The default value is 1000 milliseconds.

**Related  
Documentation**

- [Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices](#) on page 1985
- [Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices](#) on page 1902

## Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDb

The Juniper Networks Junos operating system (Junos OS) implementation of the Open vSwitch Database (OVSDb) management protocol provides a means through which VMware NSX controllers and Juniper Networks devices that support OVSDb can communicate.

This topic explains how a Juniper Networks device with Virtual Extensible LAN (VXLAN) and OVSDb management protocol capabilities handles the following types of traffic:

- Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic that originates in an OVSDb-managed VXLAN and is forwarded to interfaces within the same VXLAN.
- Layer 3 multicast traffic that is received by an integrated routing and bridging (IRB) interface in an OVSDb-managed VXLAN and is forwarded to interfaces in another OVSDb-managed VXLAN.

By default, Layer 2 BUM traffic that originates in an OVSDb-managed VXLAN is handled by one or more service nodes in the same VXLAN. When this option is used, the table for remote multicast MAC addresses in the OVSDb schema for physical devices contains only one entry that has the keyword **unknown-dst** as the MAC string and a list of software virtual tunnel endpoints (VTEPs) that host the service nodes.

Given the previously described table entry, Layer 2 BUM traffic received on an interface in the OVSDB-managed VXLAN is forwarded to one of the software VTEPs. The software VTEP, and therefore, the service node to which a BUM packet is forwarded, is determined by the Juniper Networks device on which the OVSDB-managed VXLAN is configured. On receiving the BUM packet, the service node replicates the packet and forwards the replicas to all interfaces within the VXLAN.

Instead of using service nodes, you can optionally enable ingress node replication to handle Layer 2 BUM traffic on Juniper Networks devices that support OVSDB.



**NOTE:** Ingress node replication is supported on all Juniper Networks devices that support OVSDB except the QFX5100 switch.

With ingress node replication enabled, on receiving a Layer 2 BUM packet on an interface in an OVSDB-managed VXLAN, the Juniper Networks device replicates the packet and then forwards the replicas to all software VTEPs included in the unicast MACs remote table in the OVSDB schema. The software VTEPs then forward the replicas to all virtual machines (VMs), except service VMs or nodes, on the same host.



**NOTE:** When Juniper Networks devices replicate Layer 2 BUM packets to a large number of remote software VTEPs, the performance of the Juniper Networks devices can be impacted.

On IRB interfaces that forward Layer 3 multicast traffic from one OVSDB-managed VXLAN to another, ingress node replication is automatically implemented. With ingress node replication, the Juniper Networks device replicates a Layer 3 multicast packet and then the IRB interface forwards the replicas to all hardware and software VTEPs, but not to service nodes, in the other OVSDB-managed VXLAN. For the routing of Layer 3 multicast traffic from one OVSDB-managed VXLAN to another, ingress node replication is the only option and does not need to be configured.

#### Related Documentation

- [Configuring OVSDB-Managed VXLANs on page 1989](#)
- [Open vSwitch Database Schema For Physical Devices on page 1910](#)

## Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment

The Juniper Networks Junos operating system (Junos OS) implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which VMware NSX controllers and Juniper Networks devices that support OVSDB can communicate.

In a Junos OS environment, the concept of an OVSDB-managed Layer 2 broadcast domain in which data flows are limited to that domain is known as a *VXLAN*. In an NSX environment, the same concept is known as a *logical switch*. Understanding the different terminology in turn enables you to better understand the configuration tasks required for setting up OVSDB-managed VXLANs.

The following sections explain what you need to do to set up OVSDB-managed VXLANs properly for each Juniper Networks device that supports OVSDB and VXLAN:

- [Understanding How to Set Up OVSDB-Managed VXLANs On All Juniper Networks Devices Except QFX5100 Switches on page 1906](#)
- [Understanding How to Set Up OVSDB-Managed VXLANs On QFX5100 Switches on page 1907](#)
- [Understanding Automatically Created OVSDB-Managed VXLANs on a QFX5100 Switch on page 1908](#)
- [Understanding How to Determine the State of an OVSDB-Managed VXLAN on page 1909](#)

### **[Understanding How to Set Up OVSDB-Managed VXLANs On All Juniper Networks Devices Except QFX5100 Switches](#)**

---

For each VXLAN that you plan to implement, you must first configure a logical switch, using NSX Manager or the NSX API. Based on the name and the VXLAN network identifier (VNI) that you specify, NSX automatically generates a universally unique identifier (UUID) for the logical switch. You must retain the UUID of the logical switch for later use.

Next, on the Juniper Networks device, you must configure the corresponding VXLAN, including the same VNI specified for the logical switch, using the Junos OS CLI. For the name of the VXLAN, you must specify the UUID for the logical switch.

When configuring a logical switch and a corresponding VXLAN, it is important that the UUID and VNI in both configurations are the same. If these elements are not the same, the logical switch and VXLAN cannot become operational, which means they cannot exchange MAC addresses learned in the NSX and Junos OS environments, respectively.

[Table 130 on page 1907](#) provides a summary of the procedure that you must perform for each OVSDB-managed VXLAN on each Juniper Networks device, where to get more information about the configuration task, and the configuration statements that you must use to configure the VXLAN.

**Table 130: Summary of Configuration Tasks for Setting Up An OVSDb-Managed VXLAN on All Juniper Networks Devices Except QFX5100 Switches**

Juniper Networks Device That Supports OVSDb and VXLAN	Configure Logical Switch, Using NSX Manager or the NSX API?	Where to Find More Configuration Information	Configure Corresponding VXLAN on Juniper Networks Device?	Junos OS Statement to Configure the OVSDb-Managed VXLAN	Where to Find More Configuration Information
MX Series routers	Yes	See the documentation that accompanies NSX Manager or the NSX API.	Yes	<b>ovsdb-managed</b> statement in the <b>[edit bridge-domains domain-name vxlan]</b> hierarchy.  For the name of the VXLAN, specify the UUID for the logical switch configured in NSX Manager or in the NSX API.	<a href="#">"Configuring OVSDb-Managed VXLANs" on page 1989</a>

### Understanding How to Set Up OVSDb-Managed VXLANs On QFX5100 Switches

For each OVSDb-managed VXLAN that you plan to implement, we recommend a particular configuration workflow. First, you must specify a few statements in the Junos OS CLI of the QFX5100 switch. Then, you must configure a logical switch, using NSX Manager or the NSX API. Based on the name and the VNI that you specify, NSX automatically generates a UUID for the logical switch.

After you create a logical switch in NSX Manager or in the NSX API, the NSX controller pushes relevant information to the logical switch table in the OVSDb schema for physical devices. This schema resides in the QFX5100 switch. Based on the information pushed by the NSX controller, the switch then automatically creates a corresponding VXLAN, thereby eliminating the need for you to perform this task, using the Junos OS CLI. For the name of the VXLAN, the switch uses the UUID of the logical switch. For more information about automatically created VXLANs, see ["Understanding Automatically Created OVSDb-Managed VXLANs on a QFX5100 Switch" on page 1908](#).

[Table 131 on page 1908](#) provides a summary of the procedure that you must perform for each OVSDb-managed VXLAN on a QFX5100 switch, where to get more information about the configuration task, and the configuration statements that you must use to configure the VXLAN.

Table 131: Summary of Configuration Tasks for Setting Up An OVSDb-Managed VXLAN on QFX5100 Switches

Required Junos OS Configuration	Configure Logical Switch, Using NSX Manager or the NSX API?	Where to Find More Configuration Information	Configure Corresponding VXLAN on Juniper Networks Device?	Junos OS Statement to Configure the OVSDb-Managed VXLAN
<ul style="list-style-type: none"> <li>Enter the <b>set switch-options ovssdb-managed</b> configuration mode command.</li> <li>Use the <b>set interfaces interface-name unit logical-unit-number family ethernet-switching interface-mode access</b> configuration mode command to configure each interface that you want to associate with the VXLAN.</li> </ul>	Yes	See the documentation that accompanies NSX Manager or the NSX API.	No. The QFX5100 switch automatically creates a corresponding OVSDb-managed VXLAN.	—

As described in [Table 131 on page 1908](#), you must issue the **set switch-options ovssdb-managed** statement in the Junos OS CLI of the switch. Issuing this statement and committing the configuration enable the QFX5100 switch to automatically configure an OVSDb-managed VXLAN after the NSX controller pushes relevant information about the corresponding logical switch to the QFX5100 switch. Upon receipt of a logical switch information, if the QFX5100 switch does not detect the presence of this statement, it cannot start the automatic configuration of the corresponding VXLAN.

### Understanding Automatically Created OVSDb-Managed VXLANs on a QFX5100 Switch

After a QFX5100 switch creates a VXLAN, you can issue the **show configuration** operational mode command in the Junos OS CLI, and a configuration similar to the following sample appears:

```
set vlans 28805c1d-0122-495d-85df-19abd647d772 vxlan vni 100
set vlans 28805c1d-0122-495d-85df-19abd647d772 vlan-id 75
```

Also, if you specified the following command in the Junos OS CLI of the switch to configure interface ge-1/0/0 unit 0:

```
set interfaces ge-1/0/0 unit 0 family ethernet-switching interface-mode access
```

The switch automatically associates the interface with the VXLAN as shown:

```
set interfaces ge-1/0/0 unit 0 family ethernet-switching interface-mode access vlan
members 28805c1d-0122-495d-85df-19abd647d772
```

Note the following about this sample configuration:

- The name of the VXLAN is 28805c1d-0122-495d-85df-19abd647d772, which assumes that the logical switch UUID generated by NSX is 28805c1d-0122-495d-85df-19abd647d772.
- The QFX5100 switch automatically generated the VLAN ID of 75 for the VXLAN. In general, when generating a VLAN ID, the switch uses a VLAN ID that is not currently used either by NSX Manager or the NSX API, or by the QFX5100 switch.

Furthermore, if the switch detects an already used VLAN ID in a configuration performed in the Junos OS CLI, the switch displays an error message.

- After the switch automatically associates interface ge-1/0/0 unit 0 with VXLAN 28805c1d-0122-495d-85df-19abd647d772, keep in mind that you must also configure the interface as OVSDB-managed in the Junos OS CLI. For information about configuring an interface as OVSDB-managed, see [“Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices” on page 1985](#).
- The QFX5100 switch can automatically associate access interfaces with the VXLAN. It cannot automatically associate trunk interfaces with the VXLAN.

If you need to modify or delete an OVSDB-managed VXLAN that was automatically created by a QFX5100 switch, you must modify or delete the corresponding logical switch configuration in NSX Manager or in the NSX API. After you update the logical switch configuration, the NSX controller pushes the update to the QFX5100 switch, and the switch modifies or deletes its configuration accordingly. Modifying or deleting the VXLAN configuration by using the Junos OS CLI on the switch does not work.

### Understanding How to Determine the State of an OVSDB-Managed VXLAN

Regardless of the Juniper Networks device and the procedure that you follow to set up OVSDB-managed VXLANs, after configuring one or more logical switches in NSX Manager or in the NSX API, the NSX controller pushes relevant information to the logical switch table in the OVSDB schema, which resides on the respective devices.

To determine the state of a VXLAN and corresponding logical switch, you can use the **show ovbdb logical-switch** command. The following are possible states:

- Created by Controller**—A logical switch was configured in NSX Manager or in the NSX API, but the corresponding VXLAN is not yet created on the Juniper Networks device. In this state, the VXLAN and corresponding logical switch are not yet operational.
- Created by L2ALD**—A VXLAN was created, but the corresponding logical switch is not yet configured in NSX Manager or in the NSX API. In this state, the VXLAN and corresponding logical switch are not yet operational.
- Created by both**—A logical switch was configured in NSX Manager or in the NSX API, and a corresponding VXLAN was created. In this state, the VXLAN and corresponding logical switch are operational.
- Tunnel key mismatch**—The VNIs specified in the logical switch and corresponding VXLAN configurations do not match. In this state, the VXLAN and corresponding logical switch are not yet operational.

## Related Documentation

- [show ovssdb logical-switch on page 2019](#)
- [Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSSDB-Managed VXLAN on page 2045](#)

## Open vSwitch Database Schema For Physical Devices

An Open vSwitch Database (OVSSDB) server runs on a Juniper Networks device that supports the OVSSDB management protocol. When this device is connected to one or more VMware NSX controllers, the connections provide a means through which the Juniper Networks device and the controllers can communicate.

In an NSX multi-hypervisor environment, Juniper Networks devices that support OVSSDB and NSX controllers exchange control and statistical data. This data is stored in the OVSSDB database schema, which was defined for physical devices, and the schema resides in the OVSSDB server. The schema includes several tables. Juniper Networks devices and NSX controllers, which are both OVSSDB clients, can add rows to the tables as well as monitor for the addition, deletion, and modification of rows.

For example, the OVSSDB client on a Juniper Networks device or NSX controller can collect MAC routes learned by entities in the physical or virtual networks, respectively, and publish the routes to the appropriate table in the schema. By using the MAC routes and other information provided in the table, Juniper Networks devices in the physical network and entities in the virtual network can determine where to forward virtual machine (VM) traffic.

Some of the OVSSDB table names include the words *local* or *remote*, for example, the *unicast MACs local table* and the *unicast MACs remote table*. Information in *local* tables is learned by a Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP), while information in *remote* tables is learned from other software or hardware VTEPs.

[Table 132 on page 1910](#) describes the tables in the schema, the physical or virtual entity that is the source of the data provided in the table, and the command that you can enter in the CLI of the Juniper Networks device to get similar information.

**Table 132: OVSSDB Schema Tables**

Table Name	Description	Source of Information	Command
Global table	Includes the top-level configuration for the Juniper Networks device.	Juniper Networks device	None
Manager table	Includes information for each NSX controller that is connected to the Juniper Networks device.	<ul style="list-style-type: none"> <li>• Juniper Networks device</li> <li>• NSX controller</li> </ul>	<a href="#">show ovssdb controller</a>



Table 132: OVSDB Schema Tables (*continued*)

Table Name	Description	Source of Information	Command
Physical switch table	Includes information about the Juniper Networks device on which a hardware VTEP is implemented. This table includes information only for the device on which the table resides.	Juniper Networks device	None
Physical port table	Includes information about OVSDB-managed interfaces.	Juniper Networks device	<a href="#">show ovssdb interface</a>
Logical switch table	Includes information about logical switches, which you configured in NSX Manager or the NSX API, and the corresponding Virtual Extensible LANs (VXLANs), which was configured on the Juniper Networks device.	Juniper Networks device	<a href="#">show ovssdb logical-switch</a>
Logical binding statistics table	Includes statistics for OVSDB-managed interfaces.	Juniper Networks device	<a href="#">show ovssdb statistics interface</a>
Physical locator table	Includes information about Juniper Networks devices configured as hardware VTEPs, software VTEPs, and service nodes.	Juniper Networks device	<a href="#">show ovssdb virtual-tunnel-end-point</a>
Physical locator set table	Includes a list of service nodes for a logical switch.	Juniper Networks device	None
Unicast MACs remote table	Reachability information, including unicast MAC addresses, for entities in the virtual network.	NSX controller	<a href="#">show ovssdb mac</a>
Unicast MACs local table	Reachability information, including unicast MAC addresses, for entities in the physical network.	Juniper Networks device that is configured as a hardware VTEP.	<a href="#">show ovssdb mac</a>
Multicast MACs remote table	Includes only one row. In this row, the MAC column includes the keyword <b>unknown dst</b> along with a list of software VTEPs that host a cluster of service nodes, which handle multicast traffic.	NSX controller	<a href="#">show ovssdb mac</a>

**Related Documentation** • [Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices on page 1902](#)

- [Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903](#)
- [Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment on page 1905](#)

## VXLAN Overview

---

- [Understanding VXLANs on page 1912](#)

### Understanding VXLANs

- [VXLAN Benefits on page 1912](#)
- [What is a VXLAN? on page 1913](#)
- [Using a QFX5100 Switch with VXLANs on page 1913](#)
- [Using an MX Series Routers as a VTEP on page 1914](#)
- [Manual VXLANs Require PIM on page 1914](#)
- [Load Balancing VXLAN Traffic on page 1915](#)

#### VXLAN Benefits

---

Virtual Extensible Local Area Network (VXLAN) is a technology that allows you to segment your networks (as VLANs do) but that also solves the scaling limitation of VLANs and provides benefits that VLANs cannot. Here are the most important benefits of using VXLANs:

- You can theoretically create as many as 16 million VXLANs in an administrative domain (as opposed to 4094 VLANs on a Juniper Networks device).
- MX Series routers support as many as 32K VXLANs. This means that VXLANs provide network segmentation at the scale required by cloud builders to support very large numbers of tenants.
- QFX 5100 switches support 4K VXLANs.
- You can enable migration of virtual machines between servers that exist in separate Layer 2 domains by tunneling the traffic over Layer 3 networks. This functionality allows you to dynamically allocate resources within or between data centers without being constrained by Layer 2 boundaries or being forced to create large or geographically stretched Layer 2 domains.

Using VXLANs to create smaller Layer 2 domains that are connected over a Layer 3 network means that you don't need to use STP to converge the topology but can use more-robust routing protocols in the Layer 3 network instead. In the absence of STP, none of your links are blocked, which means you can get full value from all the ports that you purchase. Using routing protocols to connect your Layer 2 domains also allows you to load balance the traffic to ensure that you get the best use of your available bandwidth. Given the amount of east-west traffic that often flows within or between data centers, maximizing your network performance for that traffic is very important.

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of using VXLANs.



Video: [Why Use an Overlay Network in a Data Center?](#)


### What is a VXLAN?

VXLAN is often described as an overlay technology because it allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. The devices that encapsulate traffic that must be transported over a VXLAN and de-encapsulate traffic when it must leave the VXLAN tunnel are virtual tunnel endpoints (VTEPs), which can be end hosts or network switches or routers. To encapsulate an Ethernet frame, VTEPs add a number of fields, including the following:

- Outer MAC destination address (MAC address of the tunnel endpoint VTEP)
- Outer MAC source address (MAC address of the tunnel source VTEP)
- Outer IP destination address (IP address of the tunnel endpoint VTEP)
- Outer IP source address (IP address of the tunnel source VTEP)
- Outer UDP header
- A VXLAN header that includes a 24-bit field—called the VXLAN network identifier (VNI)—that is used to uniquely identify the VXLAN. The VNI is similar to a VLAN ID, but having 24 bits allows you to create many more VXLANs than VLANs.

Figure 30 on page 1913 shows the VXLAN packet format.

**Figure 30: VXLAN Packet Format**



### Using a QFX5100 Switch with VXLANs

You can configure a QFX5100 switch to perform all of the following roles:

- Act as a transit Layer 3 switch for downstream hosts acting as VTEPs. In this configuration, you do not need to configure any VXLAN functionality on the switch. You do need to configure IGMP and PIM so that the switch can form the multicast trees for the VXLAN multicast groups. (See [Manual VXLANs Require PIM on page 1914](#) for more information.)

- Act as a Layer 2 gateway between virtualized and non-virtualized networks in the same data center or between data centers. For example, you can use a QFX5100 switch to connect a network that uses VXLANs to one that uses VLANs.
- Act as a Layer 2 gateway between virtualized networks in the same or different data centers and allow virtual machines to move (VMotion) between those networks and data centers. For example, if you want to allow VMotion between devices in two different networks, you can create the same VLAN in both networks and put both devices on that VLAN. The QFX5100 switches connected to these devices, acting as VTEPs, can map that VLAN to the same VXLAN, and the VXLAN traffic can then be routed between the two networks.



**NOTE:** A QFX 5100 switch cannot route traffic between different VXLANs. To connect devices in different VXLANs you need a VXLAN-capable Layer 3 gateway, such as a Juniper Networks MX Series router.

---

### Using an MX Series Routers as a VTEP

---

You can configure an MX Series router to act as a VTEP and perform all of the following roles:

- Act as a Layer 2 gateway between virtualized and non-virtualized networks in the same data center or between data centers. For example, you can use an MX Series router to connect a network that uses VXLANs to one that uses VLANs.
- Act as a Layer 2 gateway between virtualized networks in the same or different data centers and allow virtual machines to move (VMotion) between those networks and data centers.
- Act as a Layer 3 gateway to route traffic between different VXLANs in the same data center.
- Act as a Layer 3 gateway to route traffic between different VXLANs in different data centers over a WAN or the Internet using standard routing protocols or VPLS tunnels.



**NOTE:** If you want an MX Series router to be a VXLAN Layer 3 gateway, you must configure integrated routing and bridging (IRB) interfaces to connect the VXLANs, just as you do if you want to route traffic between VLANs.

---

### Manual VXLANs Require PIM

---

You can use a controller (such as VMware's NSX) to provision VXLANs on a Juniper Networks device. A controller also provides a control plane that VTEPs use to advertise their reachability and learn about the reachability of other VTEPs. You can also manually create VXLANs on Juniper Networks devices instead of using a controller. If you use this approach, you must also configure PIM on the VTEPs so that they can create VXLAN tunnels between themselves.

You must also configure each VTEP in a given VXLAN to be a member of the same multicast group. (If possible, you should assign a different multicast group address to each VXLAN.) The VTEPs can then forward ARP requests they receive from their connected hosts to the multicast group. The other VTEPs in the group de-encapsulate the VXLAN information, and (assuming they are members of the same VXLAN) they forward the ARP request to their connected hosts. When the target host receives the ARP request, it responds with its MAC address, and its VTEP forwards this ARP reply back to the source VTEP. Through this process, the VTEPs learn the IP addresses of the other VTEPs in the VXLAN and the MAC addresses of the hosts connected to the other VTEPs.

The multicast groups and trees are also used to forward broadcast, unknown unicast, and multicast (BUM) traffic between VTEPs. This prevents BUM traffic from being unnecessarily flooded outside the VXLAN.



**NOTE:** Multicast traffic that is forwarded through a VXLAN tunnel is sent only to the remote VTEPs in the VXLAN. That is, the encapsulating VTEP does not copy and send copies of the packets according to the multicast tree—it only forwards the received multicast packets to the remote VTEPs. The remote VTEPs de-encapsulate the encapsulated multicast packets and forward them the appropriate Layer 2 interfaces. The remote VTEPs also do not copy and send copies of the packets according to the multicast tree.

### Load Balancing VXLAN Traffic

On QFX5100 switches, the Layer 3 routes that form VXLAN tunnels use per-packet load balancing by default, which means that load balancing is implemented if there are ECMP paths to the remote VTEP. This is different from normal routing behavior in which per-packet load balancing is not used by default. (Normal routing uses per-prefix load balancing by default.)

The source port field in the UDP header is used to enable ECMP load balancing of the VXLAN traffic in the Layer 3 network. This field is set to a hash of the inner packet fields, which results in a variable that ECMP can use to distinguish between tunnels (flows). (None of the other fields that flow-based ECMP normally uses are suitable for use with VXLANs. All tunnels between the same two VTEPs have the same outer source and destination IP addresses, and the UDP destination port is set to port 4789 by definition. Therefore, none of these fields provide a sufficient way for ECMP to differentiate flows.)

#### Related Documentation

- [Examples: Configuring VXLANs on QFX Series Switches on page 1944](#)
- [Example: Configuring VXLAN on MX Series Routers on page 1934](#)
- [Open vSwitch Database Support on Juniper Networks Devices on page 1899](#)



## CHAPTER 19

# Configuration

- [Configuration Examples on page 1917](#)
- [Configuration Tasks on page 1983](#)
- [OVSDB Configuration Statements on page 1992](#)
- [VXLAN Configuration Statements on page 2002](#)

### Configuration Examples

---

- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections Between Virtual and Physical Entities in a Data Center on page 1917](#)
- [Example: Setting Up Inter-VXLAN Routing and OVSDB Connections in a Data Center on page 1925](#)
- [Example: Configuring VXLAN on MX Series Routers on page 1934](#)
- [Examples: Configuring VXLANs on QFX Series Switches on page 1944](#)
- [Example: Configuring VXLAN to VPLS Stitching with OVSDB on page 1952](#)

#### Example: Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections Between Virtual and Physical Entities in a Data Center

In a physical network, a Juniper Networks device that supports Virtual Extensible LAN (VXLAN) can function as a hardware virtual tunnel endpoint (VTEP), which enables software applications running directly on physical servers to communicate with virtual machines (VMs) in a virtual network.

In this environment, you can also include VMware NSX controllers and implement the Open vSwitch Database (OVSDB) management protocol on the Juniper Networks device that functions as a hardware VTEP. The Junos OS implementation of OVSDB provides a means through which VMware NSX controllers and Juniper Networks devices that support OVSDB can communicate. These components serve the following purposes:

- Centralized configuration (QFX5100 switch only)—After you configure a logical switch, using NSX Manager or the NSX API, the NSX controller pushes relevant information about the configuration to the switch (the QFX5100 switch in this example) that functions as a hardware VTEP. Using the relevant configuration information, the switch automatically creates the configuration of a VXLAN, which is the Junos OS-equivalent of the logical switch.

- Centralized storage and exchange of MAC route information—Availability of MAC routes enables the hardware VTEP in the physical network and the software VTEP in the virtual network to forward VM traffic between entities in the physical and virtual networks.

This example explains how to configure a QFX5100 switch as a hardware VTEP, which serves as a Layer 2 gateway, and set up this device with an OVSDb connection to an NSX controller.

- [Requirements on page 1918](#)
- [Overview and Topology on page 1919](#)
- [Configuration on page 1921](#)
- [Verification on page 1923](#)

### Requirements

---

The topology for this example includes the following hardware and software components:

- A physical server on which software applications directly run.
- A QFX5100 switch that is running Junos OS release 14.1X53-D10 or later and an OVSDb software package. The OVSDb software package release must be the same as the Junos OS release running on the device. This switch is configured to function as a hardware VTEP.
- A cluster of five NSX controllers, each of which is running NSX software version 4.0.3. (In this example, you explicitly configure a connection with one NSX controller.)
- NSX Manager version 4.0.3.
- A service node that handles the replication and forwarding of Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic in the VXLAN used in this example.
- A host that includes VMs managed by a hypervisor, which includes a software VTEP.

Before you begin the configuration, you need to perform the following tasks:

- Using NSX Manager version 4.0.3, specify the IP address of the service node.
- Create an SSL private key and certificate, and install them in the `/var/db/certs` directory of the QFX5100 switch. For more information, see [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers” on page 1984](#).
- Issue the **set switch-options ovbdb-managed** configuration mode command in the Junos OS CLI on the QFX5100 switch. Issuing this command and committing the configuration enable the QFX5100 switch to automatically create OVSDb-managed VXLANs.
- For each interface that you want to associate with the OVSDb-managed VXLAN, issue the **set interfaces interface-name unit logical-unit-number family ethernet-switching interface-mode access** configuration mode command in the Junos OS CLI on the QFX5100 switch. In this example, interface ge-1/0/0.0 is associated with the VXLAN,



so the command is issued once with `ge-1/0/0` specified as the interface name and 0 specified as the logical unit number.

- Using NSX Manager version 4.0.3, configure a logical switch for each VXLAN that OVSDb will manage. This example implements one OVSDb-managed VXLAN; therefore, you must configure one logical switch. After you configure the logical switch, the QFX5100 switch creates the configuration of a Junos OS-equivalent VXLAN. The name of the VXLAN is derived from the universally unique identifier (UUID) that the NSX controller automatically generated for the logical switch. A sample UUID is 28805c1d-0122-495d-85df-19abd647d772.

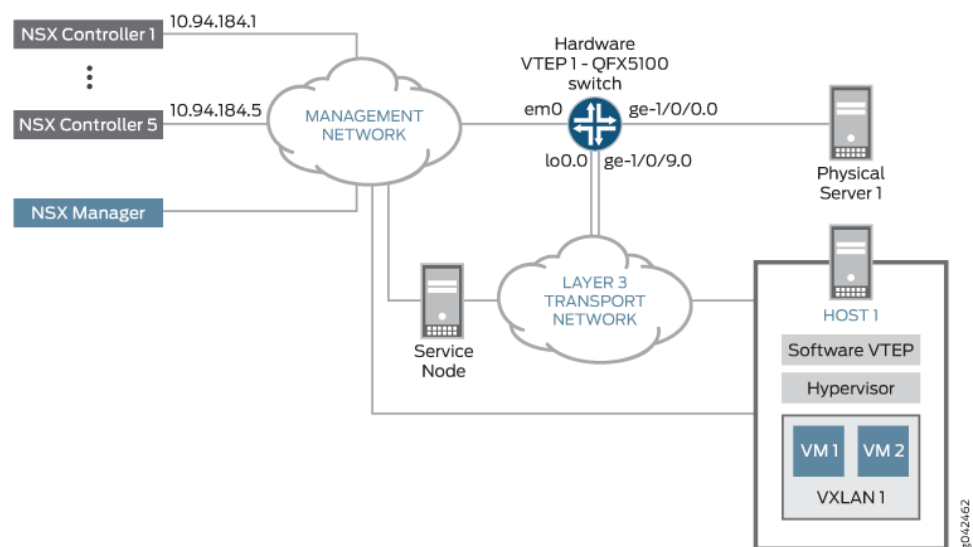
For more information about logical switches, VXLANs, and the automatic creation of VXLANs by the QFX5100 switch, see [“Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment”](#) on page 1905.

For information about using NSX Manager, see the documentation that accompanies these products.

### Overview and Topology

Figure 31 on page 1919 shows a topology in which a software application running directly on Physical Server 1 in the physical network needs to communicate with virtual machine VM 1 in VXLAN 1 and vice versa. To enable this communication, a QFX5100 switch is configured as Hardware VTEP 1.

Figure 31: VXLAN/OVSDb Layer 2 Gateway Topology



Based on the configuration of a logical switch in NSX Manager, the QFX5100 switch automatically creates a configuration for a Junos OS-equivalent VXLAN. A sample configuration using the UUID of 28805c1d-0122-495d-85df-19abd647d772 is as follows:

```
set vlans 28805c1d-0122-495d-85df-19abd647d772 vxlan vni 100
set vlans 28805c1d-0122-495d-85df-19abd647d772 vlan-id 75
```



**NOTE:** In the sample configuration, the switch automatically generates a VLAN ID of 75. In general, when generating a VLAN ID, the switch uses a VLAN ID that is not currently used by either NSX Manager or the NSX API, or by the QFX5100 switch.

Also, the following command was specified in the Junos OS CLI to configure interface ge-1/0/0 unit 0:

```
set interfaces ge-1/0/0 unit 0 family ethernet-switching interface-mode access
```

The switch automatically associates the interface with the VXLAN as shown:

```
set interfaces ge-1/0/0 unit 0 family ethernet-switching interface-mode access vlan  
members 28805c1d-0122-495d-85df-19abd647d772
```



**NOTE:** Keep in mind that you must also configure this interface as OVSDB-managed in the Junos OS CLI. This example provides information about performing this step.

The purpose of VXLAN 28805c1d-0122-495d-85df-19abd647d772 is to provide a means of mapping physical server 1 to VXLAN 1 using VXLAN 1's VNI of 100.

On the management interface **em0** or **em1** of the QFX5100 switch, a connection with an NSX controller is explicitly configured, by using the Junos OS CLI.

Each VXLAN-encapsulated packet must include a source IP address, which identifies the source hardware or software VTEP, in the outer IP header. In this example, the IP address of the loopback interface (lo0) on the QFX5100 switch is used for hardware VTEP 1.

Within VXLAN 28805c1d-0122-495d-85df-19abd647d772, Layer 2 BUM packets are replicated by the service node, which then forwards the replicas to all interfaces in the VXLAN. Having the service node handle the Layer 2 BUM traffic is the default behavior, and no configuration is required on the QFX5100 switch.

In this example, the tracing of all OVSDB events is configured. The output of the OVSDB events is placed in a file named **ovsdb**, which is stored in the **/var/log** directory. By default, a maximum of 10 trace files can exist, and the configured maximum size of each file is 10 MB.

The components of the topology for this example are shown in [Table 133 on page 1921](#).

Table 133: Components of the Topology for Setting Up a VXLAN Layer 2 Gateway and OVSDB Connections

Property	Settings
OVSDB-managed VXLAN	<p><b>NOTE:</b> The QFX5100 switch automatically creates this VXLAN configuration, which is based on the NSX-equivalent logical switch configuration in NSX Manager. Therefore, no configuration is required.</p> <p>VXLAN name: 28805c1d-0122-495d-85df-19abd647d772</p> <p>VLAN ID: 100</p> <p>VNI: 100 (VXLAN 1)</p>
OVSDB-managed interface	<p>Interface name: ge-1/0/0.0</p> <p>Interface type: access</p> <p><b>NOTE:</b> We recommend that you configure this interface in the Junos OS CLI <i>before</i> you create the logical switch in NSX Manager. After you configure the logical switch, the NSX controller pushes relevant information to the QFX5100 switch, and the switch automatically associates the interface with the VXLAN. You must then specify this interface as an OVSDB-managed interface in the Junos OS CLI.</p>
NSX controller	IP address: 10.94.184.1
Handling of Layer 2 BUM traffic in VXLAN 28805c1d-0122-495d-85df-19abd647d772	<p>Service node</p> <p><b>NOTE:</b> By default, one or more service nodes handle Layer 2 BUM traffic in a VXLAN; therefore, no configuration is required.</p>
Hardware VTEP source identifier	<p>Source interface: loopback (lo0.0)</p> <p>Source IP address: 10.17.17.17/32</p>
OVSDB tracing operations	<p>Filename: /var/log/ovsdb</p> <p>File size: 10 MB</p> <p>Flag: All</p>

### Configuration

Perform this task:

- [Configuring the QFX5100 Switch as Hardware VTEP 1 on page 1922](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis network-services enhanced-ip
set interfaces ge-1/0/9 unit 0 family inet address 10.40.40.1/24
set routing-options static route 10.19.19.19/32 next-hop 10.40.40.2
```

```
set routing-options router-id 10.17.17.17
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-1/0/9.0
set interfaces lo0 unit 0 family inet address 10.17.17.17/32 primary
set interfaces lo0 unit 0 family inet address 10.17.17.17/32 preferred
set switch-options vtep-source-interface lo0.0
set protocols ovssdb traceoptions file ovssdb
set protocols ovssdb traceoptions file size 10m
set protocols ovssdb traceoptions flag all
set protocols ovssdb interfaces ge-1/0/0.0
set protocols ovssdb controller 10.94.184.1
```



**NOTE:** After completing this configuration, you must configure a gateway, which is the NSX-equivalent of a hardware VTEP. This example implements one hardware VTEP; therefore, you must configure one gateway, a gateway service, and a logical switch port by using NSX Manager or the NSX API. For more information about the tasks you must perform and key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints”](#) on page 1986.

---

### *Configuring the QFX5100 Switch as Hardware VTEP 1*

#### **Step-by-Step Procedure**

To configure a QFX5100 switch as hardware VTEP 1 and with an OVSSDB connection to an NSX controller, follow these steps:

1. Configure the Layer 3 network.

```
[edit chassis]
user@switch# set network-services enhanced-ip
[edit interfaces]
user@switch# set ge-1/0/9 unit 0 family inet address 10.40.40.1/24
[edit routing-options]
user@switch# set static route 10.19.19.19/32 next-hop 10.40.40.2
user@switch# set router-id 10.17.17.17
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-1/0/9.0
```

2. Specify an IP address for the loopback interface. This IP address serves as the source IP address in the outer header of any VXLAN-encapsulated packets.

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 10.17.17.17/32 primary
user@switch# set lo0 unit 0 family inet address 10.17.17.17/32 preferred
```

3. Set the loopback interface as the interface that identifies hardware VTEP 1.

```
[edit switch-options]
user@switch# set vtep-source-interface lo0.0
```

4. Set up OVSSDB tracing operations.

```
[edit protocols]
user@switch# set ovssdb traceoptions file ovssdb
```

```
user@switch# set ovssdb traceoptions file size 10m
user@switch# set ovssdb traceoptions flag all
```

5. Specify that the interface between hardware VTEP 1 and physical server 1 is managed by OVSSDB.

```
[edit protocols]
user@switch# set ovssdb interfaces ge-1/0/0.0
```

6. Configure a connection with an NSX controller.

```
[edit protocols]
user@switch# set ovssdb controller 10.94.184.1
```



**NOTE:** After completing this configuration, you must configure a gateway, which is the NSX-equivalent of a hardware VTEP. This example implements one hardware VTEP; therefore, you must configure one gateway, a gateway service, and a logical switch port by using NSX Manager or the NSX API. For more information about the tasks you must perform and key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints” on page 1986](#).

## Verification

- [Verifying the Logical Switch on page 1923](#)
- [Verifying the MAC Address of VM 1 on page 1924](#)
- [Verifying the NSX Controller Connection on page 1924](#)
- [Verifying the OVSSDB-Managed Interface on page 1924](#)

### Verifying the Logical Switch

**Purpose** Verify that the configuration of the logical switch with the UUID of 28805c1d-0122-495d-85df-19abd647d772 is present in the OVSSDB schema for physical devices and that the state (**Flags**) of the logical switch is **Created by both**.

**Action** Issue the **show ovssdb logical-switch** operational mode command.

```
user@switch> show ovssdb logical-switch
Logical switch information:
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
Flags: Created by both
VNI: 100
Num of Remote MAC: 1
Num of Local MAC: 0
```

**Meaning** The output verifies that the configuration for the logical switch is present. The **Created by both** state indicates that the logical switch was configured in NSX Manager, and that

the QFX5100 switch automatically created the corresponding VXLAN. In this state, the logical switch and VXLAN are operational.

If the state of the logical switch is something other than **Created by both**, see [“Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN” on page 2045](#).

#### ***Verifying the MAC Address of VM 1***

**Purpose** Verify that the MAC address of VM 1 is present in the OVSDb schema.

**Action** Issue the **show ovssdb mac remote** operational mode command.

```
user@switch> show ovssdb mac remote
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
  Mac      IP      Encapsulation      Vtep
  Address  Address
a8:59:5e:f6:38:90    0.0.0.0      Vxlan over Ipv4    10.17.17.17
```

**Meaning** The output shows that the MAC address for VM 1 is present and is associated with the logical switch with the UUID of 28805c1d-0122-495d-85df-19abd647d772. Given that the MAC address is present, VM 1 is reachable through the QFX5100 switch, which functions as a hardware VTEP.

#### ***Verifying the NSX Controller Connection***

**Purpose** Verify that the connection with the NSX controller is up.

**Action** Issue the **show ovssdb controller** operational mode command, and verify that the controller connection state is **up**.

```
user@switch> show ovssdb controller
VTEP controller information:
Controller IP address: 10.94.184.1
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 542325
Controller seconds-since-disconnect: 542346
Controller connection status: active
```

**Meaning** The output shows that the connection state of the NSX controller is **up**, in addition to other information about the controller. By virtue of this connection being up, OVSDb is enabled on the QFX5100 switch.

#### ***Verifying the OVSDb-Managed Interface***

**Purpose** Verify that interface ge-1/0/0.0 is managed by OVSDb.

**Action** Issue the **show ovssdb interface** operational mode command, and verify that interface ge-1/0/0.0 is managed by OVSSDB.

```
user@switch> show ovssdb interface
Interface          VLAN ID          Bridge-domain
ge-1/0/0.0
```

**Meaning** The output shows that interface ge-1/0/0.0 is managed by OVSSDB and that the VLAN ID and VLAN are correctly configured.

## Example: Setting Up Inter-VXLAN Routing and OVSSDB Connections in a Data Center

This example shows a data center in which virtual machines (VMs) in different Virtual Extensible LANs (VXLANs) need to communicate. The Juniper Networks device that is integrated into this environment functions as a hardware virtual tunnel endpoint (VTEP) that can route VM traffic from one VXLAN (Layer 2) environment to another.

The Juniper Networks device implements the Open vSwitch Database (OVSSDB) management protocol and has a connection with a VMware NSX controller, both of which enable these two entities to exchange MAC routes to and from VMs in the physical and virtual networks. .

This example explains how to configure a Juniper Networks device as hardware VTEPs and set up OVSSDB connections to an NSX controller.

- [Requirements on page 1925](#)
- [Overview and Topology on page 1926](#)
- [Configuration on page 1929](#)
- [Verification on page 1932](#)

### Requirements

In this example, the topology includes the following hardware and software components:

- A cluster of five NSX controllers, each of which is running NSX software version 4.0.3.
- NSX Manager.
- A service node that handles broadcast, unknown unicast, and multicast (BUM) traffic within each of the two VXLANs.
- Two hosts, each of which includes VMs managed by a hypervisor. Each hypervisor includes a software VTEP. The VMs on each of the hosts belong to different VXLANs.
- A Juniper Networks device that routes VM traffic between the two VXLANs. For example, an MX Series router running Junos OS Release 14.1R2 or later. The Juniper Networks device must also run an OVSSDB software package, and the release of this package must be the same as the Junos OS release running on the device. This device is configured to function as a hardware VTEP.

Before you start the configuration of the Juniper Networks device, you need to perform the following tasks:

- In NSX Manager version 4.0.3 or the NSX API, specify the IP address of the service node.
- In NSX Manager version 4.0.3 or the NSX API, configure a logical switch for each VXLAN that OVSDb will manage. This example implements two OVSDb-managed VXLANs; therefore, you must configure two logical switches. After the configuration of each logical switch, NSX automatically generates a universally unique identifier (UUID) for the logical switch. If you have not already, retrieve the UUID for each logical switch. A sample UUID is 28805c1d-0122-495d-85df-19abd647d772. When configuring the equivalent VXLANs on the Juniper Networks device, you must use the UUID of the logical switch as the bridge domain or VLAN name.

For more information about logical switches and VXLANs, see [“Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment”](#) on page 1905.

- Create an SSL private key and certificate, and install them in the `/var/db/certs` directory of the Juniper Networks device. For more information, see [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers”](#) on page 1984.

For information about using NSX Manager or the NSX API to perform these configuration tasks, see the documentation that accompanies these respective products.

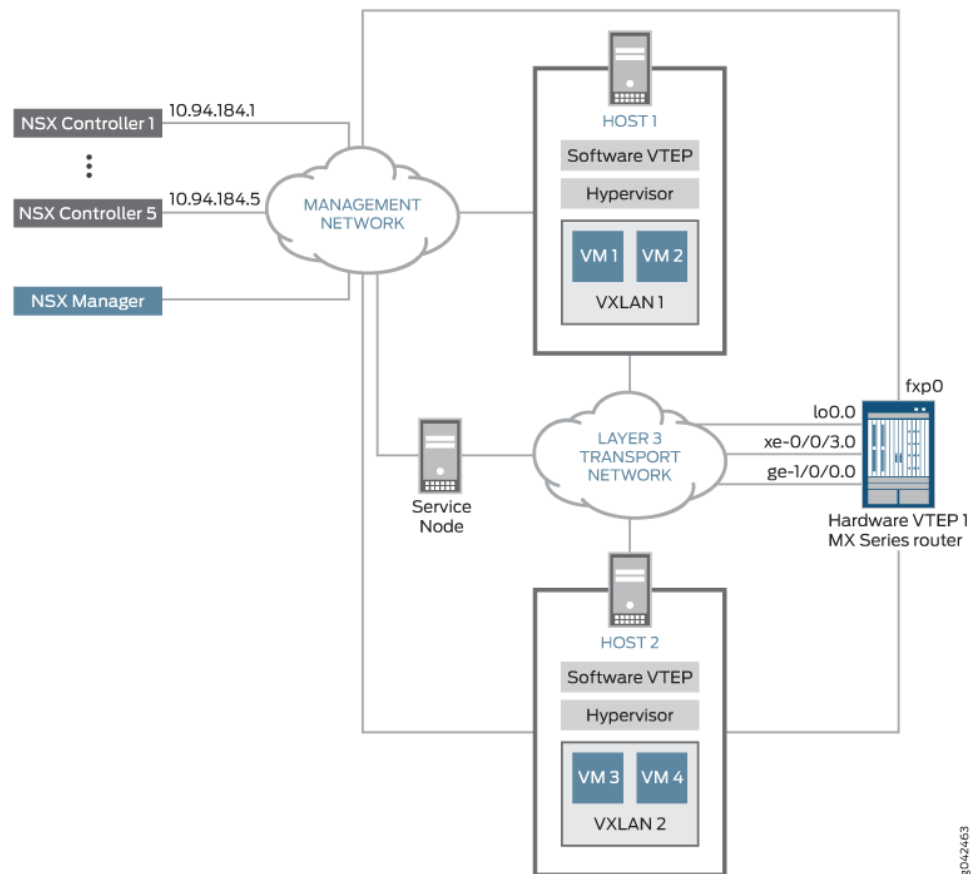
### Overview and Topology

---

In the topology shown in [Figure 32 on page 1927](#), VM 1 in VXLAN 1 needs to communicate with VM 3 in VXLAN 2. To enable this communication, hardware VTEP 1, which can be an MX Series router, is configured to route VM traffic between the two VXLANs.



Figure 32: Inter-VXLAN Routing and OVSDb Topology



On hardware VTEP 1, a routing instance (virtual switch) is set up. Within the routing instance, two VXLANs are configured: VXLAN 1 and VXLAN 2. Both of the VXLANs are associated with an integrated routing and bridging (IRB) interface, over which VM traffic is routed between the VXLANs.

On hardware VTEP 1, a connection with an NSX controller is configured on the management interface fxp0. This configuration enables the NSX controller to push MAC routes for VM 1 and VM 3 to the hardware VTEP by way of the table for remote unicast MAC addresses in the OVSDb schema for physical devices.

Each VXLAN-encapsulated packet must include a source IP address, which identifies the source hardware or software VTEP, in the outer IP header. In this example, for hardware VTEP 1, the IP address of the loopback interface (lo0) is used.

Within each of the two VXLANs, a service node replicates Layer 2 BUM packets then forwards the replicas to all interfaces in the VXLANs. Having the service node handle the Layer 2 BUM traffic is the default behavior, and no configuration is required for this Juniper Networks device.

Between the two VXLANs, ingress node replication is automatically implemented and does not need to be configured. With this feature, hardware VTEP 1 replicates a Layer 3 multicast packet, then the IRB interface associated with the VXLAN that originated the packet forwards the replicas to all hardware and software VTEPs, but not to service nodes, in the other OVSDB-managed VXLAN.

In this example, the tracing of all OVSDB events are configured. The output of the OVSDB events are placed in a file named **ovsdb**, which is stored in the **/var/log** directory. By default, a maximum of 10 trace files can exist, and the configured maximum size of each file is 50 MB.

[Table 133 on page 1921](#) describes the components of the example topology.

**Table 134: Components of the Topology for Setting Up Inter-VXLAN Routing and OVSDB Connections in a Data Center**

Property	Settings
Routing instance	Name: vx1  Type: virtual switch  OVSDB-managed VXLANs included: VXLAN 1 and VXLAN 2
VXLAN 1	Bridge domain or VLAN associated with: 28805c1d-0122-495d-85df-19abd647d772  Interface: xe-0/0/2.0  VLAN ID: 100  VNI: 100  IRB for inter-VXLAN traffic: irb.0; 10.20.20.1/24
VXLAN 2	Bridge domain or VLAN associated with: 96a382cd-a570-4ac8-a77a-8bb8b16bde70  Interface: xe-1/2/0.0  VLAN ID: 200  VNI: 200  IRB for inter-VXLAN traffic: irb.1; 10.10.10.3/24
Handling of BUM traffic in each VXLAN	Service node  <b>NOTE:</b> By default, one or more service nodes handle Layer 2 BUM traffic in a VXLAN; therefore, no configuration is required.
Handling of Layer 3 multicast traffic between VXLANs	Ingress node replication  <b>NOTE:</b> On IRB interfaces that forward Layer 3 multicast traffic from one OVSDB-managed VXLAN to another, ingress node replication is automatically implemented; therefore, no configuration is required.

Table 134: Components of the Topology for Setting Up Inter-VXLAN Routing and OVSDb Connections in a Data Center (*continued*)

Property	Settings
Hardware VTEP source identifier	Source interface: loopback (lo0.0) Source IP address: 10.19.19.19/32
NSX controller	IP address: 10.94.184.1
OVSDb tracing operations	Filename: /var/log/ovsdb File size: 50 m Flag: All

### Configuration

An MX Series router can function as hardware VTEP 1 in this example.

To configure inter-VXLAN routing and OVSDb connections in a data center topology, you need to perform this task:

- [Configuring an MX Series Router as a Hardware VTEP with an OVSDb Connection on page 1930](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.



**NOTE:** After completing this configuration, you must configure a gateway, which is the NSX-equivalent of a hardware VTEP. This example implements one hardware VTEP; therefore, you must configure one gateway, a gateway service, and a logical switch port using NSX Manager or the NSX API. For more information about the tasks you must perform and key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints” on page 1986](#).

MX Series router configuration:

```
set chassis network-services enhanced-ip
set interfaces xe-0/0/3 unit 0 family inet address 10.50.50.2/24
set interfaces ge-1/0/0 unit 0 family inet address 10.100.100.99/24
set routing-options router-id 10.19.19.19
set protocols ospf area 0.0.0.0 interface xe-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-1/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set interfaces xe-0/0/2 unit 0 family bridge interface-mode access
set interfaces xe-0/0/2 unit 0 family bridge vlan-id 100
```

```
set interfaces xe-1/2/0 unit 0 family bridge interface-mode access
set interfaces xe-1/2/0 unit 0 family bridge vlan-id 200
set interfaces irb unit 0 family inet address 10.20.20.1/24
set interfaces irb unit 1 family inet address 10.10.10.3/24
set routing-instances vx1 vtep-source-interface lo0.0
set routing-instances vx1 instance-type virtual-switch
set routing-instances vx1 interface xe-0/0/2.0
set routing-instances vx1 interface xe-1/2/0.0
set routing-instances vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
  vlan-id 100
set routing-instances vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
  routing-interface irb.0
set routing-instances vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
  vxlan ovssdb-managed
set routing-instances vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
  vxlan vni 100
set routing-instances vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
  vlan-id 200
set routing-instances vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
  routing-interface irb.1
set routing-instances vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
  vxlan ovssdb-managed
set routing-instances vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
  vxlan vni 200
set interfaces lo0 unit 0 family inet address 10.19.19.19/32 primary
set interfaces lo0 unit 0 family inet address 10.19.19.19/32 preferred
set protocols ovssdb traceoptions file ovssdb
set protocols ovssdb traceoptions file size 50m
set protocols ovssdb traceoptions flag all
set protocols ovssdb controller 10.94.184.1
set protocols ovssdb interfaces xe-0/0/2.0
set protocols ovssdb interfaces xe-1/2/0.0
```

### *Configuring an MX Series Router as a Hardware VTEP with an OVSSDB Connection*

**Step-by-Step Procedure** To configure an MX Series router as hardware VTEP 1 with an OVSSDB connection to an NSX controller, follow these steps:

1. Create the Layer 3 network.  

```
[edit chassis]
user@router# set network-services enhanced-ip
[edit interfaces]
user@router# set xe-0/0/3 unit 0 family inet address 10.50.50.2/24
user@router# set ge-1/0/0 unit 0 family inet address 10.100.100.99/24
[edit routing-options]
user@router# set router-id 10.19.19.19
[edit protocols]
user@router# set ospf area 0.0.0.0 interface xe-0/0/3.0
user@router# set ospf area 0.0.0.0 interface ge-1/0/0.0
user@router# set ospf area 0.0.0.0 interface lo0.0
```
2. Create an access interface for VXLAN 1, and associate the interface with the VXLAN.  

```
[edit interfaces]
user@router# set xe-0/0/2 unit 0 family bridge interface-mode access
```

- ```

user@router# set xe-0/0/2 unit 0 family bridge vlan-id 100

```
3. Create an access interface for VXLAN 2, and associate the interface with the VXLAN.
 

```

[edit interfaces]
user@router# set xe-1/2/0 unit 0 family bridge interface-mode access
user@router# set xe-1/2/0 unit 0 family bridge vlan-id 200

```
  4. Create an IRB interface to handle inter-VXLAN traffic for VXLAN 1.
 

```

[edit interfaces]
user@router# set irb unit 0 family inet address 10.20.20.1/24

```
  5. Create an IRB interface to handle inter-VXLAN traffic for VXLAN 2.
 

```

[edit interfaces]
user@router# set irb unit 1 family inet address 10.10.10.3/24

```
  6. Set up the virtual switch routing instance.
 

```

[edit routing-instances]
user@router# set vx1 vtep-source-interface lo0.0
user@router# set vx1 instance-type virtual-switch
user@router# set vx1 interface xe-0/0/2.0
user@router# set vx1 interface xe-1/2/0.0
user@router# set vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
vlan-id 100
user@router# set vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
routing-interface irb.0
user@router# set vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
vxlan ovsdb-managed
user@router# set vx1 bridge-domains 28805c1d-0122-495d-85df-19abd647d772
vxlan vni 100
user@router# set vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
vlan-id 200
user@router# set vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
routing-interface irb.1
user@router# set vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
vxlan ovsdb-managed
user@router# set vx1 bridge-domains 96a382cd-a570-4ac8-a77a-8bb8b16bde70
vxlan vni 200

```
  7. Specify an IP address for the loopback interface. This IP address serves as the source IP address in the outer header of any VXLAN-encapsulated packets.
 

```

[edit interfaces]
user@router# set lo0 unit 0 family inet address 10.19.19.19/32 primary
user@router# set lo0 unit 0 family inet address 10.19.19.19/32 preferred

```
  8. Set up OVSDB tracing operations.
 

```

[edit protocols]
user@router# set ovsdb traceoptions file ovsdb
user@router# set ovsdb traceoptions file size 50m
user@router# set ovsdb traceoptions flag all

```
  9. Configure a connection with an NSX controller.
 

```

[edit protocols]
user@router# set ovsdb controller 10.94.184.1

```

10. Configure interfaces xe-0/0/2.0 and xe-1/2/0.0 to be managed by OVSDb.

[edit protocols]

user@router# set **ovsdb interfaces** xe-0/0/2.0

user@router# set **ovsdb interfaces** xe-1/2/0.0



**NOTE:** After completing this configuration, you must configure a gateway, which is the NSX-equivalent of a hardware VTEP. This example implements one hardware VTEP; therefore, you must configure one gateway, a gateway service, and a logical switch port by using NSX Manager or the NSX API. For more information about the tasks you must perform and key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints”](#) on page 1986.

---

## Verification

- [Verifying the Logical Switches on page 1932](#)
- [Verifying the MAC Addresses of VM 1 and VM 3 on page 1933](#)
- [Verifying the NSX Controller Connection on page 1933](#)

### *Verifying the Logical Switches*

**Purpose** Verify that logical switches with the UUIDs of 28805c1d-0122-495d-85df-19abd647d772 and 96a382cd-a570-4ac8-a77a-8bb8b16bde70 are configured in NSX Manager or the NSX API, and that information about the logical switches is published in the OVSDb schema.

**Action** Issue the **show ovsdb logical-switch** operational mode command.

```
user@host> show ovsdb logical-switch
Logical switch information:
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
Flags: Created by both
VNI: 100
Num of Remote MAC: 1
Num of Local MAC: 0
Logical Switch Name: 96a382cd-a570-4ac8-a77a-8bb8b16bde70
Flags: Created by both
VNI: 200
Num of Remote MAC: 1
Num of Local MAC: 1
```

**Meaning** The output verifies that information about the logical switches is published in the OVSDb schema. The **Created by both** state indicates that the logical switches are configured in NSX Manager or the NSX API, and the corresponding VXLANs are configured on the Juniper Networks device. In this state, the logical switches and VXLANs are operational.

If the state of the logical switches is something other than **Created by both**, see [“Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN” on page 2045](#).

### ***Verifying the MAC Addresses of VM 1 and VM 3***

**Purpose** Verify that the MAC addresses of VM 1 and VM 3 are present in the OVSDb schema.

**Action** Issue the **show ovssdb mac remote** operational mode command, and verify that the MAC addresses for VM 1 and VM 3 are present.

```
user@host> show ovssdb mac remote
Logical Switch Name: 28805c1d-0122-495d-85df-19abd647d772
  Mac          IP          Encapsulation  Vtep
  Address      Address
  08:33:9d:5f:a7:f1  0.0.0.0        Vxlan over Ipv4  10.19.19.19
Logical Switch Name: 96a382cd-a570-4ac8-a77a-8bb8b16bde70
  Mac          IP          Encapsulation  Vtep
  Address      Address
  a8:59:5e:f6:38:90  0.0.0.0        Vxlan over Ipv4  10.19.19.10
```

**Meaning** The output shows that the MAC addresses for VM 1 and VM 3 are present and are associated with logical switches with the UUIDs of 28805c1d-0122-495d-85df-19abd647d772 and 96a382cd-a570-4ac8-a77a-8bb8b16bde70, respectively. Given that the MAC addresses are present, VM 1 and VM 3 are reachable through hardware VTEP 1.

### ***Verifying the NSX Controller Connection***

**Purpose** Verify that the connection with the NSX controller is up.

**Action** Issue the **show ovssdb controller** operational mode command, and verify that the controller connection state is **up**.

```
user@host> show ovssdb controller
VTEP controller information:
Controller IP address: 10.94.184.1
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 542325
Controller seconds-since-disconnect: 542346
Controller connection status: active
```

**Meaning** The output shows that the connection state of the NSX controller is **up**, in addition to other information about the controller. By virtue of this connection being up, OVSDb is enabled on the Juniper Networks device.

**Related Documentation**

- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections Between Virtual and Physical Entities in a Data Center on page 1917](#)

- [Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDb on page 1904](#)
- [Open vSwitch Database Schema For Physical Devices on page 1910](#)

## Example: Configuring VXLAN on MX Series Routers

Virtual Extensible Local Area Network (VXLAN) is a Layer 3 encapsulation protocol that enables MX Series routers to push Layer 2 or Layer 3 packets through a VXLAN tunnel to a virtualized data center or the Internet. Communication is established between two virtual tunnel endpoints (VTEPs). VTEPs encapsulate the virtual machine traffic into a VXLAN header and strip off the encapsulation.

This example shows how to configure VXLAN on MX Series routers using switch options in a default bridge domain.

- [Requirements on page 1934](#)
- [Overview on page 1934](#)
- [Configuring VXLAN on MX Series Routers on page 1935](#)
- [Verification on page 1941](#)

### Requirements

---

This example uses the following hardware and software components:

- An MX Series router
- A VXLAN capable peer router
- Junos OS Release 14.1

### Overview

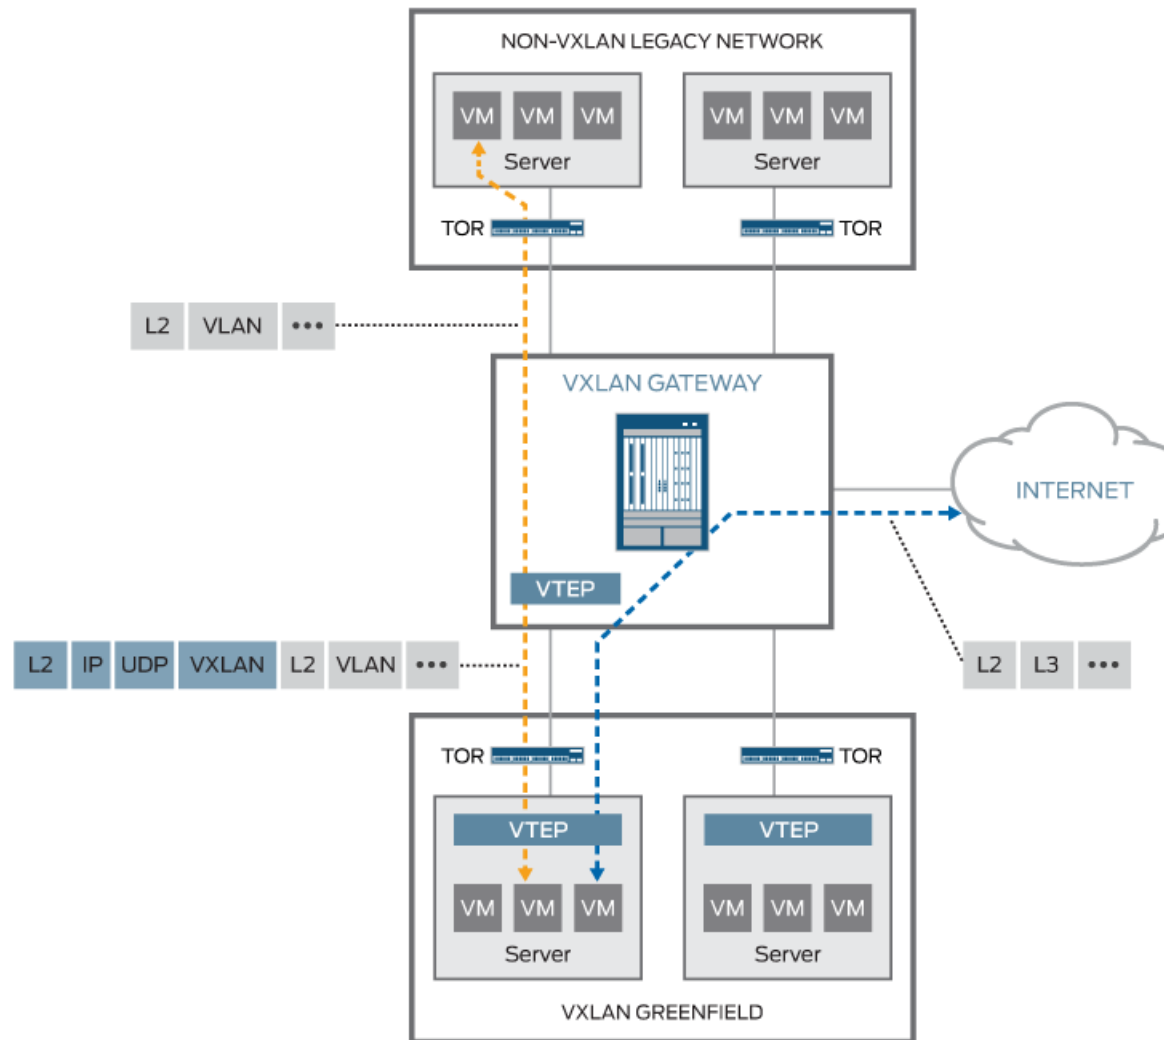
---

In this example, VXLAN is configured to run on a default bridge domain. VTEP interfaces sources are configured to the loopback address, and VLAN groups are configured under bridge domains with VXLAN enabled. Interfaces are configured for VLAN tagging and encapsulation, and IRB is enabled. OSPF and PIM protocols are configured to facilitate unicast and multicast routing. The chassis is configured for GRES and enhanced IP services.



### Topology

Figure 1: VXLAN Topology



### Configuring VXLAN on MX Series Routers

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set switch-options vtep-source-interface lo0.0
set bridge-domains vlan-5 vxlan vni 100
set bridge-domains vlan-5 vxlan multicast-group 239.1.1.1
set bridge-domains vlan-5 vlan-id 100
set bridge-domains vlan-5 routing-interface irb.0
set bridge-domains vlan-5 interface xe-1/0/0.0
```

```
set bridge-domains vlan-6 vxlan vni 200
set bridge-domains vlan-6 vxlan multicast-group 239.1.1.1
set bridge-domains vlan-6 vlan-id 200
set bridge-domains vlan-6 routing-interface irb.1
set bridge-domains vlan-6 interface xe-2/0/0.0
set interfaces xe-1/0/0 vlan-tagging
set interfaces xe-1/0/0 encapsulation flexible-ethernet-services
set interfaces xe-1/0/0 unit 0 encapsulation vlan-bridge
set interfaces xe-1/0/0 unit 0 vlan-id 100
set interfaces xe-2/0/0 vlan-tagging
set interfaces xe-2/0/0 encapsulation flexible-ethernet-services
set interfaces xe-2/0/0 unit 0 encapsulation vlan-bridge
set interfaces xe-2/0/0 unit 0 vlan-id 200
set interface irb unit 0 family inet address 5.5.5.1/24
set interface irb unit 1 family inet address 6.6.6.1/24
set interfaces lo0 unit 0 family inet address 3.3.3.3/32
set protocols ospf area 0.0.0.0 interface ge-8/3/8.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface xe-0/1/3.0
set protocols ospf area 0.0.0.0 interface ge-8/3/2.0
set protocols pim rp static address 10.2.1.3
set protocols pim interface lo0.0 mode bidirectional-sparse
set protocols pim interface ge-8/3/8.0 mode bidirectional-sparse
set protocols pim interface xe-0/1/3.0 mode bidirectional-sparse
set protocols pim interface ge-8/3/2.0 mode bidirectional-sparse
set chassis redundancy graceful-switchover
set chassis aggregated-devices ethernet device-count 10
set chassis fpc 1 pic 0 tunnel-services bandwidth 10g
set chassis network-services enhanced-ip
```

### *Configuring VXLAN*

#### **Step-by-Step Procedure**

The following example show how to set up a basic VXLAN configuration with default bridge domains and switch options. To configure VXLAN on an MX Series router, follow these steps:

1. Configure VTEP interface sources under **switch-options** for the default-switch.  
[edit]  
user@router# set switch-options vtep-source-interface lo0.0
2. Set up a VLAN group named **vlan-5** and set its VXLAN Network Identifier (VNI) to 100.  
[edit]  
user@router# set bridge-domains vlan-5 vxlan vni 100
3. Configure the **vlan-5** multicast group address for VXLAN.  
[edit]  
user@router# set bridge-domains vlan-5 vxlan multicast-group 239.1.1.1
4. Set the VLAN ID to 100 for **vlan-5**.  
[edit]  
user@router# set bridge-domains vlan-5 vlan-id 100

5. Configure integrated bridging and routing (IRB) for **vlan-5**.  

```
[edit]
user@router# set bridge-domains vlan-5 routing-interface irb.0
```
6. Assign the xe-1/0/0.0 interface to **vlan-5**.  

```
[edit]
user@router# set bridge-domains vlan-5 interface xe-1/0/0.0
```
7. Set up a VLAN group named **vlan-6** and set its VXLAN Network Identifier (VNI) to 200.  

```
[edit]
user@router# set bridge-domains vlan-6 vxlan vni 200
```
8. Configure the **vlan-6** multicast group address for VXLAN.  

```
[edit]
user@router# set bridge-domains vlan-6 vxlan multicast-group 239.1.1.1
```
9. Set the VLAN ID to 100 for **vlan-6**.  

```
[edit]
user@router# set bridge-domains vlan-6 vlan-id 200
```
10. Configure IRB for **vlan-6**.  

```
[edit]
user@router# set bridge-domains vlan-6 routing-interface irb.1
```
11. Assign the xe-2/0/0.0 interface to **vlan-6**.  

```
[edit]
user@router# set bridge-domains vlan-6 interface xe-2/0/0.0
```
12. Set up VLAN tagging for xe-1/0/0.  

```
[edit]
user@router# set interfaces xe-1/0/0 vlan-tagging
```
13. Configure flexible Ethernet service encapsulation on xe-1/0/0.  

```
[edit]
user@router# set interfaces xe-1/0/0 encapsulation flexible-ethernet-services
```
14. Set up VLAN bridging encapsulation for xe-1/0/0 unit 0.  

```
[edit]
user@router# set interfaces xe-1/0/0 unit 0 encapsulation vlan-bridge
```
15. Set the xe-1/0/0 unit 0 VLAN ID to 100.  

```
[edit]
user@router# set interfaces xe-1/0/0 unit 0 vlan-id 100
```
16. Configure VLAN tagging for xe-2/0/0  

```
[edit]
user@router# set interfaces xe-2/0/0 vlan-tagging
```

17. Set up flexible Ethernet service encapsulation on xe-2/0/0.  
[edit]  
user@router# set interfaces xe-2/0/0 encapsulation flexible-ethernet-services
18. Configure VLAN bridging encapsulation for xe-2/0/0 unit 0.  
[edit]  
user@router# set interfaces xe-2/0/0 unit 0 encapsulation vlan-bridge
19. Set the xe-2/0/0 unit 0 VLAN ID to 200.  
[edit]  
user@router# set interfaces xe-2/0/0 unit 0 vlan-id 200
20. Configure the IRB unit 0 family inet address.  
[edit]  
user@router# set interface irb unit 0 family inet address 5.5.5.1/24
21. Configure the IRB unit 1 family inet address.  
[edit]  
user@router# set interface irb unit 1 family inet address 6.6.6.1/24
22. Set the family inet address for the loopback unit 0.  
[edit]  
user@router# set interfaces lo0 unit 0 family inet address 3.3.3.3/32
23. Set up OSPF for the ge-8/3/8.0 interface.  
[edit]  
user@router# set protocols ospf area 0.0.0.0 interface ge-8/3/8.0
24. Configure OSPF for the loopback interface.  
[edit]  
user@router# set protocols ospf area 0.0.0.0 interface lo0.0
25. Set up OSPF for the xe-0/1/3.0 interface.  
[edit]  
user@router# set protocols ospf area 0.0.0.0 interface xe-0/1/3.0
26. Configure OSPF for the ge-8/3/2.0 interface.  
[edit]  
user@router# set protocols ospf area 0.0.0.0 interface ge-8/3/2.0
27. Set up the static address for the physical interface module (PIM) rendezvous point (RP).  
[edit]  
user@router# set protocols pim rp static address 10.2.1.3
28. Configure the loopback interface to bidirectional sparse mode for the PIM protocol.  
[edit]  
user@router# set protocols pim interface lo0.0 mode bidirectional-sparse

29. Set the ge-8/3/8.0 interface to bidirectional sparse mode for the PIM protocol.  

```
[edit]
user@router# set protocols pim interface ge-8/3/8.0 mode bidirectional-sparse
```
30. Configure the xe-0/1/3.0 interface to bidirectional sparse mode for the PIM protocol.  

```
[edit]
user@router# set protocols pim interface xe-0/1/3.0 mode bidirectional-sparse
```
31. Set the ge-8/3/2.0 interface to bidirectional sparse mode for the PIM protocol.  

```
[edit]
user@router# set protocols pim interface ge-8/3/2.0 mode bidirectional-sparse
```
32. Configure redundant graceful switchover on the chassis.  

```
[edit]
user@router# set chassis redundancy graceful-switchover
```
33. Set the aggregated ethernet device count to 10.  

```
[edit]
user@router# set chassis aggregated-devices ethernet device-count 10
```
34. Configure the tunnel services bandwidth for FPC 1/PIC 0.  

```
[edit]
user@router# set chassis fpc 1 pic 0 tunnel-services bandwidth 10g
```
35. Enable enhanced IP for network services on the chassis.  

```
[edit]
user@router# set chassis network-services enhanced-ip
```

### Results

From configuration mode, confirm your configuration by entering the following commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@router# show switch-options
```

```
switch-options {
  vtep-source-interface lo0.0;
}
```

```
user@router# show bridge-domains
```

```
bridge-domains {
  vlan-5 {
    vxlan {
      vni 100;
      multicast-group 239.1.1.1;
    }
    vlan-id 100;
    routing-interface irb.0;
    interface xe-1/0/0.0;
  }
}
```

```
vlan-6 {  
  vxlan {  
    vni 200;  
    multicast-group 239.2.1.1;  
  }  
  vlan-id 200;  
  routing-interface irb.1;  
  interface xe-2/0/0.0;  
}  
}
```

user@router# show interfaces

```
interfaces {  
  xe-1/0/0 {  
    vlan-tagging;  
    encapsulation flexible-ethernet-services;  
    unit 0 {  
      encapsulation vlan-bridge;  
      vlan-id 100;  
    }  
  }  
  xe-2/0/0 {  
    vlan-tagging;  
    encapsulation flexible-ethernet-services;  
    unit 0 {  
      encapsulation vlan-bridge;  
      vlan-id 200;  
    }  
  }  
  irb {  
    unit 0 {  
      family inet {  
        address 5.5.5.1/24;  
      }  
    }  
    unit 1 {  
      family inet {  
        address 6.6.6.1/24;  
      }  
    }  
  }  
  lo0 {  
    unit 0 {  
      family inet {  
        address 3.3.3.3/32;  
      }  
    }  
  }  
}
```

user@router# show protocols ospf

```
area 0.0.0.0 {  
  interface ge-8/3/8.0;  
  interface lo0.0;  
  interface xe-0/1/3.0;  
  interface ge-8/3/2.0;
```

```
}  
user@router# show protocols pim  
  
rp {  
  static {  
    address 10.2.1.3;  
  }  
}  
  
user@router# show chassis  
  
redundancy {  
  graceful-switchover;  
}  
aggregated-devices {  
  ethernet {  
    device-count 10;  
  }  
}  
fpc 1 {  
  pic 0 {  
    tunnel-services {  
      bandwidth 10g;  
    }  
  }  
}  
network-services enhanced-ip;
```

---

## Verification

Confirm that the configuration is working properly.

- [Verifying Reachability on page 1941](#)
- [Verifying VXLAN on page 1942](#)

### *Verifying Reachability*

**Purpose** Verify that the network is up and running with the proper interfaces and routes installed.

**Action** user@router> show interfaces terse irb

| Interface | Admin | Link | Proto | Local        | Remote |
|-----------|-------|------|-------|--------------|--------|
| irb       | up    | up   |       |              |        |
| irb.0     | up    | up   | inet  | 5.5.5.1/24   |        |
|           |       |      |       | multiservice |        |
| irb.1     | up    | up   | inet  | 6.6.6.1/24   |        |
|           |       |      |       | multiservice |        |

user@router> ping 5.5.5.1/24

```
PING 5.5.5.1 (5.5.5.1): 56 data bytes
64 bytes from 5.5.5.1: icmp_seq=0 ttl=64 time=0.965 ms
64 bytes from 5.5.5.1: icmp_seq=1 ttl=64 time=0.960 ms
64 bytes from 5.5.5.1: icmp_seq=2 ttl=64 time=0.940 ms
^C
--- 1.1.1.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.940/0.955/0.965/0.011 ms
```

**Meaning** Use the **show interfaces terse irb** command to verify that the IRB interface has been properly configured. The **irb.0** and **irb.1** interfaces should display the proper multiservice inet addresses.

Use the **ping** command to confirm that the network is connected to the IRB multiservice address.

#### *Verifying VXLAN*

**Purpose** Verify that VXLAN is working and the proper protocols are enabled.



**Action** `user@router> show interfaces vtep`

```
Physical interface: vtep, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 575
  Type: Software-Pseudo, Link-level type: VxLAN-Tunnel-Endpoint, MTU: 1600, Speed:
Unlimited
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vtep.32768 (Index 334) (SNMP ifIndex 607)
    Flags: Up SNMP-Traps Encapsulation: ENET2
    VXLAN Endpoint Type: Source, VXLAN Endpoint Address: 10.255.187.32, L2 Routing
Instance: default-switch, L3 Routing Instance: default
    Input packets : 0
    Output packets: 0

user@router> show l2-learning vxlan-tunnel-end-point remote mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Logical system : <default>
Routing instance : default-switch
  Bridging domain : vlan-5+100, VLAN : 100, VNID : 100
  Bridging domain : vlan-6+200, VLAN : 200, VNID : 200

user@router> show l2-learning vxlan-tunnel-end-point source
```

| Logical System Name | Id | SVTEP-IP      | IFL   | L3-Idx |             |
|---------------------|----|---------------|-------|--------|-------------|
| <default>           | 0  | 10.255.187.32 | lo0.0 | 0      |             |
| L2-RTT              |    | Bridge Domain |       | VNID   | MC-Group-IP |
| default-switch      |    | vlan-5+100    |       | 100    | 239.1.1.1   |
| default-switch      |    | vlan-6+200    |       | 200    | 239.1.1.1   |

**Meaning** Use the `show interface vtep` command to displays information about VXLAN endpoint configuration. Make sure the routing instance is assigned to the default-switch..

Use the `show l2-learning vxlan-tunnel-end-point remote mac-table` command to confirm that the bridging domain VLAN groups were configured correctly.

Use the `show l2-learning vxlan-tunnel-end-point source` command to confirm the multicast IP addresses for bridging domain VLAN groups.

**Related Documentation**

- [Understanding VXLANs on page 1912](#)
- [show bridge mac-table on page 2010](#)
- [show vpls mac-table on page 2029](#)

## Examples: Configuring VXLANs on QFX Series Switches

These examples show how to configure VXLANs on QFX Series Switches for several use cases.

- [Example: Configuring a VXLAN Transit Switch on page 1944](#)
- [Example: Configuring a VXLAN Layer 2 Gateway on page 1945](#)

### Example: Configuring a VXLAN Transit Switch

If a QFX switch acts as a transit switch for downstream devices acting as VTEPs, you do not need to configure any VXLAN information on the QFX switch. You do need to configure PIM on the switch so that it can form the multicast tree required so that the VTEPs can establish reachability with each other.

- [Requirements on page 1944](#)
- [Overview on page 1944](#)
- [Configuring PIM on the Transit Switches on page 1945](#)

#### **Requirements**

This example uses the following hardware and software components:

- Two QFX5100 switches
- Junos OS 14.1X53-D10

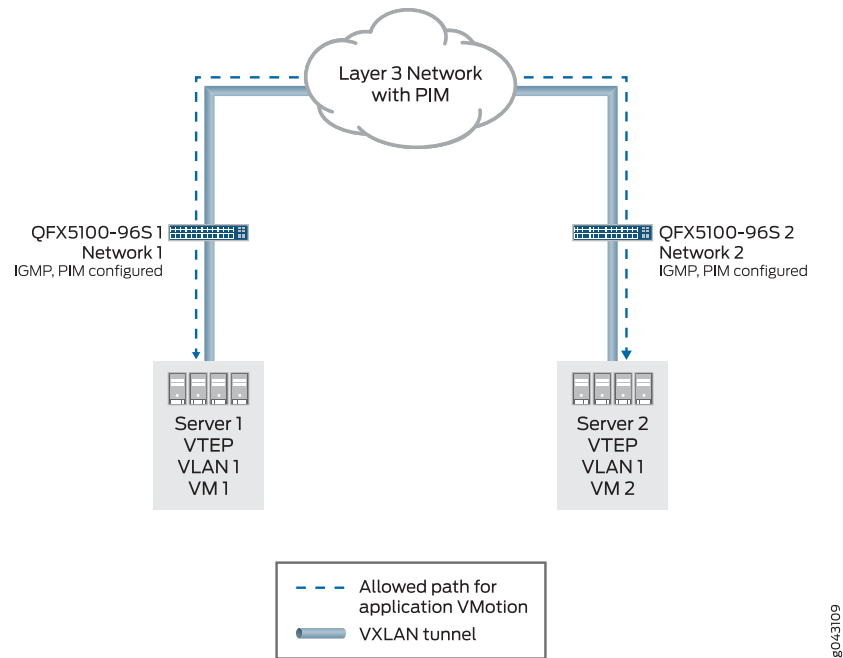
#### **Overview**

This example shows a simple use case in which QFX switches are connected to downstream servers acting as VTEPs and need to forward VXLAN packets between VM 1 on Server 1 and VM 2 on Server 2. Because this configuration allows Layer 2 connectivity between the VMs through the VXLAN tunnels, applications can VMotion between the VMs.

#### **Topology**

[Figure 33 on page 1945](#) shows a QFX 5100 switch configured to forward VXLAN packets for downstream VTEPs.

Figure 33: QFX5100 Acting as a VXLAN Transit Switch



### Configuring PIM on the Transit Switches

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols pim interface all
set protocols pim rp static address ip-address
```

#### Step-by-Step Procedure

If you are not using a controller to create a VXLAN control plane, you must enable PIM on each switch so that the VTEP can use multicast groups to advertise its existence and to learn about other VTEPs. (Configuring PIM automatically enables IGMP.) You do not need to perform any VXLAN-specific configuration. Note that you also do not need to configure VLAN 1 or 2 on either switch.

1. Enable PIM:
 

```
[edit]
user@switch# set protocols pim interface all
```
2. Configure the address of a PIM rendezvous point:
 

```
[edit]
user@switch# set protocols pim rp static address ip-address
```

### Example: Configuring a VXLAN Layer 2 Gateway

If a QFX switch is connected to a downstream server that hosts a VM that needs Layer 2 connectivity with another VM that is reachable only through a Layer 3 network, you

must configure the switch to act as a VTEP—that is, a Layer 2 gateway for downstream Layer 2 devices. You also need to configure PIM on the switch so that it can form the multicast tree required for reachability with other VTEPs and to allow BUM traffic to be forwarded between the VTEPs.

- [Requirements on page 1946](#)
- [Overview on page 1946](#)
- [Configuring the Switches on page 1947](#)
- [Verification on page 1950](#)

### **Requirements**

This example uses the following hardware and software components:

- Two QFX5100 switches
- Junos OS 14.1X53-D10

### **Overview**

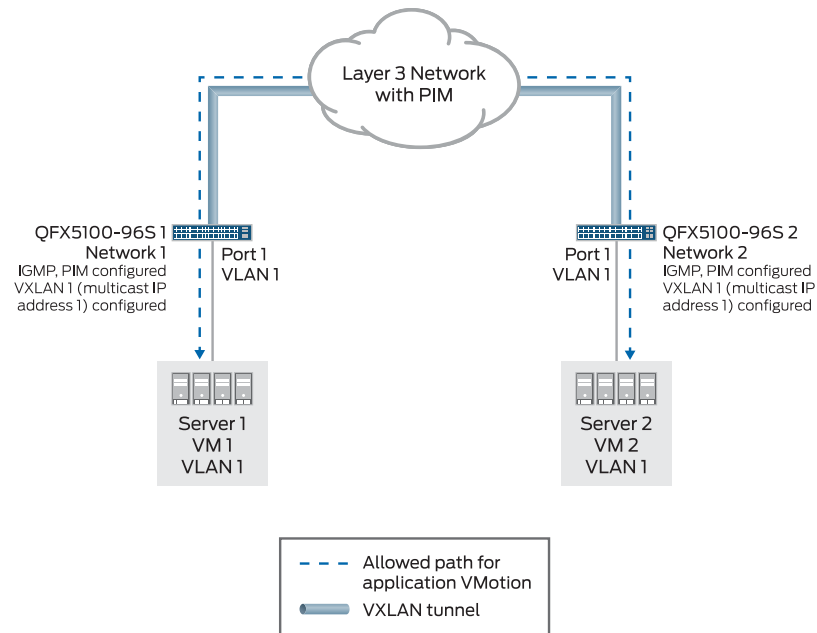
This example shows a use case in which QFX switches are connected to downstream VTEPs and need to allow Layer 2 connectivity between VM 1 on Server 1 and VM 2 on Server 2 so that VMotion can occur between the VMs. The servers in this example can be in the same or different data centers—the only constraint is that there must be Layer 3 connectivity between the QFX switches. This allows your network to be very agile in response to demand for server usage or changes in bandwidth requirements.

Note that because the same VLAN exists in both Layer 2 domains and both switches encapsulate the VLAN traffic into the same VXLAN, you do not need a gateway for the VXLAN traffic in the Layer 3 network. The Layer 3 VXLAN packets are routed normally and no de-encapsulation or re-encapsulation is required..

### **Topology**

[Figure 34 on page 1947](#) shows a QFX 5100 switch configured to act as a VTEP.

Figure 34: QFX5100 Acting as a VTEP



8043108

### Configuring the Switches

#### CLI Quick Configuration

To quickly configure the QFX5100-96S 1 in this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces lo0 unit 0 family inet address 10.1.1.1
set switch-options vtep-source-interface lo0.0
set protocols pim interface all
set protocols pim rp static address 10.2.2.2
set vlans VLAN1 vlan-id 100 vxlan vni 100 multicast-group 232.1.1.1
set vlans VLAN1 vxlan encapsulate-inner-vlan
set vlans VLAN1 vxlan decapsulate-inner-vlan
set vlans VLAN1 vxlan unreachable-vtep-aging-timer 600
set interfaces xe-0/0/0 unit 0 family inet address 10.2.2.100/24
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members all
```

The configuration for QFX5100-96S 2 is almost identical. The only changes a few of the addresses:

```
set interfaces lo0 unit 0 family inet address 10.1.1.2
set switch-options vtep-source-interface lo0.0
set protocols pim interface all
set protocols pim rp static address 10.2.2.2
set vlans VLAN1 vlan-id 100 vxlan vni 100 multicast-group 232.1.1.1
set vlans VLAN1 vxlan encapsulate-inner-vlan
set vlans VLAN1 vxlan decapsulate-inner-vlan
set vlans VLAN1 vxlan unreachable-vtep-aging-timer 600
```

```
set interfaces xe-0/0/0 unit 0 family inet address 10.2.2.200/24
set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members all
```



**NOTE:** You must configure the same multicast group address for VLAN1 on both switches.

**Step-by-Step Procedure** Perform the following procedure on both switches to set up the example configuration. You do not need to perform any VXLAN-specific configuration. Note that you also do not need to configure VLAN 1 or 2 on either switch.

1. Create a reachable IPv4 address on the loopback interface.  

```
[edit]
user@switch# set interfaces lo0.0 unit 0 family inet address 10.1.1.1
```

For switch QFX5100-96S 2, use address 10.1.1.2.
2. Configure the loopback interface—and therefore, its associated address—to be used as the tunnel source address:  

```
[edit]
user@switch# set switch-options vtep-source-interface lo0.0
```
3. Enable PIM:  

```
[edit]
user@switch# set protocols pim interface all
```
4. Configure the address of a PIM rendezvous point:  

```
[edit]
user@switch# set protocols pim rp static address 10.2.2.2
```
5. Create a VLAN, map it to a VXLAN, and assign a multicast group address to the VXLAN. All members of a VXLAN must use the same multicast group address:  

```
[edit]
user@switch# set vlans VLAN1 vlan-id 100 vxlan vni 100 multicast-group 232.1.1.1
```

In this example, the **vlan-id** and **vni** are both set to 100. This is done only for simplicity and clarity. You do not need to set the **vlan-id** and **vni** to the same value.
6. (Optional) Configure the switch to retain the original VLAN tag (in the inner Ethernet packet) after VXLAN encapsulation. By default, the original tag is dropped when the packet is encapsulated:  

```
[edit]
user@switch# set vlans VLAN1 vxlan encapsulate-inner-vlan
```
7. (Optional) Configure the switch to de-encapsulate and accept original VLAN tags in VXLAN packets. By default, a preserved VLAN tag is dropped when the packet is de-encapsulated:  

```
[edit]
user@switch# set vlans VLAN1 vxlan decapsulate-accept-inner-vlan
```

(Optional) Configure the system to age out the address for the remote VTEP (the other QFX5100 switch) if all the MAC addresses learned from that VTEP age out. The address for the remote VTEP expires the configured number of seconds after the last learned MAC address expires.

- ```
[edit]
user@switch# set vlans VLAN1 vxlan unreachable-vtep-aging-timer 600
```
8. Configure the interface that connects to the Layer 3 network:
- ```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.2.2.100/24
For switch QFX5100-96S 2, use address 10.2.2.200.
```
9. Configure the server-facing interface to support multiple VLANs:
- ```
[edit]
user@switch# set interfaces xe-0/0/1 unit 0 family ethernet-switching interface-mode
trunk
[edit]
user@switch# set interfaces xe-0/0/1 unit 0 family ethernet-switching vlan members all
```



**NOTE:** Because this example shows only one VLAN, this step is not required for the example. In a real-world configuration, however, it would be required in order to support multiple VMs connected to multiple VLANs. In this case you would also need to configure additional VLAN to VXLAN mappings.

### Results

From configuration mode, confirm your configuration by entering the following commands on QFX5100-96S 1. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch# show switch-options
```

```
vtep-source-interface lo0.0;
```

```
user@switch# show vlans
```

```
VLAN1 {
  vlan-id 100;
  vxlan {
    vni 100;
    multicast-group 232.1.1.1;
    encapsulate-inner-vlan;
  }
}
```

```
user@switch# show interfaces
```

```
xe-0/0/0 {
  unit 0 {
    family inet {
      address 10.2.2.100/24;
    }
  }
}
xe-0/0/1 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
    }
  }
}
```

```
        vlan {
            members all;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.1.1.1/32;
        }
    }
}

user@switch# show protocols pim

rp {
    static {
        address 10.2.2.2;
    }
}
interface xe-0/0/1.0 {
    mode sparse;
}
```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying VXLAN Reachability on page 1950](#)
- [Verifying That the Local VTEP is Configured Correctly on page 1951](#)
- [Verifying MAC Learning from the Remote VTEP on page 1951](#)
- [Monitor the Remote Interface on page 1951](#)

### **Verifying VXLAN Reachability**

**Purpose** On QFX5100-96S 1, verify that there is connectivity with the remote VTEP (QFX5100-96S 2).

**Action** user@switch> show ethernet-switching vxlan-tunnel-end-point remote

Logical System Name	Id	SVTEP-IP	IFL	L3-Idx
<default>	0	10.1.1.2	1o0.0	0
RVTEP-IP	IFL-Idx	NH-Id		
10.1.1.2	559	1728		
VNID	MC-Group-IP			
100	232.1.1.1			

**Meaning** The VTEP on QFX5100-96S 2 is reachable because its IP address (the address assigned to the loopback interface) appears in the output. The output also shows that the VXLAN (VNI 100) and corresponding multicast group are configured correctly on the remote VTEP.



**Verifying That the Local VTEP is Configured Correctly**

**Purpose** On QFX5100-96S 1, verify that the tunnel endpoint is correct..

**Action** user@switch> show ethernet-switching vxlan-tunnel-end-point source

Logical System Name	Id	SVTEP-IP	IFL	L3-Idx
<default>	0	10.1.1.1	lo0.0	0
L2-RTT	Bridge Domain		VNID	MC-Group-IP
default-switch	VLAN1+100		100	232.1.1.1

**Meaning** The VTEP on QFX5100-96S 1 shows the correct tunnel source IP address (assigned to the loopback interface), VLAN, and multicast group for the VXLAN.

**Verifying MAC Learning from the Remote VTEP**

**Purpose** On QFX5100-96S 1, verify that it is learning MAC addresses from the remote VTEP.

**Action** user@switch> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

Vlan name	MAC address	MAC flags	Age	Logical interface
VLAN1	00:00:00:ff:ff:ff	D	-	vtep.12345
VLAN1	00:10:94:00:00:02	D	-	xe-0/0/0.0

**Meaning** This shows the MAC addresses learned from the remote VTEP (in addition to those learned on the normal Layer 2 interfaces). It also shows the logical name of the remote VTEP interface (**vtep.12345** in the above output).

**Monitor the Remote Interface**

**Purpose** On QFX5100-96S 1, monitor traffic details for the remote VTEP interface.

**Action** user@switch> show interface vtep.12345 detail

```
M   Flags: Up SNMP-Traps Encapsulation: ENET2
      VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 10.1.1.2, L2 Routing
Instance: default-switch, L3 Routing Instance: default
      Traffic statistics:
        Input bytes :          228851738624
        Output bytes :              0
        Input packets:          714162415
        Output packets:              0
      Local statistics:
        Input bytes :              0
        Output bytes :              0
        Input packets:              0
        Output packets:              0
      Transit statistics:
        Input bytes :          228851738624          0 bps
        Output bytes :              0          0 bps
        Input packets:          714162415          0 pps
        Output packets:              0          0 pps
      Protocol eth-switch, MTU: 1600, Generation: 277, Route table: 5
```

**Meaning** This shows traffic details for the remote VTEP interface. To get this information, you must supply the logical name of the remote VTEP interface (vtep.12345 in the above output), which you can learn by using the **show ethernet-switching table** command.

**Related Documentation**

- [Understanding VXLANs on page 1912](#)

## Example: Configuring VXLAN to VPLS Stitching with OVSDb

Virtual Extensible LAN (VXLAN) can be utilized with the Open vSwitch Database (OVSDb) management protocol in a VPLS-enabled network to stitch a virtualized data center into a Layer 2 VPN network. This configuration allows for seamless interconnection between different data centers using Layer 2 VPN regardless of whether it is virtualized, physical, or both.

- [Requirements on page 1952](#)
- [Overview on page 1953](#)
- [Configuration on page 1954](#)
- [Verification on page 1967](#)

### Requirements

---

This example uses the following hardware and software components:

- Two MX Series routers running Junos OS 14.1R2 or later
- Two MX Series routers running Junos OS 14.1R2 or later with an OVSDb software package. The release of this package must be the same as the Junos OS release running on the device.
- One EX9200 switch

- One VMware NSX controller running NSX software version 4.0.3
- NSX Manager version 4.0.3

Before you start the configuration, you must perform the following tasks:

- In NSX Manager version 4.0.3 or the NSX API, configure a logical switch for each VXLAN that OVSDb will manage. This example implements two OVSDb-managed VXLANs, so you must configure two logical switches. After the configuration of each logical switch, NSX automatically generates a universally unique identifier (UUID) for the logical switch. If you have not done so already, retrieve the UUID for each logical switch. A sample UUID is 28805c1d-0122-495d-85df-19abd647d772. When configuring the equivalent VXLANs on the Juniper Networks device, you must use the UUID of the logical switch as the bridge domain name.

For more information about logical switches and VXLANs, see [“Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment” on page 1905](#).


- Create an SSL private key and certificate, and install them in the `/var/db/certs` directory of the Juniper Networks device. For more information, see [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers” on page 1984](#).

---

## Overview

In this example, four MX Series routers are configured to function together for VXLAN to virtual private LAN service (VPLS) stitching. Each router performs a different role in the configuration. The following diagram shows the topology of these MX Series routers. MX1 is the core router that handles Layer 3 traffic and protocols. MX2 is the VXLAN gateway router that functions as a virtual tunnel endpoint (VTEP) and handles switching for Layer 2, VPLS, and VXLAN. The MX3 router is configured to handle VPLS traffic. The MX4 router is configured as a VTEP to accept and decapsulate VXLAN packets.

### Topology



### Configuration

To configure VXLAN to VPLS stitching with OVSDDB:

- [Configuring MX1 on page 1959](#)
- [Configuring MX2 on page 1960](#)
- [Configuring MX3 on page 1963](#)
- [Configuring MX4 on page 1963](#)
- [Results on page 1967](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
MX1 set groups global interfaces lo0 unit 0 family inet address 127.0.0.1/32
set groups global interfaces lo0 unit 0 family inet address 10.255.181.13/32primary
set groups global interfaces lo0 unit 0 family iso
47.0005.80ff.f800.0000.0108.0001.0102.5518.1013.00
set groups global interfaces lo0 unit 0 family inet6 address abcd::10:255:181:13/128primary
set interfaces apply-groups LAG-options
set interfaces xe-0/0/2 mtu 9000
set interfaces xe-0/0/2 unit 0 family inet address 80.80.0.250/24
set interfaces xe-0/0/3 unit 0 family inet address 30.30.30.6/30
set interfaces ge-1/0/5 mtu 1600
set interfaces ge-1/0/5 unit 0 family inet address 20.20.20.2/30
set interfaces ge-1/0/6 unit 0 family inet address 20.20.20.10/30
set interfaces ge-1/0/7 gigether-options 802.3ad ae2
set interfaces ge-1/0/8 gigether-options 802.3ad ae2
set interfaces ge-1/1/2 unit 0 family inet address 30.30.30.2/30
set interfaces ae2 mtu 1600
```

```

set interfaces ae2 unit 0 family inet address 20.20.20.6/30
set protocols ospf area 0.0.0.0 interface xe-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-1/0/5.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ae2.0
set protocols ospf area 0.0.0.0 interface ge-1/0/6.0
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.181.13
set protocols pim interface all mode sparse-dense

```

```

MX2 set interfaces xe-1/2/0 gigether-options 802.3ad ae1
set interfaces xe-5/0/0 gigether-options 802.3ad ae1
set interfaces ge-7/0/0 gigether-options 802.3ad ae2
set interfaces ge-7/0/1 mtu 1600
set interfaces ge-7/0/1 unit 0 family inet address 20.20.20.1/30
set interfaces ge-7/1/3 gigether-options 802.3ad ae2
set interfaces xe-10/3/0 vlan-tagging
set interfaces xe-10/3/0 encapsulation flexible-ethernet-services
set interfaces xe-10/3/0 unit 100 family bridge interface-mode trunk
set interfaces xe-10/3/0 unit 100 family bridge vlan-id-list 100-101
set interfaces xe-10/3/0 unit 102 family bridge interface-mode trunk
set interfaces xe-10/3/0 unit 102 family bridge vlan-id-list 102-103
set interfaces ae1 vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 unit 100 family bridge interface-mode trunk
set interfaces ae1 unit 100 family bridge vlan-id-list 100-101
set interfaces ae1 unit 102 family bridge interface-mode trunk
set interfaces ae1 unit 102 family bridge vlan-id-list 102-103
set interfaces ae2 mtu 1600
set interfaces ae2 unit 0 family inet address 20.20.20.5/30
set interfaces irb unit 1 family inet address 2.2.1.1/24
set interfaces irb unit 2 family inet address 2.2.2.1/24
set interfaces irb unit 3 family inet address 2.2.3.1/24
set interfaces irb unit 4 family inet address 2.2.4.1/24
set interfaces irb unit 5 family inet address 2.2.5.1/24
set interfaces irb unit 11 family inet address 2.2.11.1/24
set interfaces irb unit 12 family inet address 2.2.12.1/24
set interfaces irb unit 13 family inet address 2.2.13.1/24
set interfaces irb unit 14 family inet address 2.2.14.1/24
set interfaces irb unit 15 family inet address 2.2.15.1/24
set interfaces lo0 unit 1 family inet address 200.1.1.1/32
set protocols ospf area 0.0.0.0 interface ge-7/0/1.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols l2-learning traceoptions file vxlan-l2ald.log
set protocols l2-learning traceoptions file size 100m
set protocols l2-learning traceoptions file files 10
set protocols l2-learning traceoptions level all
set protocols l2-learning traceoptions flag all
set protocols layer2-control nonstop-bridging
set protocols ovssdb traceoptions file ovssdb.log
set protocols ovssdb traceoptions file size 100m
set protocols ovssdb traceoptions file files 10
set protocols ovssdb traceoptions level all
set protocols ovssdb traceoptions flag all

```

```

set protocols ovssdb interfaces xe-10/3/0.1
set protocols ovssdb interfaces xe-10/3/0.0
set protocols ovssdb interfaces ae1.0
set protocols ovssdb interfaces ae1.1
set protocols ovssdb controller 192.168.182.45 protocol ssl port 6632
set routing-instances default-VS1 vtep-source-interface lo0.1
set routing-instances default-VS1 instance-type virtual-switch
set routing-instances default-VS1 interface xe-10/3/0.102
set routing-instances default-VS1 interface ae1.102
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vlan-id 102
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab routing-interface irb.102
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ovssdb-managed
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan vni 102
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ingress-node-replication
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vlan-id 103
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 routing-interface irb.103
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ovssdb-managed
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan vni 103
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ingress-node-replication
set routing-instances vrf1 instance-type vrf
set routing-instances vrf1 interface ae2.0
set routing-instances vrf1 interface lo0.1
set routing-instances vrf1 route-distinguisher 100:100
set routing-instances vrf1 vrf-target target:100:100
set routing-instances vrf1 protocols ospf area 0.0.0.0 interface ae2.0
set routing-instances vrf1 protocols ospf area 0.0.0.0 interface lo0.1 passive
set routing-instances vrf1 protocols pim rp static address 10.255.181.13
set routing-instances vrf1 protocols pim interface all
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vlan-id 100
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 routing-interface irb.100
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan ovssdb-managed
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan vni 100
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan
    ingress-node-replication
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vlan-id 101
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 routing-interface irb.101
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan ovssdb-managed
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan vni 101
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan
    ingress-node-replication
set switch-options vtep-source-interface lo0.0

MX3 set groups global interfaces lo0 unit 0 family inet address 127.0.0.1/32
set groups global interfaces lo0 unit 0 family inet address 10.255.181.98/32 primary
set groups global interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5518.1098.00

```

```

set groups global interfaces lo0 unit 0 family inet6 address abcd::10:255:181:98/128
  primary
set interfaces xe-0/0/0 gigether-options 802.3ad ae1
set interfaces xe-0/0/1 gigether-options 802.3ad ae1
set interfaces xe-0/0/3 vlan-tagging
set interfaces xe-0/0/3 encapsulation flexible-ethernet-services
set interfaces xe-0/0/3 unit 0 family bridge interface-mode trunk
set interfaces xe-0/0/3 unit 0 family bridge vlan-id-list 1-40
set interfaces ae1 vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 unit 0 family bridge interface-mode trunk vlan-id-list 1-40
set protocols pim dense-groups 224.0.1.39/32
set protocols pim dense-groups 224.0.1.40/32
set protocols pim rp auto-rp discovery
set protocols pim interface all mode sparse-dense
set routing-instances vs1 instance-type virtual-switch
set routing-instances vs1 interface xe-0/0/3.0
set routing-instances vs1 interface ae1.0
set routing-instances vs1 bridge-domains v1 vlan-id-list 1-40

```

MX4

```

set groups global interfaces lo0 unit 0 family inet address 127.0.0.1/32
set groups global interfaces lo0 unit 0 family inet address 10.255.181.43/32 primary
set groups global interfaces lo0 unit 0 family iso address
  47.0005.80ff.f800.0000.0108.0001.0102.5518.1043.00
set groups global interfaces lo0 unit 0 family inet6 address abcd::10:255:181:43/128
  primary
set interfaces xe-0/0/0 vlan-tagging
set interfaces xe-0/0/0 encapsulation flexible-ethernet-services
set interfaces xe-0/0/0 unit 0 family bridge interface-mode trunk
set interfaces xe-0/0/0 unit 0 family bridge vlan-id-list 1-10
set interfaces xe-0/0/0 unit 1 family bridge interface-mode trunk
set interfaces xe-0/0/0 unit 1 family bridge vlan-id-list 11-15
set interfaces xe-0/0/0 unit 2 family bridge interface-mode trunk
set interfaces xe-0/0/0 unit 2 family bridge vlan-id-list 21-30
set interfaces xe-0/0/0 unit 3 family bridge interface-mode trunk
set interfaces xe-0/0/0 unit 3 family bridge vlan-id-list 31-40
set interfaces ge-0/2/6 unit 0 family inet address 30.30.30.1/30
set interfaces ge-0/3/0 unit 0 family inet address 3.3.3.2/30
set interfaces irb unit 1 family inet address 2.2.1.2/24
set interfaces irb unit 2 family inet address 2.2.2.2/24
set interfaces irb unit 3 family inet address 2.2.3.2/24
set interfaces irb unit 4 family inet address 2.2.4.2/24
set interfaces irb unit 5 family inet address 2.2.5.2/24
set interfaces irb unit 11 family inet address 2.2.11.2/24
set interfaces irb unit 12 family inet address 2.2.12.2/24
set interfaces irb unit 13 family inet address 2.2.13.2/24
set interfaces irb unit 14 family inet address 2.2.14.2/24
set interfaces irb unit 15 family inet address 2.2.15.2/24
set protocols ospf area 0.0.0.0 interface ge-0/2/6.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/3/0.0
set protocols l2-learning traceoptions file vxlan-l2ald.log size 100m files 10
set protocols l2-learning traceoptions level all
set protocols l2-learning traceoptions flag all
set protocols layer2-control nonstop-bridging

```

```
set protocols pim rp auto-rp discovery
set protocols pim rp static address 10.255.181.13
set protocols pim rp interface all mode sparse-dense
set protocols ovssdb traceoptions file ovssdb.log size 100m files 10
set protocols ovssdb traceoptions level all
set protocols ovssdb traceoptions flag all
set protocols ovssdb interfaces xe-0/0/0.1
set protocols ovssdb interfaces xe-0/0/0.0
set protocols ovssdb controller 192.168.182.45 protocol ssl port 6632
set routing-instances default-vs1 vtep-source-interface lo0.0
set routing-instances default-vs1 instance-type virtual-switch
set routing-instances default-vs1 interface xe-0/0/0.1
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vlan-id 11
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab routing-interface irb.11
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ovssdb-managed
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan vni 16777214
set routing-instances default-VS1 bridge-domains
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ingress-node-replication
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vlan-id 12
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 routing-interface irb.12
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ovssdb-managed
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan vni 12
set routing-instances default-VS1 bridge-domains
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ingress-node-replication
set routing-instances default-VS1 bridge-domains v13 vlan-id 13
set routing-instances default-VS1 bridge-domains v13 routing-interface irb.13
set routing-instances default-VS1 bridge-domains v13 vxlan vni 13
set routing-instances default-VS1 bridge-domains v13 vxlan multicast-group 228.1.1.13
set routing-instances default-VS1 bridge-domains v14 vlan-id 14
set routing-instances default-VS1 bridge-domains v14 routing-interface irb.14
set routing-instances default-VS1 bridge-domains v14 vxlan vni 14
set routing-instances default-VS1 bridge-domains v14 vxlan multicast-group 228.1.1.14
set routing-instances default-VS1 bridge-domains v15 vlan-id 15
set routing-instances default-VS1 bridge-domains v15 routing-interface irb.15
set routing-instances default-VS1 bridge-domains v15 vxlan vni 15
set routing-instances default-VS1 bridge-domains v15 vxlan multicast-group 228.1.1.15
set routing-instances default-VS2 bridge-domains v21 vlan-id 21
set routing-instances default-VS2 bridge-domains v21 vxlan vni 21
set routing-instances default-VS2 bridge-domains v21 vxlan multicast-group 228.1.1.21
set routing-instances default-VS2 bridge-domains v22 vlan-id 22
set routing-instances default-VS2 bridge-domains v22 vxlan vni 22
set routing-instances default-VS2 bridge-domains v22 vxlan multicast-group 228.1.1.22
set routing-instances default-VS2 bridge-domains v23 vlan-id 23
set routing-instances default-VS2 bridge-domains v23 vxlan vni 23
set routing-instances default-VS2 bridge-domains v23 vxlan multicast-group 228.1.1.23
set routing-instances default-VS2 bridge-domains v24 vlan-id 24
set routing-instances default-VS2 bridge-domains v24 vxlan vni 24
set routing-instances default-VS2 bridge-domains v24 vxlan multicast-group 228.1.1.24
```



```

set routing-instances default-VS2 bridge-domains v25 vlan-id 25
set routing-instances default-VS2 bridge-domains v25 vxlan vni 25
set routing-instances default-VS2 bridge-domains v25 vxlan multicast-group 228.1.1.25
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vlan-id 3
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 routing-interface irb.3
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan ovsdb-managed
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan vni 3
set bridge-domains 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan
    ingress-node-replication
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vlan-id 2
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 routing-interface irb.2
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan ovsdb-managed
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan vni 2
set bridge-domains cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan
    ingress-node-replication
set bridge-domains v1 vlan-id 1
set bridge-domains v1 routing-interface irb.1
set bridge-domains v1 vxlan vni 1
set bridge-domains v1 vxlan multicast-group 228.1.1.1
set bridge-domains v4 vlan-id 4
set bridge-domains v4 routing-interface irb.4
set bridge-domains v4 vxlan vni 4
set bridge-domains v4 vxlan multicast-group 228.1.1.4
set bridge-domains v5 vlan-id 5
set bridge-domains v5 routing-interface irb.5
set bridge-domains v5 vxlan vni 5
set bridge-domains v5 vxlan multicast-group 228.1.1.5
set switch-options vtep-source-interface lo0.0

```

### Configuring MX1

**Step-by-Step Procedure** The first router to be configured is the core router. This MX Series router handles Layer 3 traffic and protocols for the rest of the network.

To configure the MX1 router:

1. Specify the IPv4, IPv6, and ISO addresses for the loopback interface.
 

```

[edit groups global interfaces]
user@MX1# set lo0 unit 0 family inet address 127.0.0.1/32
user@MX1# set lo0 unit 0 family inet address 10.255.181.13/32primary
user@MX1# set lo0 unit 0 family iso
    47.0005.80ff.f800.0000.0108.0001.0102.5518.1013.00
user@MX1# set lo0 unit 0 family inet6 address abcd::10:255:181:13/128primary

```
2. Configure the Layer 3 network.
 

```

[edit interfaces]
user@MX1# set apply-groups LAG-options
user@MX1# set xe-0/0/2 mtu 9000
user@MX1# set xe-0/0/2 unit 0 family inet address 80.80.0.250/24
user@MX1# set xe-0/0/3 unit 0 family inet address 30.30.30.6/30
user@MX1# set ge-1/0/5 mtu 1600
user@MX1# set ge-1/0/5 unit 0 family inet address 20.20.20.2/30
user@MX1# set ge-1/0/6 unit 0 family inet address 20.20.20.10/30
user@MX1# set ge-1/0/7 gigether-options 802.3ad ae2
user@MX1# set ge-1/0/8 gigether-options 802.3ad ae2

```

```
user@MX1# set ge-1/1/2 unit 0 family inet address 30.30.30.2/30
user@MX1# set ae2 mtu 1600
user@MX1# set ae2 unit 0 family inet address 20.20.20.6/30
```

3. Enable OSPF and PIM.

```
[edit protocols]
user@MX1# set ospf area 0.0.0.0 interface xe-0/0/2.0
user@MX1# set ospf area 0.0.0.0 interface ge-1/0/5.0
user@MX1# set ospf area 0.0.0.0 interface lo0.0 passive
user@MX1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@MX1# set ospf area 0.0.0.0 interface ae2.0
user@MX1# set ospf area 0.0.0.0 interface ge-1/0/6.0
user@MX1# set ospf area 0.0.0.0 interface all
user@MX1# set pim rp local address 10.255.181.13
user@MX1# set pim interface all mode sparse-dense
```

### *Configuring MX2*

**Step-by-Step Procedure** The second router to be configured is the VXLAN gateway router. This MX Series router is configured as a VTEP, and it handles switching for Layer 2, VPLS, and VXLAN.

To configure the MX2 router:

1. Configure interfaces for the VXLAN gateway.

```
[edit interfaces]
user@MX2# set xe-1/2/0 gigether-options 802.3ad ae1
user@MX2# set xe-5/0/0 gigether-options 802.3ad ae1
user@MX2# set ge-7/0/0 gigether-options 802.3ad ae2
user@MX2# set ge-7/0/1 mtu 1600
user@MX2# set ge-7/0/1 unit 0 family inet address 20.20.20.1/30
user@MX2# set ge-7/1/3 gigether-options 802.3ad ae2
user@MX2# set xe-10/3/0 vlan-tagging
user@MX2# set xe-10/3/0 encapsulation flexible-ethernet-services
user@MX2# set xe-10/3/0 unit 100 family bridge interface-mode trunk
user@MX2# set xe-10/3/0 unit 100 family bridge vlan-id-list 100-101
user@MX2# set xe-10/3/0 unit 102 family bridge interface-mode trunk
user@MX2# set xe-10/3/0 unit 102 family bridge vlan-id-list 102-103
user@MX2# set ae1 vlan-tagging
user@MX2# set ae1 encapsulation flexible-ethernet-services
user@MX2# set ae1 unit 100 family bridge interface-mode trunk
user@MX2# set ae1 unit 100 family bridge vlan-id-list 100-101
user@MX2# set ae1 unit 102 family bridge interface-mode trunk
user@MX2# set ae1 unit 102 family bridge vlan-id-list 102-103
user@MX2# set ae2 mtu 1600
user@MX2# set ae2 unit 0 family inet address 20.20.20.5/30
user@MX2# set irb unit 1 family inet address 2.2.1.1/24
user@MX2# set irb unit 2 family inet address 2.2.2.1/24
user@MX2# set irb unit 3 family inet address 2.2.3.1/24
user@MX2# set irb unit 4 family inet address 2.2.4.1/24
user@MX2# set irb unit 5 family inet address 2.2.5.1/24
user@MX2# set irb unit 11 family inet address 2.2.11.1/24
user@MX2# set irb unit 12 family inet address 2.2.12.1/24
user@MX2# set irb unit 13 family inet address 2.2.13.1/24
user@MX2# set irb unit 14 family inet address 2.2.14.1/24
user@MX2# set irb unit 15 family inet address 2.2.15.1/24
```

```
user@MX2# set lo0 unit 1 family inet address 200.1.1/32
```

2. Configure OSPF interface settings and the Layer 2 learning traceoption file.

```
[edit protocols]
user@MX2# set ospf area 0.0.0.0 interface ge-7/0/1.0
user@MX2# set ospf area 0.0.0.0 interface fxp0.0 disable
user@MX2# set ospf area 0.0.0.0 interface lo0.0 passive
user@MX2# set l2-learning traceoptions file vxlan-l2ald.log
user@MX2# set l2-learning traceoptions file size 100m
user@MX2# set l2-learning traceoptions file files 10
user@MX2# set l2-learning traceoptions level all
user@MX2# set l2-learning traceoptions flag all
user@MX2# set layer2-control nonstop-bridging
```

3. Set up OVSDb tracing operations.

```
[edit protocols]
user@MX2# set ovssdb traceoptions file ovssdb.log
user@MX2# set ovssdb traceoptions file size 100m
user@MX2# set ovssdb traceoptions file files 10
user@MX2# set ovssdb traceoptions level all
user@MX2# set ovssdb traceoptions flag all
```

4. Specify that interfaces xe-10/3/0.1, xe-10/3/0.0, ae1.0, and ae1.1 are managed by OVSDb.

```
[edit protocols]
user@MX2# set ovssdb interfaces xe-10/3/0.1
user@MX2# set ovssdb interfaces xe-10/3/0.0
user@MX2# set ovssdb interfaces ae1.0
user@MX2# set ovssdb interfaces ae1.1
```

5. Configure a connection with an NSX controller.

```
[edit protocols]
user@MX2# set ovssdb controller 192.168.182.45 protocol ssl port 6632
```

6. Configure the **default-VS1** virtual switch instance as a VTEP.

```
[edit routing-instances]
user@MX2# set default-VS1 vtep-source-interface lo0.1
user@MX2# set default-VS1 instance-type virtual-switch
user@MX2# set default-VS1 interface xe-10/3/0.102
user@MX2# set default-VS1 interface ae1.102
```

7. Configure a set of bridge domains that are associated with VXLAN under the virtual switch instance.

```
[edit routing-instances]
user@MX2# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vlan-id 102
user@MX2# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab routing-interface irb.102
user@MX2# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ovssdb-managed
user@MX2# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan vni 102
user@MX2# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ingress-node-replication
```

```

user@MX2# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vlan-id 103
user@MX2# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 routing-interface irb.103
user@MX2# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ovsdb-managed
user@MX2# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan vni 103
user@MX2# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ingress-node-replication

```

8. Set up VPN routing and forwarding.

```

[edit routing- instances]
user@MX2# set vrf1 instance-type vrf
user@MX2# set vrf1 interface ae2.0
user@MX2# set vrf1 interface lo0.1
user@MX2# set vrf1 route-distinguisher 100:100
user@MX2# set vrf1 vrf-target target:100:100
user@MX2# set vrf1 protocols ospf area 0.0.0.0 interface ae2.0
user@MX2# set vrf1 protocols ospf area 0.0.0.0 interface lo0.1 passive
user@MX2# set vrf1 protocols pim rp static address 10.255.181.13
user@MX2# set vrf1 protocols pim interface all

```

9. Configure bridge domains with VXLAN information.

```

[edit bridge-domains]
user@MX2# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vlan-id 100
user@MX2# set 24a76aff-7e61-4520-a78d-3eca26ad7510 routing-interface irb.100
user@MX2# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan ovsdb-managed
user@MX2# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan vni 100
user@MX2# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan
ingress-node-replication
user@MX2# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vlan-id 101
user@MX2# set cadbc185-f60f-48a6-93fd-dc14a6420c60 routing-interface irb.101
user@MX2# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan ovsdb-managed
user@MX2# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan vni 101
user@MX2# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan
ingress-node-replication

```

10. Configure the loopback interface to be used as the tunnel source address.

```

[edit switch-options]
user@MX2# set vtep-source-interface lo0.0

```



**NOTE:** After completing this configuration, you must configure a gateway, which is the NSX equivalent of a hardware VTEP. This configuration implements one hardware VTEP, so you must configure one gateway, a gateway service, and a logical switch port using NSX Manager or the NSX API. For more information about the tasks you must perform as well as key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints” on page 1986](#).

**Configuring MX3**

**Step-by-Step Procedure** The third MX Series router must be configured to handle VPLS traffic.

To configure the MX3 router:

1. Specify the IPv4, IPv6, and ISO addresses for the loopback interface.
 

```
[edit groups global interfaces]
user@MX3# set lo0 unit 0 family inet address 127.0.0.1/32
user@MX3# set lo0 unit 0 family inet address 10.255.181.98/32 primary
user@MX3# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5518.1098.00
user@MX3# set lo0 unit 0 family inet6 address abcd::10:255:181:98/128 primary
```
2. Configure the interfaces.
 

```
[edit interfaces]
user@MX3# set fxp0 unit 0 family inet address 192.168.181.97/25
user@MX3# set xe-0/0/0 gigether-options 802.3ad ae1
user@MX3# set xe-0/0/1 gigether-options 802.3ad ae1
user@MX3# set xe-0/0/3 vlan-tagging
user@MX3# set xe-0/0/3 encapsulation flexible-ethernet-services
user@MX3# set xe-0/0/3 unit 0 family bridge interface-mode trunk
user@MX3# set xe-0/0/3 unit 0 family bridge vlan-id-list 1-40
user@MX3# set ae1 vlan-tagging
user@MX3# set ae1 encapsulation flexible-ethernet-services
user@MX3# set ae1 unit 0 family bridge interface-mode trunk vlan-id-list 1-40
```
3. Enable PIM.
 

```
[edit protocols]
user@MX3# set pim dense-groups 224.0.1.39/32
user@MX3# set pim dense-groups 224.0.1.40/32
user@MX3# set pim rp auto-rp discovery
user@MX3# set pim interface all mode sparse-dense
```
4. Configure the VPLS bridge domain and interfaces.
 

```
[edit routing-instances]
user@MX3# set vs1 instance-type virtual-switch
user@MX3# set vs1 interface xe-0/0/3.0
user@MX3# set vs1 interface ae1.0
user@MX3# set vs1 bridge-domains v1 vlan-id-list 1-40
```

**Configuring MX4**

**Step-by-Step Procedure** The fourth MX Series router is configured as a VTEP to accept and decapsulate VXLAN packets.

To configure the MX4 router:

1. Specify the IPv4, IPv6, and ISO addresses for the loopback interface.
 

```
[edit groups global interfaces]
user@MX4# set lo0 unit 0 family inet address 127.0.0.1/32
user@MX4# set lo0 unit 0 family inet address 10.255.181.43/32 primary
```

```
user@MX4# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5518.1043.00
user@MX4# set lo0 unit 0 family inet6 address abcd::10:255:181:43/128 primary
```

2. Configure the interfaces.

```
[edit interfaces]
user@MX4# set xe-0/0/0 vlan-tagging
user@MX4# set xe-0/0/0 encapsulation flexible-ethernet-services
user@MX4# set xe-0/0/0 unit 0 family bridge interface-mode trunk
user@MX4# set xe-0/0/0 unit 0 family bridge vlan-id-list 1-10
user@MX4# set xe-0/0/0 unit 1 family bridge interface-mode trunk
user@MX4# set xe-0/0/0 unit 1 family bridge vlan-id-list 11-15
user@MX4# set xe-0/0/0 unit 2 family bridge interface-mode trunk
user@MX4# set xe-0/0/0 unit 2 family bridge vlan-id-list 21-30
user@MX4# set xe-0/0/0 unit 3 family bridge interface-mode trunk
user@MX4# set xe-0/0/0 unit 3 family bridge vlan-id-list 31-40
user@MX4# set ge-0/2/6 unit 0 family inet address 30.30.30.1/30
user@MX4# set ge-0/3/0 unit 0 family inet address 3.3.3.2/30
user@MX4# set irb unit 1 family inet address 2.2.1.2/24
user@MX4# set irb unit 2 family inet address 2.2.2.2/24
user@MX4# set irb unit 3 family inet address 2.2.3.2/24
user@MX4# set irb unit 4 family inet address 2.2.4.2/24
user@MX4# set irb unit 5 family inet address 2.2.5.2/24
user@MX4# set irb unit 11 family inet address 2.2.11.2/24
user@MX4# set irb unit 12 family inet address 2.2.12.2/24
user@MX4# set irb unit 13 family inet address 2.2.13.2/24
user@MX4# set irb unit 14 family inet address 2.2.14.2/24
user@MX4# set irb unit 15 family inet address 2.2.15.2/24
```

3. Configure OSPF interface settings and the Layer 2 learning traceoption file.

```
[edit protocols]
user@MX4# set ospf area 0.0.0.0 interface ge-0/2/6.0
user@MX4# set ospf area 0.0.0.0 interface fxp0.0 disable
user@MX4# set ospf area 0.0.0.0 interface lo0.0 passive
user@MX4# set ospf area 0.0.0.0 interface ge-0/3/0.0
user@MX4# set l2-learning traceoptions file vxlan-l2ald.log size 100m files 10
user@MX4# set l2-learning traceoptions level all
user@MX4# set l2-learning traceoptions flag all
user@MX4# set layer2-control nonstop-bridging
```

4. Enable PIM.

```
[edit protocols]
user@MX4# set pim rp auto-rp discovery
user@MX4# set pim rp static address 10.255.181.13
user@MX4# set pim rp interface all mode sparse-dense
```

5. Set up OVSDB tracing operations.

```
[edit protocols]
user@MX4# set ovsdb traceoptions file ovsdb.log size 100m files 10
user@MX4# set ovsdb traceoptions level all
user@MX4# set ovsdb traceoptions flag all
```

6. Specify that interfaces xe-0/0/0.1 and xe-0/0/0.0 are managed by OVSDB.

```
[edit protocols]
```

- ```

user@MX4# set ovsdb interfaces xe-0/0/0.1
user@MX4# set ovsdb interfaces xe-0/0/0.0

```
7. Configure a connection with an NSX controller.
 

```

[edit protocols]
user@MX4# set ovsdb controller 192.168.182.45 protocol ssl port 6632

```
  8. Configure the VPLS interface.
 

```

[edit routing-instances]
user@MX4# set default-vs1 vtep-source-interface lo0.0
user@MX4# set default-vs1 instance-type virtual-switch
user@MX4# set default-vs1 interface xe-0/0/0.1

```
  9. Configure the **default-VS1** instance with a set of bridge domains that are associated with VXLAN.
 

```

[edit routing-instances]
user@MX4# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vlan-id 11
user@MX4# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab routing-interface irb.11
user@MX4# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ovsdb-managed
user@MX4# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan vni 16777214
user@MX4# set default-VS1 bridge-domains
bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab vxlan ingress-node-replication
user@MX4# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vlan-id 12
user@MX4# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 routing-interface irb.12
user@MX4# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ovsdb-managed
user@MX4# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan vni 12
user@MX4# set default-VS1 bridge-domains
f293dd5b-a901-4dba-bcbf-18a9979cf9d3 vxlan ingress-node-replication
user@MX4# set default-VS1 bridge-domains v13 vlan-id 13
user@MX4# set default-VS1 bridge-domains v13 routing-interface irb.13
user@MX4# set default-VS1 bridge-domains v13 vxlan vni 13
user@MX4# set default-VS1 bridge-domains v13 vxlan multicast-group 228.1.1.13
user@MX4# set default-VS1 bridge-domains v14 vlan-id 14
user@MX4# set default-VS1 bridge-domains v14 routing-interface irb.14
user@MX4# set default-VS1 bridge-domains v14 vxlan vni 14
user@MX4# set default-VS1 bridge-domains v14 vxlan multicast-group 228.1.1.14
user@MX4# set default-VS1 bridge-domains v15 vlan-id 15
user@MX4# set default-VS1 bridge-domains v15 routing-interface irb.15
user@MX4# set default-VS1 bridge-domains v15 vxlan vni 15
user@MX4# set default-VS1 bridge-domains v15 vxlan multicast-group 228.1.1.15

```
  10. Configure the **default-VS2** instance with a set of bridge domains that are associated with VXLAN.
 

```

[edit routing-instances]
user@MX4# set default-VS2 bridge-domains v21 vlan-id 21
user@MX4# set default-VS2 bridge-domains v21 vxlan vni 21

```

```

user@MX4# set default-VS2 bridge-domains v21 vxlan multicast-group 228.1.1.21
user@MX4# set default-VS2 bridge-domains v22 vlan-id 22
user@MX4# set default-VS2 bridge-domains v22 vxlan vni 22
user@MX4# set default-VS2 bridge-domains v22 vxlan multicast-group 228.1.1.22
user@MX4# set default-VS2 bridge-domains v23 vlan-id 23
user@MX4# set default-VS2 bridge-domains v23 vxlan vni 23
user@MX4# set default-VS2 bridge-domains v23 vxlan multicast-group 228.1.1.23
user@MX4# set default-VS2 bridge-domains v24 vlan-id 24
user@MX4# set default-VS2 bridge-domains v24 vxlan vni 24
user@MX4# set default-VS2 bridge-domains v24 vxlan multicast-group 228.1.1.24
user@MX4# set default-VS2 bridge-domains v25 vlan-id 25
user@MX4# set default-VS2 bridge-domains v25 vxlan vni 25
user@MX4# set default-VS2 bridge-domains v25 vxlan multicast-group 228.1.1.25

```

11. Configure a set of VXLAN-enabled bridge domains.

```

[edit bridge-domains]
user@MX4# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vlan-id 3
user@MX4# set 24a76aff-7e61-4520-a78d-3eca26ad7510 routing-interface irb.3
user@MX4# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan ovsdb-managed
user@MX4# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan vni 3
user@MX4# set 24a76aff-7e61-4520-a78d-3eca26ad7510 vxlan
    ingress-node-replication
user@MX4# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vlan-id 2
user@MX4# set cadbc185-f60f-48a6-93fd-dc14a6420c60 routing-interface irb.2
user@MX4# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan ovsdb-managed
user@MX4# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan vni 2
user@MX4# set cadbc185-f60f-48a6-93fd-dc14a6420c60 vxlan
    ingress-node-replication
user@MX4# set v1 vlan-id 1
user@MX4# set v1 routing-interface irb.1
user@MX4# set v1 vxlan vni 1
user@MX4# set v1 vxlan multicast-group 228.1.1.1
user@MX4# set v4 vlan-id 4
user@MX4# set v4 routing-interface irb.4
user@MX4# set v4 vxlan vni 4
user@MX4# set v4 vxlan multicast-group 228.1.1.4
user@MX4# set v5 vlan-id 5
user@MX4# set v5 routing-interface irb.5
user@MX4# set v5 vxlan vni 5
user@MX4# set v5 vxlan multicast-group 228.1.1.5

```

12. Configure the loopback interface to be used as the tunnel source address.

```

[edit switch-options]
user@MX4# set vtep-source-interface lo0.0

```



**NOTE:** After completing this configuration, you must configure a gateway, which is the NSX equivalent of a hardware VTEP. This configuration implements one hardware VTEP, so you must configure one gateway, a gateway service, and a logical switch port using NSX Manager or the NSX API. For more information about the tasks you must perform as well as key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints” on page 1986](#).



**Results**

From configuration mode, confirm your configuration by entering the following commands on each router. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

**Verification**

---

Confirm that the configuration is working properly.

- [Verifying MX1 on page 1967](#)
- [Verifying MX2 on page 1970](#)
- [Verifying MX3 on page 1975](#)
- [Verifying MX4 on page 1977](#)

**Verifying MX1**

**Purpose** Verify your configuration on MX1.

**Action** Verify that the interfaces are configured properly.

```
user@MX1# show interface
```

```
apply-groups LAG-options;
xe-0/0/2 {
  mtu 9000;
  unit 0 {
    family inet {
      address 80.80.0.250/24;
    }
  }
}
xe-0/0/3 {
  unit 0 {
    family inet {
      address 30.30.30.6/30;
    }
  }
}
ge-1/0/5 {
  mtu 1600;
  unit 0 {
    family inet {
      address 20.20.20.2/30;
    }
  }
}
ge-1/0/6 {
  unit 0 {
    family inet {
      address 20.20.20.10/30;
    }
  }
}
ge-1/0/7 {
  gigether-options {
    802.3ad ae2;
  }
}
ge-1/0/8 {
  gigether-options {
    802.3ad ae2;
  }
}
ge-1/1/2 {
  unit 0 {
    family inet {
      address 30.30.30.2/30;
    }
  }
}
ae2 {
  mtu 1600;
  unit 0 {
```

```

        family inet {
            address 20.20.20.6/30;
        }
    }
}

```

Verify the loopback addresses.

**user@MX1# show groups global interfaces**

```

lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.181.13/32 {
        primary;
      }
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.5518.1013.00;
    }
    family inet6 {
      address abcd::10:255:181:13/128 {
        primary;
      }
    }
  }
}

```

Verify that OSPF and PIM are configured correctly.

**user@MX1# show protocols**

```

ospf {
  area 0.0.0.0 {
    interface xe-0/0/2.0;
    interface ge-1/0/5.0;
    interface lo0.0 {
      passive;
    }
    interface fxp0.0 {
      disable;
    }
    interface ae2.0;
    interface ge-1/0/6.0;
    interface all;
  }
}
pim {
  rp {
    local {
      address 10.255.181.13;
    }
  }
  interface all {
    mode sparse-dense;
  }
}

```

}

### *Verifying MX2*

**Purpose** Verify your configuration on MX2.

**Action** Verify that the interfaces are configured properly.

```
user@MX2# show interfaces
```

```

xe-1/2/0 {
  gigeether-options {
    802.3ad ae1;
  }
}
xe-5/0/0 {
  gigeether-options {
    802.3ad ae1;
  }
}
ge-7/0/0 {
  gigeether-options {
    802.3ad ae2;
  }
}
ge-7/0/1 {
  mtu 1600;
  unit 0 {
    family inet {
      address 20.20.20.1/30;
    }
  }
}
ge-7/1/3 {
  gigeether-options {
    802.3ad ae2;
  }
}
xe-10/3/0 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 100 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 100-101;
    }
  }
  unit 102 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 102-103;
    }
  }
}
ae1 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 100 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 100-101;
    }
  }
}

```

```
    }  
  }  
  unit 102 {  
    family bridge {  
      interface-mode trunk;  
      vlan-id-list 102-103;  
    }  
  }  
}  
ae2 {  
  mtu 1600;  
  unit 0 {  
    family inet {  
      address 20.20.20.5/30;  
    }  
  }  
}  
irb {  
  unit 1 {  
    family inet {  
      address 2.2.1.1/24;  
    }  
  }  
  unit 2 {  
    family inet {  
      address 2.2.2.1/24;  
    }  
  }  
  unit 3 {  
    family inet {  
      address 2.2.3.1/24;  
    }  
  }  
  unit 4 {  
    family inet {  
      address 2.2.4.1/24;  
    }  
  }  
  unit 5 {  
    family inet {  
      address 2.2.5.1/24;  
    }  
  }  
  unit 11 {  
    family inet {  
      address 2.2.11.1/24;  
    }  
  }  
  unit 12 {  
    family inet {  
      address 2.2.12.1/24;  
    }  
  }  
  unit 13 {  
    family inet {  
      address 2.2.13.1/24;  
    }  
  }  
}
```

```

    }
  }
  unit 14 {
    family inet {
      address 2.2.14.1/24;
    }
  }
  unit 15 {
    family inet {
      address 2.2.15.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 200.1.1.1/32;
    }
  }
}
}

```

Verify that OSPF is configured properly.

```
user@MX2# show protocols ospf
```

```

area 0.0.0.0 {
  interface ge-7/0/1.0;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    passive;
  }
}

```

Verify that Layer 2 learning is configured properly.

```
user@MX2# show protocols l2-learning
```

```

l2-learning {
  traceoptions {
    file vxlan-l2ald.log size 100m files 10;
    level all;
    flag all;
  }
}

```

Verify that Layer 2 control is configured properly.

```
user@MX2# show protocols layer2-control
```

```

layer2-control {
  nonstop-bridging;
}

```

Verify that OVSDb is configured properly.

```
user@MX2# show protocols ovssdb
```

```

ovssdb {

```

```
    traceoptions {
      file ovssdb.log size 100m files 10;
      level all;
      flag all;
    }
    interfaces {
      xe-10/3/0.1;
      xe-10/3/0.0;
      ae1.0;
      ae1.1;
    }
    controller 192.168.182.45 {
      protocol {
        ssl port 6632;
      }
    }
  }
}
```

Verify the **default-VS1** routing instance configuration.

user@MX2# show routing-instances

```
default-VS1 {
  vtep-source-interface lo0.0;
  instance-type virtual-switch;
  interface xe-10/3/0.102;
  interface ae1.102;
  bridge-domains {
    bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab {
      vlan-id 102;
      routing-interface irb.102;
      vxlan {
        ovssdb-managed;
        vni 102;
        ingress-node-replication;
      }
    }
    f293dd5b-a901-4dba-bcbf-18a9979cf9d3 {
      vlan-id 103;
      routing-interface irb.103;
      vxlan {
        ovssdb-managed;
        vni 103;
        ingress-node-replication;
      }
    }
  }
}
```

Verify the **vrf1** routing instance configuration.

user@MX2# show routing-instances

```
vrf1 {
  instance-type vrf;
  interface ae2.0;
  interface lo0.1;
  route-distinguisher 100:100;
```



```

    vrf-target target:100:100;
  protocols {
    ospf {
      area 0.0.0.0 {
        interface ae2.0;
        interface lo0.0 {
          passive;
        }
      }
    }
    pim {
      rp {
        static {
          address 10.255.181.13;
        }
      }
      interface all ;
    }
  }
}

```

Verify the bridge domains configuration.

```
user@MX2# show bridge-domains
```

```

24a76aff-7e61-4520-a78d-3eca26ad7510 {
  vlan-id 100;
  routing-interface irb.100;
  vxlan {
    ovsdb-managed;
    vni 100;
    ingress-node-replication;
  }
}
cadbc185-f60f-48a6-93fd-dc14a6420c60 {
  vlan-id 101;
  routing-interface irb.101;
  vxlan {
    ovsdb-managed;
    vni 101;
    ingress-node-replication;
  }
}

```

Verify that the loopback interface is used as the tunnel source address.

```

user@MX2# show switch-options
  vtep-source-interface lo0.0;

```

### Verifying MX3

**Purpose** Verify your configuration on MX3.

**Action** Verify that the global group interfaces are configured properly.

```
user@MX3# show groups global interfaces
```

```
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.181.98/32 {
        primary;
      }
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.5518.1098.00;
    }
    family inet6 {
      address abcd::10:255:181:98/128 {
        primary;
      }
    }
  }
}
```

Verify that the interfaces are configured properly.

```
user@MX3# show interfaces
```

```
xe-0/0/0 {
  gigether-options {
    802.3ad ae1;
  }
}
xe-0/0/1 {
  gigether-options {
    802.3ad ae1;
  }
}
xe-0/0/3 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-40;
    }
  }
}
ae1 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-40;
    }
  }
}
```

```
}
```

Verify that the PIM is configured properly.

```
user@MX3# show protocols pim
  dense-groups {
    224.0.1.39/32;
    224.0.1.40/32;
  }
  rp {
    auto-rp discovery;
  }
  interface all {
    mode sparse-dense;
  }
```

Verify the VPLS bridge domain and interfaces configuration.

```
user@MX3# show protocols pim
  instance-type virtual-switch;
  interface xe-0/0/3.0;
  interface ae1.0;
  bridge-domains {
    v1 {
      vlan-id-list 1-40;
    }
  }
```

#### ***Verifying MX4***

**Purpose** Verify your configuration on MX4.

**Action** Verify that the global group interfaces are configured properly.

```
user@MX4# show groups global interfaces
```

```
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.181.43/32 {
        primary;
      }
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.5518.1043.00;
    }
    family inet6 {
      address abcd::10:255:181:43/128 {
        primary;
      }
    }
  }
}
```

Verify that the interfaces are configured properly.

```
user@MX4# show interfaces
```

```
xe-0/0/0 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-10;
    }
  }
  unit 1 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 11-15;
    }
  }
  unit 2 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 21-30;
    }
  }
  unit 3 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 31-40;
    }
  }
}
ge-0/2/6 {
  unit 0 {
    family inet {
      address 30.30.30.1/30;
    }
  }
}
```

```
    }  
  }  
  ge-0/3/0 {  
    unit 0 {  
      family inet {  
        address 3.3.3.2/30;  
      }  
    }  
  }  
  irb {  
    unit 1 {  
      family inet {  
        address 2.2.1.2/24;  
      }  
    }  
    unit 2 {  
      family inet {  
        address 2.2.2.2/24;  
      }  
    }  
    unit 3 {  
      family inet {  
        address 2.2.3.2/24;  
      }  
    }  
    unit 4 {  
      family inet {  
        address 2.2.4.2/24;  
      }  
    }  
    unit 5 {  
      family inet {  
        address 2.2.5.2/24;  
      }  
    }  
    unit 11 {  
      family inet {  
        address 2.2.11.2/24;  
      }  
    }  
    unit 12 {  
      family inet {  
        address 2.2.12.2/24;  
      }  
    }  
    unit 13 {  
      family inet {  
        address 2.2.13.2/24;  
      }  
    }  
    unit 14 {  
      family inet {  
        address 2.2.14.2/24;  
      }  
    }  
    unit 15 {  
      family inet {  
        address 2.2.15.2/24;  
      }  
    }  
  }  
}
```

Verify that the OSPF interface settings are configured properly.

```
user@MX4# show protocols ospf
  area 0.0.0.0 {
    interface ge-0/2/6.0;
    interface fxp0.0 {
      disable;
    }
    interface lo0.0 {
      passive;
    }
    interface ge-0/3/0.0;
  }
```

Verify that the Layer 2 learning traceoption file is configured properly.

```
user@MX4# show protocols l2-learning
  traceoptions {
    file vxlan-l2ald.log size 100m files 10;
    level all;
    flag all;
  }
```

Verify that the PIM protocol is configured properly.

```
user@MX4# show protocols PIM
  rp {
    auto-rp discovery;
    static {
      address 10.255.181.13;
    }
  }
  interface all {
    mode sparse-dense;
  }
```

Verify that Layer 2 control is configured properly.

```
user@MX4# show protocols layer2-control
  nonstop-bridging;
```

Verify that OVSDB is configured properly.

```
user@MX4# show protocols ovsdb
  traceoptions {
    file ovsdb.log size 100m files 10;
    level all;
    flag all;
  }
  interfaces {
    xe-0/0/0.1;
    xe-0/0/0.0;
  }
  controller 192.168.182.45 {
    protocol {
      ssl port 6632;
    }
  }
```

Verify the **default-VS1** routing instance configuration and bridge domains.

```
user@MX4# show routing-instances default-VS1
vtep-source-interface lo0.0;
instance-type virtual-switch;
interface xe-0/0/0.1;
bridge-domains {
  bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab {
    vlan-id 11;
    routing-interface irb.11;
    vxlan {
      ovsdb-managed;
      vni 16777214;
      ingress-node-replication;
    }
  }
  f293dd5b-a901-4dba-bcbf-18a9979cf9d3 {
    vlan-id 12;
    routing-interface irb.12;
    vxlan {
      ovsdb-managed;
      vni 12;
      ingress-node-replication;
    }
  }
  v13 {
    vlan-id 13;
    routing-interface irb.13;
    vxlan {
      vni 13;
      multicast-group 228.1.1.13;
    }
  }
  v14 {
    vlan-id 14;
    routing-interface irb.14;
    vxlan {
      vni 14;
      multicast-group 228.1.1.14;
    }
  }
  v15 {
    vlan-id 15;
    routing-interface irb.15;
    vxlan {
      vni 15;
      multicast-group 228.1.1.15;
    }
  }
}
```

Verify the **default-VS2** routing instance configuration and bridge domains.

```
user@MX4# show routing-instances default-VS2
bridge-domains {
  v21 {
    vlan-id 21;
    vxlan {
      vni 21;
      multicast-group 228.1.1.21;
    }
  }
}
```

```
    }  
  }  
  v22 {  
    vlan-id 22;  
    vxlan {  
      vni 22;  
      multicast-group 228.1.1.22;  
    }  
  }  
  v23 {  
    vlan-id 23;  
    vxlan {  
      vni 23;  
      multicast-group 228.1.1.23;  
    }  
  }  
  v24 {  
    vlan-id 24;  
    vxlan {  
      vni 24;  
      multicast-group 228.1.1.24;  
    }  
  }  
  v25 {  
    vlan-id 25;  
    vxlan {  
      vni 25;  
      multicast-group 228.1.1.25;  
    }  
  }  
}
```

Verify that the bridge domains are configured properly.

```
user@MX4# show bridge-domains  
24a76aff-7e61-4520-a78d-3eca26ad7510 {  
  vlan-id 3;  
  routing-interface irb.3;  
  vxlan {  
    ovsdb-managed;  
    vni 3;  
    ingress-node-replication;  
  }  
}  
cadbc185-f60f-48a6-93fd-dc14a6420c60 {  
  vlan-id 2;  
  routing-interface irb.2;  
  vxlan {  
    ovsdb-managed;  
    vni 2;  
    ingress-node-replication;  
  }  
}  
v1 {  
  vlan-id 1;  
  routing-interface irb.1;  
  vxlan {  
    vni 1;  
    multicast-group 228.1.1.1;  
  }  
}
```



```

v4 {
  vlan-id 4;
  routing-interface irb.4;
  vxlan {
    vni 4;
    multicast-group 228.1.1.4;
  }
}
v5 {
  vlan-id 5;
  routing-interface irb.5;
  vxlan {
    vni 5;
    multicast-group 228.1.1.5;
  }
}

```

Verify that the loopback interface is used as the tunnel source address.

```

user@MX4# show switch-options
vtep-source-interface lo0.0;

```

#### Related Documentation

- [Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment on page 1905](#)
- [Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers on page 1984](#)

## Configuration Tasks

- [Installing Open vSwitch Database Components on Juniper Networks Devices on page 1983](#)
- [Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers on page 1984](#)
- [Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985](#)
- [VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints on page 1986](#)
- [Configuring OVSDb-Managed VXLANs on page 1989](#)
- [Configuring VXLANs on a QFX5100 Switch on page 1991](#)

### Installing Open vSwitch Database Components on Juniper Networks Devices

To install Open vSwitch Database (OVSDb) components on a Juniper Networks device, you must copy the OVSDb software package to the Juniper Networks device and then install the package. The OVSDb software package name uses the following format:

`jsdn-packageID-release`

where:

- *packageID* identifies the package that should run on each Juniper Networks device.

- *release* identifies the OVSDB release; for example, 14.1R2. The OVSDB software release and the Junos OS release running on the device must be the same.

For information about OVSDB support on Juniper Networks devices and the software package for each device, see [“Open vSwitch Database Support on Juniper Networks Devices” on page 1899](#).

To install the OVSDB software package on a Juniper Networks device:

1. Download the software package to the Juniper Networks device.
2. If an OVSDB software package already exists on the Juniper Networks device, remove the package by issuing the **request system software delete** operational mode command.

```
user@device> request system software delete existing-ovsdb-package
```

3. Install the new software package by using the **request system software add** operational mode command.

```
user@device> request system software add path-to-ovsdb-package
```

**Related Documentation** • [Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices on page 1902](#)

## Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers

To secure a connection between a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol and one or more VMware NSX controllers, the following Secure Sockets Layer (SSL) files must be present in the **/var/db/certs** directory on the device:

- vtep-privkey.pem
- vtep-cert.pem
- ca-cert.pem

You must create the vtep-privkey.pem and vtep-cert.pem files for the device, and then install the two files in the **/var/db/certs** directory on the device.

Upon the initial connection between a Juniper Networks device with OVSDB implemented and an NSX controller, the ca-cert.pem file is automatically generated, and then installed in the **/var/db/certs** directory on the device.

The procedure provided in this topic uses the OpenFlow public key infrastructure (PKI) management utility ovs-pki on a Linux computer to initialize a public key infrastructure (PKI) and create the vtep-privkey.pem and vtep-cert.pem files. (If you have an existing PKI on your Linux computer, you can skip the step to initialize a new one.) By default, the utility initializes the PKI and places these files in the **/usr/local/share/openvswitch/pki** directory of the Linux computer.

To create and install an SSL key and certificate on a Juniper Networks device:

1. Initialize a PKI if one does not already exist on your Linux computer.

```
# ovs-pki init
```

2. On the same Linux computer on which the PKI exists, create a new key and certificate for the Juniper Networks device.

```
# ovs-pki req+sign vtep
```

3. Copy only the vtep-privkey.pem and vtep-cert.pem files from the Linux computer to the `/var/db/certs` directory on the Juniper Networks device.

#### Related Documentation

- [Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903](#)
- [Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985](#)

## Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices

To implement the Open vSwitch Database (OVSDB) management protocol on a Juniper Networks device, you must explicitly configure a connection to at least one VMware NSX controller, using the Junos OS CLI.

All NSX controller connections are made on the management interface (fxp0, em0, or em1) of the Juniper Networks device. This connection is secured by using the Secure Sockets Layer (SSL) protocol. The default port number over which the connection is made is 6632.

You must also specify that any interface with a physical server is managed by OVSDB. By performing this configuration, you are essentially disabling the Juniper Networks device from learning about other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs) and the MAC addresses learned by the hardware VTEPs, and enabling OVSDB to learn about these elements.

Before setting up OVSDB on a Juniper Networks device, you must do the following:

- Ensure that the Juniper Networks device has an OVSDB software package installed, and that the OVSDB software package release is the same as the Junos OS release running on the device.
- Determine the IP address of the NSX controller.
- Create an SSL private key and certificate, and install them in the `/var/db/certs` directory of the Juniper Networks device. For more information, see [“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers” on page 1984](#).

To set up OVSDb on a Juniper Networks device:

1. Specify the IP address of the NSX controller.

```
[edit protocols ovssdb]
user@host# set controller ip-address
```

2. Specify SSL as the protocol that secures the connection.

```
[edit protocols ovssdb controller ip-address]
user@host# set protocol ssl
```

3. Set the number of the port over which the connection to the NSX controller is made.

```
[edit protocols ovssdb controller ip-address protocol ssl]
user@host# set port number
```

4. (Optional) Specify (in milliseconds) how long the connection can be inactive before an inactivity probe is sent.

```
[edit protocols ovssdb controller ip-address]
user@host# set inactivity-probe-duration milliseconds
```

5. (Optional) Specify (in milliseconds) how long the device must wait before it can try to connect to the NSX controller again if the previous attempt failed.

```
[edit protocols ovssdb controller ip-address]
user@host# set maximum-backoff-duration milliseconds
```

6. (Optional) Repeat steps 1 through 5 to explicitly configure a connection to an additional NSX controller in the same cluster.

7. Specify the interfaces that you want OVSDb to manage.

```
[edit protocols ovssdb]
user@host# set interfaces interface-name unit logical-unit-number
```



**NOTE:** After completing this procedure, you must set up OVSDb-managed VXLANs. For more information, see [“Configuring OVSDb-Managed VXLANs” on page 1989](#).

#### Related Documentation

- [Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903](#)

## VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints

For each Juniper Networks Junos operating system (Junos OS) network device that you plan to deploy as a hardware virtual tunnel endpoint (VTEP) in a physical network, you must create a VMware NSX-equivalent entity, which is known as a *gateway*, in NSX Manager version 4.0.3 or in the NSX API. You must also map the gateway to a logical switch, which is the NSX equivalent of an Open vSwitch Database (OVSDb)-managed Virtual Extensible LAN (VXLAN) in the physical network. Performing this configuration enables connectivity between physical servers in the physical network and virtual machines (VMs) in the virtual network.

This topic provides a high-level summary of the tasks that you must perform to create a gateway. Although you can create a gateway either in NSX Manager or in the NSX API, this topic describes the necessary tasks from the perspective of NSX Manager. Also, this

topic does not include a complete procedure for each task. Rather, it includes key NSX Manager configuration details for ensuring the correct configuration of the virtual entities so that they function properly with the physical entities. For complete information about performing the tasks described in this topic, see the documentation that accompanies NSX Manager.

For each hardware VTEP that you deploy in the physical network, you must perform the following tasks:

- [Creating a Gateway on page 1987](#)
- [Creating a Gateway Service on page 1987](#)
- [Creating a Logical Switch Port on page 1988](#)

### Creating a Gateway

In NSX Manager, you must create a gateway for each hardware VTEP that you implement in the physical network. [Table 135 on page 1987](#) provides a summary of key configuration fields in NSX Manager and how to configure them when creating a gateway.

Before you begin this task, you must configure a logical switch in NSX Manager or in the NSX API for each OVSDB-managed VXLAN that you plan to implement in the physical network. For information about configuring a logical switch, see the documentation that accompanies NSX Manager or the NSX API.

**Table 135: Create a Gateway in NSX Manager: Key Configurations**

| NSX Manager Configuration Page/Dialog Box | NSX Manager Configuration Field | How to Configure   |
|---|---------------------------------|--|
| Type                                      | Transport Node Type             | Select <b>Gateway</b> .  |
| Properties                                | VTEP Enabled                    | Select <b>VTEP Enabled</b> .   |
| Credential                                | Type                            | Select <b>Management Address</b> .   |
| Credential                                | Management Address              | Specify the management IP address of the Juniper Networks device.                      |
| Connections/Create Transport Connector    | Transport Type                  | Select <b>VXLAN</b> .  |
| Connections/Create Transport Connector    | Transport Zone UUID             | Select the UUID of an existing transport zone, or create a new transport zone.         |
| Connections/Create Transport Connector    | IP Address                      | Specify the IP address of the loopback interface (lo0) of the Juniper Networks device. |

### Creating a Gateway Service

In NSX Manager, you must create a gateway service for each hardware VTEP that you implement in the physical network. [Table 136 on page 1988](#) provides a summary of key

configuration fields in NSX Manager and how to configure them when creating a gateway service.

Before you start this task, make sure that you have configured the OVSDDB-managed interfaces on the hardware VTEP. For information about configuring OVSDDB-managed interfaces, see [“Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices” on page 1985](#).

**Table 136: Create a Gateway Service in NSX Manager: Key Configurations**

| NSX Manager Configuration Page/Dialog Box | NSX Manager Configuration Field | How to Configure  |
|---|---------------------------------|---|
| Type                                      | Gateway Service Type            | Select <b>VTEP L2 Gateway Service</b> .                             |
| Transport Nodes/Edit Gateway              | Transport Node                  | Select the gateway that you created for the hardware VTEP.          |
| Transport Nodes/Edit Gateway              | Port ID                         | Select an OVSDDB-managed interface configured on the hardware VTEP. |

#### Creating a Logical Switch Port

In NSX Manager, you must create a logical switch port for each hardware VTEP that you implement in the physical network. [Table 137 on page 1988](#) provides a summary of key configuration fields in NSX Manager and how to configure them when creating a logical switch port.

Before you start this task, you must configure a logical switch in NSX Manager or in the NSX API for each OVSDDB-managed VXLAN that you plan to implement in the physical network. For information about configuring a logical switch, see the documentation that accompanies NSX Manager or the NSX API.

**Table 137: Create a Logical Switch Port in NSX Manager: Key Configurations**

| NSX Manager Configuration Page/Dialog Box | NSX Manager Configuration Field | How to Configure   |
|---|---------------------------------|--|
| Logical Switch                            | Logical Switch UUID             | Select the UUID of the logical switch that corresponds to the hardware VTEP. |
| Attachment                                | Attachment Type                 | Select <b>VTEP L2 Gateway</b> .  |
| Attachment                                | VTEP L2 Gateway Service UUID    | Select the UUID of the gateway service you created for the hardware VTEP.    |

**Related Documentation**

- [Configuring OVSDDB-Managed VXLANs on page 1989](#)

## Configuring OVSDB-Managed VXLANs



**NOTE:** This topic does not apply to QFX5100 switches that support Open vSwitch Database (OVSDB) and Virtual Extensible LAN (VXLAN). The QFX5100 switch automatically creates OVSDB-managed VXLANs, thereby eliminating the need for you to configure them, using the Junos OS CLI. In addition, the QFX5100 switch does not support ingress node replication, which is described in this topic. However, there are other configuration tasks that must be performed to set up OVSDB on a QFX5100 switch. For more information, see [“Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment”](#) on page 1905.

To implement the OVSDB management protocol on a Juniper Networks device, you must configure OVSDB-managed VXLANs.

For Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic that originates in an OVSDB-managed VXLAN and is forwarded to interfaces within the same VXLAN, you can optionally enable ingress node replication. With this feature enabled, the Juniper Networks device handles the replication of these packets and the forwarding of the replicas to interfaces within the same OVSDB-managed VXLAN. For more information about using ingress node replication or a service node, which is the default way to handle Layer 2 BUM traffic, see [“Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDB”](#) on page 1904.



**NOTE:** When Juniper Networks devices replicate Layer 2 BUM packets to a large number of remote software virtual tunnel endpoints (VTEPs), the performance of the Juniper Networks devices can be impacted.

Before you configure VXLANs on a Juniper Networks device, using the Junos OS CLI:

- For each OVSDB-managed VXLAN that you plan to configure on a Juniper Networks device, you must configure a logical switch in VMware NSX Manager version 4.0.3 or the NSX API. (For information about configuring a logical switch, see the documentation that accompanies NSX Manager or the NSX API.) Based on the name and VXLAN network identifier (VNI) that you configure for the logical switch, NSX automatically generates a universally unique identifier (UUID) for the logical switch. You must retain the UUID of the logical switch for use when configuring a corresponding VXLAN on the Juniper Networks device as described in the following procedure.
- You must perform the configuration described in [“Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices”](#) on page 1985.

To configure an OVSDB-managed VXLAN on a Juniper Networks device:

1. Configure the VXLANs that you want OVSDB to manage. You can configure the VXLANs in the context of a bridge domain, routing instance, or switching instance.



**NOTE:** For the name of the bridge domain, you must specify the UUID for the logical switch configured in NSX Manager or the NSX API.

Bridge domains:

```
[edit bridge-domains bridge-domain-name vxlan]
user@host# set ovsdb-managed
```

Bridge domains within the specified routing instance:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name
vxlan]
user@host# set ovsdb-managed
```

VLANs within the specified routing instance:

```
[edit routing-instances routing-instance-name vlans vlan-name vxlan]
user@device# set ovsdb-managed
```

Default switching instance within the specified routing instance:

```
[edit routing-instances routing-instance-name switch-options]
user@host# set ovsdb-managed
```

All VXLAN entities within the specified routing instance:

```
[edit routing-instances routing-instance-name vxlan]
user@host# set ovsdb-managed
```

2. (Optional) Enable ingress node replication to handle Layer 2 BUM traffic on interfaces in the same VXLAN in which the traffic originated. You can configure ingress node replication in the context of a bridge domain or routing instance.

Bridge domains:

```
[edit bridge-domains bridge-domain-name vxlan]
user@host# set ingress-node-replication
```

Bridge domains or all VXLAN entities, respectively, within the specified routing instance:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name
vxlan]
[edit routing-instances routing-instance-name vlans vlan-name vxlan]
[edit routing-instances routing-instance-name vxlan]
user@host# set ingress-node-replication
```

3. For each Juniper Networks device that you plan to implement as a hardware VTEP, you must perform some configuration tasks in NSX Manager or in the NSX API.

For more information about the tasks you must perform and key NSX Manager configuration details, see [“VMware NSX Configuration for Juniper Networks Devices That Function as Virtual Tunnel Endpoints” on page 1986](#).

#### Related Documentation

- [Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment on page 1905](#)
- [Example: Setting Up a VXLAN Layer 2 Gateway and OVSDDB Connections Between Virtual and Physical Entities in a Data Center on page 1917](#)
- [Example: Setting Up Inter-VXLAN Routing and OVSDDB Connections in a Data Center on page 1925](#)



## Configuring VXLANs on a QFX5100 Switch

Follow these steps to configure a QFX5100 switch to act as a VTEP. (If the switch is acting as a transit Layer 3 switch for downstream VTEPs, you do not need to perform these steps. No special configuration is needed in this case.)

- [Configuring a Source IP Address on page 1991](#)
- [Configuring PIM for VXLANs on page 1991](#)
- [Configuring VXLANs on page 1991](#)

### Configuring a Source IP Address

On a switch that will act as a VTEP, you must configure an IP address that will be used as the source address in the outer IP header of the VXLAN packet. This is the VXLAN tunnel source address.

1. Create a reachable IPv4 address on the loopback interface and configure it to be used as the tunnel source address:

```
[edit]
user@switch# set interfaces lo0.0 unit 0 family inet address ip-address
[edit]
user@switch# set switch-options vtep-interface-source lo0.0
```

### Configuring PIM for VXLANs

If you are not using a controller to create a VXLAN control plane, you must enable PIM on the switch so that the VTEP can use multicast groups to establish reachability with other VTEPs and forward BUM traffic.

1. Enable PIM:

```
[edit]
user@switch# set protocols pim interface all
```

2. Configure the address of a PIM rendezvous point:

```
[edit]
user@switch# set protocols pim rp static address ip-address
```

### Configuring VXLANs

You configure VXLANs under the **vlan** stanza (which is why a QFX5100 switch supports 4K VLANs). You must also configure the server-facing interfaces to be VLAN members.

1. Create a VLAN to VXLAN mapping and assign a multicast group address to the VXLAN. All members of a VXLAN must use the same multicast group address:

```
[edit]
user@switch# set vlans name vlan-id ID vxlan vni ID multicast-group multicast-group-address
```

2. (Optional) Configure the switch to retain the original VLAN tag (in the inner Ethernet packet) after VXLAN encapsulation. By default, the original tag is dropped when the packet is encapsulated:

```
[edit]
user@switch# set vlans name vxlan encapsulate-inner-vlan
```

3. (Optional) Configure the switch to de-encapsulate and accept original VLAN tags in VXLAN packets. By default, the original tag is dropped when the packet is encapsulated:

```
[edit]
user@switch# set protocols l2-learning decapsulate-accept-inner-vlan
```

4. Configure server-facing interfaces to support multiple VLANs:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching interface-mode trunk
```

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching vlan members all
```

You must create a VLAN to VXLAN mapping for each VLAN that will need Layer 2 connectivity over the Layer 3 network.

**Related Documentation**

- [Understanding VXLANs on page 1912](#)

---

## OVSDB Configuration Statements

---

- [controller \(OVSDB\) on page 1993](#)
- [inactivity-probe-duration on page 1994](#)
- [ingress-node-replication on page 1995](#)
- [interfaces \(OVSDB\) on page 1996](#)
- [maximum-backoff-duration on page 1996](#)
- [ovsdb on page 1997](#)
- [ovsdb-managed on page 1998](#)
- [port \(OVSDB\) on page 1999](#)
- [protocol \(OVSDB\) on page 2000](#)
- [traceoptions \(OVSDB\) on page 2001](#)

## controller (OVSDB)


|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> controller <i>ip-address</i> {     <i>inactivity-probe-duration</i> <i>milliseconds</i>;     <i>maximum-backoff-duration</i> <i>milliseconds</i>;     protocol <i>protocol</i> {         port <i>number</i>;     } } </pre>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb</a> ]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>   |
| <b>Description</b>              | <p>Configure a connection between a Juniper Networks device and a VMware NSX controller. The Junos OS device must be running a release that supports the Open vSwitch Database (OVSDB) management protocol and have the OVSDB software package installed. The OVSDB software package release must be the same as the Junos OS release running on the device.</p> <p>The Junos OS implementation of OVSDB supports one cluster of NSX controllers, which includes three or five controllers as per VMware recommendations.</p> <p>To implement OVSDB on a Junos OS device, you must explicitly configure a connection to at least one NSX controller, using the Junos OS CLI. If the NSX controller to which you explicitly configure a connection is in a cluster, the controller pushes information about other controllers in the same cluster to the device, and the device establishes connections with the other controllers. However, you can also explicitly configure connections with the other controllers in the cluster, using the Junos OS CLI.</p> |
| <b>Options</b>                  | <p><b><i>ip-address</i></b> —IPv4 address of the NSX controller.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985</a></li> <li>• <a href="#">Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903</a></li> </ul>   |

## inactivity-probe-duration

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>inactivity-probe-duration <i>milliseconds</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb controller</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Configure the maximum amount of time, in milliseconds, that the connection between a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol and a VMware NSX controller can be inactive before an inactivity probe is sent.  |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Number of milliseconds that the connection can be inactive before an inactivity probe is sent.<br><b>Range:</b> 0 through 4,294,967,295<br><b>Default:</b> 0. This value indicates that an inactivity probe is never sent.  |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985</a></li><li>• <a href="#">Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903</a></li></ul> |

## ingress-node-replication

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ingress-node-replication;  |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> vxlan],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> vxlan],<br>[edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> vxlan],<br>[edit routing-instances <i>routing-instance-name</i> vxlan]<br>[edit vlans <i>vlan-name</i> vxlan]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.   |
| <b>Description</b>              | <p>Enable ingress node replication for a specified Virtual Extensible LAN (VXLAN) that is managed by the Open vSwitch Database (OVSDB) management protocol.</p> <p>With this feature enabled, instead of service nodes, Juniper Networks devices with OVSDB implemented handle incoming broadcast, unknown unicast, or multicast (BUM) traffic. For more information about the scenarios in which you can use ingress node replication and how it works, see <a href="#">“Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDB”</a> on page 1904.</p> |
|                                 | <div>  <p><b>NOTE:</b> When Juniper Networks devices replicate Layer 2 BUM packets to a large number of remote software VTEPs, the performance of the Juniper Networks devices can be impacted.</p> </div>   |
| <b>Default</b>                  | If you do not include the <b>ingress-node-replication</b> statement, one or more service nodes handle BUM traffic.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring OVSDB-Managed VXLANs on page 1989</a></li> </ul>  |

## interfaces (OVSDB)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>interfaces <i>interface-name</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Specify the interfaces to be managed by the Open vSwitch Database (OVSDB) management protocol. Typically, the only interfaces that need to be managed by OVSDB are interfaces with physical servers. |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface, including the logical unit number—for example, xe-1/1/0.0.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring OVSDB-Managed VXLANs on page 1989</a></li></ul>  |

## maximum-backoff-duration

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>maximum-backoff-duration <i>milliseconds</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb controller</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Specify (in milliseconds) how long a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol waits before it re-attempts to connect with a VMware NSX controller if a previous attempt failed.  |
| <b>Options</b>                  | <i>milliseconds</i> —Number of milliseconds a Juniper Networks device waits before it re-attempts to connect with an NSX controller.<br><b>Range:</b> 1000 through 4,294,967,295<br><b>Default:</b> 1000  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985</a></li><li>• <a href="#">Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903</a></li></ul> |

## ovsdb

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> ovsdb {   controller <i>ip-address</i> {     inactivity-probe-duration <i>milliseconds</i>;     maximum-backoff-duration <i>milliseconds</i>;     protocol <i>protocol</i> {       port <i>number</i>;     }   }   interfaces <i>interface-name</i>;   traceoptions {     file &lt;<i>filename</i>&gt; &lt;<i>files number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;no-world-readable         world-readable&gt; &lt;<i>size size</i>&gt;;     flag <i>flag</i>;     no-remote-trace;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>  |
| <b>Description</b>              | <p>Configure support for the Open vSwitch Database (OVSDB) management protocol on a Juniper Networks device. The Juniper Networks device must be running a release that supports OVSDB and have the OVSDB software package installed. The OVSDB software package release must be the same as the Junos OS release that is running on the device.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Default</b>                  | The OVSDB management protocol is disabled on Juniper Networks devices.  |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding the Open vSwitch Database Management Protocol Running on Juniper Networks Devices on page 1902</a></li> <li>• <a href="#">Configuring OVSDB-Managed VXLANs on page 1989</a></li> </ul>   |

## ovsdb-managed

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ovsdb-managed;  |
| <b>Hierarchy Level</b>          | [edit bridge-domains <i>bridge-domain-name</i> vxlan],<br>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> vxlan],<br>[edit routing-instances <i>routing-instance-name</i> switch-options],<br>[edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> vxlan],<br>[edit routing-instances <i>routing-instance-name</i> vxlan],<br>[edit switch-options]<br>[edit vlans <i>vlan-name</i> vxlan]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | <p>Disable a Junos OS device from learning about other Junos OS devices that function as hardware virtual tunnel endpoints (VTEPs) in a specified Virtual Extensible LAN (VXLAN) and the MAC addresses learned by the hardware VTEPs. Instead, the Junos OS device uses the Open vSwitch Database (OVSDB) management protocol to learn about the hardware VTEPs in the VXLAN and the MAC addresses learned by the hardware VTEPs.</p> <p>The specified VXLAN must have a VXLAN Network Identifier (VNI) configured, using the <b>vni</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instance <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy.</p> <p>Also, this implementation of OVSDB uses the multicast scheme described in <a href="#">“Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDB” on page 1904</a>. Therefore, specifying the <b>multicast-group</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instances <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy has no effect.</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring OVSDB-Managed VXLANs on page 1989</a></li></ul>   |



---


## port (OVSDB)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>port number;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">ovsdb controller protocol</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Specify the VMware NSX controller port to which a Juniper Networks device that supports the Open vSwitch Database (OVSDb) management protocol connects.   |
| <b>Options</b>                  | <b>number</b> —Port number of NSX controller port.<br><b>Range:</b> 1024 through 65,535<br><b>Default:</b> 6632   |
| <b>Required Privilege Level</b> | <b>admin</b> —To view this statement in the configuration.<br><b>admin-control</b> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985</a></li><li>• <a href="#">Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903</a></li></ul> |

## protocol (OVSDB)

---

|  |   |
|--|---|
| <b>Syntax</b>  | <code>protocol protocol {<br/>    port number;<br/>}</code>   |
| <b>Hierarchy Level</b>   | [edit protocols <a href="#">ovsdb controller</a> ]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 14.1R2.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>   | <p>Configure the security protocol that protects the connection between a Juniper Networks device that supports the Open vSwitch Database (OVSDB) management protocol and a VMware NSX controller.</p> <p>The Secure Sockets Layer (SSL) connection requires a private key and certificates, which must be stored in the <code>/var/db/certs</code> directory of the Juniper Networks device. For more information about the files, including actions you must take to create and install some of the files, see <a href="#">“Creating and Installing an SSL Key and Certificate on a Juniper Networks Device for Connection with VMware NSX Controllers” on page 1984</a>.</p> |
| <b>Options</b>   | <b>protocol</b> —Establish a secure connection to the NSX controller, using SSL or the Transmission Control Protocol (TCP).   |
| <hr/> <div> <b>NOTE:</b> SSL is the only supported connection protocol.</div> <hr/> |   |
| <b>Default:</b>  | ssl   |
|  | The remaining statement is explained separately.  |
| <b>Required Privilege Level</b>  | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985</a></li><li>• <a href="#">Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903</a></li></ul>   |

## traceoptions (OVSDB)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;match regular-expression&gt; &lt;no-world-readable       world-readable&gt; &lt;size size&gt;;     flag flag;     no-remote-trace; } </pre>   |
| <b>Hierarchy Level</b>     | [edit protocols <a href="#">ovsdb</a> ]   |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>  |
| <b>Description</b>         | Define tracing operations for the Open vSwitch Database (OVSDB) management protocol, which is supported on Juniper Networks devices.  |
| <b>Default</b>             | If you do not include this statement, OVSDB-specific tracing operations are not performed.  |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of file in which the system places the output of the tracing operations. By default, the system places all files in the <code>/var/log</code> directory.</p> <p><b>Default:</b> <code>/var/log/vgd</code></p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file reaches the size specified by the <b>size</b> option, the filename is appended with 0 and compressed. For example, a trace file named <b>trace-file.gz</b> would be renamed <b>trace-file.0.gz</b>. When <b>trace-file.0.gz</b> reaches the specified size, it is renamed <b>trace-file.1.gz</b> and its contents are compressed to <b>trace-file.0.gz</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option and a filename.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. You can include one or more of the following flags:</p> <ul style="list-style-type: none"> <li><b>all</b>—All OVSDB events.</li> <li><b>configuration</b>—OVSDB configuration events.</li> <li><b>core</b>—OVSDB core events.</li> <li><b>function</b>—OVSDB function events.</li> <li><b>interface</b>—OVSDB interface events.</li> <li><b>l2-client</b>—OVSDB Layer 2 client events.</li> <li><b>ovs-client</b>—OVSDB client events.</li> </ul> |

**match *regular-expression***—(Optional) Only log lines that match the regular expression.

**no-remote-trace**—(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or passed through the Juniper Networks device.

**no-world-readable**—Restrict access to the trace files to the owner.

**Default:** no-world-readable

**size *size***—(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files by using the **files** option and a filename by using the **file** option.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB.

**Range:** 10,240 through 1,073,741,824 bytes

**Default:** 128 KB

**world-readable**—Enable any user to access the trace files.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up a VXLAN Layer 2 Gateway and OVSDb Connections Between Virtual and Physical Entities in a Data Center on page 1917</a></li><li>• <a href="#">Example: Setting Up Inter-VXLAN Routing and OVSDb Connections in a Data Center on page 1925</a></li></ul> |
|------------------------------|---|

---

## VXLAN Configuration Statements

---

- [decapsulate-accept-inner-vlan on page 2003](#)
- [encapsulate-inner-vlan on page 2003](#)
- [multicast-group on page 2004](#)
- [ovsdb-managed on page 2005](#)
- [unreachable-vtep-aging-timer on page 2006](#)
- [vni on page 2006](#)
- [vtep-source-interface on page 2007](#)
- [vxlan on page 2007](#)

## decapsulate-accept-inner-vlan

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | decapsulate-accept-inner-vlan  |
| <b>Hierarchy Level</b>          | [edit protocols l2-learning]   |
| <b>Release Information</b>      | Statement modified in Junos OS 14.1X53 for the QFX Series.   |
| <b>Description</b>              | Configure the switch to de-encapsulate and accept original VLAN tags in VXLAN packets.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding VXLANs on page 1912</a></li> <li>• <a href="#">encapsulate-inner-vlan on page 2003</a></li> </ul> |

## encapsulate-inner-vlan

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | encapsulate-inner-vlan   |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.  |
| <b>Description</b>              | Configure the switch to preserve the original VLAN tag (in the inner Ethernet packet) when performing VXLAN encapsulation.   |
| <b>Default</b>                  | The original tag is dropped when the packet is encapsulated.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding VXLANs on page 1912</a></li> <li>• <a href="#">Configuring VXLANs on a QFX5100 Switch on page 1991</a></li> <li>• <a href="#">Examples: Configuring VXLANs on QFX Series Switches on page 1944</a></li> <li>• <a href="#">decapsulate-accept-inner-vlan on page 2003</a></li> </ul> |

## multicast-group

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | multicast-group  |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.  |
| <b>Description</b>              | Assign a multicast group address to a VXLAN. All members of a VXLAN must use the same multicast group address  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 1912</a></li><li>• <a href="#">Configuring VXLANs on a QFX5100 Switch on page 1991</a></li><li>• <a href="#">Examples: Configuring VXLANs on QFX Series Switches on page 1944</a></li></ul> |

## ovsdb-managed

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ovsdb-managed;  |
| <b>Hierarchy Level</b>          | <p>[edit bridge-domains <i>bridge-domain-name</i> vxlan],<br/> [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> vxlan],<br/> [edit routing-instances <i>routing-instance-name</i> switch-options],<br/> [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i> vxlan],<br/> [edit routing-instances <i>routing-instance-name</i> vxlan],<br/> [edit switch-options]<br/> [edit vlans <i>vlan-name</i> vxlan]</p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 14.1R2.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>  |
| <b>Description</b>              | <p>Disable a Junos OS device from learning about other Junos OS devices that function as hardware virtual tunnel endpoints (VTEPs) in a specified Virtual Extensible LAN (VXLAN) and the MAC addresses learned by the hardware VTEPs. Instead, the Junos OS device uses the Open vSwitch Database (OVSDB) management protocol to learn about the hardware VTEPs in the VXLAN and the MAC addresses learned by the hardware VTEPs.</p> <p>The specified VXLAN must have a VXLAN Network Identifier (VNI) configured, using the <b>vni</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instance <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy.</p> <p>Also, this implementation of OVSDB uses the multicast scheme described in <a href="#">“Understanding How Layer 2 BUM Traffic and Layer 3 Routed Multicast Traffic Are Handled in VXLANs Managed by OVSDB” on page 1904</a>. Therefore, specifying the <b>multicast-group</b> statement in the [edit bridge-domains <i>bridge-domain-name</i> vxlan], [edit routing-instances <i>routing-instance-name</i> vxlan], or [edit vlans <i>vlan-name</i> vxlan] hierarchy has no effect.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring OVSDB-Managed VXLANs on page 1989</a></li> </ul>   |

## unreachable-vtep-aging-timer

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | unreachable-vtep-aging-timer [300–1800]  |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.  |
| <b>Description</b>              | Configure the system to age out the address for the remote VTEP if all the MAC addresses learned from that VTEP age out. The address for the remote VTEP expires the configured number of seconds after the last learned MAC address expires.                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 1912</a></li><li>• <a href="#">Configuring VXLANs on a QFX5100 Switch on page 1991</a></li><li>• <a href="#">Examples: Configuring VXLANs on QFX Series Switches on page 1944</a></li></ul> |

## vni

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | vni [1–16777214]   |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>VLAN</i> <a href="#">vxlan</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.  |
| <b>Description</b>              | Assign a numeric value to identify a VXLAN. All members of a VXLAN must use the same VNI.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding VXLANs on page 1912</a></li><li>• <a href="#">Configuring VXLANs on a QFX5100 Switch on page 1991</a></li><li>• <a href="#">Examples: Configuring VXLANs on QFX Series Switches on page 1944</a></li></ul> |



## vtep-source-interface

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>vtep-source-interface <i>logical-interface</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">switch-options</a> ,   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.  |
| <b>Description</b>              | Configure a source interface for a VXLAN tunnel. You must provide the name of a logical interface configured on the loopback interface.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding VXLANs on page 1912</a></li> <li>• <a href="#">Configuring VXLANs on a QFX5100 Switch on page 1991</a></li> <li>• <a href="#">Examples: Configuring VXLANs on QFX Series Switches on page 1944</a></li> </ul> |

## vxlan

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>vxlan {   encapsulate-inner-vlan   multicast-group   ovsdb-managed   unreachable-vtep-aging-timer   vni }</pre>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> ],   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10.  |
| <b>Description</b>              |  |
| <b>Options</b>                  | The remaining statements are explained separately.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding VXLANs on page 1912</a></li> <li>• <a href="#">Configuring VXLANs on a QFX5100 Switch on page 1991</a></li> <li>• <a href="#">Examples: Configuring VXLANs on QFX Series Switches on page 1944</a></li> </ul> |



## CHAPTER 20

# Administration

- [OVSDB Monitoring Commands on page 2009](#)
- [VXLAN Monitoring Commands on page 2033](#)

### OVSDB Monitoring Commands

---

- [show bridge mac-table](#)
- [show ovssdb controller](#)
- [show ovssdb interface](#)
- [show ovssdb logical-switch](#)
- [show ovssdb mac](#)
- [show ovssdb statistics interface](#)
- [show ovssdb virtual-tunnel-end-point](#)
- [show vpls mac-table](#)

## show bridge mac-table

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>show bridge mac-table &lt;brief   count   detail   extensive&gt; &lt;bridge-domain (all   <i>bridge-domain-name</i>)&gt; &lt;global-count&gt; &lt;interface <i>interface-name</i>&gt; &lt;mac-address&gt; &lt;vlan-id (all-vlan   <i>vlan-id</i>)&gt;</pre>   |
| Release Information      | Command introduced in Junos OS Release 8.4.  |
| Description              | (MX Series routers only) Display Layer 2 MAC address information.  |
| Options                  | <p><b>none</b>—Display all learned Layer 2 MAC address information.</p> <p><b>brief   count   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain (all   <i>bridge-domain-name</i>)</b>—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p><b>global-count</b>—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned Layer 2 MAC addresses for the specified interface.</p> <p><b>mac-address</b>—(Optional) Display the specified learned Layer 2 MAC address information.</p> <p><b>vlan-id (all-vlan   <i>vlan-id</i>)</b>—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p> |
| Additional Information   | When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.   |
| Required Privilege Level | view   |
| List of Sample Output    | <p><a href="#">show bridge mac-table on page 2011</a></p> <p><a href="#">show bridge mac-table (with VXLAN enabled) on page 2012</a></p> <p><a href="#">show bridge mac-table count on page 2012</a></p> <p><a href="#">show bridge mac-table detail on page 2013</a></p>  |
| Output Fields            | <p><a href="#">Table 138 on page 2011</a> describes the output fields for the <b>show bridge mac-table</b> command. Output fields are listed in the approximate order in which they appear.</p>  |

Table 138: show bridge mac-table Output fields

| Field Name                | Field Description   |
|---------------------------|---|
| <b>Routing instance</b>   | Name of the routing instance.   |
| <b>Bridging domain</b>    | Name of the bridging domain.  |
| <b>MAC address</b>        | MAC address or addresses learned on a logical interface.  |
| <b>MAC flags</b>          | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address is configured.</li> <li>• <b>D</b>—Dynamic MAC address is configured.</li> <li>• <b>L</b>—Locally learned MAC address is configured.</li> <li>• <b>C</b>—Control MAC address is configured.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Non-configured MAC.</li> <li>• <b>R</b>—Remote PE MAC address is configured.</li> </ul> |
| <b>Logical interface</b>  | Name of the logical interface.  |
| <b>MAC count</b>          | Number of MAC addresses learned on the specific routing instance or interface.  |
| <b>Learning interface</b> | Name of the logical interface on which the MAC address was learned.   |
| <b>Learning VLAN</b>      | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.  |
| <b>VXLAN ID/VXLAN</b>     | VXLAN Network Identifier (VNI)  |
| <b>Layer 2 flags</b>      | Debugging flags signifying that the MAC address is present in various lists.  |
| <b>Epoch</b>              | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.   |
| <b>Sequence number</b>    | Sequence number assigned to this MAC address. Used for debugging.   |
| <b>Learning mask</b>      | Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.   |
| <b>IPC generation</b>     | Creation time of the logical interface when this MAC address was learned. Used for debugging.   |

## Sample Output

### show bridge mac-table

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : default-switch
Bridging domain : test1, VLAN : 1
MAC          MAC      Logical   NH      RTR
address      flags    interface Index   ID
01:00:0c:cc:cc:cc S,NM    NULL
01:00:0c:cc:cc:cd S,NM    NULL
01:00:0c:cd:cd:d0 S,NM    NULL
64:87:88:6a:17:d0 D        ae0.1
64:87:88:6a:17:f0 D        ae0.1

```

### show bridge mac-table (with VXLAN enabled)

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : default-switch
Bridging domain : vlan-1, VLAN : 1
VXLAN: Id : 100, Multicast group: 226.1.1.1
MAC          MAC      Logical   NH      RTR
address      flags    interface Index   ID
00:01:01:00:01:f7 D,SE    vtep.1052010
00:03:00:32:01:f7 D,SE    vtep.1052011
00:00:21:11:11:10 DL        ge-1/0/0.0
00:00:21:11:11:11 DL        ge-1/1/0.0

```

```

Routing instance : default-switch
Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
VXLAN: Id : 200, Multicast group: 226.1.1.2
MAC          MAC      Logical   NH      RTR
address      flags    interface Index   ID
00:02:01:33:01:f7 D,SE    vtep.1052010
00:04:00:14:01:f7 D,SE    vtep.1052011
00:00:21:11:21:10 DL        ge-1/0/0.1
00:00:21:11:21:11 DL        ge-1/1/0.1

```

### show bridge mac-table count

```

user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

```

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-11/0/3.0       | 1         |
| ge-11/1/4.100     | 0         |
| ge-11/1/1.100     | 0         |
| ge-11/1/0.100     | 0         |
| xe-10/2/0.100     | 1         |
| xe-10/0/0.100     | 0         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 2         |

```

0 MAC address learned in routing instance vs1 bridge domain vlan200

```

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
|-------------------|-----------|

|               |   |
|---------------|---|
| ge-11/1/0.200 | 0 |
| ge-11/1/1.200 | 0 |
| ge-11/1/4.200 | 0 |
| xe-10/0/0.200 | 0 |
| xe-10/2/0.200 | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 0         |

### show bridge mac-table detail

```
user@host> show bridge mac-table detail
```

MAC address: 00:00:00:19:1c:db

Routing instance: vs1

Bridging domain: vlan100

Learning interface: ge-11/0/3.0      Learning VLAN: 0

Layer 2 flags: in\_ifd, in\_ifl, in\_vlan, kernel

Epoch: 4      Sequence number: 0

Learning mask: 0x800      IPC generation: 0

MAC address: 00:00:00:59:3a:2f

Routing instance: vs1

Bridging domain: vlan100

Learning interface: xe-10/2/0.100      Learning VLAN: 0

Layer 2 flags: in\_ifd, in\_ifl, in\_vlan, kernel

Epoch: 7      Sequence number: 0

Learning mask: 0x400      IPC generation: 0

## show ovsdb controller

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show ovsdb controller</code><br><code>&lt;address ip-address&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | <p>Display information and connection status for VMware NSX controllers to which the Juniper Networks device is connected. This command displays information for NSX controllers with connections to a Juniper Networks device that are made in the following ways:</p> <ul style="list-style-type: none"> <li>• With explicit configuration—The connection is explicitly configured using the Junos OS CLI.</li> <li>• Without explicit configuration—An NSX controller to which the Juniper Networks device is connected pushes information about other controllers in the same cluster to the device. With this method, the Juniper Networks device learns about the other controllers in the same cluster and connections to these controllers are established without explicit configuration.</li> </ul> |
| <b>Options</b>                  | <p><b>none</b>—Display information about all NSX controllers to which the Juniper Networks device is connected.</p> <p><b>address ip-address</b>—Display information about the NSX controller at the specified IP address.</p>  |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Setting Up the Open vSwitch Database Management Protocol on Juniper Networks Devices on page 1985</a></li> <li>• <a href="#">Understanding How to Set Up Open vSwitch Database Connections Between Juniper Networks Devices and Controllers on page 1903</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show ovsdb controller on page 2015</a><br><a href="#">show ovsdb controller address on page 2015</a>  |
| <b>Output Fields</b>            | Table 139 on page 2014 lists the output fields for the <b>show ovsdb controller</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 139: show ovsdb controller Output Fields

| Field Name            | Field Descriptions  |
|-----------------------|---|
| Controller IP address | IP address of an NSX controller to which the Juniper Networks device is connected.            |
| Controller protocol   | Protocol used by the Juniper Networks device to initiate a connection with an NSX controller. |



Table 139: show ovssdb controller Output Fields (*continued*)

| Field Name                          | Field Descriptions   |
|-------------------------------------|--|
| Controller port                     | NSX controller port to which the Juniper Networks device is connected.   |
| Controller connection               | State of the connection between the Juniper Networks device and an NSX controller.                             |
| Controller seconds-since-connect    | Number of seconds since the connection between the Juniper Networks device and NSX controller was established. |
| Controller seconds-since-disconnect | Number of seconds since the connection between the Juniper Networks device and NSX controller was dropped.     |
| Controller connection status        | Status of the connection between the Juniper Networks device and an NSX controller.                            |

## Sample Output

### show ovssdb controller

```

user@host> show ovssdb controller
VTEP controller information:
Controller IP address: 10.168.66.189
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56290
Controller seconds-since-disconnect: 0
Controller connection status: active

Controller IP address: 10.168.181.54
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56292
Controller seconds-since-disconnect: 0
Controller connection status: active

Controller IP address: 10.168.182.45
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56292
Controller seconds-since-disconnect: 0
Controller connection status: active

```

### show ovssdb controller address

```

user@host> show ovssdb controller address 10.168.182.45
VTEP controller information:
Controller IP address: 192.168.182.45
Controller protocol: ssl
Controller port: 6632
Controller connection: up
Controller seconds-since-connect: 56347
Controller seconds-since-disconnect: 0
Controller connection status: active

```



## show ovssdb interface

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show ovssdb interface</code><br><code>&lt;interface-name&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Display information about Open vSwitch Database (OVSSDB)-managed interfaces configured by using the <b>interfaces interface-name</b> statement in the <b>[edit protocols ovssdb]</b> hierarchy. |
| <b>Options</b>                  | <b>none</b> —Display information about all OVSSDB-managed interfaces.<br><br><b>interface-name</b> —Display information about the specified OVSSDB-managed interface.                           |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring OVSSDB-Managed VXLANs on page 1989</a></li> <li>• <a href="#">show ovssdb statistics interface on page 2025</a></li> </ul>     |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb interface on page 2017</a><br><a href="#">show ovssdb (Specific Interface) on page 2018</a>   |
| <b>Output Fields</b>            | Table 140 on page 2017 lists the output fields for the <b>show ovssdb interface</b> command. Output fields are listed in the approximate order in which they appear.                            |

Table 140: show ovssdb interface Output Fields

| Field Name            | Field Description   |
|-----------------------|---|
| Interface             | Name of interface.  |
| VLAN ID               | ID of Virtual Extensible LAN (VXLAN) with which the interface is associated.<br><br><b>NOTE:</b> This field is not supported by MX Series routers.. |
| Bridge domain or VLAN | Bridge domain or VLAN under which the VXLAN is created.<br><br><b>NOTE:</b> This field is not supported by MX Series routers.                       |

## Sample Output

### show ovssdb interface

```

user@host> show ovssdb interface
Interface          VLAN ID          Bridge-domain
ge-7/0/9.0
ge-7/0/9.1
irb.11
irb.12

```

```
irb.2
irb.3
xe-10/3/0.0
xe-10/3/0.1
```

#### show ovssdb (Specific Interface)

```
user@host> show ovssdb interface ge-7/0/9.0
Interface          VLAN ID      Bridge-domain
ge-7/0/9.0
```

## show ovssdb logical-switch

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show ovssdb logical-switch</code><br><code>&lt;logical-switch-name&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | <p>Display information about logical switches, which you configured in NSX Manager or the NSX API, and the corresponding Virtual Extensible LANs (VXLANs), which were configured on the Juniper Networks device.</p> <p>In the command output, each logical switch is identified by a universally unique identifier (UUID), which in the context of this command, is also known as a logical switch name.</p> <p>The <b>show ovssdb logical-switch</b> command displays the state of the logical switch (<b>Flags</b>), which can be one of the following:</p> <p><b>Created by Controller</b>—A logical switch was configured in NSX Manager. In this state, the logical switch and corresponding VXLAN are not yet operational.</p> <p><b>Created by L2ALD</b>—A VXLAN was configured on a Juniper Networks device. In this state, the logical switch-VXLAN are not yet operational.</p> <p><b>Created by both</b>—A logical switch was configured in NSX Manager, and a corresponding VXLAN was configured on a Juniper Networks device. In this state, the logical switch-VXLAN are operational.</p> <p><b>Tunnel key mismatch</b>—The VNIs specified in the logical switch and corresponding VXLAN configurations do not match. In this state, the logical switch-VXLAN are not yet operational.</p> <p>For more information about configuring the logical switch and corresponding VXLAN, see <a href="#">“Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment” on page 1905</a>.</p> |
| <b>Options</b>                  | <p><b>none</b>—Display information about all logical switches that are present in the OVSSDB schema for physical devices.</p> <p><b>logical-switch-name</b>—Display information about the specified logical switch.</p>   |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Open vSwitch Database Schema For Physical Devices on page 1910</a></li> <li>• <a href="#">Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSSDB-Managed VXLAN on page 2045</a></li> </ul>   |
| <b>List of Sample Output</b>    | <p><a href="#">show ovssdb logical-switch on page 2020</a></p> <p><a href="#">show ovssdb logical-switch (Specific Logical Switch) on page 2020</a></p>   |

**Output Fields** [Table 141 on page 2020](#) lists the output fields for the **show ovssdb logical-switch** command. Output fields are listed in the approximate order in which they appear.

**Table 141: show ovssdb logical-switch Output Fields**

| Field Name          | Field Description   |
|---------------------|---|
| Logical Switch Name | UUID that is automatically generated by NSX and assigned to the logical switch after you configure it in NSX Manager or the NSX API. When configuring the corresponding VXLAN in the Junos OS CLI, the same UUID must be specified as the VXLAN name. |
| Flags               | State of the logical switch. For possible states, see the Description section of this topic.  |
| VNI                 | VNI that is configured for the logical switch and corresponding VXLAN.  |
| Num of Remote MAC   | The total number of remote MAC addresses associated with the logical switch. These addresses are learned by software and hardware virtual tunnel endpoints (VTEPs) in the NSX environment.  |
| Num of Local MAC    | The total number of local MAC addresses associated with the logical switch. <i>Local MAC addresses</i> are addresses learned on the local physical ports.   |

## Sample Output

### show ovssdb logical-switch

```

user@host> show ovssdb logical-switch
Logical switch information:
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Flags: Created by both
VNI: 3
Num of Remote MAC: 13
Num of Local MAC: 12
Logical Switch Name: 9b4f880e-dac8-4612-a832-97ad9dec270f
Flags: Created by Controller
VNI: 50
Num of Remote MAC: 0
Num of Local MAC: 0
Logical Switch Name: bc0da2da-6c16-44bf-b655-442484294ded
Flags: Created by Controller
VNI: 51
Num of Remote MAC: 0
Num of Local MAC: 0

```

### show ovssdb logical-switch (Specific Logical Switch)

```

user@host> show ovssdb logical-switch 24a76aff-7e61-4520-a78d-3eca26ad7510
Logical switch information:
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
Flags: Created by both
VNI: 3
Num of Remote MAC: 13
Num of Local MAC: 12

```

## show ovssdb mac

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>show ovssdb mac &lt;address <i>mac-address</i>&gt; &lt;local&gt; &lt;logical-switch <i>logical-switch-uuid</i>&gt; &lt;multicast&gt; &lt;remote&gt; &lt;unicast&gt;</pre>  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 14.1R2.</p> <p>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>  |
| <b>Description</b>              | <p>Display MAC addresses, as well as information about the MAC addresses, learned by a Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP). Using the Open vSwitch Database (OVSSDB) management protocol, this hardware VTEP can learn about MAC addresses directly or from other software or hardware VTEPs. The MAC addresses learned directly by the hardware VTEP are known as local addresses, while the addresses learned from other software or hardware VTEPs are known as remote addresses.</p>  |
| <b>Options</b>                  | <p>Use one or more of the following options to display a more specific list of MAC addresses and information about the MAC addresses. For example, to display a list of local unicast MAC addresses, you can issue the <b>show ovssdb mac local unicast</b> command.</p> <p><b>none</b>—Display all MAC addresses, which includes all local, remote, unicast, and multicast addresses associated with all logical switches.</p> <p><b>address <i>mac-address</i></b>—Display the specified MAC address.</p> <p><b>local</b>—Display all local MAC addresses.</p> <p><b>logical-switch <i>logical-switch-uuid</i></b>—Display all MAC addresses associated with the specified logical switch.</p> <p><b>multicast</b>—Display all multicast MAC addresses.</p> <p><b>remote</b>—Display all remote MAC addresses.</p> <p><b>unicast</b>—Display all unicast MAC addresses.</p> |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Open vSwitch Database Schema For Physical Devices on page 1910</a></li> </ul>  |
| <b>List of Sample Output</b>    | <p><a href="#">show ovssdb mac on page 2022</a></p> <p><a href="#">show ovssdb mac address on page 2023</a></p> <p><a href="#">show ovssdb mac logical-switch on page 2023</a></p> <p><a href="#">show ovssdb mac local unicast on page 2023</a></p>  |

**Output Fields** Table 142 on page 2022 lists the output fields for the **show ovssdb mac** command. Output fields are listed in the approximate order in which they appear.

**Table 142: show ovssdb mac Output Fields**

| Field Name          | Field Description  |
|---------------------|--|
| Logical Switch Name | Universally unique identifier (UUID) of the logical switch. For more information about logical switches and UUIDs, see <a href="#">“Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment”</a> on page 1905. |
| MAC Address         | MAC addresses of virtual machines (VMs).   |
| IP Address          | IP address of VMs.<br><br><b>NOTE:</b> If the IP addresses of VMs are not published by the NSX controller, this field displays 0.0.0.0.  |
| Encapsulation       | Encapsulation type.  |
| VTEP Address        | IP address of the hardware or software VTEP from which the MAC address was learned. Further, this VTEP can forward VM traffic to the associated host.  |

## Sample Output

### show ovssdb mac

```

user@host> show ovssdb mac
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
  Mac      IP      Encapsulation      Vtep
  Address  Address
02:00:00:00:03:01  0.0.0.0      Vxlan over Ipv4      10.255.18.22
02:00:00:00:03:02  0.0.0.0      Vxlan over Ipv4      10.255.18.22
02:00:00:00:03:03  0.0.0.0      Vxlan over Ipv4      10.255.18.22
02:00:00:00:03:04  0.0.0.0      Vxlan over Ipv4      10.255.18.22
02:00:00:00:03:05  0.0.0.0      Vxlan over Ipv4      10.255.18.22
04:00:00:00:03:05  0.0.0.0      Vxlan over Ipv4      10.255.18.22
06:00:00:00:03:01  0.0.0.0      Vxlan over Ipv4      10.255.18.22
06:00:00:00:03:02  0.0.0.0      Vxlan over Ipv4      10.255.18.22
06:00:00:00:03:03  0.0.0.0      Vxlan over Ipv4      10.255.18.22
06:00:00:00:03:04  0.0.0.0      Vxlan over Ipv4      10.255.18.22
06:00:00:00:03:05  0.0.0.0      Vxlan over Ipv4      10.255.18.22
40:b4:f0:06:6f:f0  0.0.0.0      Vxlan over Ipv4      10.255.18.22
ff:ff:ff:ff:ff:ff  0.0.0.0      Vxlan over Ipv4      10.100.100.1

Logical Switch Name: bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab
  Mac      IP      Encapsulation      Vtep
  Address  Address
02:00:00:00:11:01  0.0.0.0      Vxlan over Ipv4      10.1.1.29
02:00:00:00:11:02  0.0.0.0      Vxlan over Ipv4      10.1.1.29
02:00:00:00:11:03  0.0.0.0      Vxlan over Ipv4      10.1.1.29
02:00:00:00:11:04  0.0.0.0      Vxlan over Ipv4      10.1.1.29
02:00:00:00:11:05  0.0.0.0      Vxlan over Ipv4      10.1.1.29
04:00:00:00:11:05  0.0.0.0      Vxlan over Ipv4      10.1.1.29
06:00:00:00:11:01  0.0.0.0      Vxlan over Ipv4      10.1.1.29
06:00:00:00:11:02  0.0.0.0      Vxlan over Ipv4      10.1.1.29
06:00:00:00:11:03  0.0.0.0      Vxlan over Ipv4      10.1.1.29

```



```

06:00:00:00:11:04      0.0.0.0      Vxlan over Ipv4      10.1.1.29
06:00:00:00:11:05      0.0.0.0      Vxlan over Ipv4      10.1.1.29
40:b4:f0:06:6f:f0      0.0.0.0      Vxlan over Ipv4      10.1.1.29
00:23:9c:5e:a7:f0      0.0.0.0      Vxlan over Ipv4      10.255.18.22
08:00:00:00:11:01      0.0.0.0      Vxlan over Ipv4      10.255.18.22
08:00:00:00:11:02      0.0.0.0      Vxlan over Ipv4      10.255.18.22
08:00:00:00:11:03      0.0.0.0      Vxlan over Ipv4      10.255.18.22
08:00:00:00:11:04      0.0.0.0      Vxlan over Ipv4      10.255.18.22
08:00:00:00:11:05      0.0.0.0      Vxlan over Ipv4      10.255.18.22
ff:ff:ff:ff:ff:ff      0.0.0.0      Vxlan over Ipv4      10.110.110.1
...

```

### show ovssdb mac address

```
user@host> show ovssdb mac address 02:00:00:00:03:01
```

| Mac<br>Address    | IP<br>Address | Encapsulation   | Vtep<br>Address |
|-------------------|---------------|-----------------|-----------------|
| 02:00:00:00:03:01 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |

### show ovssdb mac logical-switch

```
user@host> show ovssdb mac logical-switch bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab
```

```
Logical Switch Name: bf6d4fd4-f5f6-430c-8c37-4033ef1c55ab
```

| Mac<br>Address    | IP<br>Address | Encapsulation   | Vtep<br>Address |
|-------------------|---------------|-----------------|-----------------|
| 02:00:00:00:11:01 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 02:00:00:00:11:02 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 02:00:00:00:11:03 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 02:00:00:00:11:04 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 02:00:00:00:11:05 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 04:00:00:00:11:05 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 06:00:00:00:11:01 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 06:00:00:00:11:02 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 06:00:00:00:11:03 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 06:00:00:00:11:04 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 06:00:00:00:11:05 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 40:b4:f0:06:6f:f0 | 0.0.0.0       | Vxlan over Ipv4 | 10.1.1.29       |
| 00:23:9c:5e:a7:f0 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |
| 08:00:00:00:11:01 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |
| 08:00:00:00:11:02 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |
| 08:00:00:00:11:03 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |
| 08:00:00:00:11:04 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |
| 08:00:00:00:11:05 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.18.22    |
| ff:ff:ff:ff:ff:ff | 0.0.0.0       | Vxlan over Ipv4 | 10.110.110.1    |

### show ovssdb mac local unicast

```
user@host> show ovssdb mac local unicast
```

```
Logical Switch Name: 24a76aff-7e61-4520-a78d-3eca26ad7510
```

| Mac<br>Address    | IP<br>Address | Encapsulation   | Vtep<br>Address |
|-------------------|---------------|-----------------|-----------------|
| 02:00:00:00:03:01 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 02:00:00:00:03:02 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 02:00:00:00:03:03 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 02:00:00:00:03:04 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 02:00:00:00:03:05 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 04:00:00:00:03:05 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 06:00:00:00:03:01 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 06:00:00:00:03:02 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 06:00:00:00:03:03 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |
| 06:00:00:00:03:04 | 0.0.0.0       | Vxlan over Ipv4 | 10.255.181.72   |

|                   |         |                 |               |
|-------------------|---------|-----------------|---------------|
| 06:00:00:00:03:05 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.181.72 |
| 40:b4:f0:06:6f:f0 | 0.0.0.0 | Vxlan over Ipv4 | 10.255.181.72 |
| ...               |         |                 |               |

## show ovssdb statistics interface

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>show ovssdb statistics interface</code><br><code>&lt;interface-name&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display statistics for Open vSwitch Database (OVSSDB)-managed interfaces configured by using the <b>interfaces</b> <i>interface-name</i> statement in the <b>[edit protocols ovssdb]</b> hierarchy.<br><br>When an interface is configured as OVSSDB-managed, the collection of statistics for that interface begins, and the statistics displayed at any given time reflects the data collected up to that point. |
| <b>Options</b>                  | <b>none</b> —Display statistics for all configured OVSSDB-managed interfaces.<br><br><i>interface-name</i> —Display statistics for the specified interface.  |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">interfaces on page 1996</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb statistics interface on page 2025</a><br><a href="#">show ovssdb statistics interface (Specific Interface) on page 2026</a>  |
| <b>Output Fields</b>            | <a href="#">Table 143 on page 2025</a> lists the output fields for the <b>show ovssdb statistics interface</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 143: show ovssdb statistics interface Output Fields

| Field Name      | Field Descriptions                           |
|-----------------|--|
| Num of rx pkts  | Number of packets received by the interface. |
| Num of tx pkts  | Number of packets sent by the interface.     |
| Num of rx bytes | Number of bytes received by the interface.   |
| Num of tx bytes | Number of bytes sent by the interface.       |

## Sample Output

### show ovssdb statistics interface

```

user@host> show ovssdb statistics interface
Interface Name: ge-7/0/9.0
Num of rx pkts: 945                      Num of tx pkts: 113280890
Num of rx bytes: 56700                   Num of tx bytes: 57531319540
Interface Name: ge-7/0/10.0

```

|                             |                              |
|-----------------------------|------------------------------|
| Num of rx pkts: 459         | Num of tx pkts: 473840856    |
| Num of rx bytes: 84747      | Num of tx bytes: 45830738532 |
| Interface Name: ge-7/0/11.0 |                              |
| Num of rx pkts: 305         | Num of tx pkts: 367483456    |
| Num of rx bytes: 98974      | Num of tx bytes: 33495468092 |

#### show ovsdb statistics interface (Specific Interface)

```
user@host> show ovsdb statistics interface ge-7/0/9.0
```

|                            |                              |
|----------------------------|------------------------------|
| Interface Name: ge-7/0/9.0 |                              |
| Num of rx pkts: 945        | Num of tx pkts: 113280890    |
| Num of rx bytes: 56700     | Num of tx bytes: 57531319540 |

## show ovssdb virtual-tunnel-end-point

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | show ovssdb virtual-tunnel-end-point<br>address <ip-address><br>encapsulation <encapsulation-type>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1R2.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display information about the following entities that the Juniper Networks device has learned: <ul style="list-style-type: none"> <li>• Other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs)</li> <li>• Software VTEPs</li> <li>• Service nodes</li> </ul>  |
| <b>Options</b>                  | <b>none</b> —Display information about all VTEPs and service nodes that the Juniper Networks device has learned.<br><br><b>address ip-address</b> —Display information about the entity with specified IP address.<br><br><b>encapsulation encapsulation-type</b> —Display information about all entities with the specified encapsulation type.   |
| <b>Required Privilege Level</b> | admin  |
| <b>List of Sample Output</b>    | <a href="#">show ovssdb virtual-tunnel-end-point on page 2028</a><br><a href="#">show ovssdb virtual-tunnel-end-point address (Specific Address) on page 2028</a><br><a href="#">show ovssdb virtual-tunnel-end-point encapsulation (Specific Encapsulation) on page 2028</a><br><a href="#">show ovssdb virtual-tunnel-end-point address (Specific Address) encapsulation (Specific Encapsulation) on page 2028</a> |
| <b>Output Fields</b>            | Table 144 on page 2027 lists the output fields for the <b>show ovssdb virtual-tunnel-end-point</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 144: show ovssdb virtual-tunnel-end-point Output Fields

| Field Name    | Field Description                              |
|---------------|--|
| Encapsulation | Encapsulation type of entity.                  |
| IP Address    | IP address of entity.                          |
| Num of MACs   | Number of MAC addresses learned by the entity. |

## Sample Output

### show ovssdb virtual-tunnel-end-point

```
user@host> show ovssdb virtual-tunnel-end-point
Encapsulation      Ip Address      Num of MAC's
VXLAN over IPv4    10.255.181.43   24
VXLAN over IPv4    10.255.181.50   12
VXLAN over IPv4    10.255.181.72   24
```

### show ovssdb virtual-tunnel-end-point address (Specific Address)

```
user@host> show ovssdb virtual-tunnel-end-point address 10.255.181.43
Encapsulation      Ip Address      Num of MAC's
VXLAN over IPv4    10.255.181.43   24
```

### show ovssdb virtual-tunnel-end-point encapsulation (Specific Encapsulation)

```
user@host> show ovssdb virtual-tunnel-end-point encapsulation vxlan-over-ipv4
Encapsulation      Ip Address      Num of MAC's
VXLAN over IPv4    10.255.181.43   24
VXLAN over IPv4    10.255.181.50   12
VXLAN over IPv4    10.255.181.72   24
```

### show ovssdb virtual-tunnel-end-point address (Specific Address) encapsulation (Specific Encapsulation)

```
user@host> show ovssdb virtual-tunnel-end-point address 10.255.181.43 encapsulation
vxlan-over-ipv4
Encapsulation      Ip Address      Num of MAC's
VXLAN over IPv4    10.255.181.43   24
```

## show vpls mac-table

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>show vpls mac-table &lt;brief   detail   extensive   summary&gt; &lt;bridge-domain <i>bridge-domain-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;mac-address&gt; &lt;vlan-id <i>vlan-id-number</i>&gt;</pre>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.   |
| <b>Description</b>              | (MX960 routers only) Display learned VPLS MAC address information.  |
| <b>Options</b>                  | <p><b>none</b>—Display all learned VPLS MAC address information.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p><b>mac-address</b>—(Optional) Display the specified learned VPLS MAC address information..</p> <p><b>vlan-id <i>vlan-id-number</i></b>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view  |
| <b>List of Sample Output</b>    | <p><a href="#">show vpls mac-table on page 2030</a></p> <p><a href="#">show vpls mac-table (with VXLAN enabled) on page 2031</a></p> <p><a href="#">show vpls mac-table count on page 2031</a></p> <p><a href="#">show vpls mac-table detail on page 2032</a></p> <p><a href="#">show vpls mac-table extensive on page 2032</a></p>   |
| <b>Output Fields</b>            | <p><a href="#">Table 145 on page 2029</a> describes the output fields for the <b>show bridge mac-table</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

Table 145: show vpls mac-table Output fields

| Field Name       | Field Description             |
|------------------|-------------------------------|
| Routing instance | Name of the routing instance. |

Table 145: show vpls mac-table Output fields (*continued*)

| Field Name                | Field Description  |
|---------------------------|--|
| <b>Bridging domain</b>    | Name of the bridging domain.   |
| <b>MAC address</b>        | MAC address or addresses learned on a logical interface.   |
| <b>MAC flags</b>          | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address configured.</li> <li>• <b>D</b>—Dynamic MAC address learned.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Nonconfigured MAC.</li> </ul> |
| <b>Logical interface</b>  | Name of the logical interface.   |
| <b>MAC count</b>          | Number of MAC addresses learned on a specific routing instance or interface.   |
| <b>Learning interface</b> | Logical interface or logical Label Switched Interface (LSI) the address is learned on.   |
| <b>Learn VLAN ID/VLAN</b> | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.   |
| <b>VXLAN ID/VXLAN</b>     | VXLAN Network Identifier (VNI)   |
| <b>Layer 2 flags</b>      | Debugging flags signifying that the MAC address is present in various lists.   |
| <b>Epoch</b>              | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.  |
| <b>Sequence number</b>    | Sequence number assigned to this MAC address. Used for debugging.  |
| <b>Learning mask</b>      | Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.  |
| <b>IPC generation</b>     | Creation time of the logical interface when this MAC address was learned. Used for debugging.  |

## Sample Output

### show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
  MAC          MAC      Logical
  address      flags    interface
  00:90:69:9c:1c:5d  D      ge-0/2/5.400

MAC flags (S -static MAC, D -dynamic MAC,
           SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_red

```



```

VLAN : 401
MAC          MAC      Logical
address      flags    interface
00:00:aa:12:12:12 D      lsi.1051138
00:05:85:74:9f:f0 D      lsi.1051138

```

### show vpls mac-table (with VXLAN enabled)

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
           SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
VXLAN: Id : 300, Multicast group: 226.1.1.3
MAC          MAC      Logical
address      flags    interface
00:01:01:00:01:f4 D,SE    ge-4/2/0.1000
00:02:01:33:01:f4 D,SE    lsi.1052004
00:03:00:32:01:f4 D,SE    lsi.1048840
00:04:00:14:01:f4 D,SE    lsi.1052005
00:02:01:33:02:f7 D,SE    vtep.1052010
00:04:00:14:02:f7 D,SE    vtep.1052011

```

### show vpls mac-table count

```

user@host> show vpls mac-table count
0 MAC address learned in routing instance __juniper_private1__

```

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lc-0/0/0.32769    | 0         |
| lc-0/1/0.32769    | 0         |
| lc-0/2/0.32769    | 0         |
| lc-2/0/0.32769    | 0         |
| lc-0/3/0.32769    | 0         |
| lc-2/1/0.32769    | 0         |
| lc-9/0/0.32769    | 0         |
| lc-11/0/0.32769   | 0         |
| lc-2/2/0.32769    | 0         |
| lc-9/1/0.32769    | 0         |
| lc-11/1/0.32769   | 0         |
| lc-2/3/0.32769    | 0         |
| lc-9/2/0.32769    | 0         |
| lc-11/2/0.32769   | 0         |
| lc-11/3/0.32769   | 0         |
| lc-9/3/0.32769    | 0         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 0         |

1 MAC address learned in routing instance vpls\_ldp1

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lsi.1051137       | 0         |
| ge-0/2/5.400      | 1         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
|---------------|-----------|

|   | 0         | 1 |
|---|-----------|---|
| 1 MAC address learned in routing instance vpls_red        |           |   |
| MAC address count per interface within routing instance:  |           |   |
| Logical interface   | MAC count |   |
| ge-0/2/5.300  | 1         |   |
| MAC address count per learn VLAN within routing instance: |           |   |
| Learn VLAN ID   | MAC count |   |
| 0   | 1         |   |

### show vpls mac-table detail

```

user@host> show vpls mac-table detail
MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_red
Learning interface: ge-0/2/5.300
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

```

### show vpls mac-table extensive

```

user@host> show vpls mac-table extensive
MAC address: 00:00:aa:12:12:12
Routing instance: vpls_ldp1
Learning interface: lsi.1051137
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0
Routing instance: vpls_ldp1
Learning interface: lsi.1051137
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:00:aa:12:12:12
Routing instance: vpls_red
Learning interface: lsi.1051138
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 0
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0

```

```
Routing instance: vpls_red
Learning interface: lsi.1051138
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 0
Learning mask: 0x1                       IPC generation: 0
```

---

## VXLAN Monitoring Commands

---

- [show bridge mac-table](#)
- [show vpls mac-table](#)
- [Verifying VXLAN Reachability on page 2042](#)
- [Verifying That a Local VXLAN VTEP is Configured Correctly on page 2042](#)
- [Verifying MAC Learning from a Remote VTEP on page 2042](#)
- [Monitor a Remote VTEP Interface on page 2043](#)

## show bridge mac-table

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>show bridge mac-table &lt;brief   count   detail   extensive&gt; &lt;bridge-domain (all   <i>bridge-domain-name</i>)&gt; &lt;global-count&gt; &lt;interface <i>interface-name</i>&gt; &lt;mac-address&gt; &lt;vlan-id (all-vlan   <i>vlan-id</i>)&gt;</pre>   |
| Release Information      | Command introduced in Junos OS Release 8.4.  |
| Description              | (MX Series routers only) Display Layer 2 MAC address information.  |
| Options                  | <p><b>none</b>—Display all learned Layer 2 MAC address information.</p> <p><b>brief   count   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain (all   <i>bridge-domain-name</i>)</b>—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p><b>global-count</b>—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned Layer 2 MAC addresses for the specified interface.</p> <p><b>mac-address</b>—(Optional) Display the specified learned Layer 2 MAC address information.</p> <p><b>vlan-id (all-vlan   <i>vlan-id</i>)</b>—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p> |
| Additional Information   | When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.   |
| Required Privilege Level | view   |
| List of Sample Output    | <p><a href="#">show bridge mac-table on page 2035</a></p> <p><a href="#">show bridge mac-table (with VXLAN enabled) on page 2036</a></p> <p><a href="#">show bridge mac-table count on page 2036</a></p> <p><a href="#">show bridge mac-table detail on page 2037</a></p>  |
| Output Fields            | <p><a href="#">Table 138 on page 2011</a> describes the output fields for the <b>show bridge mac-table</b> command. Output fields are listed in the approximate order in which they appear.</p>  |

Table 146: show bridge mac-table Output fields

| Field Name                | Field Description   |
|---------------------------|---|
| <b>Routing instance</b>   | Name of the routing instance.   |
| <b>Bridging domain</b>    | Name of the bridging domain.  |
| <b>MAC address</b>        | MAC address or addresses learned on a logical interface.  |
| <b>MAC flags</b>          | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address is configured.</li> <li>• <b>D</b>—Dynamic MAC address is configured.</li> <li>• <b>L</b>—Locally learned MAC address is configured.</li> <li>• <b>C</b>—Control MAC address is configured.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Non-configured MAC.</li> <li>• <b>R</b>—Remote PE MAC address is configured.</li> </ul> |
| <b>Logical interface</b>  | Name of the logical interface.  |
| <b>MAC count</b>          | Number of MAC addresses learned on the specific routing instance or interface.  |
| <b>Learning interface</b> | Name of the logical interface on which the MAC address was learned.   |
| <b>Learning VLAN</b>      | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.  |
| <b>VXLAN ID/VXLAN</b>     | VXLAN Network Identifier (VNI)  |
| <b>Layer 2 flags</b>      | Debugging flags signifying that the MAC address is present in various lists.  |
| <b>Epoch</b>              | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.   |
| <b>Sequence number</b>    | Sequence number assigned to this MAC address. Used for debugging.   |
| <b>Learning mask</b>      | Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.   |
| <b>IPC generation</b>     | Creation time of the logical interface when this MAC address was learned. Used for debugging.   |

## Sample Output

### show bridge mac-table

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : default-switch
Bridging domain : test1, VLAN : 1
  MAC      MAC      Logical  NH      RTR
  address   flags    interface Index  ID
01:00:0c:cc:cc:cc S,NM    NULL
01:00:0c:cc:cc:cd S,NM    NULL
01:00:0c:cd:cd:d0 S,NM    NULL
64:87:88:6a:17:d0 D        ae0.1
64:87:88:6a:17:f0 D        ae0.1

```

### show bridge mac-table (with VXLAN enabled)

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : default-switch
Bridging domain : vlan-1, VLAN : 1
VXLAN: Id : 100, Multicast group: 226.1.1.1
  MAC      MAC      Logical
  address   flags    interface
00:01:01:00:01:f7 D,SE    vtep.1052010
00:03:00:32:01:f7 D,SE    vtep.1052011
00:00:21:11:11:10 DL        ge-1/0/0.0
00:00:21:11:11:11 DL        ge-1/1/0.0

```

```

Routing instance : default-switch
Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
VXLAN: Id : 200, Multicast group: 226.1.1.2
  MAC      MAC      Logical
  address   flags    interface
00:02:01:33:01:f7 D,SE    vtep.1052010
00:04:00:14:01:f7 D,SE    vtep.1052011
00:00:21:11:21:10 DL        ge-1/0/0.1
00:00:21:11:21:11 DL        ge-1/1/0.1

```

### show bridge mac-table count

```

user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

```

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| ge-11/0/3.0       | 1         |
| ge-11/1/4.100     | 0         |
| ge-11/1/1.100     | 0         |
| ge-11/1/0.100     | 0         |
| xe-10/2/0.100     | 1         |
| xe-10/0/0.100     | 0         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 2         |

```

0 MAC address learned in routing instance vs1 bridge domain vlan200

```

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
|-------------------|-----------|

|               |   |
|---------------|---|
| ge-11/1/0.200 | 0 |
| ge-11/1/1.200 | 0 |
| ge-11/1/4.200 | 0 |
| xe-10/0/0.200 | 0 |
| xe-10/2/0.200 | 0 |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 0         |

### show bridge mac-table detail

```
user@host> show bridge mac-table detail
```

MAC address: 00:00:00:19:1c:db

Routing instance: vs1

Bridging domain: vlan100

Learning interface: ge-11/0/3.0      Learning VLAN: 0

Layer 2 flags: in\_ifd, in\_ifl, in\_vlan, kernel

Epoch: 4      Sequence number: 0

Learning mask: 0x800      IPC generation: 0

MAC address: 00:00:00:59:3a:2f

Routing instance: vs1

Bridging domain: vlan100

Learning interface: xe-10/2/0.100      Learning VLAN: 0

Layer 2 flags: in\_ifd, in\_ifl, in\_vlan, kernel

Epoch: 7      Sequence number: 0

Learning mask: 0x400      IPC generation: 0

## show vpls mac-table

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>show vpls mac-table &lt;brief   detail   extensive   summary&gt; &lt;bridge-domain <i>bridge-domain-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;mac-address&gt; &lt;vlan-id <i>vlan-id-number</i>&gt;</pre>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.   |
| <b>Description</b>              | (MX960 routers only) Display learned VPLS MAC address information.  |
| <b>Options</b>                  | <p><b>none</b>—Display all learned VPLS MAC address information.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bridge-domain <i>bridge-domain-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p><b>mac-address</b>—(Optional) Display the specified learned VPLS MAC address information..</p> <p><b>vlan-id <i>vlan-id-number</i></b>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p> |
| <b>Required Privilege Level</b> | view  |
| <b>List of Sample Output</b>    | <p><a href="#">show vpls mac-table on page 2039</a></p> <p><a href="#">show vpls mac-table (with VXLAN enabled) on page 2040</a></p> <p><a href="#">show vpls mac-table count on page 2040</a></p> <p><a href="#">show vpls mac-table detail on page 2041</a></p> <p><a href="#">show vpls mac-table extensive on page 2041</a></p>   |
| <b>Output Fields</b>            | <p><a href="#">Table 145 on page 2029</a> describes the output fields for the <b>show bridge mac-table</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

Table 147: show vpls mac-table Output fields

| Field Name       | Field Description             |
|------------------|-------------------------------|
| Routing instance | Name of the routing instance. |



Table 147: show vpls mac-table Output fields (*continued*)

| Field Name                | Field Description  |
|---------------------------|--|
| <b>Bridging domain</b>    | Name of the bridging domain.   |
| <b>MAC address</b>        | MAC address or addresses learned on a logical interface.   |
| <b>MAC flags</b>          | Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> <li>• <b>S</b>—Static MAC address configured.</li> <li>• <b>D</b>—Dynamic MAC address learned.</li> <li>• <b>SE</b>—MAC accounting is enabled.</li> <li>• <b>NM</b>—Nonconfigured MAC.</li> </ul> |
| <b>Logical interface</b>  | Name of the logical interface.   |
| <b>MAC count</b>          | Number of MAC addresses learned on a specific routing instance or interface.   |
| <b>Learning interface</b> | Logical interface or logical Label Switched Interface (LSI) the address is learned on.   |
| <b>Learn VLAN ID/VLAN</b> | VLAN ID of the routing instance or bridge domain in which the MAC address was learned.   |
| <b>VXLAN ID/VXLAN</b>     | VXLAN Network Identifier (VNI)   |
| <b>Layer 2 flags</b>      | Debugging flags signifying that the MAC address is present in various lists.   |
| <b>Epoch</b>              | Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.  |
| <b>Sequence number</b>    | Sequence number assigned to this MAC address. Used for debugging.  |
| <b>Learning mask</b>      | Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.  |
| <b>IPC generation</b>     | Creation time of the logical interface when this MAC address was learned. Used for debugging.  |

## Sample Output

### show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
  MAC          MAC      Logical
  address      flags    interface
  00:90:69:9c:1c:5d  D      ge-0/2/5.400

MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_red

```

```

VLAN : 401
MAC          MAC      Logical
address      flags    interface
00:00:aa:12:12:12 D      lsi.1051138
00:05:85:74:9f:f0 D      lsi.1051138

```

### show vpls mac-table (with VXLAN enabled)

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
VXLAN: Id : 300, Multicast group: 226.1.1.3
MAC          MAC      Logical
address      flags    interface
00:01:01:00:01:f4 D,SE ge-4/2/0.1000
00:02:01:33:01:f4 D,SE lsi.1052004
00:03:00:32:01:f4 D,SE lsi.1048840
00:04:00:14:01:f4 D,SE lsi.1052005
00:02:01:33:02:f7 D,SE vtep.1052010
00:04:00:14:02:f7 D,SE vtep.1052011

```

### show vpls mac-table count

```

user@host> show vpls mac-table count
0 MAC address learned in routing instance __juniper_private1__

```

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lc-0/0/0.32769    | 0         |
| lc-0/1/0.32769    | 0         |
| lc-0/2/0.32769    | 0         |
| lc-2/0/0.32769    | 0         |
| lc-0/3/0.32769    | 0         |
| lc-2/1/0.32769    | 0         |
| lc-9/0/0.32769    | 0         |
| lc-11/0/0.32769   | 0         |
| lc-2/2/0.32769    | 0         |
| lc-9/1/0.32769    | 0         |
| lc-11/1/0.32769   | 0         |
| lc-2/3/0.32769    | 0         |
| lc-9/2/0.32769    | 0         |
| lc-11/2/0.32769   | 0         |
| lc-11/3/0.32769   | 0         |
| lc-9/3/0.32769    | 0         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
| 0             | 0         |

1 MAC address learned in routing instance vpls\_ldp1

MAC address count per interface within routing instance:

| Logical interface | MAC count |
|-------------------|-----------|
| lsi.1051137       | 0         |
| ge-0/2/5.400      | 1         |

MAC address count per learn VLAN within routing instance:

| Learn VLAN ID | MAC count |
|---------------|-----------|
|---------------|-----------|

|   | 0         | 1 |
|---|-----------|---|
| 1 MAC address learned in routing instance vpls_red        |           |   |
| MAC address count per interface within routing instance:  |           |   |
| Logical interface   | MAC count |   |
| ge-0/2/5.300  | 1         |   |
| MAC address count per learn VLAN within routing instance: |           |   |
| Learn VLAN ID   | MAC count |   |
| 0   | 1         |   |

### show vpls mac-table detail

```

user@host> show vpls mac-table detail
MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_red
Learning interface: ge-0/2/5.300
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

```

### show vpls mac-table extensive

```

user@host> show vpls mac-table extensive
MAC address: 00:00:aa:12:12:12
Routing instance: vpls_ldp1
Learning interface: lsi.1051137
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0
Routing instance: vpls_ldp1
Learning interface: lsi.1051137
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:00:aa:12:12:12
Routing instance: vpls_red
Learning interface: lsi.1051138
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 0
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0

```

```
Routing instance: vpls_red
Learning interface: lsi.1051138
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 0
Learning mask: 0x1                       IPC generation: 0
```

## Verifying VXLAN Reachability

**Purpose** On the local VTEP, verify that there is connectivity with the remote VTEP.

**Action** `user@switch> show ethernet-switching vxlan-tunnel-end-point remote`

| Logical System Name | Id          | SVTEP-IP | IFL   | L3-Idx |
|---------------------|-------------|----------|-------|--------|
| <default>           | 0           | 10.1.1.2 | 1o0.0 | 0      |
| RVTEP-IP            | IFL-Idx     | NH-Id    |       |        |
| 10.1.1.2            | 559         | 1728     |       |        |
| VNID                | MC-Group-IP |          |       |        |
| 100                 | 232.1.1.1   |          |       |        |

**Meaning** The remote VTEP is reachable because its IP address appears in the output. The output also shows that the VXLAN (VNI 100) and corresponding multicast group are configured correctly on the remote VTEP.

- Related Documentation**
- [Understanding VXLANs on page 1912](#)
  - [Configuring VXLANs on a QFX5100 Switch on page 1991](#)
  - [Examples: Configuring VXLANs on QFX Series Switches on page 1944](#)

## Verifying That a Local VXLAN VTEP is Configured Correctly

**Purpose** Verify that a local VTEP is correct..

**Action** `user@switch> show ethernet-switching vxlan-tunnel-end-point source`

| Logical System Name | Id            | SVTEP-IP | IFL   | L3-Idx      |
|---------------------|---------------|----------|-------|-------------|
| <default>           | 0             | 10.1.1.1 | 1o0.0 | 0           |
| L2-RTT              | Bridge Domain |          | VNID  | MC-Group-IP |
| default-switch      | VLAN1+100     |          | 100   | 232.1.1.1   |

**Meaning** The output should show the correct tunnel source IP address (loopback address), VLAN, and multicast group for the VXLAN.

- Related Documentation**
- [Understanding VXLANs on page 1912](#)
  - [Configuring VXLANs on a QFX5100 Switch on page 1991](#)
  - [Examples: Configuring VXLANs on QFX Series Switches on page 1944](#)

## Verifying MAC Learning from a Remote VTEP

**Purpose** Verify that a local VTEP is learning MAC addresses from a remote VTEP.

**Action** user@switch> show ethernet-switching table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static

SE - statistics enabled, NM - non configured MAC, R - remote PE MAC)

Ethernet switching table : 2 entries, 2 learned

Routing instance : default-switch

| Vlan<br>name | MAC<br>address    | MAC<br>flags | Age | Logical<br>interface |
|--------------|-------------------|--------------|-----|----------------------|
| VLAN1        | 00:00:00:ff:ff:ff | D            | -   | vtep.12345           |
| VLAN1        | 00:10:94:00:00:02 | D            | -   | xe-0/0/0.0           |

**Meaning** This shows the MAC addresses learned from the remote VTEP (in addition to those learned on the normal Layer 2 interfaces). It also shows the logical name of the remote VTEP interface (**vtep.12345** in the above output).

- Related Documentation**
- [Understanding VXLANs on page 1912](#)
  - [Configuring VXLANs on a QFX5100 Switch on page 1991](#)
  - [Examples: Configuring VXLANs on QFX Series Switches on page 1944](#)

## Monitor a Remote VTEP Interface

**Purpose** Monitor traffic details for a remote VTEP interface.

**Action** user@switch> show interface *logical-name* detail

```

M   Flags: Up SNMP-Traps Encapsulation: ENET2
      VXLAN Endpoint Type: Remote, VXLAN Endpoint Address: 10.1.1.2, L2 Routing
Instance: default-switch, L3 Routing Instance: default
      Traffic statistics:
        Input bytes :          228851738624
        Output bytes :              0
        Input packets:          714162415
        Output packets:              0
      Local statistics:
        Input bytes :              0
        Output bytes :              0
        Input packets:              0
        Output packets:              0
      Transit statistics:
        Input bytes :          228851738624          0 bps
        Output bytes :              0          0 bps
        Input packets:          714162415          0 pps
        Output packets:              0          0 pps
      Protocol eth-switch, MTU: 1600, Generation: 277, Route table: 5

```

**Meaning** This shows traffic details for the remote VTEP interface. To get this information, you must supply the logical name of the remote VTEP interface (vtep.12345 in the above output), which you can learn by using the **show ethernet-switching table** command.

- Related Documentation**
- [Understanding VXLANs on page 1912](#)
  - [Configuring VXLANs on a QFX5100 Switch on page 1991](#)

- [Examples: Configuring VXLANs on QFX Series Switches on page 1944](#)

# Troubleshooting

- [Troubleshooting Procedures on page 2045](#)

## Troubleshooting Procedures

---

- [Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN on page 2045](#)

### Troubleshooting a Nonoperational VMware NSX Logical Switch and Corresponding Junos OS OVSDb-Managed VXLAN

**Problem**    **Description:** A logical switch configured by using VMware NSX Manager or the NSX API, and the corresponding Open vSwitch Database (OVSDb)-Managed Virtual Extensible LAN (VXLAN), which is configured on a Juniper Networks device, are not exchanging MAC addresses learned in the NSX and Junos OS environments, respectively. Also, the **Flags** field in the **show ovssdb logical-switch** operational mode command output is one of the following:

- **Created by Controller**
- **Created by L2ALD**
- **Tunnel key mismatch**

**Cause**

- If the **Flags** field displays **Created by Controller**, a logical switch is configured in NSX Manager or in the NSX API, but a corresponding VXLAN is not configured or is improperly configured on the Juniper Networks device.
- If the **Flags** field displays **Created by L2ALD**, a VXLAN is configured on the Juniper Networks device, but a corresponding logical switch is not configured in NSX Manager or in the NSX API.
- If the **Flags** field displays **Tunnel key mismatch**, the VXLAN network identifiers (VNIs) in the logical switch and corresponding VXLAN configurations do not match.

**Solution**    If the **Flags** field displays **Created by Controller**, take the following action:

- On a QFX5100 switch, verify that the **set switch-options ovssdb-managed** configuration command was issued in the Junos OS CLI. Issuing this command and committing the configuration enable the switch to automatically create OVSDb-managed VXLANs.

If this command was already issued, another possible cause is that the L2ALD daemon has become non-functional. If this is the case, wait for a few seconds, reissue the **show ovssdb logical-switch** operational mode command, and recheck the setting of the **Flags** field.

- On all other Juniper Networks devices that support VXLAN and OVSSDB, determine whether a corresponding VXLAN is configured on the device. If the VXLAN is not configured, configure it using the procedure in [“Configuring OVSSDB-Managed VXLANs” on page 1989](#). If a VXLAN is configured, check the VXLAN name to make sure that it is the same as the universally unique identifier (UUID) of the logical switch. Also, check the VNI to make sure that the value is the same as the value in the logical switch configuration.

If the **Flags** field displays **Created by L2ALD**, take the following action:

- On a QFX5100 switch, two issues exist. First, despite the fact that the switch automatically creates OVSSDB-managed VXLANs, this VXLAN was manually configured by using the Junos OS CLI. Second, a corresponding logical switch was not configured. To resolve both issues, configure a logical switch in NSX Manager or in the NSX API. After the NSX controller pushes relevant logical switch information to the switch, the switch automatically creates a corresponding VXLAN and deletes the manually configured VXLAN.
- On all other Juniper Networks devices that support VXLAN and OVSSDB, determine whether a corresponding logical switch is configured in NSX Manager or in the NSX API. If a logical switch is not configured, configure one, keeping in mind that NSX automatically generates a UUID for the logical switch and that this UUID must be used as the name of the VXLAN. That is, the VXLAN name must be reconfigured with the logical switch UUID. Another possibility is that the logical switch might exist, but the logical switch UUID might not be the VXLAN name. In NSX Manager or in the NSX API, check for a logical switch that has the same configuration as the VXLAN but has a different UUID.

If the **Flags** field displays **Tunnel key mismatch**, take the following action:

- For a QFX5100 switch, check the configuration of the VNI in NSX Manager or in the NSX API to see whether it was changed after the switch created the corresponding VXLAN. If it was changed, update the VNI on the QFX5100 switch, using the Junos OS CLI.
- On all other Juniper Networks devices that support VXLAN and OVSSDB, check the values of the VNI in NSX Manager or in the NSX API and the Junos OS CLI, and change the incorrect value.

#### Related Documentation

- [Understanding How to Set Up Virtual Extensible LANs in an Open vSwitch Database Environment on page 1905](#)
- [show ovssdb logical-switch on page 2019](#)



## PART 7

# OpenFlow

- [Overview on page 2049](#)
- [Installing Support for OpenFlow on page 2097](#)
- [OpenFlow Basic Configuration on page 2099](#)
- [Configuring OpenFlow Hybrid Interfaces on page 2127](#)
- [Configuring OpenFlow Traffic Steering Across MPLS Networks on page 2149](#)
- [Configuration Statements on page 2169](#)
- [Operational Commands on page 2183](#)



## CHAPTER 22

# Overview

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS on page 2051](#)
- [Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS on page 2056](#)
- [Understanding the OpenFlow Version Negotiation Between the Controller and Junos OS Devices on page 2057](#)
- [Understanding OpenFlow Flows and Filters on Devices Running Junos OS on page 2058](#)
- [Understanding OpenFlow Flow Instructions on Devices Running Junos OS on page 2060](#)
- [Understanding How the OpenFlow Group Action Works on page 2061](#)
- [Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS on page 2062](#)
- [Understanding OpenFlow Barrier Messages on Devices Running Junos OS on page 2064](#)
- [Understanding OpenFlow Multipart Messages on Devices Running Junos OS on page 2064](#)
- [OpenFlow v1.0 Compliance Matrix for Devices Running Junos OS on page 2065](#)
- [OpenFlow v1.0 Compliance Matrix for QFX5100 Switches on page 2071](#)
- [OpenFlow v1.0 Compliance Matrix for EX4550 Switches on page 2078](#)
- [OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS on page 2087](#)

## Understanding Support for OpenFlow on Devices Running Junos OS

---

- [OpenFlow Overview on page 2049](#)
- [OpenFlow Virtual Switches on page 2050](#)
- [OpenFlow Interfaces on page 2050](#)

### OpenFlow Overview

OpenFlow is an open standard that enables you to control traffic and run experimental protocols in an existing network using a remote controller. The OpenFlow components consist of a controller, an OpenFlow or OpenFlow-enabled switch, and the OpenFlow protocol. The OpenFlow protocol is a Layer 2 protocol that permits an OpenFlow controller access to the data plane of an OpenFlow-enabled switch over an SSL or TCP/IP connection.

Using OpenFlow, you can control traffic paths in a network by creating, deleting, and modifying flows in each device along a path. Flow entries specify match conditions against which packets are compared, and a set of actions (OpenFlow v1.0) or instructions (OpenFlow v1.3.1) that are applied to matching packets.

You can configure certain devices running the Junos operating system (Junos OS) as OpenFlow-enabled switches. The Junos OS process, `openflowd (ofd)`, handles OpenFlow functionality on these devices. When implementing OpenFlow in an existing network, you must isolate experimental flows from production flows so that normal network traffic is not impacted. On devices running Junos OS, you isolate OpenFlow traffic by configuring one or more virtual switches that act as logically separate flood domains. The virtual switch and controller communicate by exchanging OpenFlow protocol messages, which the controller uses to add, delete, and modify flows on the switch.

For more information about OpenFlow, see the Open Networking Foundation website at <https://www.opennetworking.org/>.

## OpenFlow Virtual Switches

To isolate and control OpenFlow traffic on devices running Junos OS, you configure virtual switches. Each virtual switch configuration contains the controller connection information, the set of logical interfaces participating in OpenFlow, and the default action executed when a packet does not match any existing flow entry. You configure the OpenFlow protocol and OpenFlow virtual switches at the **[edit protocols openflow]** hierarchy level.

Depending on the platform, a default VLAN or bridge domain is assigned to each virtual switch. This VLAN or bridge domain acts as a logically separate flood domain, which isolates OpenFlow traffic from normal traffic. On certain platforms, you must also configure a separate virtual switch routing instance at the **[edit routing-instances]** hierarchy level.

You can configure a single OpenFlow virtual switch on devices running Junos OS that support OpenFlow, and you can configure one controller connection per virtual switch. By default, if you configure a virtual switch with a single controller, the controller is in active mode. If a controller is in active mode, the switch automatically initiates a connection to the controller.

## OpenFlow Interfaces

When you configure an OpenFlow virtual switch on a device running Junos OS, you must specify the logical interfaces that are participating in OpenFlow for that virtual switch instance. OpenFlow traffic can only either enter or exit from OpenFlow-enabled interfaces, and MAC address learning is disabled on these interfaces.

Interfaces participating in OpenFlow must be configured as Layer 2 interfaces. To configure the interface as OpenFlow-enabled, you add the logical interface to the OpenFlow virtual switch configuration at the **[edit protocols openflow switch switch-name interfaces]** hierarchy level. An OpenFlow interface can only be configured under a single virtual switch. On platforms that require a separate virtual switch routing instance for OpenFlow traffic, you must also configure the OpenFlow interfaces under the OpenFlow virtual switch routing instance.

On certain devices running Junos OS, you can only configure a single logical unit using logical unit number 0 on an OpenFlow interface. However, on certain platforms that support OpenFlow, a single physical interface can be configured as a hybrid interface that supports both OpenFlow and non-OpenFlow logical interfaces. For example, you could configure interface ge-1/0/1 to have two logical interfaces ge-1/0/1.0 and ge-1/0/1.1, where ge-1/0/1.0 does not participate in OpenFlow, and ge-1/0/1.1 is an OpenFlow-enabled interface.

#### Related Documentation

- [OpenFlow Support on Devices Running Junos OS](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
- [Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS on page 2056](#)
- [Understanding OpenFlow Flows and Filters on Devices Running Junos OS on page 2058](#)
- [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS on page 2051](#)
- [OpenFlow v1.0 Compliance Matrix for QFX5100 Switches on page 2071](#)
- [OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS on page 2087](#)

## Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS

This topic summarizes the OpenFlow features and forwarding actions supported on devices running Junos OS. For detailed information about support for specific OpenFlow v1.0 messages and fields, match conditions, wildcards, flow actions, statistics, and features, see “[OpenFlow v1.0 Compliance Matrix for Devices Running Junos OS](#)” on [page 2065](#). For a detailed list of supported OpenFlow v1.3.1 messages and fields, port structure flags and numbering, match conditions, flow actions, multipart messages, flow instructions, and group types, see “[OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS](#)” on [page 2087](#).

- [OpenFlow Operation and Support on page 2051](#)
- [OpenFlow Forwarding Actions on page 2054](#)

### OpenFlow Operation and Support

To isolate and control OpenFlow traffic on devices running Junos OS, you configure virtual switches. You can configure one OpenFlow virtual switch and one active OpenFlow controller on each device running Junos OS that supports OpenFlow. You configure the OpenFlow protocol, virtual switch, and controller connection information at the **[edit protocols openflow]** hierarchy level.

OpenFlow traffic can only either enter or exit from OpenFlow-enabled ports. If a flow modification message is sent to an ingress port that is not enabled for OpenFlow, the device sends an ofp\_error\_msg with an OFPET\_FLOW\_MOD\_FAILED error type and OFPFMFC\_UNKNOWN code to the controller. If a flow modification action is requested

for a port that is not enabled for OpenFlow, the device sends an ofp\_error\_msg with an OFPET\_BAD\_ACTION error type and OFPBAC\_BAD\_OUT\_PORT code to the controller.

[Table 148 on page 2052](#) summarizes the general features supported on devices running Junos OS that support OpenFlow v1.0. For information about support on specific platforms, see *OpenFlow Support on Devices Running Junos OS*.

**Table 148: OpenFlow v1.0 Support on Devices Running Junos OS**

| Feature   | Support  |
|---|--|
| OpenFlow v1.0   | Supported.   |
| OpenFlow virtual switch                                       | 1 OpenFlow virtual switch.   |
| Controller  | 1 active OpenFlow controller per virtual switch. Tested controllers include Floodlight and OESS.   |
| Controller connection   | TCP/IP connection. Only passive connections are accepted. The controller cannot actively connect to the OpenFlow switch.<br><br>SSL connections are not supported.   |
| Emergency mode  | Not supported as defined in OpenFlow Switch Specification v1.0. If the controller connection is lost and cannot be reestablished, the switch maintains all flow states in the control and data planes.   |
| Flow classification and mapping as a Layer 2 or Layer 3 route | Not supported.   |
| Flow priority   | Supported as per OpenFlow Switch Specification v1.3 in which there is no prioritization of exact match entries over wildcard entries.  |
| Flow table  | Single flow table.   |
| Forwarding actions  | <ul style="list-style-type: none"> <li>Forward to an OpenFlow-enabled physical port</li> <li>ALL, CONTROLLER, NORMAL, and FLOOD for normal flow actions</li> <li>ALL and FLOOD for Send Packet flow actions</li> </ul> <p><b>NOTE:</b> The QFX5100 switch does not support NORMAL for normal flow actions.</p> |
| Hybrid interfaces   | Supported on some devices. OpenFlow-enabled devices that support hybrid interfaces permit a physical interface to concurrently support logical interfaces for normal traffic and logical interfaces for OpenFlow traffic.  |
| Interfaces  | You can configure Ethernet interfaces only as OpenFlow interfaces.   |
| Multi-VLAN actions  | Supported on some devices. OpenFlow-enabled devices that support multi-VLAN actions have the ability to associate a different VLAN and different VLAN action with each egress port.  |

**Table 148: OpenFlow v1.0 Support on Devices Running Junos OS (continued)**

| Feature                                    | Support  |
|--|--|
| Port modification                          | Not supported. OpenFlow-enabled devices ignore all OpenFlow controller OFPT_PORT_MOD requests. |
| Queues, queue messages, or enqueue actions | Not supported.   |

[Table 149 on page 2053](#) summarizes the general features supported on devices running Junos OS that support OpenFlow v1.3.1. For information about support on specific platforms, see *OpenFlow Support on Devices Running Junos OS*.

**Table 149: OpenFlow v1.3.1 Support on Devices Running Junos OS**

| Feature   | Support  |
|---|--|
| OpenFlow v1.3.1   | Supported.   |
| OpenFlow virtual switch                                       | 1 OpenFlow virtual switch.   |
| Controller  | 1 active OpenFlow controller per virtual switch. Tested controllers include NEC and Ixia. .  |
| Controller connection   | TCP/IP connection. Only passive connections are accepted. The controller cannot actively connect to the OpenFlow switch.<br><br>SSL connections are not supported.   |
| Flow classification and mapping as a Layer 2 or Layer 3 route | Not supported.   |
| Flow priority   | Supported as per OpenFlow Switch Specification v1.3 in which there is no prioritization of exact match entries over wildcard entries.  |
| Flow instructions   | For each flow entry, 1 flow instruction is supported. A flow instruction can be one of the following: <ul style="list-style-type: none"> <li>• OFPIT_APPLY_ACTIONS</li> <li>• OFPIT_WRITE_ACTIONS</li> </ul> |
| Flow table  | Single flow table.   |

**Table 149: OpenFlow v1.3.1 Support on Devices Running Junos OS (continued)**

| Feature                                    | Support  |
|--|--|
| Forwarding actions                         | <ul style="list-style-type: none"> <li>Forward to an OpenFlow-enabled physical port</li> <li>ALL, CONTROLLER, NORMAL, and FLOOD for normal flow actions</li> <li>ALL and FLOOD for Send Packet flow actions</li> </ul> <p><b>NOTE:</b> The QFX5100 switch does not support NORMAL for normal flow actions.</p> |
| Group action                               | <p>Supported. A group can include 1 to 32 buckets, and a bucket can have a set of actions (set, pop, or output).</p> <p>Group types OFPGT_ALL and OFPGT_INDIRECT are supported.</p>  |
| Interfaces                                 | You can configure Ethernet interfaces only as OpenFlow interfaces.   |
| Multi-VLAN actions                         | Supported on some devices. OpenFlow-enabled devices that support multi-VLAN actions have the ability to associate a different VLAN and different VLAN action with each egress port.  |
| Multipart messages                         | <p>Supported for requesting and returning the following information:</p> <ul style="list-style-type: none"> <li>Switch, group, or port descriptions</li> <li>Single-flow, aggregate-flow, flow table, port, or group statistics</li> <li>Group or table features</li> </ul>                                    |
| OpenFlow version negotiation               | Supported for OpenFlow controller and Junos OS device OpenFlow version negotiation.  |
| Port modification                          | Not supported. OpenFlow-enabled devices ignore all OpenFlow controller OFPT_PORT_MOD requests.   |
| Queues, queue messages, or enqueue actions | Not supported.   |

## OpenFlow Forwarding Actions



**NOTE:** The information in this section applies to both OpenFlow v1.0 and v1.3.1 except where noted.



OpenFlow-enabled devices running Junos OS support several flow actions for forwarding OpenFlow packets. For normal flow actions, the following forwarding actions are supported:

- **physical port**—Forward unicast or multicast packets out the specified OpenFlow-enabled interfaces.
- **ALL**—Flood the packet out all OpenFlow interfaces configured for that virtual switch instance except the ingress interface.
- **CONTROLLER**—Send the packet to the OpenFlow controller for processing.
- **FLOOD**—Flood the packet along the minimum spanning tree, which includes all OpenFlow interfaces configured for that virtual switch instance except the ingress interface and any interfaces that are disabled by the Spanning Tree Protocol (STP). Because devices running Junos OS do not support 802.1D STP capabilities for OpenFlow, the FLOOD forwarding action behaves like the ALL forwarding action.
- **NORMAL**—Process the packet, using traditional Layer 2 or Layer 3 processing.



**NOTE:** The QFX5100 switch does not support NORMAL for normal flow actions.

The OpenFlow controller can also use a Send Packet message (OFPT\_PACKET\_OUT) to direct the OpenFlow virtual switch to send a packet out of a specified port. The Send Packet message includes the packet to be forwarded and the forwarding action indicating the interface out of which the packet must be forwarded. Supported forwarding actions for the Send Packet message include ALL and FLOOD.

Each OpenFlow virtual switch is a logically separate flood domain. Therefore, the OpenFlow ALL and FLOOD actions only flood packets out OpenFlow interfaces configured under that specific virtual switch excluding the ingress OpenFlow interface.

#### Related Documentation

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [OpenFlow v1.0 Compliance Matrix for QFX5100 Switches on page 2071](#)
- [OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS on page 2087](#)
- [Understanding How the OpenFlow Group Action Works on page 2061](#)
- [Understanding OpenFlow Flow Instructions on Devices Running Junos OS on page 2060](#)
- [Understanding OpenFlow Multipart Messages on Devices Running Junos OS on page 2064](#)
- [Understanding the OpenFlow Version Negotiation Between the Controller and Junos OS Devices on page 2057](#)
- [Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS on page 2056](#)
- [Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS on page 2062](#)

## Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS

---

On devices running the Junos operating system (Junos OS), each OpenFlow virtual switch establishes an independent connection with the controller and is represented by a unique runtime datapath ID consisting of the management port MAC address and an internally assigned virtual switch ID. The controller and virtual switch connect to each other using a TCP/IP connection on the management plane. Thus, OpenFlow-enabled devices that are managed by a controller must be connected to the management network (for example, connected using the `me0`, `fxp0`, `em0`, or `em1` management port) and must be reachable from the controller IP address.

Upon establishing a connection with the controller, the switch and the controller exchange hello messages that specify the highest OpenFlow protocol version supported by the sender. If the first packet received by the switch is not an `OFPT_HELLO` message, the switch tears down the connection and attempts to establish a new connection with the controller. Additionally, if the controller and the switch negotiate an OpenFlow protocol version that one of the parties does not support, the connection is terminated with an error message indicating an `OFPET_HELLO_FAILED` error type and an `OFPHFC_INCOMPATIBLE` code.

The session is established when the switch and controller successfully exchange Hello messages and negotiate the OpenFlow protocol version. After session establishment, the controller sends the switch a feature request message to request the capabilities supported by the switch. The switch responds with a feature reply message, which includes the local MAC address in the virtual switch datapath ID field. If the local MAC address is unavailable, the switch terminates the connection.

After establishing the session, the controller and virtual switch exchange echo request and reply messages as a keepalive mechanism. The keepalive timer is reset if the virtual switch or controller receives either an echo reply or a packet. Echo requests are sent every 10 seconds during idle windows in the absence of other messages. If the switch receives no echo reply or other message from the controller for 120 seconds, the connection is considered lost, and the switch attempts to reestablish the connection for 10 seconds. If the connection cannot be established, the switch enters emergency mode as defined in the OpenFlow v1.3 specification. In emergency mode, the switch deletes normal flow entries, and after 30 seconds, purges flow entries that are installed in hardware.

If at any point after the session is established the recipient receives an OpenFlow message that specifies the wrong OpenFlow version, the recipient responds with an error message indicating an `OFPET_BAD_REQUEST` type and `OFPBRC_BAD_VERSION` code. If the switch cannot process the version and type of an OpenFlow packet in the TCP buffer, or if the switch fails sending OpenFlow messages to the controller, the switch terminates the connection.

Modifying, deleting, or deactivating the virtual switch configuration also impacts the connection to the controller. If you modify an existing virtual switch configuration, the virtual switch tears down the existing connection to the controller and establishes a new session with the updated configuration information. If you delete or deactivate an existing

virtual switch configuration, the virtual switch automatically disconnects from the controller.

To summarize, the switch disconnects from the controller under the following circumstances:

- The first packet the switch receives from the controller is not a hello message.
- The switch receives a hello message with an unsupported OpenFlow version.
- The local MAC address is not available for inclusion in the feature reply message.
- The switch receives no echo reply or other message from the controller for 120 seconds.
- The existing virtual switch configuration is deleted or deactivated.
- The existing virtual switch configuration is modified. In this case, after disconnecting from the controller, the switch attempts to establish a new connection and session.
- The switch cannot process the version and type of an OpenFlow packet in the TCP buffer.
- The switch fails to send OpenFlow messages to the controller, which is treated as a dead TCP socket connection.

**Related  
Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding the OpenFlow Version Negotiation Between the Controller and Junos OS Devices on page 2057](#)

## Understanding the OpenFlow Version Negotiation Between the Controller and Junos OS Devices

Upon establishing an initial connection, an OpenFlow controller and a Junos OS device negotiate the OpenFlow version to be used. In general, the OpenFlow controller must support at least one of the versions run on the Junos OS device. Otherwise, a connection cannot be established.



**NOTE:** The Junos OS implementation of OpenFlow 1.3.1 does not support the OFPHET\_VERSIONBITMAP Hello message element.

[Table 150 on page 2057](#) outlines the OpenFlow versions run by the Junos OS device and controller, the negotiated version, and the *numerical value* associated with each version.

**Table 150: OpenFlow Versions Negotiated Between the Controller and a Junos OS Device and the Numerical Value Associated with Each Version**

| OpenFlow Version Run by Junos OS Device | OpenFlow Version Supported by Controller | Negotiated Version | Numerical Value Associated with Negotiated OpenFlow Version |
|---|--|--------------------|---|
| 1.0                                     | 1.0                                      | 1.0                | 1   |

**Table 150: OpenFlow Versions Negotiated Between the Controller and a Junos OS Device and the Numerical Value Associated with Each Version (*continued*)**

| OpenFlow Version Run by Junos OS Device | OpenFlow Version Supported by Controller   | Negotiated Version  | Numerical Value Associated with Negotiated OpenFlow Version |
|---|--|---------------------|---|
| 1.3.1                                   | 1.3.1  | 1.3.1               | 4   |
| 1.0 and 1.3.1                           | 1.0 and 1.3.1  | 1.3.1               | 4   |
| 1.0 and 1.3.1                           | 1.0  | 1.0                 | 1   |
| 1.0 and 1.3.1                           | 1.3.1  | 1.3.1               | 4   |
| 1.0 and/or 1.3.1                        | <ul style="list-style-type: none"> <li>Neither 1.0 nor 1.3.1</li> <li>Connection with Junos OS device is down</li> </ul> | None; no connection | 0   |

To determine the negotiated version running on a Junos OS device, you enter the [show openflow controller](#) command. The output of this command includes a **Negotiated version** field and a numerical value that represents the negotiated version number. Use [Table 150 on page 2057](#) to correlate the numerical values shown in this field with the negotiated versions.

- Related Documentation**
- [Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS on page 2056](#)

## Understanding OpenFlow Flows and Filters on Devices Running Junos OS

OpenFlow flows are defined by various elements. [Table 151 on page 2058](#) outlines the flow elements supported in OpenFlow v1.0 and v1.3.1. The elements supported by the OpenFlow versions uniquely identify a flow.

**Table 151: OpenFlow Flow Elements**

| Flow Element                | Supported In OpenFlow v1.0? | Supported In OpenFlow v1.3.1? |
|-----------------------------|-----------------------------|-------------------------------|
| Match conditions            | Yes                         | Yes                           |
| Set of actions              | Yes                         | No                            |
| Flow instructions           | No                          | Yes                           |
| Flow priority               | Yes                         | Yes                           |
| Flow timeout information    | Yes                         | Yes                           |
| Flow cookie and cookie mask | No                          | Yes                           |

Flow entries specify wildcard match conditions for fields that do not require an exact match. If a flow entry contains wildcards for all match conditions, then all packets match that flow entry.

To implement OpenFlow flow-based forwarding, devices running Junos OS use filters. For each logical interface configured to participate in OpenFlow, a single filter is created and applied to the logical interface in the input direction. The filter name is the concatenation of the interface name, including the logical unit number, and an internally assigned virtual switch ID, for example `ge-1/1/0.0_0`.



**NOTE:** If you manually configure a filter name or a filter term name that conflicts with an autogenerated OpenFlow filter name or filter term name, Junos OS does not generate an error during a **commit** check. If there is a conflict, the commit succeeds, but one of the filters or filter terms is rejected based on the order in which they were received.

A filter consists of one or more terms with match conditions, and actions (for OpenFlow v1.0) or instructions (for OpenFlow v1.3.1). OpenFlow flows are mapped to filter terms, and OpenFlow controller requests to add, delete, and modify flows result in the addition, deletion, or modification of terms in the filter. When the OpenFlow controller sends a flow modification request, the flow entry match condition for the ingress port determines which logical interface filter is updated. The OpenFlow flow priority determines the order of the terms in the filter, where higher priority terms are installed above lower priority terms. Flow match conditions are mapped to the filter term match conditions, and flow actions or instructions are mapped to the filter term **then** statement. Depending on the flow action or instruction, the **then** statement might include actions for forwarding the packet to the next hop or OpenFlow controller, or discarding the packet.



**NOTE:** If the OpenFlow controller sends a request to modify a flow, but no flow entries match the conditions, OpenFlow v1.0 adds an entry for the flow to the flow table. However, in the same situation, OpenFlow v1.3.1 does not add this flow to the flow table, nor is an error logged.

Each OpenFlow flow entry corresponds to a filter term. However, each flow entry might map to a term in one or more filters depending on the match condition for the ingress port. If the ingress port is a wildcard match, the flow entry appears as a term in all of the interface filters for that OpenFlow virtual switch. For example, suppose that the OpenFlow controller sends a request to add a new flow entry with a wildcard match for the ingress port field. In this case, the flow is added as a new filter term for all OpenFlow logical interfaces configured under that virtual switch.

Devices running Junos OS support both strict and non-strict flow mod commands for modifying and deleting flows. OpenFlow controller strict modify and strict delete flow mod requests only modify or delete flows that exactly match the description for all header fields including wildcards and priorities. Non-strict modify and delete flow mod requests modify or delete flows that exactly match or are more specific than the request.

In addition to the functionality already described, OpenFlow v1.3.1 supports a flow cookie, which is an identifier that the OpenFlow controller can specify when a flow is installed in the flow table. This cookie enables OpenFlow to filter flows selected for flow modification and delete operations.

You can configure the default action for packets that do not match on any flow entry as either **drop**, which discards the packet, or **packet-in**, which accepts the packet and forwards it to the controller. The default action is specific to the OpenFlow virtual switch and is the same across all filters associated with that virtual switch. If you do not explicitly configure the default action, the default is **packet-in**.

In the event that a logical interface becomes unavailable, the filter associated with that logical interface is removed from the Packet Forwarding Engine. Although the filter is removed, the Routing Engine retains flows that match on the logical interface as the ingress port until such time as the flows are purged in response to OpenFlow timers. For information about OpenFlow timers, see [“Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS” on page 2062](#). If the logical interface becomes available before the flows are purged, the filter and any flows retained by the Routing Engine at that point are reinstalled in hardware.

Similarly, when a logical interface becomes unavailable, flows that have that logical interface as the only active egress interface in their action set or instruction are considered invalid. The invalid flows are removed from the Packet Forwarding Engine but are indefinitely retained by the Routing Engine until the flows are purged in response to various OpenFlow timers. Alternatively, flows that include the logical interface as one of several active egress interfaces in their action set or instruction are still valid. In that case, the flow remains in the Packet Forwarding Engine, but the multicast next hop is updated to remove that logical interface as a valid egress interface.

**Related  
Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Flow Instructions on Devices Running Junos OS on page 2060](#)

---

## Understanding OpenFlow Flow Instructions on Devices Running Junos OS

---



**NOTE:** Flow instructions are supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.

When a packet matches a particular OpenFlow flow, a Juniper Networks device running OpenFlow v1.0 applies a set of actions to the packet. Instead of applying a set of actions, starting with OpenFlow v1.3.1, the Juniper Networks device applies a flow instruction to a matching packet.

In the Junos OS implementation of OpenFlow v1.3.1, a flow entry can include only one flow instruction, which can be one of the following:

- Apply actions (OFPIT\_APPLY\_ACTIONS)
- Write actions (OFPIT\_WRITE\_ACTIONS)

Each of the instructions mentioned above includes a list of actions that the device applies immediately in the order in which they appear the list.

**Related  
Documentation**

- [OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS on page 2087](#)

## Understanding How the OpenFlow Group Action Works



**NOTE:** The group action is supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.

OpenFlow uses flow entries as a means to match flows and specify an action for incoming packets on logical OpenFlow interfaces. The action specified in one or more flow entries can direct packets to, or reference, a base action called a *group* action. The purpose of the group action is to further process these packets and assign a more specific forwarding action to them.

A group can include 1 to 32 buckets, and in turn, a bucket can have a set of actions (set, pop, or output).

For information about the specific actions that are supported for each base type, see the [“OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS” on page 2087](#).

Junos OS devices support the following *group types*, which define how buckets are implemented:

- All—Multiple buckets are implemented for the handling of multicast and broadcast packets. Each incoming packet is replicated and processed by each bucket in the group.
- Indirect—One bucket is implemented. An indirect group is typically referenced by multiple flow entries, thereby allowing each of these entities to have a centralized action that can be easily updated.

For example, an all group type with a unique OpenFlow-controller-assigned identifier, say, 50 can have two buckets: bucket 1 and bucket 2. The action associated with bucket 1 might be to set the VLAN ID field in the packet to 3022 and to output the packet to an OpenFlow port—for example, 118. The action associated with bucket 2 might be to set the VLAN ID field in the packet to—for example, 2022—and to output the packet to an OpenFlow port—for example, 117.

You can add a group with one or more buckets on the OpenFlow controller, and the controller pushes the group to the Junos OS devices with which it is connected. Each Junos OS device checks to see whether the group already exists. If it does not, the group is added to the group table on the Junos OS devices. After the group is in the group table, you can modify or delete it from the table by way of the OpenFlow controller.

**Related  
Documentation**

- [show openflow groups on page 2201](#)
- [show openflow statistics groups on page 2211](#)

## Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS

- [OpenFlow Flow Entry Timer Overview on page 2062](#)
- [Idle Timeout and Hard Timeout on page 2062](#)
- [Purge Flow Timer on page 2063](#)

### OpenFlow Flow Entry Timer Overview

For each logical interface participating in OpenFlow on a device running Junos OS, a single filter is created and applied to the logical interface in the input direction. OpenFlow flows are mapped to the filter as filter terms. Each flow has a number of timers associated with it, some of which are configured through the OpenFlow controller while others are configured through the Junos OS CLI. OpenFlow flow entry timers include the idle timeout, the hard timeout, and the purge flow timer. [Table 152 on page 2062](#) summarizes the various OpenFlow flow timers.

**Table 152: OpenFlow Flow Entry Timers**

| Timer            | Configured Through  | Range (Seconds)      |
|------------------|---|----------------------|
| Idle timeout     | Controller  | 0, 11 through 65,535 |
| Hard timeout     | Controller  | 0 through 65,535     |
| Purge flow timer | Junos OS CLI by using the <b>purge-flow-timer</b> configuration statement | 0 through 300        |

### Idle Timeout and Hard Timeout

Each flow entry has an idle timeout and a hard timeout associated with it, both of which are configured through the OpenFlow controller. The idle timeout is the number of seconds after which the flow is removed from the flow table and the hardware provided there are no matching packets. The hard timeout is the number of seconds after which the flow is removed from the flow table and the hardware regardless of the number of matching packets.

If a flow entry has both an idle timer and a hard timer associated with it, the first timer to expire causes the flow entry to be removed. If the idle timer expires first, the flow is removed at that point only if there are no matching packets. Otherwise, the flow is removed when the hard timer expires.

When the controller sends a flow entry modification message (OFPT\_FLOW\_MOD) to the switch, it specifies the idle timeout and hard timeout for that flow entry. On devices running Junos OS, the idle timeout value can be 0, or it can range from 11 through 65,535 seconds. If the controller sets the idle timeout to 0, the flow entry does not experience an idle time out. The hard timeout value can range from 0 through 65,535 seconds. If the controller sets the hard timeout to 0, the flow entry does not experience a hard time out. If the controller requests an invalid timeout value, the switch rejects the flow modification message and sends an error message back to the controller.



## Purge Flow Timer

On devices running Junos OS, you can configure a purge flow timer, which is the number of seconds after which an invalid OpenFlow flow entry is deleted from the flow table. The **purge-flow-timer** statement is configured through the Junos OS CLI at the **[edit protocols openflow switch switch-name]** hierarchy level. The **purge-flow-timer** value is specific to the OpenFlow virtual switch under which it is configured, and it is the same for all flow entries associated with that virtual switch.

If you do not configure the **purge-flow-timer** statement, the device purges invalid flows from hardware, but indefinitely retains the corresponding flow entries in the flow table on the Routing Engine. If you configure the **purge-flow-timer** statement, the device purges invalid flow entries from hardware, and after the specified number of seconds, deletes the invalid flow entries from the flow table. Configuring a value of 0 causes the device to immediately delete invalid flow entries from the flow table.

For example, consider the case of an OpenFlow logical interface that becomes temporarily unavailable. When the interface becomes unavailable, flow entries that have the logical interface as the matching ingress interface or as the only active egress interface in their action set (for OpenFlow v1.0) or flow instruction (for OpenFlow v1.3.1) are marked as invalid. Although the logical interface is not available, the flow entries could still be valid. The **purge-flow-timer** configuration statement determines how to handle the flows.

In this example, if you do not configure the **purge-flow-timer** statement, then when the logical interface becomes unavailable, the device removes the invalid flows from the hardware but indefinitely retains the flow entries in the flow table. If the logical interface later becomes available, the flows are reinstalled in the hardware without any controller intervention.

On the other hand, if you configure the **purge-flow-timer** statement, then when the logical interface becomes unavailable, the device removes the flows from the hardware, and retains the flow entries in the flow table for the configured number of **purge-flow-timer** seconds. If the interface does not become available and the timer expires, the device deletes the flow entries from the flow table. After the interface comes back up, the OpenFlow controller must send new flow entry modification messages to the OpenFlow switch in order to restore the flow entries to the flow table and to the hardware.



**NOTE:** By default, if you remove an active OpenFlow logical interface from an existing OpenFlow configuration, flow entries that match on this logical interface as the ingress interface and flow entries that include this logical interface as the only active egress interface in their action list or flow instruction are invalid and are automatically purged from the flow table and from the hardware regardless of whether you configure the **purge-flow-timer** statement.

### Related Documentation

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [purge-flow-timer on page 2178](#)

## Understanding OpenFlow Barrier Messages on Devices Running Junos OS

---

OpenFlow-enabled devices running Junos OS support the OpenFlow protocol controller-to-switch Barrier Request message (OFPT\_BARRIER\_REQUEST). The OpenFlow controller sends a Barrier Request message to request that the OpenFlow-enabled switch complete processing of all messages sent before the Barrier Request message before processing any messages sent after the Barrier Request message. This ensures that the virtual switch processes all message dependencies and sends all notifications for completed operations before proceeding with new requests.

When the OpenFlow virtual switch receives a Barrier Request message, it queues all subsequent incoming messages, with the exception of echo request and reply messages, until processing of all prior messages is complete. Echo request and reply messages are required to maintain connectivity to the controller.

When the switch completes an operation, it sends a reply message back to the controller. Only after the reply is sent to the controller does the switch mark the message or operation as processed. After the switch completes processing for all operations requested prior to the Barrier Request message, the switch sends a Barrier Reply (OFPT\_BARRIER\_REPLY) message, which includes the ID of the original request message, to the OpenFlow controller. At that point, the switch resumes processing of the queued messages.

### Related Documentation

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Flows and Filters on Devices Running Junos OS on page 2058](#)

## Understanding OpenFlow Multipart Messages on Devices Running Junos OS

---



**NOTE:** Multipart messages are supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.

To more efficiently process large OpenFlow data responses, OpenFlow v1.3.1 introduces support for multipart messages.

The OpenFlow controller can request the following information, using a multipart request message:

- Switch, group, or port descriptions
- Single-flow, aggregate-flow, flow table, port, or group statistics
- Group or table features

In response, a Juniper Networks device can send one or more multipart response messages wherein each message includes the same request identifier. In addition, each message in the sequence, except the last message, includes a flag that indicates more messages are to follow.

**Related  
Documentation**

- [OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS on page 2087](#)

## OpenFlow v1.0 Compliance Matrix for Devices Running Junos OS

The following tables list the support for OpenFlow v1.0 messages and fields, match conditions, wild cards, flow actions, statistics, and features on the indicated platforms.

- [Table 153 on page 2065](#) lists support for message types.
- [Table 154 on page 2067](#) lists support for port structure flags.
- [Table 155 on page 2068](#) lists support for match conditions.
- [Table 156 on page 2069](#) lists support for wildcards.
- [Table 157 on page 2069](#) lists support for flow actions.
- [Table 158 on page 2070](#) lists support for flow actions in Send Packet messages (OFPT\_PACKET\_OUT).
- [Table 159 on page 2071](#) lists support for statistics.
- [Table 160 on page 2071](#) lists support for features.

[Table 153 on page 2065](#) lists the support for OpenFlow v1.0 message types.

**Table 153: Junos OS Support for OpenFlow v1.0 Message Types**

| Section | Specification         | MX Series     | EX9200        |
|---------|-----------------------|---------------|---------------|
| 5.1     | OFPT_HELLO            | Supported     | Supported     |
|         | OFPT_ERROR            | Supported     | Supported     |
|         | OFPT_ECHO_REQUEST     | Supported     | Supported     |
|         | OFPT_ECHO_REPLY       | Supported     | Supported     |
|         | OFPT_VENDOR           | Not supported | Not supported |
|         | OFPT_FEATURES_REQUEST | Supported     | Supported     |
|         | OFPT_FEATURES_REPLY:  | Supported     | Supported     |
|         | Datapath ID           | Supported     | Supported     |
|         | N_buffers             | -1            | -1            |
|         | N_tables              | 1             | 1             |
|         | OFPC_FLOW_STATS       | Supported     | Supported     |
|         | OFPC_TABLE_STATS      | Supported     | Supported     |
|         | OFPC_PORT_STATS       | Supported     | Supported     |
|         | OFPC_STP              | Not supported | Not supported |
|         | OFPC_IP_REASM         | Not supported | Not supported |
|         | OFPC_QUEUE_STATS      | Supported     | Supported     |
|         | OFPC_ARP_MATCH_IP     | Not supported | Not supported |

**Table 153: Junos OS Support for OpenFlow v1.0 Message Types (continued)**

| Section | Specification  | MX Series     | EX9200        |
|---------|--|---------------|---------------|
|         | OFPT_GET_CONFIG_REQUEST  | Supported     | Supported     |
|         | OFPT_GET_CONFIG_REPLY  | Supported     | Supported     |
|         | OFPT_SET_CONFIG  | Supported     | Supported     |
|         | OFPT_PACKET_IN   | Supported     | Supported     |
|         | OFPT_PACKET_IN with buffer_id                                  | Not supported | Not supported |
|         | OFPT_FLOW_REMOVED  | Supported     | Supported     |
|         | OFPT_PORT_STATUS   | Supported     | Supported     |
|         | OFPT_PACKET_OUT  | Supported     | Supported     |
|         | OFPT_PACKET_OUT with buffer_id                                 | Not supported | Not supported |
|         | OFPT_FLOW_MOD:   | Supported     | Supported     |
|         | OFPPC_ADD  | Supported     | Supported     |
|         | OFPPC_ADD with OFPPF_CHECK_OVERLAP                             | Supported     | Supported     |
|         | OFPPC_MODIFY   | Supported     | Supported     |
|         | OFPPC_MODIFY_STRICT  | Supported     | Supported     |
|         | OFPPC_DELETE   | Supported     | Supported     |
|         | OFPPC_DELETE_STRICT  | Supported     | Supported     |
|         | OFPT_FLOW_MOD with buffer_id                                   | Not supported | Not supported |
|         | OFPT_PORT_MOD  | Not supported | Not supported |
|         | OFPT_STATS_REQUEST   | Supported     | Supported     |
|         | OFPT_STATS_REPLY<br>See <a href="#">Table 159 on page 2071</a> | Supported     | Supported     |
|         | OFPT_BARRIER_REQUEST   | Supported     | Supported     |
|         | OFPT_BARRIER_REPLY   | Supported     | Supported     |
|         | OFPT_QUEUE_GET_CONFIG_REQUEST                                  | Not supported | Not supported |
|         | OFPT_QUEUE_GET_CONFIG_REPLY                                    | Not supported | Not supported |

[Table 154 on page 2067](#) lists the support for OpenFlow v1.0 port structure flags.

Table 154: Junos OS Support for OpenFlow v1.0 Port Structure Flags

| Section | Specification      | MX Series     | EX9200        |
|---------|--------------------|---------------|---------------|
| 5.2.1   | OFPPC_PORT_DOWN    | Not supported | Not supported |
|         | OFPPC_NO_STP       | Not supported | Not supported |
|         | OFPPC_NO_RECV      | Not supported | Not supported |
|         | OFPPC_NO_RECV_STP  | Not supported | Not supported |
|         | OFPPC_NO_FLOOD     | Not supported | Not supported |
|         | OFPPC_NO_FWD       | Not supported | Not supported |
|         | OFPPC_NO_PACKET_IN | Not supported | Not supported |
|         | OFPPS_LINK_DOWN    | Supported     | Supported     |
|         | OFPPS_STP_LISTEN   | Not supported | Not supported |
|         | OFPPS_STP_LEARN    | Not supported | Not supported |
|         | OFPPS_STP_FORWARD  | Not supported | Not supported |
|         | OFPPS_STP_BLOCK    | Not supported | Not supported |
|         | OFPPS_STP_MASK     | Not supported | Not supported |
|         | OFPPF_10MB_HD      | Supported     | Supported     |
|         | OFPPF_10MB_FD      | Supported     | Supported     |
|         | OFPPF_100MB_HD     | Supported     | Supported     |
|         | OFPPF_100MB_FD     | Supported     | Supported     |
|         | OFPPF_1GB_HD       | Supported     | Supported     |
|         | OFPPF_1GB_FD       | Supported     | Supported     |
|         | OFPPF_10GB_FD      | Supported     | Supported     |
|         | OFPPF_COPPER       | Supported     | Supported     |
|         | OFPPF_FIBER        | Supported     | Supported     |
|         | OFPPF_AUTONEG      | Supported     | Supported     |
|         | OFPPF_PAUSE        | Not supported | Not supported |

**Table 154: Junos OS Support for OpenFlow v1.0 Port Structure Flags (*continued*)**

| Section | Specification    | MX Series     | EX9200        |
|---------|------------------|---------------|---------------|
|         | OFPPF_PAUSE_ASYM | Not supported | Not supported |

[Table 155 on page 2068](#) lists the support for OpenFlow v1.0 match conditions.

**Table 155: Junos OS Support for OpenFlow v1.0 Match Conditions**

| Section | Specification   | MX Series | EX9200    |
|---------|---|-----------|-----------|
| 5.2.3   | dL_src (Ethernet source address)  | Supported | Supported |
|         | dL_dst (Ethernet destination address)   | Supported | Supported |
|         | dL_vlan (Input VLAN ID)   | Supported | Supported |
|         | Note: The flow match condition for the VLAN ID must be less than 4096. Otherwise, the flow is not installed. The only exception is VLAN ID 65535, which corresponds to untagged frames. |           |           |
|         | dL_vlan_pcp (Input VLAN priority)   | Supported | Supported |
|         | Note: The flow match condition for the VLAN priority must be in accordance with 802.1p. Otherwise, the flow is not installed.   |           |           |
|         | dL_type (Ethernet frame type)   | Supported | Supported |
|         | nw_tos (IP TOS (6 bits DSCP))   | Supported | Supported |
|         | nw_proto (IP Protocol or lower 8 bits of ARP opcode)  | Supported | Supported |
|         | nw_src (IP source address)  | Supported | Supported |
|         | nw_dst (IP destination address)   | Supported | Supported |
|         | tp_src (TCP/UDP source port)  | Supported | Supported |
|         | tp_dst (TCP/UDP destination port)   | Supported | Supported |
|         | Match all 12 tuples or a combination of tuples  | Supported | Supported |

[Table 156 on page 2069](#) lists the support for OpenFlow v1.0 wildcards.

Table 156: Junos OS Support for OpenFlow v1.0 Wildcards

| Section | Specification                             | MX Series | EX9200    |
|---------|---|-----------|-----------|
| 5.2.3   | OFFFW_IN_PORT                             | Supported | Supported |
|         | OFFFW_DL_VLAN                             | Supported | Supported |
|         | OFFFW_DL_SRC                              | Supported | Supported |
|         | OFFFW_DL_DST                              | Supported | Supported |
|         | OFFFW_DL_TYPE                             | Supported | Supported |
|         | OFFFW_NW_PROTO                            | Supported | Supported |
|         | OFFFW_TP_SRC                              | Supported | Supported |
|         | OFFFW_TP_DST                              | Supported | Supported |
|         | No wild cards set. Match entire 12 tuple. | Supported | Supported |

Table 157 on page 2069 lists the support for OpenFlow v1.0 flow actions.

Table 157: Junos OS Support for OpenFlow v1.0 Flow Actions

| Section | Specification       | MX Series     | EX9200        |
|---------|---------------------|---------------|---------------|
| 5.2.4   | OFFPAT_OUTPUT:      |               |               |
|         | OFPP_IN_PORT        | Not supported | Not supported |
|         | OFPP_TABLE          | Not supported | Not supported |
|         | OFPP_NORMAL         | Supported     | Supported     |
|         | OFPP_FLOOD          | Supported     | Supported     |
|         | OFPP_ALL            | Supported     | Supported     |
|         | OFPP_CONTROLLER     | Supported     | Supported     |
|         | OFPP_LOCAL          | Not supported | Not supported |
|         | OFFPAT_SET_VLAN_VID | Supported     | Supported     |
|         | OFFPAT_SET_VLAN_PCP | Not supported | Not supported |
|         | OFFPAT_STRIP_VLAN   | Supported     | Supported     |
|         | OFFPAT_SET_DL_SRC   | Not supported | Not supported |
|         | OFFPAT_SET_DL_DST   | Not supported | Not supported |
|         | OFFPAT_SET_NW_SRC   | Not supported | Not supported |
|         | OFFPAT_SET_NW_DST   | Not supported | Not supported |

**Table 157: Junos OS Support for OpenFlow v1.0 Flow Actions (*continued*)**

| Section | Specification     | MX Series     | EX9200        |
|---------|-------------------|---------------|---------------|
|         | OFFPAT_SET_NW_TOS | Not supported | Not supported |
|         | OFFPAT_SET_TP_SRC | Not supported | Not supported |
|         | OFFPAT_SET_TP_DST | Not supported | Not supported |
|         | OFFPAT_ENQUEUE    | Not supported | Not supported |

Table 158 on page 2070 lists the support for OpenFlow v1.0 flow actions in Send Packet messages (OFPT\_PACKET\_OUT).

**Table 158: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT\_PACKET\_OUT)**

| Section | Specification       | MX Series     | EX9200        |
|---------|---------------------|---------------|---------------|
| 5.2.4   | OFFPAT_OUTPUT:      |               |               |
|         | OFPP_IN_PORT        | Not supported | Not supported |
|         | OFPP_TABLE          | Not supported | Not supported |
|         | OFPP_NORMAL         | Not supported | Not supported |
|         | OFPP_FLOOD          | Supported     | Supported     |
|         | OFPP_ALL            | Supported     | Supported     |
|         | OFPP_CONTROLLER     | Not supported | Not supported |
|         | OFPP_LOCAL          | Not supported | Not supported |
|         | OFFPAT_SET_VLAN_VID | Not supported | Not supported |
|         | OFFPAT_SET_VLAN_PCP | Not supported | Not supported |
|         | OFFPAT_STRIP_VLAN   | Not supported | Not supported |
|         | OFFPAT_SET_DL_SRC   | Not supported | Not supported |
|         | OFFPAT_SET_DL_DST   | Not supported | Not supported |
|         | OFFPAT_SET_NW_SRC   | Not supported | Not supported |
|         | OFFPAT_SET_NW_DST   | Not supported | Not supported |
|         | OFFPAT_SET_NW_TOS   | Not supported | Not supported |
|         | OFFPAT_SET_TP_SRC   | Not supported | Not supported |
|         | OFFPAT_SET_TP_DST   | Not supported | Not supported |
|         | OFFPAT_ENQUEUE      | Not supported | Not supported |



[Table 159 on page 2071](#) lists the support for OpenFlow v1.0 statistics.

**Table 159: Junos OS Support for OpenFlow v1.0 Statistics**

| Section | Specification  | MX Series          | EX9200             |
|---------|----------------|--------------------|--------------------|
| 5.3.5   | OFST_DESC      | Supported          | Supported          |
|         | OFST_FLOW      | Supported          | Supported          |
|         | OFST_AGGREGATE | Supported          | Supported          |
|         | OFST_TABLE     | Supported          | Supported          |
|         | OFST_PORT      | Supported          | Supported          |
|         | OFST_QUEUE     | Supported          | Supported          |
|         | OFST_VENDOR    | Gracefully ignored | Gracefully ignored |

[Table 160 on page 2071](#) lists the support for OpenFlow v1.0 features.

**Table 160: Junos OS Support for OpenFlow v1.0 Features**

| Section | Specification  | MX Series     | EX9200        |
|---------|--|---------------|---------------|
| 4.4     | Encryption. Controller and switch communicate through a TLS connection | Not supported | Not supported |
| 5.3.3   | Flow Idle Timeout  | Supported     | Supported     |
|         | Flow Hard Timeout  | Supported     | Supported     |
|         | Flow Priority  | Supported     | Supported     |

**Related Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS on page 2051](#)
- [OpenFlow Operational Mode Commands on page 2183](#)

## OpenFlow v1.0 Compliance Matrix for QFX5100 Switches

[Table 153 on page 2065](#) through [Table 160 on page 2071](#) list the OpenFlow v1.0 support for QFX5100 switches.

- [Table 153 on page 2065](#) lists support for message types.
- [Table 154 on page 2067](#) lists support for port structure flags.

- [Table 155 on page 2068](#) lists support for match conditions.
- [Table 156 on page 2069](#) lists support for wildcards.
- [Table 157 on page 2069](#) lists support for flow actions.
- [Table 158 on page 2070](#) lists support for flow actions in Send Packet messages (OFPT\_PACKET\_OUT).
- [Table 159 on page 2071](#) lists support for statistics.
- [Table 160 on page 2071](#) lists support for features.

[Table 153 on page 2065](#) lists the OpenFlow v1.0 message type support.

**Table 161: Junos OS Support for OpenFlow v1.0 Message Types**

| Section | Specification                 | QFX5100       |
|---------|-------------------------------|---------------|
| 5.1     | OFPT_HELLO                    | Supported     |
|         | OFPT_ERROR                    | Supported     |
|         | OFPT_ECHO_REQUEST             | Supported     |
|         | OFPT_ECHO_REPLY               | Supported     |
|         | OFPT_VENDOR                   | Not supported |
|         | OFPT_FEATURES_REQUEST         | Supported     |
|         | OFPT_FEATURES_REPLY:          | Supported     |
|         | Datapath ID                   | Supported     |
|         | N_buffers                     | -1            |
|         | N_tables                      | 1             |
|         | OFPC_FLOW_STATS               | Supported     |
|         | OFPC_TABLE_STATS              | Supported     |
|         | OFPC_PORT_STATS               | Supported     |
|         | OFPC_STP                      | Not supported |
|         | OFPC_IP_REASM                 | Not supported |
|         | OFPC_QUEUE_STATS              | Supported     |
|         | OFPC_ARP_MATCH_IP             | Not supported |
|         | OFPT_GET_CONFIG_REQUEST       | Supported     |
|         | OFPT_GET_CONFIG_REPLY         | Supported     |
|         | OFPT_SET_CONFIG               | Supported     |
|         | OFPT_PACKET_IN                | Supported     |
|         | OFPT_PACKET_IN with buffer_id | Not supported |
|         | OFPT_FLOW_REMOVED             | Supported     |

Table 161: Junos OS Support for OpenFlow v1.0 Message Types (*continued*)

| Section | Specification  | QFX5100       |
|---------|--|---------------|
|         | OFPT_PORT_STATUS   | Supported     |
|         | OFPT_PACKET_OUT  | Supported     |
|         | OFPT_PACKET_OUT with buffer_id                                 | Not supported |
|         | OFPT_FLOW_MOD:   | Supported     |
|         | OFPPC_ADD  | Supported     |
|         | OFPPC_ADD with OFPFF_CHECK_OVERLAP                             | Supported     |
|         | OFPPC_MODIFY   | Supported     |
|         | OFPPC_MODIFY_STRICT  | Supported     |
|         | OFPPC_DELETE   | Supported     |
|         | OFPPC_DELETE_STRICT  | Supported     |
|         | OFPT_FLOW_MOD with buffer_id                                   | Not supported |
|         | OFPT_PORT_MOD  | Not supported |
|         | OFPT_STATS_REQUEST   | Supported     |
|         | OFPT_STATS_REPLY<br>See <a href="#">Table 159 on page 2071</a> | Supported     |
|         | OFPT_BARRIER_REQUEST   | Supported     |
|         | OFPT_BARRIER_REPLY   | Supported     |
|         | OFPT_QUEUE_GET_CONFIG_REQUEST                                  | Not supported |
|         | OFPT_QUEUE_GET_CONFIG_REPLY                                    | Not supported |

[Table 154 on page 2067](#) lists the OpenFlow v1.0 port structure flag support

Table 162: Junos OS Support for OpenFlow v1.0 Port Structure Flags

| Section | Specification     | QFX5100       |
|---------|-------------------|---------------|
| 5.2.1   | OFPPC_PORT_DOWN   | Not supported |
|         | OFPPC_NO_STP      | Not supported |
|         | OFPPC_NO_RECV     | Not supported |
|         | OFPPC_NO_RECV_STP | Not supported |
|         | OFPPC_NO_FLOOD    | Not supported |
|         | OFPPC_NO_FWD      | Not supported |

**Table 162: Junos OS Support for OpenFlow v1.0 Port Structure Flags (continued)**

| Section | Specification      | QFX5100       |
|---------|--------------------|---------------|
|         | OFPPC_NO_PACKET_IN | Not supported |
|         | OFPPS_LINK_DOWN    | Supported     |
|         | OFPPS_STP_LISTEN   | Not supported |
|         | OFPPS_STP_LEARN    | Not supported |
|         | OFPPS_STP_FORWARD  | Not supported |
|         | OFPPS_STP_BLOCK    | Not supported |
|         | OFPPS_STP_MASK     | Not supported |
|         | OFPPF_10MB_HD      | Supported     |
|         | OFPPF_10MB_FD      | Supported     |
|         | OFPPF_100MB_HD     | Supported     |
|         | OFPPF_100MB_FD     | Supported     |
|         | OFPPF_1GB_HD       | Supported     |
|         | OFPPF_1GB_FD       | Supported     |
|         | OFPPF_10GB_FD      | Supported     |
|         | OFPPF_COPPER       | Supported     |
|         | OFPPF_FIBER        | Supported     |
|         | OFPPF_AUTONEG      | Supported     |
|         | OFPPF_PAUSE        | Not supported |
|         | OFPPF_PAUSE_ASYM   | Not supported |

[Table 155 on page 2068](#) lists OpenFlow v1.0 match condition support.

**Table 163: Junos OS Support for OpenFlow v1.0 Match Conditions**

| Section | Specification                    | QFX5100   |
|---------|----------------------------------|-----------|
| 5.2.3   | dl_src (Ethernet source address) | Supported |

**Table 163: Junos OS Support for OpenFlow v1.0 Match Conditions** (*continued*)

| Section | Specification  | QFX5100   |
|---------|--|-----------|
|         | dl_dst (Ethernet destination address)  | Supported |
|         | dl_vlan (Input VLAN ID)  | Supported |
|         | <b>NOTE:</b> The flow match condition for the VLAN ID must be less than 4096. Otherwise, the flow is not installed. The only exception is VLAN ID 65535, which corresponds to untagged frames. |           |
|         | dl_vlan_pcp (Input VLAN priority)  | Supported |
|         | <b>NOTE:</b> The flow match condition for the VLAN priority must be in accordance with 802.1p specifications. Otherwise, the flow is not installed.  |           |
|         | dl_type (Ethernet frame type)  | Supported |
|         | nw_tos (IP TOS (6-bit DSCP))   | Supported |
|         | nw_proto (IP Protocol or lower 8 bits of ARP opcode)   | Supported |
|         | nw_src (IP source address)   | Supported |
|         | nw_dst (IP destination address)  | Supported |
|         | tp_src (TCP/UDP source port/ICMPv4 type)   | Supported |
|         | tp_dst (TCP/UDP destination port/ICMPv4 code)  | Supported |
|         | Match all 12 tuples or a combination of tuples   | Supported |

Table 156 on page 2069 lists the OpenFlow v1.0 wildcard support.

**Table 164: Junos OS Support for OpenFlow v1.0 Wildcards**

| Section | Specification  | QFX5100   |
|---------|----------------|-----------|
| 5.2.3   | OFPFW_IN_PORT  | Supported |
|         | OFPFW_DL_VLAN  | Supported |
|         | OFPFW_DL_SRC   | Supported |
|         | OFPFW_DL_DST   | Supported |
|         | OFPFW_DL_TYPE  | Supported |
|         | OFPFW_NW_PROTO | Supported |

Table 164: Junos OS Support for OpenFlow v1.0 Wildcards (*continued*)

| Section | Specification                             | QFX5100   |
|---------|---|-----------|
|         | OFPFW_TP_SRC                              | Supported |
|         | OFPFW_TP_DST                              | Supported |
|         | No wild cards set. Match entire 12 tuple. | Supported |

Table 157 on page 2069 lists the OpenFlow v1.0 flow action support.

Table 165: Junos OS Support for OpenFlow v1.0 Flow Actions

| Section | Specification      | QFX5100       |
|---------|--------------------|---------------|
| 5.2.4   | OFPAT_OUTPUT:      |               |
|         | OFPP_IN_PORT       | Not supported |
|         | OFPP_TABLE         | Not supported |
|         | OFPP_NORMAL        | Not supported |
|         | OFPP_FLOOD         | Supported     |
|         | OFPP_ALL           | Supported     |
|         | OFPP_CONTROLLER    | Supported     |
|         | OFPP_LOCAL         | Not supported |
|         | OFPAT_SET_VLAN_VID | Supported     |
|         | OFPAT_SET_VLAN_PCP | Not supported |
|         | OFPAT_STRIP_VLAN   | Supported     |
|         | OFPAT_SET_DL_SRC   | Not supported |
|         | OFPAT_SET_DL_DST   | Not supported |
|         | OFPAT_SET_NW_SRC   | Not supported |
|         | OFPAT_SET_NW_DST   | Not supported |
|         | OFPAT_SET_NW_TOS   | Not supported |
|         | OFPAT_SET_TP_SRC   | Not supported |
|         | OFPAT_SET_TP_DST   | Not supported |
|         | OFPAT_ENQUEUE      | Not supported |

Table 158 on page 2070 lists the OpenFlow v1.0 flow action support in Send Packet messages (OFPT\_PACKET\_OUT).

**Table 166: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT\_PACKET\_OUT)**

| Section | Specification   | QFX5100   |
|---------|---|---|
| 5.2.4   | OFPAT_OUTPUT:<br>OFPP_IN_PORT<br>OFPP_TABLE<br>OFPP_NORMAL<br>OFPP_FLOOD<br>OFPP_ALL<br>OFPP_CONTROLLER<br>OFPP_LOCAL | Not supported<br>Not supported<br>Not supported<br>Supported<br>Supported<br>Not supported<br>Not supported |
|         | OFPAT_SET_VLAN_VID  | Supported   |
|         | OFPAT_SET_VLAN_PCP  | Not supported   |
|         | OFPAT_STRIP_VLAN  | Supported   |
|         | OFPAT_SET_DL_SRC  | Not supported   |
|         | OFPAT_SET_DL_DST  | Not supported   |
|         | OFPAT_SET_NW_SRC  | Not supported   |
|         | OFPAT_SET_NW_DST  | Not supported   |
|         | OFPAT_SET_NW_TOS  | Not supported   |
|         | OFPAT_SET_TP_SRC  | Not supported   |
|         | OFPAT_SET_TP_DST  | Not supported   |
|         | OFPAT_ENQUEUE   | Not supported   |

[Table 159 on page 2071](#) lists the OpenFlow v1.0 statistics support.

**Table 167: Junos OS Support for OpenFlow v1.0 Statistics**

| Section | Specification   | QFX5100   |
|---------|-----------------|-----------|
| 5.3.5   | OFPST_DESC      | Supported |
|         | OFPST_FLOW      | Supported |
|         | OFPST_AGGREGATE | Supported |
|         | OFPST_TABLE     | Supported |
|         | OFPST_PORT      | Supported |

**Table 167: Junos OS Support for OpenFlow v1.0 Statistics (*continued*)**

| Section | Specification | QFX5100            |
|---------|---------------|--------------------|
|         | OFPST_QUEUE   | Not supported      |
|         | OFPST_VENDOR  | Gracefully ignored |

[Table 160 on page 2071](#) lists the OpenFlow v1.0 feature support.

**Table 168: Junos OS Support for OpenFlow v1.0 Features**

| Section | Specification   | QFX5100       |
|---------|---|---------------|
| 4.4     | Encryption. Controller and switch communicate through a TLS connection. | Not supported |
| 5.3.3   | Flow Idle Timeout   | Supported     |
|         | Flow Hard Timeout   | Supported     |
|         | Flow Priority   | Supported     |

**Related  
Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS on page 2051](#)
- [OpenFlow Operational Mode Commands on page 2183](#)

## OpenFlow v1.0 Compliance Matrix for EX4550 Switches

[Table 153 on page 2065](#) through [Table 160 on page 2071](#) list the OpenFlow v1.0 support for the EX4550 switch.

- [Table 153 on page 2065](#) lists support for message types.
- [Table 154 on page 2067](#) lists support for port structure flags.
- [Table 155 on page 2068](#) lists match condition support.
- [Table 156 on page 2069](#) lists wildcard support.
- [Table 157 on page 2069](#) lists flow action support.
- [Table 158 on page 2070](#) lists flow action in Send Packet messages (OFPT\_PACKET\_OUT) support.
- [Table 159 on page 2071](#) lists statistics support.
- [Table 160 on page 2071](#) lists feature support.

[Table 153 on page 2065](#) lists the OpenFlow v1.0 message type support.



Table 169: Junos OS Support for OpenFlow v1.0 Message Types

| Section | Specification                  | EX4550        |
|---------|--------------------------------|---------------|
| 5.1     | OFPT_HELLO                     | Supported     |
|         | OFPT_ERROR                     | Supported     |
|         | OFPT_ECHO_REQUEST              | Supported     |
|         | OFPT_ECHO_REPLY                | Supported     |
|         | OFPT_VENDOR                    | Not supported |
|         | OFPT_FEATURES_REQUEST          | Supported     |
|         | OFPT_FEATURES_REPLY:           | Supported     |
|         | Datapath ID                    | Supported     |
|         | N_buffers                      | -1            |
|         | N_tables                       | 1             |
|         | OFPC_FLOW_STATS                | Not supported |
|         | OFPC_TABLE_STATS               | Supported     |
|         | OFPC_PORT_STATS                | Supported     |
|         | OFPC_STP                       | Not supported |
|         | OFPC_IP_REASM                  | Not supported |
|         | OFPC_QUEUE_STATS               | Supported     |
|         | OFPC_ARP_MATCH_IP              | Not supported |
|         | OFPT_GET_CONFIG_REQUEST        | Supported     |
|         | OFPT_GET_CONFIG_REPLY          | Supported     |
|         | OFPT_SET_CONFIG                | Supported     |
|         | OFPT_PACKET_IN                 | Supported     |
|         | OFPT_PACKET_IN with buffer_id  | Not supported |
|         | OFPT_FLOW_REMOVED              | Supported     |
|         | OFPT_PORT_STATUS               | Supported     |
|         | OFPT_PACKET_OUT                | Supported     |
|         | OFPT_PACKET_OUT with buffer_id | Not supported |

**Table 169: Junos OS Support for OpenFlow v1.0 Message Types (continued)**

| Section | Specification  | EX4550        |
|---------|--|---------------|
|         | OFPT_FLOW_MOD:   | Supported     |
|         | OFPPC_ADD  | Supported     |
|         | OFPPC_ADD with OFPFF_CHECK_OVERLAP                             | Supported     |
|         | OFPPC_MODIFY   | Supported     |
|         | OFPPC_MODIFY_STRICT  | Supported     |
|         | OFPPC_DELETE   | Supported     |
|         | OFPPC_DELETE_STRICT  | Supported     |
|         | OFPT_FLOW_MOD with buffer_id                                   | Not supported |
|         | OFPT_PORT_MOD  | Not supported |
|         | OFPT_STATS_REQUEST   | Supported     |
|         | OFPT_STATS_REPLY<br>See <a href="#">Table 159 on page 2071</a> | Supported     |
|         | OFPT_BARRIER_REQUEST   | Supported     |
|         | OFPT_BARRIER_REPLY   | Supported     |
|         | OFPT_QUEUE_GET_CONFIG_REQUEST                                  | Not supported |
|         | OFPT_QUEUE_GET_CONFIG_REPLY                                    | Not supported |

[Table 154 on page 2067](#) lists the OpenFlow v1.0 port structure flag support.

**Table 170: Junos OS Support for OpenFlow v1.0 Port Structure Flags**

| Section | Specification      | EX4550        |
|---------|--------------------|---------------|
| 5.2.1   | OFPPC_PORT_DOWN    | Not supported |
|         | OFPPC_NO_STP       | Not supported |
|         | OFPPC_NO_RECV      | Not supported |
|         | OFPPC_NO_RECV_STP  | Not supported |
|         | OFPPC_NO_FLOOD     | Not supported |
|         | OFPPC_NO_FWD       | Not supported |
|         | OFPPC_NO_PACKET_IN | Not supported |
|         | OFPPS_LINK_DOWN    | Supported     |

**Table 170: Junos OS Support for OpenFlow v1.0 Port Structure Flags (continued)**

| Section | Specification     | EX4550        |
|---------|-------------------|---------------|
|         | OFPPS_STP_LISTEN  | Not supported |
|         | OFPPS_STP_LEARN   | Not supported |
|         | OFPPS_STP_FORWARD | Not supported |
|         | OFPPS_STP_BLOCK   | Not supported |
|         | OFPPS_STP_MASK    | Not supported |
|         | OFPPF_10MB_HD     | Supported     |
|         | OFPPF_10MB_FD     | Supported     |
|         | OFPPF_100MB_HD    | Supported     |
|         | OFPPF_100MB_FD    | Supported     |
|         | OFPPF_1GB_HD      | Supported     |
|         | OFPPF_1GB_FD      | Supported     |
|         | OFPPF_10GB_FD     | Supported     |
|         | OFPPF_COPPER      | Supported     |
|         | OFPPF_FIBER       | Supported     |
|         | OFPPF_AUTONEG     | Supported     |
|         | OFPPF_PAUSE       | Not supported |
|         | OFPPF_PAUSE_ASYM  | Not supported |

[Table 155 on page 2068](#) lists OpenFlow v1.0 match condition support.

**Table 171: Junos OS Support for OpenFlow v1.0 Match Conditions**

| Section | Specification                         | EX4550    |
|---------|---------------------------------------|-----------|
| 5.2.3   | dl_src (Ethernet source address)      | Supported |
|         | dl_dst (Ethernet destination address) | Supported |

Table 171: Junos OS Support for OpenFlow v1.0 Match Conditions (*continued*)

| Section | Specification   | EX4550        |
|---------|---|---------------|
|         | dl_vlan (Input VLAN ID)<br><br><b>NOTE:</b> The flow match condition for the VLAN ID must be less than 4096. Otherwise, the flow is not installed. The only exception is VLAN ID 65535, which corresponds to untagged frames. | Supported     |
|         | dl_vlan_pcp (Input VLAN priority)<br><br><b>NOTE:</b> The flow match condition for the VLAN priority must be in accordance with 802.1p. Otherwise, the flow is not installed.   | Supported     |
|         | dl_type (Ethernet frame type)   | Supported     |
|         | nw_tos (IP TOS (6 bits DSCP))   | Supported     |
|         | nw_proto (IP Protocol or lower 8 bits of ARP opcode)  | Supported     |
|         | nw_src (IP source address)  | Supported     |
|         | nw_dst (IP destination address)   | Supported     |
|         | tp_src (TCP/UDP source port)  | Supported     |
|         | tp_dst (TCP/UDP destination port)   | Supported     |
|         | Match all 12 tuples or a combination of tuples  | Supported     |
|         | OFPXMT_OFB_IN_PORT  | Not supported |
|         | OFPXMT_OFB_IN_PHY_PORT  | Not supported |
|         | OFPXMT_OFB_METADATA   | Not supported |
|         | OFPXMT_OFB_ETH_DST  | Not supported |
|         | OFPXMT_OFB_ETH_SRC  | Not supported |
|         | OFPXMT_OFB_ETH_TYPE   | Not supported |
|         | OFPXMT_OFB_VLAN_VID   | Not supported |
|         | OFPXMT_OFB_VLAN_PCP   | Not supported |
|         | OFPXMT_OFB_IP_DSCP  | Not supported |
|         | OFPXMT_OFB_IP_ECN   | Not supported |

**Table 171: Junos OS Support for OpenFlow v1.0 Match Conditions (*continued*)**

| Section | Specification             | EX4550        |
|---------|---------------------------|---------------|
|         | OFPXMT_OFB_IP_PROTO       | Not supported |
|         | OFPXMT_OFB_IPV4_SRC       | Not supported |
|         | OFPXMT_OFB_IPV4_DST       | Not supported |
|         | OFPXMT_OFB_TCP_SRC        | Not supported |
|         | OFPXMT_OFB_TCP_DST        | Not supported |
|         | OFPXMT_OFB_UDP_SRC        | Not supported |
|         | OFPXMT_OFB_UDP_DST        | Not supported |
|         | OFPXMT_OFB_SCTP_SRC       | Not supported |
|         | OFPXMT_OFB_SCTP_DST       | Not supported |
|         | OFPXMT_OFB_ICMPV4_TYPE    | Not supported |
|         | OFPXMT_OFB_ICMPV4_CODE    | Not supported |
|         | OFPXMT_OFB_ARP_OP         | Not supported |
|         | OFPXMT_OFB_ARP_SPA        | Not supported |
|         | OFPXMT_OFB_ARP_TPA        | Not supported |
|         | OFPXMT_OFB_ARP_SHA        | Not supported |
|         | OFPXMT_OFB_ARP_THA        | Not supported |
|         | OFPXMT_OFB_IPV6_SRC       | Not supported |
|         | OFPXMT_OFB_IPV6_DST       | Not supported |
|         | OFPXMT_OFB_IPV6_FLABEL    | Not supported |
|         | OFPXMT_OFB_ICMPV6_TYPE    | Not supported |
|         | OFPXMT_OFB_ICMPV6_CODE    | Not supported |
|         | OFPXMT_OFB_IPV6_ND_TARGET | Not supported |
|         | OFPXMT_OFB_IPV6_ND_SLL    | Not supported |

**Table 171: Junos OS Support for OpenFlow v1.0 Match Conditions** (*continued*)

| Section | Specification          | EX4550        |
|---------|------------------------|---------------|
|         | OFPXMT_OFB_IPV6_ND_TLL | Not supported |
|         | OFPXMT_OFB_MPLS_LABEL  | Not supported |
|         | OFPXMT_OFB_MPLS_TC     | Not supported |
|         | OFPXMT_OFB_MPLS_BOS    | Not supported |
|         | OFPXMT_OFB_PBB_ISID    | Not supported |
|         | OFPXMT_OFB_TUNNEL_ID   | Not supported |
|         | OFPXMT_OFB_IPV6_EXTHDR | Not supported |

[Table 156 on page 2069](#) lists the OpenFlow v1.0 wildcard support.

**Table 172: Junos OS Support for OpenFlow v1.0 Wildcards**

| Section | Specification                             | EX4550    |
|---------|---|-----------|
| 5.2.3   | OFPFW_IN_PORT                             | Supported |
|         | OFPFW_DL_VLAN                             | Supported |
|         | OFPFW_DL_SRC                              | Supported |
|         | OFPFW_DL_DST                              | Supported |
|         | OFPFW_DL_TYPE                             | Supported |
|         | OFPFW_NW_PROTO                            | Supported |
|         | OFPFW_TP_SRC                              | Supported |
|         | OFPFW_TP_DST                              | Supported |
|         | No wild cards set. Match entire 12 tuple. | Supported |

[Table 157 on page 2069](#) lists the OpenFlow v1.0 flow action support.

Table 173: Junos OS Support for OpenFlow v1.0 Flow Actions

| Section | Specification   | EX4550  |
|---------|---|---|
| 5.2.4   | OFPAT_OUTPUT:<br><br>OFPP_IN_PORT<br>OFPP_TABLE<br>OFPP_NORMAL<br>OFPP_FLOOD<br>OFPP_ALL<br>OFPP_CONTROLLER<br>OFPP_LOCAL | Not supported<br>Not supported<br>Supported<br>Supported<br>Supported<br>Supported<br>Not supported |
|         | OFPAT_SET_VLAN_VID  | Not supported   |
|         | OFPAT_SET_VLAN_PCP  | Not supported   |
|         | OFPAT_STRIP_VLAN  | Not supported   |
|         | OFPAT_SET_DL_SRC  | Not supported   |
|         | OFPAT_SET_DL_DST  | Not supported   |
|         | OFPAT_SET_NW_SRC  | Not supported   |
|         | OFPAT_SET_NW_DST  | Not supported   |
|         | OFPAT_SET_NW_TOS  | Not supported   |
|         | OFPAT_SET_TP_SRC  | Not supported   |
|         | OFPAT_SET_TP_DST  | Not supported   |
|         | OFPAT_ENQUEUE   | Not supported   |

Table 158 on page 2070 lists the OpenFlow v1.0 flow action support in Send Packet messages (OFPT\_PACKET\_OUT).

Table 174: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT\_PACKET\_OUT)

| Section | Specification   | EX4550  |
|---------|---|---|
| 5.2.4   | OFPAT_OUTPUT:<br><br>OFPP_IN_PORT<br>OFPP_TABLE<br>OFPP_NORMAL<br>OFPP_FLOOD<br>OFPP_ALL<br>OFPP_CONTROLLER<br>OFPP_LOCAL | Not supported<br>Not supported<br>Not supported<br>Supported<br>Supported<br>Not supported<br>Not supported |

**Table 174: Junos OS Support for OpenFlow v1.0 Flow Actions in Send Packet Messages (OFPT\_PACKET\_OUT) (continued)**

| Section | Specification      | EX4550        |
|---------|--------------------|---------------|
|         | OFPAT_SET_VLAN_VID | Not supported |
|         | OFPAT_SET_VLAN_PCP | Not supported |
|         | OFPAT_STRIP_VLAN   | Not supported |
|         | OFPAT_SET_DL_SRC   | Not supported |
|         | OFPAT_SET_DL_DST   | Not supported |
|         | OFPAT_SET_NW_SRC   | Not supported |
|         | OFPAT_SET_NW_DST   | Not supported |
|         | OFPAT_SET_NW_TOS   | Not supported |
|         | OFPAT_SET_TP_SRC   | Not supported |
|         | OFPAT_SET_TP_DST   | Not supported |
|         | OFPAT_ENQUEUE      | Not supported |

[Table 159 on page 2071](#) lists the OpenFlow v1.0 statistics support.

**Table 175: Junos OS Support for OpenFlow v1.0 Statistics**

| Section | Specification   | EX4550             |
|---------|-----------------|--------------------|
| 5.3.5   | OFPST_DESC      | Supported          |
|         | OFPST_FLOW      | Not supported      |
|         | OFPST_AGGREGATE | Not supported      |
|         | OFPST_TABLE     | Supported          |
|         | OFPST_PORT      | Supported          |
|         | OFPST_QUEUE     | Supported          |
|         | OFPST_VENDOR    | Gracefully ignored |

[Table 160 on page 2071](#) lists the OpenFlow v1.0 feature support.



Table 176: Junos OS Support for OpenFlow v1.0 Features

| Section | Specification   | EX4550        |
|---------|---|---------------|
| 4.4     | Encryption. Controller and switch communicate through a TLS connection. | Not supported |
| 5.3.3   | Flow Idle Timeout   | Not supported |
|         | Flow Hard Timeout   | Supported     |
|         | Flow Priority   | Supported     |

**Related Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS on page 2051](#)
- [OpenFlow Operational Mode Commands on page 2183](#)

## OpenFlow v1.3.1 Compliance Matrix for Devices Running Junos OS

The following tables list the support for OpenFlow v1.3.1 features on the indicated platform.

- [Table 153 on page 2065](#) lists support for message types.
- [Table 178 on page 2090](#) lists support for features reply messages.
- [Table 154 on page 2067](#) lists support for port structure flags.
- [Table 180 on page 2091](#) lists support for port numbering.
- [Table 155 on page 2068](#) lists support for match conditions.
- [Table 157 on page 2069](#) lists support for flow actions.
- [Table 183 on page 2094](#) lists support for multipart messages.
- [Table 184 on page 2095](#) lists support for flow instructions.
- [Table 185 on page 2096](#) lists support for group types.

[Table 153 on page 2065](#) lists the support for OpenFlow v1.3.1 message types.

Table 177: Junos OS Support for OpenFlow v1.3.1 Message Types

| Specification     | QFX5100   |
|-------------------|-----------|
| OFPT_HELLO        | Supported |
| OFPT_ERROR        | Supported |
| OFPT_ECHO_REQUEST | Supported |
| OFPT_ECHO_REPLY   | Supported |

**Table 177: Junos OS Support for OpenFlow v1.3.1 Message Types (*continued*)**

| Specification                                | QFX5100       |
|--|---------------|
| OFPT_EXPERIMENTER                            | Not supported |
| OFPT_FEATURES_REQUEST                        | Supported     |
| OFPT_FEATURES_REPLY                          | Supported     |
| See <a href="#">Table 178 on page 2090</a> . |               |
| OFPT_GET_CONFIG_REQUEST                      | Supported     |
| OFPT_GET_CONFIG_REPLY                        | Supported     |
| OFPT_SET_CONFIG                              | Supported     |
| OFPT_PACKET_IN                               | Supported     |
| OFPT_PACKET_IN with buffer_id                | Not supported |
| OFPT_FLOW_REMOVED                            | Supported     |
| OFPT_PORT_STATUS                             | Supported     |
| OFPT_PACKET_OUT                              | Supported     |
| OFPT_PACKET_OUT with buffer_id               | Not supported |
| OFPT_FLOW_MOD                                | Supported     |
| OFPT_FLOW_MOD with buffer_id                 | Not supported |
| OFFPC_ADD                                    | Supported     |
| OFFPC_ADD with OFFPF_CHECK_OVERLAP           |               |
| OFFPC_MODIFY                                 | Supported     |
| OFFPC_MODIFY_STRICT                          | Supported     |
| OFFPC_DELETE                                 | Supported     |
| OFFPC_DELETE_STRICT                          | Supported     |
| Flow Modification Flags:                     | Supported     |
| OFFPF_SEND_FLOW_REM                          | Supported     |
| OFFPF_CHECK_OVERLAP                          | Supported     |
| OFFPF_RESET_COUNTS                           | Supported     |
| OFFPF_NO_PKT_COUNTS                          | Supported     |
| OFFPF_NO_BYT_COUNTS                          | Supported     |

**Table 177: Junos OS Support for OpenFlow v1.3.1 Message Types (*continued*)**

| Specification  | QFX5100       |
|--|---------------|
| OFPT_GROUP_MOD:  | Supported     |
| OFPGC_ADD  | Supported     |
| OFPGC_MODIFY   | Supported     |
| OFPGC_DELETE   | Supported     |
| OFPT_PORT_MOD  | Not supported |
| OFPT_TABLE_MOD   | Not supported |
| OFPT_MULTIPART_REQUEST<br>See <a href="#">Table 183 on page 2094</a> | Supported     |
| OFPT_MULTIPART_REPLY<br>See <a href="#">Table 183 on page 2094</a>   | Supported     |
| OFPT_BARRIER_REQUEST   | Supported     |
| OFPT_BARRIER_REPLY   | Supported     |
| OFPT_QUEUE_GET_CONFIG_REQUEST  | Not supported |
| OFPT_QUEUE_GET_CONFIG_REPLY  | Not supported |
| OFPT_ROLE_REQUEST  | Not supported |
| OFPT_ROLE_REPLY  | Not supported |
| OFPT_GET_ASYNC_REQUEST   | Not supported |
| OFPT_GET_ASYNC_REPLY   | Not supported |
| OFPT_SET_ASYNC   | Not supported |
| OFPT_METER_MOD   | Not supported |
| OFPT_VENDOR  | Not supported |

[Table 178 on page 2090](#) lists the support for OpenFlow v1.3.1 features reply messages.

Table 178: Junos OS Support for OpenFlow v1.3.1 Features Reply Messages

| Specification        | QFX5100       |
|----------------------|---------------|
| OFPT_FEATURES_REPLY: |               |
| Datapath ID          | Supported     |
| N_buffers            | -1            |
| N_tables             | 1             |
| Auxiliary ID         | 0             |
| OFPC_FLOW_STATS      | Supported     |
| OFPC_TABLE_STATS     | Supported     |
| OFPC_PORT_STATS      | Supported     |
| OFPC_GROUP_STATS     | Supported     |
| OFPC_IP_REASM        | Not supported |
| OFPC_QUEUE_STATS     | Supported     |
| OFPC_PORT_BLOCKED    | Not supported |

[Table 154 on page 2067](#) lists the support for OpenFlow v1.3.1 port structure flags.

Table 179: Junos OS Support for OpenFlow v1.3.1 Port Structure Flags

| Specification      | QFX5100       |
|--------------------|---------------|
| OFPPC_PORT_DOWN    | Not supported |
| OFPPC_NO_STP       | Not supported |
| OFPPC_NO_RECV      | Not supported |
| OFPPC_NO_RECV_STP  | Not supported |
| OFPPC_NO_FLOOD     | Not supported |
| OFPPC_NO_FWD       | Not supported |
| OFPPC_NO_PACKET_IN | Not supported |
| OFPPS_LINK_DOWN    | Supported     |
| OFPPS_BLOCKED      | Not supported |
| OFPPS_LIVE         | Not supported |
| OFPPF_10MB_HD      | Supported     |
| OFPPF_10MB_FD      | Supported     |
| OFPPF_100MB_HD     | Supported     |
| OFPPF_100MB_FD     | Supported     |
| OFPPF_1GB_HD       | Supported     |

**Table 179: Junos OS Support for OpenFlow v1.3.1 Port Structure Flags (*continued*)**

| Specification    | QFX5100       |
|------------------|---------------|
| OFPPF_1GB_FD     | Supported     |
| OFPPF_10GB_FD    | Supported     |
| OFPPF_40GB-FD    | Supported     |
| OFPPF_100GB-FD   | Not supported |
| OFPPF_1TB-FD     | Not supported |
| OFPPF_COPPER     | Not supported |
| OFPPF_FIBER      | Supported     |
| OFPPF_AUTONEG    | Supported     |
| OFPPF_PAUSE      | Not supported |
| OFPPF_PAUSE_ASYM | Not supported |

[Table 180 on page 2091](#) lists the support for OpenFlow v1.3.1 port numbering.

**Table 180: Junos OS Support for OpenFlow v1.3.1 Port Numbering**

| Specification  | QFX5100       |
|--|---------------|
| OFPP_IN_PORT   | Not supported |
| OFPP_TABLE   | Not supported |
| OFPP_NORMAL  | Not supported |
| OFPP_FLOOD (all except input and STP disabled port) (Flood and All are same) | Supported     |
| OFPP_ALL (all except input)  | Supported     |
| OFPP_CONTROLLER  | Supported     |
| OFPP_LOCAL   | Not supported |

[Table 155 on page 2068](#) lists the support for OpenFlow v1.3.1 match conditions.

Table 181: Junos OS Support for OpenFlow v1.3.1 Match Conditions

| Specification          | QFX5100       |
|------------------------|---------------|
| OFPXMT_OFB_IN_PORT     | Supported     |
| OFPXMT_OFB_IN_PHY_PORT | Not supported |
| OFPXMT_OFB_METADATA    | Not supported |
| OFPXMT_OFB_ETH_SRC     | Supported     |
| OFPXMT_OFB_ETH_DST     | Supported     |
| OFPXMT_OFB_VLAN_VID    | Supported     |
| OFPXMT_OFB_VLAN_PCP    | Supported     |
| OFPXMT_OFB_ETH_TYPE    | Supported     |
| OFPXMT_OFB_IP_DSCP     | Supported     |
| OFPXMT_OFB_IP_ECN      | Not supported |
| OFPXMT_OFB_IP_PROTO    | Supported     |
| OFPXMT_OFB_IPV4_SRC    | Supported     |
| OFPXMT_OFB_IPV4_DST    | Supported     |
| OFPXMT_OFB_TCP_SRC     | Supported     |
| OFPXMT_OFB_TCP_DST     | Supported     |
| OFPXMT_OFB_UDP_SRC     | Supported     |
| OFPXMT_OFB_UDP_DST     | Supported     |
| OFPXMT_OFB_SCTP_SRC    | Not supported |
| OFPXMT_OFB_SCTP_DST    | Not supported |
| OFPXMT_OFB_ICMPV4_TYPE | Supported     |
| OFPXMT_OFB_ICMPV4_CODE | Supported     |
| OFPXMT_OFB_ARP_OP      | Not supported |
| OFPXMT_OFB_ARP_SPA     | Not supported |
| OFPXMT_OFB_ARP_TPA     | Not supported |

**Table 181: Junos OS Support for OpenFlow v1.3.1 Match Conditions** (*continued*)

| Specification          | QFX5100       |
|------------------------|---------------|
| OFPXMT_OFB_ARP_SHA     | Not supported |
| OFPXMT_OFB_ARP_THA     | Not supported |
| OFPXMT_OFB_IPV6_SRC    | Not supported |
| OFPXMT_OFB_IPV6_DST    | Not supported |
| OFPXMT_OFB_IPV6_FLABEL | Not supported |
| OFPXMT_OFB_ICMPV6_TYPE | Not supported |
| OFPXMT_OFB_ICMPV6_CODE | Not supported |
| OXM_OF_IPV6_ND_TARGET  | Not supported |
| OXM_OF_IPV6_ND_SLL     | Not supported |
| OXM_OF_IPV6_ND_TLL     | Not supported |
| OXM_OF_IPV6_EXTHDR     | Not supported |
| OFPXMT_OFB_MPLS_LABEL  | Not supported |
| OFPXMT_OFB_MPLS_TC     | Not supported |
| OFPXMT_OFB_MPLS_BOS    | Not supported |
| OFPXMT_OFB_PBB_ISID    | Not supported |
| OFPXMT_OFB_TUNNEL_ID   | Not supported |



**NOTE:** The Junos OS implementation of OpenFlow v1.3.1 supports wildcards for all match conditions.

Also, this implementation of OpenFlow v1.3.1 does not support arbitrary bit masks for any fields. This implementation supports only continuous masks for IPv4 source and destination addresses.

Table 157 on page 2069 lists the support for OpenFlow v1.3.1 flow actions.

**Table 182: Junos OS Support for OpenFlow v1.3.1 Flow Actions**

| Specification       | QFX5100       |
|---------------------|---------------|
| OFFPAT_SET_VLAN_VID | Supported     |
| OFFPAT_SET_VLAN_PCP | Not supported |
| OFFPAT_POP_VLAN     | Supported     |
| OFFPAT_GROUP        | Supported     |
| OFFPAT_COPY_TTL_OUT | Not supported |
| OFFPAT_COPY_TTL_IN  | Not supported |
| OFFPAT_SET_MPLS_TTL | Not supported |
| OFFPAT_DEC_MPLS_TTL | Not supported |
| OFFPAT_PUSH_VLAN    | Not supported |
| OFFPAT_PUSH_MPLS    | Not supported |
| OFFPAT_POP_MPLS     | Not supported |
| OFFPAT_SET_QUEUE    | Not supported |
| OFFPAT_SET_NW_TTL   | Not supported |
| OFFPAT_DEC_NW_TTL   | Not supported |
| OFFPAT_PUSH_PBB     | Not supported |
| OFFPAT_POP_PBB      | Not supported |
| OFFPAT_EXPERIMENTER | Not supported |

[Table 183 on page 2094](#) lists the support for OpenFlow v1.3.1 multipart messages.

**Table 183: Junos OS Support for OpenFlow v1.3.1 Multipart Messages**

| Specification    | QFX5100   |
|------------------|-----------|
| OFFPMP_DESC      | Supported |
| OFFPMP_FLOW      | Supported |
| OFFPMP_AGGREGATE | Supported |
| OFFPMP_TABLE     | Supported |



Table 183: Junos OS Support for OpenFlow v1.3.1 Multipart Messages (*continued*)

| Specification        | QFX5100       |
|----------------------|---------------|
| OFPMP_PORT_STATS     | Supported     |
| OFPMP_QUEUE          | Supported     |
| OFPMP_GROUP          | Supported     |
| OFPMP_GROUP_DESC     | Supported     |
| OFPMP_GROUP_FEATURES | Supported     |
| OFPMP_METER          | Not supported |
| OFPMP_METER_CONFIG   | Not supported |
| OFPMP_METER_FEATURES | Not supported |
| OFPMP_TABLE_FEATURES | Supported     |
| OFPMP_PORT_DESC      | Supported     |
| OFPMP_EXPERIMENTER   | Not supported |

Table 184 on page 2095 lists the support for OpenFlow v1.3.1 flow instructions.



**NOTE:** A flow can have a maximum of one of the supported flow instructions listed in Table 184 on page 2095.

Table 184: Junos OS Support for OpenFlow v1.3.1 Flow Instructions

| Specification        | QFX5100       |
|----------------------|---------------|
| OFPIT_GOTO_TABLE     | Not supported |
| OFPIT_WRITE_METADATA | Not supported |
| OFPIT_WRITE_ACTIONS  | Supported     |
| OFPIT_APPLY_ACTIONS  | Supported     |
| OFPIT_CLEAR_ACTIONS  | Not supported |
| OFPIT_METER          | Not supported |
| OFPIT_EXPERIMENTER   | Not supported |

Table 185 on page 2096 lists the support for OpenFlow v1.3.1 group types.

**Table 185: Junos OS Support for OpenFlow v1.3.1 Group Types**

| Specification   | QFX5100       |
|-----------------|---------------|
| OFFPGT_ALL      | Supported     |
| OFFPGT_SELECT   | Not supported |
| OFFPGT_INDIRECT | Supported     |
| OFFPGT_FF       | Not supported |

**Related  
Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Understanding OpenFlow Operation and Forwarding Actions on Devices Running Junos OS on page 2051](#)
- [Understanding How the OpenFlow Group Action Works on page 2061](#)
- [OpenFlow Operational Mode Commands on page 2183](#)

## CHAPTER 23

# Installing Support for OpenFlow

- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)

## Installing Support for OpenFlow on Devices Running Junos OS

---

You can add support for OpenFlow on a device running Junos OS by copying the software package to your device and then installing the package. The software package is identified by the jsdn prefix, and the filename string begins with the following format:

*jsdn-packageID-release*

where:

- *packageID* identifies the devices running Junos OS on which you can install the package.
- *release* identifies the release, for example, 13.3. The jsdn software release and the Junos OS release of the device on which it is installed must match.

For information about OpenFlow support on devices running Junos OS and the corresponding installation package for that device, see *OpenFlow Support on Devices Running Junos OS*.

To install the jsdn software package on a device running Junos OS:

1. Download the software package to the device.
2. If you previously installed the jsdn software package, remove the existing package using the **request system software delete** operational mode command.

```
user@host> request system software delete existing-jsdn-package
```

3. Install the new software package using the **request system software add** operational mode command.

```
user@h> request system software add path-to-jsdn-package
```

### Related Documentation

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)



# OpenFlow Basic Configuration

- [Configuring Support for OpenFlow on MX Series Routers on page 2099](#)
- [Example: Enabling OpenFlow on MX Series Routers on page 2102](#)
- [Configuring Support for OpenFlow on EX9200 Switches on page 2106](#)
- [Example: Enabling OpenFlow on EX9200 Switches on page 2109](#)
- [Configuring Support for OpenFlow on QFX5100 Switches on page 2113](#)
- [Example: Enabling OpenFlow on QFX5100 Switches on page 2115](#)
- [Configuring Support for OpenFlow on EX4550 Switches on page 2120](#)
- [Example: Enabling OpenFlow on EX4550 Switches on page 2121](#)

## Configuring Support for OpenFlow on MX Series Routers

---

The following sections configure MX Series routers to support OpenFlow using interfaces that participate solely in OpenFlow. For information about configuring hybrid interfaces, which concurrently support OpenFlow logical interfaces and non-OpenFlow logical interfaces, see [“Configuring OpenFlow Hybrid Interfaces on MX Series Routers” on page 2128](#).

Before configuring support for OpenFlow, ensure that the router meets the following requirements:

- MX Series router running Junos OS Release 13.3 or a later release
- OpenFlow software package with a software package release that matches the Junos OS release of the device on which it is installed
- TCP connection between the router and an OpenFlow controller
- Connection between the management interface of the router and the management network, which is reachable from the controller IP address

Configuration tasks are described in detail in the following sections:

- [Configuring the OpenFlow Interfaces on page 2100](#)
- [Configuring the OpenFlow Protocol on page 2100](#)
- [Configuring the OpenFlow Routing Instance on page 2101](#)

## Configuring the OpenFlow Interfaces

You must configure interfaces participating in OpenFlow as Layer 2 interfaces. On MX Series routers, you configure the interfaces with encapsulation **ethernet-bridge** and protocol family **bridge**.

To configure the OpenFlow Interfaces:

- Configure the physical link-layer encapsulation type and the logical interface and protocol family.

```
[edit interfaces interface-name]  
user@host# set encapsulation ethernet-bridge  
user@host# set unit unit family bridge
```

## Configuring the OpenFlow Protocol

To configure support for OpenFlow, create a virtual switch instance, and specify a switch name, which must be 60 characters or less. For the virtual switch instance, configure the OpenFlow controller information and the participating logical interfaces. Optionally, configure the default action for packets that do not match a flow entry, the purge timer for invalid flows, and any OpenFlow traceoptions.

To configure the OpenFlow protocol:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch switch-name]  
user@host# set controller address address  
user@host# set controller protocol tcp
```

2. Specify the logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch switch-name]  
user@host# set interfaces interface-name1.unit1  
user@host# set interfaces interface-name2.unit1
```

3. (Optional) Configure the **default-action** statement for packets that do not match on an existing flow entry.

If you do not configure the **default-action** statement, the default is **packet-in**, which indicates that packets with no matching flow entry must be sent to the controller for processing.

```
[edit protocols openflow switch switch-name]  
user@host# set default-action (drop | packet-in)
```

4. (Optional) Configure the **purge-flow-timer** statement, which is the number of seconds after which an invalid flow is purged from the flow table.

```
[edit protocols openflow switch switch-name]  
user@host# set purge-flow-timer seconds
```

5. (Optional) Configure OpenFlow traceoptions.

If you do not configure a log filename, OpenFlow trace messages are logged in the default OpenFlow log file `/var/log/ofd`.

```
[edit protocols openflow]
user@host# set traceoptions flag flag
user@host# set traceoptions file file-name
```

## Configuring the OpenFlow Routing Instance

To configure the virtual switch routing instance for OpenFlow traffic:

1. Configure the routing instance type as **virtual-switch**.

```
[edit routing-instances routing-instance-name]
user@host# set instance-type virtual-switch
```

2. Configure the bridge domain name and type.

```
[edit routing-instances routing-instance-name]
user@host# set bridge-domains name domain-type bridge
```

3. Configure the VLAN ID as **none**.

```
[edit routing-instances routing-instance-name]
user@host# set bridge-domains name vlan-id none
```

4. Configure the OpenFlow logical interfaces that will be bound to the routing instance.

```
[edit routing-instances routing-instance-name]
user@host# set bridge-domains name interface interface-name1.unit1
user@host# set bridge-domains name interface interface-name2.unit1
```

5. (Optional) If you use the NORMAL forward action to forward OpenFlow traffic using traditional Layer 2 and Layer 3 processing, configure an integrated routing and bridging (IRB) interface, and include the appropriate logical interface in the bridge domain configuration.

```
[edit routing-instances routing-instance-name]
user@host# set bridge-domains name routing-interface irb.unit
```

### Related Documentation

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
- [Example: Enabling OpenFlow on MX Series Routers on page 2102](#)
- [OpenFlow Operational Mode Commands on page 2183](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)
- [\[edit protocols openflow\] Hierarchy Level on page 2169](#)

## Example: Enabling OpenFlow on MX Series Routers

---

OpenFlow is an open standard that allows you to control traffic paths in a network by creating, deleting, and modifying flows in each device along a path. This example shows how to configure OpenFlow support on an MX240 router running Junos OS.

- [Requirements on page 2102](#)
- [Overview on page 2102](#)
- [Configuration on page 2103](#)
- [Verification on page 2105](#)

### Requirements

This example uses the following hardware and software components:

- MX240 router running Junos OS Release 13.3 or a later release
- OpenFlow software package with a software package release that matches the Junos OS release of the device on which it is installed
- TCP connection between the router and an OpenFlow controller
- Connection between the management interface of the router and the management network, which is reachable from the OpenFlow controller IP address

### Overview

In this example, you configure support for OpenFlow on an MX240 router. The router has three interfaces that participate solely in OpenFlow: ge-1/0/0.0, ge-1/1/0.0, and xe-0/0/0.0. You first configure the interfaces as Layer 2 interfaces using physical link-layer encapsulation type **ethernet-bridge** and protocol family **bridge**.

MX Series routers require a separate virtual switch routing instance to isolate the OpenFlow traffic from the normal network traffic. This example configures a virtual switch routing instance, `rt-bd-1`, using instance type **virtual-switch** at the **[edit routing-instances]** hierarchy level. Within the routing instance, the bridge domain **of-bridge** includes all of the logical interfaces participating in OpenFlow.

You configure the OpenFlow virtual switch and OpenFlow protocol statements at the **[edit protocols openflow]** hierarchy level. In this example, the virtual switch, `OFswitch1`, connects to the controller over a TCP connection at IP address 172.16.1.1. The virtual switch configuration must include all of the logical interfaces participating in OpenFlow, and OpenFlow traffic will only enter or exit from these interfaces.

Within the OpenFlow configuration, the **default-action** statement indicates the action the switch must take for packets that do not have a matching flow entry. If you omit the **default-action** statement, the default action is **packet-in**, which indicates that packets with no matching flow entry must be sent to the controller for processing. This example explicitly configures the default action for packets that do not have a matching flow entry as **packet-in**.



This example also configures OpenFlow traceoptions. In this case, the **flag all** statement indicates that all OpenFlow trace events should be captured and logged. Since the example does not configure a specific filename for the log file, OpenFlow trace messages are logged in the default OpenFlow log file `/var/log/ofd`.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-1/0/0 encapsulation ethernet-bridge unit 0 family bridge
set interfaces ge-1/1/0 encapsulation ethernet-bridge unit 0 family bridge
set interfaces xe-0/0/0 encapsulation ethernet-bridge unit 0 family bridge
set routing-instances rt-bd-1 instance-type virtual-switch
set routing-instances rt-bd-1 bridge-domains of-bridge vlan-id none
set routing-instances rt-bd-1 bridge-domains of-bridge interface ge-1/0/0.0
set routing-instances rt-bd-1 bridge-domains of-bridge interface ge-1/1/0.0
set routing-instances rt-bd-1 bridge-domains of-bridge interface xe-0/0/0.0
set protocols openflow switch OFswitch1 controller address 172.16.1.1
set protocols openflow switch OFswitch1 controller protocol tcp
set protocols openflow switch OFswitch1 interfaces ge-1/0/0.0
set protocols openflow switch OFswitch1 interfaces ge-1/1/0.0
set protocols openflow switch OFswitch1 interfaces xe-0/0/0.0
set protocols openflow switch OFswitch1 default-action packet-in
set protocols openflow traceoptions flag all
```

**Step-by-Step Procedure** To configure support for OpenFlow:

1. Configure the OpenFlow interfaces as Layer 2 interfaces.

```
[edit interfaces]
user@host# set ge-1/0/0 encapsulation ethernet-bridge unit 0 family bridge
user@host# set ge-1/1/0 encapsulation ethernet-bridge unit 0 family bridge
user@host# set xe-0/0/0 encapsulation ethernet-bridge unit 0 family bridge
```

2. Configure the virtual switch routing instance.

```
[edit routing-instances]
user@host# set rt-bd-1 instance-type virtual-switch
user@host# set rt-bd-1 bridge-domains of-bridge vlan-id none
user@host# set rt-bd-1 bridge-domains of-bridge interface ge-1/0/0.0
user@host# set rt-bd-1 bridge-domains of-bridge interface ge-1/1/0.0
user@host# set rt-bd-1 bridge-domains of-bridge interface xe-0/0/0.0
```

3. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch OFswitch1]
user@host# set controller address 172.16.1.1
user@host# set controller protocol tcp
```

4. Configure the logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch OFswitch1]
user@host# set interfaces ge-1/0/0.0
```

```
user@host# set interfaces ge-1/1/0.0
user@host# set interfaces xe-0/0/0.0
```

5. Configure the default action for packets that do not have a matching flow entry.

```
[edit protocols openflow switch OFswitch1]
user@host# set default-action packet-in
```

6. Configure OpenFlow traceoptions.

```
[edit protocols openflow]
user@host# set traceoptions flag all
```

7. Commit the configuration.

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

---

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols openflow**, and **show routing-instances** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-1/0/0 {
  encapsulation ethernet-bridge;
  unit 0 {
    family bridge;
  }
}
ge-1/1/0 {
  encapsulation ethernet-bridge;
  unit 0 {
    family bridge;
  }
}
xe-0/0/0 {
  encapsulation ethernet-bridge;
  unit 0 {
    family bridge;
  }
}

user@host# show protocols openflow
switch OFswitch1 {
  default-action packet-in;
  interfaces {
    ge-1/0/0.0;
    ge-1/1/0.0;
    xe-0/0/0.0;
  }
  controller {
    address 172.16.1.1;
    protocol tcp;
  }
}
```

```

}
traceoptions {
  flag all;
}

user@host# show routing-instances
rt-bd-1 {
  instance-type virtual-switch;
  bridge-domains {
    of-bridge {
      vlan-id none;
      interface ge-1/0/0.0;
      interface ge-1/1/0.0;
      interface xe-0/0/0.0;
    }
  }
}
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying that the OpenFlow Controller Connection is Up on page 2105](#)
- [Verifying that the OpenFlow Interfaces Are Up on page 2105](#)

### Verifying that the OpenFlow Controller Connection is Up

**Purpose** Verify that the OpenFlow controller connection is up.

**Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**. Because the virtual switch configuration has only a single controller, the virtual switch should automatically initiate a connection to the controller after you commit the configuration.

```

user@host> show openflow controller
Openflowd controller information:
Controller socket: 11
Controller IP address: 172.16.1.1
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 1
Controller role: equal

```

**Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying that the OpenFlow Interfaces Are Up

**Purpose** Verify that the OpenFlow interfaces are up.

**Action** Issue the **show openflow interfaces** operational mode command, and verify that the state of each OpenFlow interface is **Up**.

```

user@host> show openflow interfaces

```

```
Switch name: OFswitch1
Interface Name: ge-1/0/0.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:b1
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: xe-0/0/0.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:b3
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

**Meaning** The output shows that the state of each OpenFlow interface is **Up**, in addition to other information about the interfaces.

**Related  
Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
- [Configuring Support for OpenFlow on MX Series Routers on page 2099](#)
- [OpenFlow Operational Mode Commands on page 2183](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)
- [\[edit protocols openflow\] Hierarchy Level on page 2169](#)

---

## Configuring Support for OpenFlow on EX9200 Switches

The following sections detail one method to configure EX9200 switches to support OpenFlow using interfaces that participate solely in OpenFlow. For information about configuring hybrid interfaces, which concurrently support OpenFlow logical interfaces and non-OpenFlow logical interfaces, see [“Configuring OpenFlow Hybrid Interfaces on EX9200 Switches” on page 2138](#).

Before configuring support for OpenFlow, ensure that the switch meets the following requirements:

- EX9200 switch running Junos OS Release 13.3 or a later release.
- OpenFlow software package with a software package release that matches the Junos OS release of the device on which it is installed

- TCP connection between the switch and an OpenFlow controller
- Connection between the management interface of the switch and the management network, which is reachable from the controller IP address

Configuration tasks are described in detail in the following sections:

- [Configuring the OpenFlow Interfaces on page 2107](#)
- [Configuring the OpenFlow Protocol on page 2107](#)
- [Configuring the OpenFlow Routing Instance on page 2108](#)

## Configuring the OpenFlow Interfaces

To configure the OpenFlow interfaces:

1. Specify the desired VLAN tagging and configure the encapsulation type.

```
[edit interfaces interface-name]
user@host# set flexible-vlan-tagging
user@host# set encapsulation flexible-ethernet-services
```

2. Configure the logical interface and the protocol family.

```
[edit interfaces interface-name]
user@host# set unit unit family ethernet-switching
```

3. Configure the interface as a trunk interface and specify the VLAN members associated with OpenFlow.

```
[edit interfaces interface-name]
user@host# set unit unit family ethernet-switching interface-mode trunk
user@host# set unit unit family ethernet-switching vlan members openflow-vlan-ids
```

## Configuring the OpenFlow Protocol

To configure support for OpenFlow, create a virtual switch instance, and specify a switch name, which must be 60 characters or less. For the virtual switch instance, configure the OpenFlow controller information and the participating logical interfaces. Optionally, configure the default action for packets that do not match a flow entry, the purge timer for invalid flows, and any OpenFlow traceoptions.

To configure the OpenFlow protocol:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch switch-name]
user@host# set controller address address
user@host# set controller protocol tcp
```

2. Specify the logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch switch-name]
user@host# set interfaces interface-name1.unit1
user@host# set interfaces interface-name2.unit1
```

3. (Optional) Configure the **default-action** statement for packets that do not match on an existing flow entry.

If you do not configure the **default-action** statement, the default is **packet-in**, which indicates that packets with no matching flow entry must be sent to the controller for processing.

```
[edit protocols openflow switch switch-name]  
user@host# set default-action (drop | packet-in)
```

4. (Optional) Configure the **purge-flow-timer** statement, which is the number of seconds after which an invalid flow is purged from the flow table.

```
[edit protocols openflow switch switch-name]  
user@host# set purge-flow-timer seconds
```

5. (Optional) Configure OpenFlow traceoptions.

If you do not configure a log filename, OpenFlow trace messages are logged in the default OpenFlow log file `/var/log/ofd`.

```
[edit protocols openflow]  
user@host# set traceoptions flag flag  
user@host# set traceoptions file file-name
```

## Configuring the OpenFlow Routing Instance

To configure the virtual switch routing instance for OpenFlow traffic:

1. Configure the routing instance type as **virtual-switch**.

```
[edit routing-instances routing-instance-name]  
user@host# set instance-type virtual-switch
```

2. Configure the OpenFlow logical interfaces that will be bound to the routing instance.

```
[edit routing-instances routing-instance-name]  
user@host# set interface interface-name1.unit1  
user@host# set interface interface-name2.unit1
```

3. Configure the OpenFlow VLAN members under the **vlan**s hierarchy.

```
[edit routing-instances routing-instance-name]  
user@host# set vlans name (vlan-id | vlan-id-list) openflow-vlan-ids
```

4. (Optional) If you use the NORMAL forward action to forward OpenFlow traffic using traditional Layer 2 and Layer 3 processing, configure an integrated routing and bridging (IRB) interface, and bind the appropriate logical interface to the VLAN.

```
[edit routing-instances routing-instance-name]  
user@host# set vlans name l3-interface irb.unit
```

### Related Documentation

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
- [OpenFlow Operational Mode Commands on page 2183](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)

- [\[edit protocols openflow\] Hierarchy Level on page 2169](#)

## Example: Enabling OpenFlow on EX9200 Switches

OpenFlow is an open standard that enables you to control traffic paths in a network by creating, deleting, and modifying flows in each device, including EX9200 switches that have an OpenFlow software package installed, along a path. This example shows how to configure OpenFlow support on an EX9200 switch.

- [Requirements on page 2109](#)
- [Overview on page 2109](#)
- [Configuration on page 2110](#)
- [Verification on page 2112](#)

### Requirements

This example uses the following hardware and software components:

- An EX9200 switch running Junos OS Release 13.3 or a later release.
- An OpenFlow software package is installed on the switch, and the software package release matches the Junos OS release running on the switch.
- The switch has a TCP connection to an OpenFlow controller, which needs to access the data plane of the switch.
- The switch is connected to the management network through the fxp0 interface and is reachable from the OpenFlow controller IP address.

### Overview

In this example, you configure support for OpenFlow on an EX9200 switch. The switch has three interfaces that are dedicated to handling OpenFlow traffic: ge-1/0/0.0, ge-1/1/0.0, and xe-0/0/0.0.

EX9200 switches require a separate routing instance for a virtual switch. This routing instance isolates the experimental OpenFlow traffic from the normal network traffic. In this example, you configure a routing instance for the virtual switch, **OF-ri**, by using the instance type **virtual-switch** at the **[edit routing-instances]** hierarchy level. The routing instance **OF-ri** includes all of the logical interfaces participating in OpenFlow.

The virtual switch, **OFswitch1**, connects to the OpenFlow controller over a TCP connection at the IP address 198.51.100.174. The virtual switch configuration must include all of the logical interfaces participating in OpenFlow, and OpenFlow traffic only will either enter or exit these interfaces.

A flow entry consists of a match condition against which packets entering an OpenFlow interface are compared, and the action that is applied to packets that match the condition. Each OpenFlow interface can have one or more flow entries. The **default-action** statement in the OpenFlow configuration indicates the action the switch applies for packets that

do not have a matching flow entry. If you do not explicitly configure the **default-action** statement, the default action is **packet-in**, which indicates that packets that have no matching flow entry are sent to the OpenFlow controller for processing. In this example, you explicitly configure **packet-in** as the default action for packets that do not have a matching flow entry.

In this example, you configure OpenFlow traceoptions also. When traceoptions are configured with the **flag all** statement, all OpenFlow events are captured and logged. In this example, a specific filename is not configured for the log file. Therefore, OpenFlow events are logged in the default OpenFlow log file at **/var/log/ofd**.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-1/0/0 unit 0 family ethernet-switching
set interfaces ge-1/1/0 unit 0 family ethernet-switching
set interfaces xe-0/0/0 unit 0 family ethernet-switching
set routing-instances OF-ri instance-type virtual-switch
set routing-instances OF-ri interface ge-1/0/0.0
set routing-instances OF-ri interface ge-1/1/0.0
set routing-instances OF-ri interface xe-0/0/0.0
set routing-instances OF-ri vlans of-bridge vlan-id none
set protocols openflow switch OFswitch1 controller address 198.51.100.174
set protocols openflow switch OFswitch1 controller protocol tcp port 6633
set protocols openflow switch OFswitch1 interfaces ge-1/0/0.0
set protocols openflow switch OFswitch1 interfaces ge-1/1/0.0
set protocols openflow switch OFswitch1 interfaces xe-0/0/0.0
set protocols openflow switch OFswitch1 default-action packet-in
set protocols openflow traceoptions flag all
```

### Step-by-Step Procedure

To configure support for OpenFlow:

1. Configure the OpenFlow interfaces as Layer 2 interfaces.

```
[edit interfaces]
user@switch# set ge-1/0/0 unit 0 family ethernet-switching
user@switch# set ge-1/1/0 unit 0 family ethernet-switching
user@switch# set xe-0/0/0 unit 0 family ethernet-switching
```

2. Configure the virtual switch routing instance.

```
[edit routing-instances]
user@switch# set OF-ri instance-type virtual-switch
user@switch# set OF-ri interface ge-1/0/0.0
user@switch# set OF-ri interface ge-1/1/0.0
user@switch# set OF-ri interface xe-0/0/0.0
user@switch# set OF-ri vlans of-bridge vlan-id none
```

3. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch OFswitch1]
user@switch# set controller address 198.51.100.174
```



```
user@switch# set controller protocol tcp port 6633
```

4. Configure the logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch OFswitch1]
user@switch# set interfaces ge-1/0/0.0
user@switch# set interfaces ge-1/1/0.0
user@switch# set interfaces xe-0/0/0.0
```

5. Configure the default action for packets that do not have a matching flow entry.

```
[edit protocols openflow switch OFswitch1]
user@switch# set default-action packet-in
```

6. Configure OpenFlow traceoptions.

```
[edit protocols openflow]
user@switch# set traceoptions flag all
```

7. Commit the configuration.

```
[edit]
user@switch# commit
```

## Results

From operational mode, display the results of your configuration by entering the **show configuration interfaces**, **show configuration protocols openflow**, and **show configuration routing-instances** commands. If the output does not display the specified configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show configuration interfaces
ge-1/0/0 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-1/1/0 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/0 {
  unit 0 {
    family ethernet-switching;
  }
}

user@switch> show configuration protocols openflow
switch OFswitch1 {
  default-action {
    packet-in;
  }
  interfaces {
    ge-1/0/0.0;
    ge-1/1/0.0;
```

```
        xe-0/0/0.0;
    }
    controller {
        address 198.51.100.174;
        protocol tcp {
            port 6633;
        }
    }
    traceoptions {
        flag all;
    }
}

user@switch> show configuration routing-instances
OF-ri {
    instance-type virtual-switch;
    interface ge-1/0/0.0;
    interface ge-1/1/0.0;
    interface xe-0/0/0.0;
    vlans {
        of-bridge {
            vlan-id none;
        }
    }
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying the OpenFlow Controller Connection on page 2112](#)
- [Verifying the OpenFlow Interfaces on page 2113](#)

### Verifying the OpenFlow Controller Connection

---

**Purpose** Verify that the OpenFlow controller connection is up.

**Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**. Because the virtual switch configuration has only a single controller, the virtual switch automatically initiates a connection to the controller after you commit the configuration.

```
user@switch> show openflow controller
Openflowd controller information:
Controller socket: 11
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 5
Controller role: equal
```

**Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying the OpenFlow Interfaces

|                              |  |
|------------------------------|--|
| <b>Purpose</b>               | Verify that the OpenFlow interfaces are up.  |
| <b>Action</b>                | <p>Issue the <b>show openflow interfaces</b> operational mode command, and verify that the state of each OpenFlow interface is <b>Up</b>.</p> <pre> user@switch&gt; show openflow interfaces Switch name: OFswitch1 Interface Name: ge-1/0/0.0 Interface port number: 41507 Interface Hardware Address: 00:00:5E:00:53:b1 Interface speed: 1Gb Full-duplex Interface Auto-Negotiation: Disabled Interface media type: Fiber <b>Interface state: Up</b>  Switch name: OFswitch1 Interface Name: ge-1/1/0.0 Interface port number: 44538 Interface Hardware Address: 00:00:5E:00:53:b2 Interface speed: 1Gb Full-duplex Interface Auto-Negotiation: Disabled Interface media type: Fiber <b>Interface state: Up</b>  Switch name: OFswitch1 Interface Name: xe-0/0/0.0 Interface port number: 45549 Interface Hardware Address: 00:00:5E:00:53:b3 Interface speed: 10Gb Full-duplex Interface Auto-Negotiation: Disabled Interface media type: Fiber <b>Interface state: Up</b> </pre> |
| <b>Meaning</b>               | The output shows that the state of each OpenFlow interface is <b>Up</b> , in addition to other information about the interfaces.   |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li> <li>• <a href="#">Installing Support for OpenFlow on Devices Running Junos OS on page 2097</a></li> <li>• <a href="#">Configuring Support for OpenFlow on EX9200 Switches on page 2106</a></li> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">openflow (Protocols OpenFlow) on page 2175</a></li> <li>• <a href="#">[edit protocols openflow] Hierarchy Level on page 2169</a></li> </ul>   |

### Configuring Support for OpenFlow on QFX5100 Switches

This topic describes how to configure QFX5100 switches with interfaces that participate solely in OpenFlow.

Before configuring support for OpenFlow, ensure that the switch meets the following requirements:

- QFX5100 switch running Junos OS Release 14.1X53-D10 or later
- OpenFlow software package with a release that matches the Junos OS release running on the switch
- TCP connection between the switch and an OpenFlow controller
- Connection between the management interface (em0 or em1) of the switch and the management network

Configuration tasks are described in detail in the following sections:

- [Configuring the OpenFlow Interfaces on page 2114](#)
- [Configuring the OpenFlow Protocol on page 2114](#)

## Configuring the OpenFlow Interfaces

You must configure interfaces participating in OpenFlow as Layer 2 interfaces. On QFX5100 switches, you configure the interfaces with protocol family **ethernet-switching**. Also, you can configure only a single logical port by specifying logical unit number 0.

To configure the OpenFlow interfaces:

- Configure the logical interface and protocol family.

```
[edit interfaces interface-name]  
user@switch# set unit 0 family ethernet-switching
```

## Configuring the OpenFlow Protocol

To configure support for OpenFlow, you must create a virtual switch, and then, configure the connection with the OpenFlow controller and the logical interfaces participating in OpenFlow for the virtual switch. Optionally, configure the default action for packets that do not match a flow entry, the purge timer for invalid flows, and any OpenFlow traceoptions.

To configure the OpenFlow protocol:

1. Create an OpenFlow virtual switch, and specify a switch name, which can contain a maximum of 60 characters.

```
[edit protocols openflow]  
user@switch# set switch switch-name
```

2. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch switch-name]  
user@switch# set controller address address  
user@switch# set controller protocol tcp
```

3. Specify the logical interfaces participating in OpenFlow under this virtual switch.

```
[edit protocols openflow switch switch-name]  
user@switch# set interfaces interface-name.0
```

```
user@switch# set interfaces interface-name 2.0
```

4. (Optional) Configure the **default-action** statement for packets that do not match an existing flow entry.

If you do not configure the **default-action** statement, the default is **packet-in**, which indicates that packets with no matching flow entry are sent to the controller for processing.

```
[edit protocols openflow switch switch-name]
user@switch# set default-action (drop | packet-in)
```

5. (Optional) Configure the **purge-flow-timer** statement, which specifies the number of seconds after which an invalid flow is purged from the flow table.

```
[edit protocols openflow switch switch-name]
user@switch# set purge-flow-timer seconds
```

6. (Optional) Configure OpenFlow traceoptions.

If you do not configure a log file by specifying its filename, OpenFlow trace messages are logged in the default OpenFlow log file `/var/log/ofd`.

```
[edit protocols openflow]
user@switch# set traceoptions flag all
user@switch# set traceoptions file file-name
```

#### Related Documentation

- [Example: Enabling OpenFlow on QFX5100 Switches on page 2115](#)
- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
- [OpenFlow Operational Mode Commands on page 2183](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)
- [\[edit protocols openflow\] Hierarchy Level on page 2169](#)

## Example: Enabling OpenFlow on QFX5100 Switches

OpenFlow is an open standard that enables you to control traffic paths in a network by creating, deleting, and modifying flows in each device along a path. This example shows how to configure OpenFlow support on a QFX5100 switch.

To isolate and control OpenFlow traffic on a QFX5100 switch, you configure a virtual switch. You also configure a Secure Sockets Layer (SSL) or TCP/IP connection between the virtual switch and a remote OpenFlow controller. Using this connection, the OpenFlow controller can access the flows in the virtual switch.

- [Requirements on page 2116](#)
- [Overview on page 2116](#)
- [Configuration on page 2116](#)
- [Verification on page 2118](#)

## Requirements

This example uses the following hardware and software components:

- A QFX5100 switch running Junos OS Release 14.1X53-D10 or later.
- An OpenFlow software package is installed on the switch, and the release of this package matches the Junos OS release running on the switch.
- A TCP connection between the switch and an OpenFlow controller.
- A connection between the management interface (em0 or em1) of the switch and the management network.

## Overview

In this example, you configure support for OpenFlow on a QFX5100 switch. The switch has three interfaces that are dedicated to handling OpenFlow traffic: xe-0/0/10.0, xe-0/0/11.0, and xe-0/0/12.0. Note that on QFX5100 switches, you can configure only a single logical interface, using logical unit number 0 for each OpenFlow interface.

In an OpenFlow topology, a virtual switch is used to isolate and control OpenFlow traffic. You configure the OpenFlow virtual switch and OpenFlow protocol statements at the **[edit protocols openflow]** hierarchy level.

Virtual switch 100 also connects to an OpenFlow controller over a TCP connection at the IP address 10.51.100.174. The virtual switch configuration must include all of the logical interfaces participating in OpenFlow; OpenFlow traffic enters and exits only through these interfaces.

A flow entry consists of a match condition against which packets entering an OpenFlow interface are compared, and the action that is applied to packets that match the condition. Each OpenFlow interface can have one or more flow entries. The **default-action** statement in the OpenFlow configuration indicates the action the switch applies to packets that do not have a matching flow entry. This example uses the **drop** option, which specifies that packets that do not match a flow entry are dropped.

This example also configures OpenFlow traceoptions, along with the **flag all** statement, which captures and logs all OpenFlow events. This example does not configure a specific filename for the log file. As a result, OpenFlow events are logged in the default OpenFlow log directory **/var/log/ofd**.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-0/0/10 unit 0 family ethernet-switching
set interfaces xe-0/0/11 unit 0 family ethernet-switching
set interfaces xe-0/0/12 unit 0 family ethernet-switching
set protocols openflow switch 100 controller address 10.51.100.174
set protocols openflow switch 100 controller protocol tcp
```

```

set protocols openflow switch 100 interfaces xe-0/0/10.0
set protocols openflow switch 100 interfaces xe-0/0/11.0
set protocols openflow switch 100 interfaces xe-0/0/12.0
set protocols openflow switch 100 default-action drop
set protocols openflow traceoptions flag all

```

### Step-by-Step Procedure

To configure support for OpenFlow:

1. Configure the OpenFlow interfaces as Layer 2 interfaces.

```

[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching
user@switch# set xe-0/0/11 unit 0 family ethernet-switching
user@switch# set xe-0/0/12 unit 0 family ethernet-switching

```

2. Configure an OpenFlow virtual switch named 100.

```

[edit protocols openflow]
user@switch# set switch 100

```

3. Configure the OpenFlow controller IP address and the connection protocol.

```

[edit protocols openflow switch 100]
user@switch# set controller address 10.51.100.174
user@switch# set controller protocol tcp

```

4. Configure the logical interfaces in this virtual switch that participate in OpenFlow.

```

[edit protocols openflow switch 100]
user@switch# set interfaces xe-0/0/10.0
user@switch# set interfaces xe-0/0/11.0
user@switch# set interfaces xe-0/0/12.0

```

5. Configure the default action for packets that do not have a matching flow entry.

```

[edit protocols openflow switch 100]
user@switch# set default-action drop

```

6. Configure OpenFlow traceoptions.

```

[edit protocols openflow]
user@switch# set traceoptions flag all

```

7. Commit the configuration.

```

[edit]
user@switch# commit

```

### Results

From operational mode, confirm your configuration by entering the **show configuration interfaces** and **show configuration protocols openflow** commands.

```

user@switch> show configuration interfaces
xe-0/0/10 {
  unit 0 {
    family ethernet-switching;
  }
}

```

```
}
xe-0/0/11 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/12 {
  unit 0 {
    family ethernet-switching;
  }
}

user@switch> show configuration protocols openflow
switch 100 {
  default-action {
    drop;
  }
  interfaces {
    xe-0/0/10.0;
    xe-0/0/11.0;
    xe-0/0/12.0;
  }
  controller {
    protocol {
      tcp {
        address 10.51.100.174;
      }
    }
  }
}
traceoptions {
  flag all;
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying That the OpenFlow Controller Connection Is Up on page 2118](#)
- [Verifying that the OpenFlow Interfaces Are Up on page 2119](#)

### Verifying That the OpenFlow Controller Connection Is Up

**Purpose** Verify that the OpenFlow controller connection is up.

**Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**. Because the virtual switch configuration has only a single controller, the virtual switch automatically initiates a connection to the controller after you commit the configuration.

```
user@switch> show openflow controller
Openflowd controller information:
Controller socket: 12
Controller IP address: 10.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 4
Controller role: equal
```



Negotiated version: 4

**Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying that the OpenFlow Interfaces Are Up

**Purpose** Verify that the OpenFlow interfaces are up.

**Action** Issue the **show openflow interfaces** operational mode command, and verify that the state of each OpenFlow interface is **Up**.

```
user@switch> show openflow interfaces
Switch name: 100
Interface Name: xe-0/0/10.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:00
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up

Switch name: 100
Interface Name: xe-0/0/11.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:01
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up

Switch name: 100
Interface Name: xe-0/0/12.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:02
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

**Meaning** The output shows that the state of each OpenFlow interface is **Up**, in addition to other information about the interfaces.

- Related Documentation**
- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
  - [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
  - [Configuring Support for OpenFlow on QFX5100 Switches on page 2113](#)
  - [OpenFlow Operational Mode Commands on page 2183](#)
  - [openflow \(Protocols OpenFlow\) on page 2175](#)
  - [\[edit protocols openflow\] Hierarchy Level on page 2169](#)

## Configuring Support for OpenFlow on EX4550 Switches

---

The following sections configure EX4550 switches to support OpenFlow using interfaces that participate solely in OpenFlow.

Before configuring support for OpenFlow, ensure that the switch meets the following requirements:

- EX4550 switch running Junos OS Release 13.2X51-D20 or a later release
- OpenFlow software package with a software package release that matches the Junos OS release of the device on which it is installed
- TCP connection between the switch and an OpenFlow controller
- Connection between the management interface of the switch and the management network, which is reachable from the controller IP address

Configuration tasks are described in detail in the following sections:

- [Configuring the OpenFlow Interfaces on page 2120](#)
- [Configuring the OpenFlow Protocol on page 2120](#)

### Configuring the OpenFlow Interfaces

You must configure interfaces participating in OpenFlow as Layer 2 interfaces. On EX Series Ethernet Switches, you configure the interfaces with protocol family **ethernet-switching**. On EX4550 switches, you can configure only a single logical port using logical unit number 0.

To configure the OpenFlow interfaces:

- Configure the logical interface and the protocol family.

```
[edit interfaces interface-name]  
user@host# set unit 0 family ethernet-switching
```

### Configuring the OpenFlow Protocol

To configure support for OpenFlow, create a virtual switch instance, and specify a switch name, which must be 60 characters or less. For the virtual switch instance, configure the OpenFlow controller information and the participating logical interfaces. Optionally, configure the default action for packets that do not match a flow entry, the purge timer for invalid flows, and any OpenFlow traceoptions.

To configure the OpenFlow protocol:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch switch-name]  
user@host# set controller address address  
user@host# set controller protocol tcp
```

2. Specify the logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch switch-name]
user@host# set interfaces interface-name1.0
user@host# set interfaces interface-name2.0
```

3. (Optional) Configure the **default-action** statement for packets that do not match on an existing flow entry.

If you do not configure the **default-action** statement, the default is **packet-in**, which indicates that packets with no matching flow entry must be sent to the controller for processing.

```
[edit protocols openflow switch switch-name]
user@host# set default-action (drop | packet-in)
```

4. (Optional) Configure the **purge-flow-timer** statement, which is the number of seconds after which an invalid flow is purged from the flow table.

```
[edit protocols openflow switch switch-name]
user@host# set purge-flow-timer seconds
```

5. (Optional) Configure OpenFlow traceoptions.

If you do not configure a log filename, OpenFlow trace messages are logged in the default OpenFlow log file `/var/log/ofd`.

```
[edit protocols openflow]
user@host# set traceoptions flag all
user@host# set traceoptions file file-name
```

#### Related Documentation

- [Example: Enabling OpenFlow on EX4550 Switches on page 2121](#)
- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)
- [OpenFlow Operational Mode Commands on page 2183](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)
- [\[edit protocols openflow\] Hierarchy Level on page 2169](#)

## Example: Enabling OpenFlow on EX4550 Switches

OpenFlow is an open standard that enables you to control traffic paths in a network by creating, deleting, and modifying flows in each device, including EX4550 switches that have an OpenFlow software package installed, along a path. This example shows how to configure OpenFlow support on an EX4550 switch.

- [Requirements on page 2122](#)
- [Overview on page 2122](#)
- [Configuration on page 2123](#)
- [Verification on page 2124](#)

## Requirements

This example uses the following hardware and software components:

- An EX4550 switch running Junos OS Release 13.2X51-D20 or a later release.
- An OpenFlow software package is installed on the switch, and the software package release matches the Junos OS release running on the switch.
- A TCP connection between the switch and an OpenFlow controller, which needs to access the data plane of the switch.
- A connection between the me0 interface of the switch and the management network.

## Overview

In this example, you configure support for OpenFlow on an EX4550 switch. The switch has three interfaces that are dedicated to handling OpenFlow traffic: xe-0/0/4.0, xe-0/0/5.0, and xe-0/0/6.0. Note that on EX4550 switches, you can configure only a single logical unit by using logical unit number 0 for OpenFlow interfaces.

In an OpenFlow topology, a virtual switch is used to isolate and control OpenFlow traffic. You configure the OpenFlow virtual switch and OpenFlow protocol statements at the **[edit protocols openflow]** hierarchy level. In this example, the virtual switch, 100, is assigned a default VLAN, which acts as a logically separate flood domain. The assignment of the default VLAN to virtual switch 100 is automatic, and no configuration is required to set up the default VLAN.

Virtual switch 100 also connects to the controller over a TCP connection at the IP address 198.51.100.174. The virtual switch configuration must include all of the logical interfaces participating in OpenFlow, and OpenFlow traffic will only enter or exit from these interfaces.

A flow entry consists of a match condition against which packets entering an OpenFlow interface are compared, and the action that is applied to packets that match the condition. Each OpenFlow interface can have one or more flow entries. The **default-action** statement in the OpenFlow configuration indicates the action the switch applies to packets that do not have a matching flow entry. If you omit the **default-action** statement, the default action is **packet-in**, which means that packets that have no matching flow entry are sent to the controller for processing. This example explicitly configures **packet-in** as the default action for packets that do not have a matching flow entry.

This example also configures OpenFlow traceoptions, along with the **flag all** statement, which captures and logs all OpenFlow events. This example does not configure a specific filename for the log file. As a result, OpenFlow events are logged in the default OpenFlow log directory **/var/log/ofd**.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-0/0/4 unit 0 family ethernet-switching
set interfaces xe-0/0/5 unit 0 family ethernet-switching
set interfaces xe-0/0/6 unit 0 family ethernet-switching
set protocols openflow switch 100 controller address 198.51.100.174
set protocols openflow switch 100 controller protocol tcp
set protocols openflow switch 100 interfaces xe-0/0/4.0
set protocols openflow switch 100 interfaces xe-0/0/5.0
set protocols openflow switch 100 interfaces xe-0/0/6.0
set protocols openflow switch 100 default-action packet-in
set protocols openflow traceoptions flag all
```

**Step-by-Step Procedure** To configure support for OpenFlow:

1. Configure the OpenFlow interfaces as Layer 2 interfaces.  
  

```
[edit interfaces]
user@switch# set xe-0/0/4 unit 0 family ethernet-switching
user@switch# set xe-0/0/5 unit 0 family ethernet-switching
user@switch# set xe-0/0/6 unit 0 family ethernet-switching
```
2. Configure an OpenFlow virtual switch.  
  

```
[edit protocols openflow]
user@switch# set switch 100
```
3. Configure the OpenFlow controller IP address and the connection protocol.  
  

```
[edit protocols openflow switch 100]
user@switch# set controller address 198.51.100.174
user@switch# set controller protocol tcp
```
4. Configure the logical interfaces participating in OpenFlow under this virtual switch.  
  

```
[edit protocols openflow switch 100]
user@switch# set interfaces xe-0/0/4.0
user@switch# set interfaces xe-0/0/5.0
user@switch# set interfaces xe-0/0/6.0
```
5. Configure the default action for packets that do not have a matching flow entry.  
  

```
[edit protocols openflow switch 100]
user@switch# set default-action packet-in
```
6. Configure OpenFlow traceoptions.  
  

```
[edit protocols openflow]
user@switch# set traceoptions flag all
```
7. Commit the configuration.  
  

```
[edit]
user@switch# commit
```

## Results

---

From operational mode, confirm your configuration by entering the **show configuration interfaces** and **show configuration protocols openflow** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@switch> show configuration interfaces
xe-0/0/4 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/5 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/0/6 {
  unit 0 {
    family ethernet-switching;
  }
}

user@switch> show configuration protocols openflow
switch 100 {
  default-action packet-in;
  interfaces {
    xe-0/0/4.0;
    xe-0/0/5.0;
    xe-0/0/6.0;
  }
  controller {
    address 198.51.100.174;
    protocol tcp;
  }
}
traceoptions {
  flag all;
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying that the OpenFlow Controller Connection Is Up on page 2124](#)
- [Verifying that the OpenFlow Interfaces Are Up on page 2125](#)

### Verifying that the OpenFlow Controller Connection Is Up

---

**Purpose** Verify that the OpenFlow controller connection is up.

**Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**. Because the virtual switch configuration has only a

single controller, the virtual switch should automatically initiate a connection to the controller after you commit the configuration.

```
user@switch> show openflow controller
Openflowd controller information:
Controller socket: 12
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 4
Controller role: equal
```

**Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying that the OpenFlow Interfaces Are Up

**Purpose** Verify that the OpenFlow interfaces are up.

**Action** Issue the **show openflow interfaces** operational mode command, and verify that the state of each OpenFlow interface is **Up**.

```
user@switch> show openflow interfaces
Switch name: 100
Interface Name: xe-0/0/4.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:00
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: 100
Interface Name: xe-0/0/5.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:01
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: 100
Interface Name: xe-0/0/6.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:02
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```

**Meaning** The output shows that the state of each OpenFlow interface is **Up**, in addition to other information about the interfaces.

**Related Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Installing Support for OpenFlow on Devices Running Junos OS on page 2097](#)

- [Configuring Support for OpenFlow on EX4550 Switches on page 2120](#)
- [OpenFlow Operational Mode Commands on page 2183](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)
- [\[edit protocols openflow\] Hierarchy Level on page 2169](#)



# Configuring OpenFlow Hybrid Interfaces

- [Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS on page 2127](#)
- [Configuring OpenFlow Hybrid Interfaces on MX Series Routers on page 2128](#)
- [Example: Configuring OpenFlow Hybrid Interfaces on MX Series Routers on page 2131](#)
- [Configuring OpenFlow Hybrid Interfaces on EX9200 Switches on page 2138](#)
- [Example: Configuring OpenFlow Hybrid Interfaces on EX9200 Switches on page 2140](#)

## Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS

---

On Juniper Networks EX9200 Ethernet Switches and on MX Series 3D Universal Edge Routers that support OpenFlow, you can configure physical interfaces that support multiple logical interfaces as hybrid interfaces. A hybrid interface concurrently supports OpenFlow logical interfaces and non-OpenFlow logical interfaces.

On a hybrid interface, the OpenFlow protocol and the non-OpenFlow protocols essentially exist independently. Traffic does not get forwarded across OpenFlow and non-OpenFlow logical interfaces. Instead VLANs and VLAN tags are used to distinguish the OpenFlow traffic from the normal traffic. To accomplish this, you must enable the reception and transmission of 802.1Q VLAN-tagged frames on all interfaces, including both hybrid and non-hybrid interfaces. You must also configure separate virtual switch routing instances for OpenFlow traffic and for normal traffic, which serve to separate the VLAN ID space.

On devices using hybrid interfaces, traffic entering an interface must be VLAN-tagged. The VLAN ID differentiates the OpenFlow traffic from the normal traffic, and on the hybrid interface, the VLAN ID also determines the associated logical interface. Once the logical interface is known, the traffic is forwarded accordingly. The device forwards OpenFlow traffic according to OpenFlow flow entries, and it forwards normal traffic using traditional Layer 2 and Layer 3 processing. If you do not configure a native VLAN, untagged packets are dropped.

On a hybrid interface, you configure a logical interface as a trunk interface, which accepts and forwards tagged packets from multiple VLANs. Additionally, you can configure certain non-OpenFlow logical interfaces as Layer 3 subinterfaces that perform traditional Layer 3 or MPLS-based forwarding.

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, or a range or list of VLAN IDs, to the logical interface. OpenFlow interfaces must have a different set of VLANs from normal interfaces. On a hybrid interface,

OpenFlow traffic can only egress from an interface that has the same VLAN ID range as that of the ingress interface.

A hybrid interface configuration with multiple logical interfaces permits OpenFlow and non-OpenFlow traffic to traverse the same interface while keeping the traffic in separate routing or bridging domains. One advantage of using hybrid interfaces is that you can use fewer physical interfaces where port density is an issue. However, using hybrid interfaces requires some additional configuration, and untagged traffic entering a hybrid port cannot be forwarded according to OpenFlow flow entries. Additionally, several physical port properties such as Layer 1 statistics are reported for all logical interfaces on that physical interface. Thus, when you configure a physical interface in hybrid mode, these properties are reported for all OpenFlow and non-OpenFlow logical interfaces on that physical interface. These properties include queue drops, framing errors, CRC errors, and collisions. When using hybrid interfaces, if you use the Link Layer Discovery Protocol (LLDP) for topology discovery, you must ensure that any LLDP frames entering a hybrid interface are tagged appropriately.

**Related  
Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [Configuring OpenFlow Hybrid Interfaces on MX Series Routers on page 2128](#)
- [Configuring OpenFlow Hybrid Interfaces on EX9200 Switches on page 2138](#)
- [\*Binding VLAN IDs to Logical Interfaces\*](#)

---

## Configuring OpenFlow Hybrid Interfaces on MX Series Routers

---

On MX Series routers that support OpenFlow, you can configure a physical interface as a hybrid interface that concurrently supports OpenFlow logical interfaces and non-OpenFlow logical interfaces. If you configure an OpenFlow hybrid interface on a device running Junos OS, you must enable the reception and transmission of 802.1Q VLAN-tagged frames on all interfaces, including both hybrid and non-hybrid interfaces, and you must configure a virtual switch routing instance for the OpenFlow traffic and a separate virtual switch routing instance for the normal traffic.

The following sections detail configuring an MX Series router that supports OpenFlow with a mix of hybrid and normal interfaces:

- [Configuring the Hybrid Physical Interface on page 2129](#)
- [Configuring the Hybrid Interface Logical Units on page 2129](#)
- [Configuring the Non-Hybrid Interfaces on page 2129](#)
- [Configuring OpenFlow on page 2130](#)
- [Configuring the Virtual Switch Routing Instances on page 2130](#)

## Configuring the Hybrid Physical Interface

To configure the hybrid physical interface:

1. Enable VLAN tagging.

Configure **vlan-tagging** to support 802.1Q VLAN single-tag frames for both OpenFlow and non-OpenFlow traffic, or configure **flexible-vlan-tagging** to support both 802.1Q VLAN single-tag and dual-tag frames.

```
[edit interfaces interface-name]
user@host# set (vlan-tagging | flexible-vlan-tagging)
```

2. Configure flexible Ethernet services encapsulation to enable multiple per-unit Ethernet encapsulations.

```
[edit interfaces interface-name]
user@host# set encapsulation flexible-ethernet-services
```

## Configuring the Hybrid Interface Logical Units

On a hybrid interface, you configure an OpenFlow or non-OpenFlow logical interface as a Layer 2 trunk interface. Additionally, you can configure a non-OpenFlow logical interface as a Layer 3 subinterface that performs traditional Layer 3 or MPLS-based forwarding. To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, or a range or list of VLAN IDs, to the logical interface. Configure Layer 2 interfaces using family **bridge** on MX Series routers.

To configure the hybrid interface logical units:

1. Configure the OpenFlow logical interfaces and any non-OpenFlow Layer 2 logical interfaces, and specify the interface mode and VLAN membership.

```
[edit interfaces interface-name]
user@host# set unit unit family bridge interface-mode trunk
user@host# set unit unit family bridge vlan-id-list vlan-ids
```

2. Configure any non-OpenFlow Layer 3 logical interfaces, and specify the VLAN membership.

```
[edit interfaces interface-name]
user@host# set unit unit (vlan-id | vlan-id-list | vlan-id-range) vlan-ids
user@host# set unit unit family inet address address
```

## Configuring the Non-Hybrid Interfaces

Non-hybrid interfaces support either OpenFlow traffic or non-OpenFlow traffic, but not both simultaneously.

To configure the non-hybrid interfaces:

1. Configure interfaces that support only OpenFlow traffic as Layer 2 trunk interfaces, and specify the interface mode and VLAN membership.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
```

```
user@host# set unit unit family bridge interface-mode trunk
user@host# set unit unit family bridge vlan-id-list vlan-ids
```

2. Configure interfaces that support only normal traffic, and specify the interface mode for the Layer 2 interfaces and the VLAN membership.

For example:

```
[edit interfaces interface-name]
user@host# set vlan-tagging
user@host# set unit unit family bridge interface-mode trunk
user@host# set unit unit family bridge vlan-id-list vlan-ids
```

## Configuring OpenFlow

To configure the OpenFlow virtual switch instance:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch switch-name]
user@host# set controller address address
user@host# set controller protocol tcp port port
```

2. Specify all logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch switch-name]
user@host# set interfaces interface-name
```

## Configuring the Virtual Switch Routing Instances

Configure separate virtual switch routing instances for the OpenFlow traffic and the non-OpenFlow traffic. The configured interface names must include a logical unit number.

To configure the virtual switch routing instances:

1. Configure the virtual switch routing instance for the OpenFlow traffic, and specify the OpenFlow logical interfaces and VLANs.

```
[edit routing-instances of-routing-instance-name]
user@host# set instance-type virtual-switch
user@host# set interface of-interface-name1
user@host# set interface of-interface-name2
user@host# set bridge-domains name vlan-id-list of-vlan-id-list
```

2. Configure the virtual switch routing instance for the non-OpenFlow traffic, and specify the non-OpenFlow logical interfaces and VLANs.

```
[edit routing instances non-of-routing-instance-name]
user@host# set instance-type virtual-switch
user@host# set interface non-of-interface-name1
user@host# set interface non-of-interface-name2
user@host# set bridge-domains name vlan-id-list non-of-vlan-id-list
```

### Related Documentation

- [Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS on page 2127](#)
- [Example: Configuring OpenFlow Hybrid Interfaces on MX Series Routers on page 2131](#)

- *OpenFlow Feature Guide*

## Example: Configuring OpenFlow Hybrid Interfaces on MX Series Routers

On MX series routers that support OpenFlow, you can configure physical interfaces that support multiple logical interfaces as hybrid interfaces. A hybrid interface concurrently supports both OpenFlow logical interfaces and non-OpenFlow logical interfaces, thus allowing OpenFlow and non-OpenFlow traffic to traverse the same interface.

Hybrid interfaces enable you to use your physical interfaces more efficiently, especially in a situation where port density is an issue.

This example shows how to configure an MX Series router with OpenFlow hybrid interfaces.

- [Requirements on page 2131](#)
- [Overview on page 2131](#)
- [Configuration on page 2133](#)
- [Verification on page 2136](#)

### Requirements

This example uses the following hardware and software components:

- MX240 router running Junos OS Release 13.3 or a later release
- OpenFlow software package with a software package release that matches the Junos OS release of the device on which it is installed
- TCP connection between the router and an OpenFlow controller
- Connection between the fxp0 management interface of the router and the management network, which is reachable from the OpenFlow controller IP address

### Overview

In this example, you configure an MX240 router with a hybrid interface, ge-1/0/1, an OpenFlow interface, ge-1/0/2, and a non-OpenFlow interface, ge-1/0/3. On the hybrid interface, logical interface ge-1/0/1.0 participates in OpenFlow, and logical interfaces ge-1/0/1.1 and ge-1/0/1.2 do not participate in OpenFlow.

When using OpenFlow hybrid interfaces, you use VLANs to distinguish the OpenFlow traffic from the normal traffic. Thus, you must enable VLAN tagging on all interfaces, and traffic entering the interfaces must be vlan-tagged. Untagged traffic entering the hybrid interface is dropped. In this example, you configure the hybrid interface using **flexible-vlan-tagging**, which enables VLAN tagging and supports both 802.1Q VLAN single-tag and dual-tag frames for all traffic on the interface. You configure interfaces ge-1/0/2 and ge-1/0/3 using **vlan-tagging**.

You configure the hybrid interface encapsulation as flexible Ethernet services. Note that for interfaces with this encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511

are no longer reserved for normal VLANs. In this example, VLANs 1 through 100 are used for OpenFlow traffic, and VLANs 101 through 200 and VLAN 300 are used for normal traffic.

All logical interfaces except ge-1/0/1.2 are configured as Layer 2 trunk interfaces using family **bridge** and interface mode **trunk**. Logical interfaces ge-1/0/1.0 and ge-1/0/2.0 participate in OpenFlow and receive and forward traffic with OpenFlow VLAN IDs 1 through 100. Logical interfaces ge-1/0/1.1 and ge-1/0/3.0 do not participate in OpenFlow and receive and forward traffic with non-OpenFlow VLAN IDs 101 through 200.

ge-1/0/1.2 is a Layer 3 logical interface with IP address 198.51.100.10/24 that performs Layer 3 routing. This interface does not participate in OpenFlow and routes traffic with VLAN ID 300.

[Table 186 on page 2132](#) summarizes the logical interfaces, traffic type, and associated VLAN IDs.

**Table 186: Summary of Logical Interfaces**

| Logical Interface | Traffic Type | VLANs           |
|-------------------|--------------|-----------------|
| ge-1/0/1.0        | OpenFlow     | 1 through 100   |
| ge-1/0/1.1        | non-OpenFlow | 101 through 200 |
| ge-1/0/1.2        | non-OpenFlow | 300             |
| ge-1/0/2.0        | OpenFlow     | 1 through 100   |
| ge-1/0/3.0        | non-OpenFlow | 101 through 200 |

You configure the OpenFlow virtual switch and OpenFlow protocol statements at the **[edit protocols openflow]** hierarchy level. The virtual switch, OFswitch2, connects to the controller over a TCP connection at IP address 172.16.1.1. The virtual switch configuration must include all of the logical interfaces participating in OpenFlow, which includes ge-1/0/1.0 and ge-1/0/2.0.

When configuring OpenFlow on MX Series routers, you must configure a virtual switch routing instance for the OpenFlow traffic that isolates it from the normal network traffic. Additionally, when using hybrid interfaces, you configure both a virtual switch routing instance for the OpenFlow traffic and also a separate virtual switch routing instance for the normal traffic. In this example, you configure routing instance rt1 for the OpenFlow traffic and routing instance rt2 for the normal traffic.

Routing instance rt1 includes the interfaces participating in OpenFlow, ge-1/0/1.0 and ge-1/0/2.0. Within the routing instance you configure the bridge domain to include all OpenFlow VLANs 1 through 100. Routing instance rt2 includes the Layer 2 interfaces that do not participate in OpenFlow, ge-1/0/1.1 and ge-1/0/3.0. Within the routing instance you configure the bridge domain to include the non-OpenFlow VLANs 101 through 200.



**NOTE:** In order to direct OpenFlow traffic, the OpenFlow controller must install flow entries that select the appropriate traffic and forward it to the correct OpenFlow interface.

## Configuration

- [Configuring the Interfaces on page 2133](#)
- [Configuring OpenFlow on page 2134](#)
- [Configuring the Virtual Switch Routing Instances on page 2135](#)
- [Results on page 2135](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-1/0/1 flexible-vlan-tagging
set interfaces ge-1/0/1 encapsulation flexible-ethernet-services
set interfaces ge-1/0/1 unit 0 family bridge interface-mode trunk
set interfaces ge-1/0/1 unit 0 family bridge vlan-id-list 1-100
set interfaces ge-1/0/1 unit 1 family bridge interface-mode trunk
set interfaces ge-1/0/1 unit 1 family bridge vlan-id-list 101-200
set interfaces ge-1/0/1 unit 2 vlan-id 300
set interfaces ge-1/0/1 unit 2 family inet address 198.51.100.10/24
set interfaces ge-1/0/2 vlan-tagging
set interfaces ge-1/0/2 unit 0 family bridge interface-mode trunk
set interfaces ge-1/0/2 unit 0 family bridge vlan-id-list 1-100
set interfaces ge-1/0/3 vlan-tagging
set interfaces ge-1/0/3 unit 0 family bridge interface-mode trunk
set interfaces ge-1/0/3 unit 0 family bridge vlan-id-list 101-200
set protocols openflow switch OFswitch2 controller address 172.16.1.1
set protocols openflow switch OFswitch2 controller protocol tcp port 6633
set protocols openflow switch OFswitch2 interfaces ge-1/0/1.0
set protocols openflow switch OFswitch2 interfaces ge-1/0/2.0
set routing-instances rt1 instance-type virtual-switch
set routing-instances rt1 interface ge-1/0/1.0
set routing-instances rt1 interface ge-1/0/2.0
set routing-instances rt1 bridge-domains bd-of vlan-id-list 1-100
set routing-instances rt2 instance-type virtual-switch
set routing-instances rt2 interface ge-1/0/1.1
set routing-instances rt2 interface ge-1/0/3.0
set routing-instances rt2 bridge-domains bd-nonof vlan-id-list 101-200
```

### Configuring the Interfaces

#### Step-by-Step Procedure

To configure the interfaces:

1. On the hybrid physical interface, enable VLAN tagging and configure the encapsulation.

```
[edit interfaces ge-1/0/1]
user@host# set flexible-vlan-tagging
```

```
user@host# set encapsulation flexible-ethernet-services
```

2. Configure OpenFlow logical interface ge-1/0/1.0 as a Layer 2 trunk that supports VLANs 1-100.

```
[edit interfaces ge-1/0/1]
user@host# set unit 0 family bridge interface-mode trunk
user@host# set unit 0 family bridge vlan-id-list 1-100
```

3. Configure non-OpenFlow logical interface ge-1/0/1.1 as a Layer 2 trunk that supports VLANs 101-200.

```
[edit interfaces ge-1/0/1]
user@host# set unit 1 family bridge interface-mode trunk
user@host# set unit 1 family bridge vlan-id-list 101-200
```

4. Configure non-OpenFlow logical interface ge-1/0/1.2 as a Layer 3 subinterface.

```
[edit interfaces ge-1/0/1]
user@host# set unit 2 vlan-id 300
user@host# set unit 2 family inet address 198.51.100.10/24
```

5. On ge-1/0/2, enable VLAN tagging, and configure the logical interface as a Layer 2 trunk that supports VLANs 1-100.

```
[edit interfaces ge-1/0/2]
user@host# set vlan-tagging
user@host# set unit 0 family bridge interface-mode trunk
user@host# set unit 0 family bridge vlan-id-list 1-100
```

6. On ge-1/0/3, enable VLAN tagging, and configure the logical interface as a Layer 2 trunk that supports VLANs 101-200:

```
[edit interfaces ge-1/0/3]
user@host# set vlan-tagging
user@host# set unit 0 family bridge interface-mode trunk
user@host# set unit 0 family bridge vlan-id-list 101-200
```

---

### Configuring OpenFlow

#### Step-by-Step Procedure

To configure OpenFlow:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch OFswitch2]
user@host# set controller address 172.16.1.1
user@host# set controller protocol tcp port 6633
```

2. Specify the logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch OFswitch2]
user@host# set interfaces ge-1/0/1.0
user@host# set interfaces ge-1/0/2.0
```



## Configuring the Virtual Switch Routing Instances

### Step-by-Step Procedure

To configure the virtual switch routing instances:

1. Configure the virtual switch routing instance for the OpenFlow traffic.  
  

```
[edit]
user@host# set routing-instances rt1 instance-type virtual-switch
user@host# set routing-instances rt1 interface ge-1/0/1.0
user@host# set routing-instances rt1 interface ge-1/0/2.0
user@host# set routing-instances rt1 bridge-domains bd-of vlan-id-list 1-100
```
2. Configure the virtual switch routing instance for the non-OpenFlow traffic.  
  

```
[edit]
user@host# set routing-instances rt2 instance-type virtual-switch
user@host# set routing-instances rt2 interface ge-1/0/1.1
user@host# set routing-instances rt2 interface ge-1/0/3.0
user@host# set routing-instances rt2 bridge-domains bd-nonof vlan-id-list 101-200
```
3. Commit the configuration.  
  

```
[edit]
user@host# commit
```

## Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols openflow**, and **show routing-instances** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-1/0/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-100;
    }
  }
  unit 1 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 101-200;
    }
  }
  unit 2 {
    vlan-id 300;
    family inet {
      address 198.51.100.10/24;
    }
  }
}

ge-1/0/2 {
  vlan-tagging;
```

```
unit 0 {
    family bridge {
        interface-mode trunk;
        vlan-id-list 1-100;
    }
}

ge-1/0/3 {
    vlan-tagging;
    unit 0 {
        family bridge {
            interface-mode trunk;
            vlan-id-list 101-200;
        }
    }
}

user@host# show protocols openflow
switch OFswitch2 {
    interfaces {
        ge-1/0/1.0;
        ge-1/0/2.0;
    }
    controller {
        protocol tcp {
            port 6633;
        }
        address 172.16.1.1;
    }
}

user@host# show routing-instances
rt1 {
    instance-type virtual-switch;
    interface ge-1/0/1.0;
    interface ge-1/0/2.0;
    bridge-domains {
        bd-of {
            vlan-id-list 1-100;
        }
    }
}
rt2 {
    instance-type virtual-switch;
    interface ge-1/0/1.1;
    interface ge-1/0/3.0;
    bridge-domains {
        bd-nonof {
            vlan-id-list 101-200;
        }
    }
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying that the OpenFlow Controller Connection is Up on page 2137](#)
- [Verifying that the OpenFlow Interfaces Are Up on page 2137](#)

### Verifying that the OpenFlow Controller Connection is Up

- Purpose** Verify that the OpenFlow controller connection is up.
- Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**. Because the virtual switch configuration has only a single controller, the virtual switch should automatically initiate a connection to the controller after you commit the configuration.
- ```
user@host> show openflow controller
Openflowd controller information:
Controller socket: 11
Controller IP address: 172.16.1.1
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 1
Controller role: equal
```
- Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying that the OpenFlow Interfaces Are Up

- Purpose** Verify that the OpenFlow interfaces are up.
- Action** Issue the **show openflow interfaces** operational mode command, and verify that the state of each OpenFlow interface is **Up**.
- ```
user@host> show openflow interfaces
Switch name: OFswitch2
Interface Name: ge-1/0/1.0
Interface port number: 41500
Interface Hardware Address: 00:00:5e:00:53:a1
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch2
Interface Name: ge-1/0/2.0
Interface port number: 41501
Interface Hardware Address: 00:00:5e:00:53:00
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```
- Meaning** The output shows that the state of each OpenFlow interface is **Up**, in addition to other information about the interfaces.

- Related Documentation**
- [Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS on page 2127](#)
  - [Configuring OpenFlow Hybrid Interfaces on MX Series Routers on page 2128](#)

- *OpenFlow Feature Guide*

## Configuring OpenFlow Hybrid Interfaces on EX9200 Switches

---

On EX9200 switches that support OpenFlow, you can configure a physical interface as a hybrid interface that concurrently supports OpenFlow logical interfaces and non-OpenFlow logical interfaces. If you configure an OpenFlow hybrid interface on a device running Junos OS, you must enable the reception and transmission of 802.1Q VLAN-tagged frames on all interfaces, including both hybrid and non-hybrid interfaces, and you must configure a virtual switch routing instance for the OpenFlow traffic and a separate virtual switch routing instance for the normal traffic.

The following sections detail configuring an EX9200 switch that supports OpenFlow with a mix of hybrid and normal interfaces:

- [Configuring the Hybrid Physical Interface on page 2138](#)
- [Configuring the Hybrid Interface Logical Units on page 2139](#)
- [Configuring the Non-Hybrid Interfaces on page 2139](#)
- [Configuring OpenFlow on page 2140](#)
- [Configuring the Virtual Switch Routing Instances on page 2140](#)

### Configuring the Hybrid Physical Interface

To configure the hybrid physical interface:

1. Enable VLAN tagging.

Configure **vlan-tagging** to support 802.1Q VLAN single-tag frames for both OpenFlow and non-OpenFlow traffic, or configure **flexible-vlan-tagging** to support both 802.1Q VLAN single-tag and dual-tag frames.

```
[edit interfaces interface-name]  
user@host# set (vlan-tagging | flexible-vlan-tagging)
```

2. Configure flexible Ethernet services encapsulation to enable multiple per-unit Ethernet encapsulations.

```
[edit interfaces interface-name]  
user@host# set encapsulation flexible-ethernet-services
```

## Configuring the Hybrid Interface Logical Units

On a hybrid interface, you configure an OpenFlow or non-OpenFlow logical interface as a Layer 2 trunk interface. Additionally, you can configure a non-OpenFlow logical interface as a Layer 3 subinterface that performs traditional Layer 3 forwarding. To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, or a range or list of VLAN IDs, to the logical interface. Configure Layer 2 interfaces using family **ethernet-switching** on EX9200 switches.

To configure the hybrid interface logical units:

1. Configure the OpenFlow logical interfaces and any non-OpenFlow Layer 2 logical interfaces, and specify the interface mode and VLAN membership.

```
[edit interfaces interface-name]
user@host# set unit unit family ethernet-switching interface-mode trunk
user@host# set unit unit family ethernet-switching vlan members vlan-ids
```

2. Configure any non-OpenFlow Layer 3 logical interfaces, and specify the VLAN membership.

```
[edit interfaces interface-name]
user@host# set unit unit (vlan-id | vlan-id-list | vlan-id-range) vlan-ids
user@host# set unit unit family inet address address
```

## Configuring the Non-Hybrid Interfaces

Non-hybrid interfaces support either OpenFlow traffic or non-OpenFlow traffic, but not both simultaneously.

To configure the non-hybrid interfaces:

1. Configure interfaces that support only OpenFlow traffic as Layer 2 trunk interfaces, and specify the interface mode and VLAN membership.

```
[edit interfaces interface-name]
user@host# set vlan-tagging
user@host# set unit unit family ethernet-switching interface-mode trunk
user@host# set unit unit family ethernet-switching vlan members (vlan-id | vlan-id-list)
```

2. Configure interfaces that support only normal traffic, and specify the interface mode for the Layer 2 interfaces and the VLAN membership.

For example:

```
[edit interfaces interface-name]
user@host# set vlan-tagging
user@host# set unit unit family ethernet-switching interface-mode trunk
user@host# set unit unit family ethernet-switching vlan members (vlan-id | vlan-id-list)
```

## Configuring OpenFlow

To configure the OpenFlow virtual switch instance:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch switch-name]  
user@host# set controller address address  
user@host# set controller protocol tcp port port
```

2. Specify all logical interfaces participating in OpenFlow under this virtual switch instance.

```
[edit protocols openflow switch switch-name]  
user@host# set interfaces interface-name
```

## Configuring the Virtual Switch Routing Instances

Configure separate virtual switch routing instances for the OpenFlow traffic and the non-OpenFlow traffic. The configured interface names must include a logical unit number.

To configure the virtual switch routing instances:

1. Configure the virtual switch routing instance for the OpenFlow traffic, and specify the OpenFlow logical interfaces and VLANs.

```
[edit routing-instances of-routing-instance-name]  
user@host# set instance-type virtual-switch  
user@host# set interface of-interface-name1  
user@host# set interface of-interface-name2  
user@host# set vlans name vlan-id-list of-vlan-id-list
```

2. Configure the virtual switch routing instance for the non-OpenFlow traffic, and specify the non-OpenFlow logical interfaces and VLANs.

```
[edit routing-instances non-of-routing-instance-name]  
user@host# set instance-type virtual-switch  
user@host# set interface non-of-interface-name1  
user@host# set interface non-of-interface-name2  
user@host# set vlans name vlan-id-list non-of-vlan-id-list
```

### Related Documentation

- [Example: Configuring OpenFlow Hybrid Interfaces on EX9200 Switches on page 2140](#)
- [Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS on page 2127](#)
- [OpenFlow Feature Guide](#)

---

## Example: Configuring OpenFlow Hybrid Interfaces on EX9200 Switches

On EX9200 switches that have the OpenFlow software package installed, you can configure physical interfaces that support multiple logical interfaces as OpenFlow hybrid interfaces. A hybrid interface concurrently supports OpenFlow logical interfaces and non-OpenFlow logical interfaces. A hybrid interface enables OpenFlow and non-OpenFlow traffic to traverse the same physical interface while keeping the traffic in separate VLANs.

Hybrid interfaces enable you to use physical interfaces more efficiently, especially in a situation where having an adequate number of physical interfaces available is important.

This example shows how to configure an OpenFlow hybrid interface on an EX9200 switch.

- [Requirements on page 2141](#)
- [Overview on page 2141](#)
- [Configuration on page 2143](#)
- [Verification on page 2147](#)

## Requirements

This example uses the following hardware and software components:

- An EX9200 switch running Junos OS Release 13.3 or a later release.
- An OpenFlow software package is installed on the switch, and the software package release matches the Junos OS release running on the switch.
- The switch has a TCP connection to an OpenFlow controller, which needs to access the data plane of the switch.
- The switch is connected to the management network through the fxp0 interface and is reachable from the controller IP address.

## Overview

In this example, you configure an EX9200 switch with:

- One hybrid interface, xe-2/1/0
- One non-hybrid interface, xe-2/1/1, which handles OpenFlow traffic only
- One non-hybrid interface, xe-2/1/2, which handles non-OpenFlow traffic only

On the hybrid interface, logical interface xe-2/1/0.0 participates in OpenFlow, and logical interfaces xe-2/1/0.1 and xe-2/1/0.2 do not participate in OpenFlow.

When using hybrid interfaces, you use VLAN tagging to distinguish OpenFlow traffic from non-OpenFlow traffic. Thus, you must enable VLAN tagging on all interfaces, and traffic entering the interfaces must be VLAN-tagged. If you do not configure a native VLAN, untagged traffic entering a hybrid interface is dropped. In this example, you configure the hybrid interface by using **flexible-vlan-tagging**, which enables VLAN tagging and supports both 802.1Q VLAN single-tag and dual-tag frames for all traffic on the interface. You also configure the OpenFlow interface xe-2/1/1 and the non-OpenFlow interface xe-2/1/2 by using **vlan-tagging**, which enables VLAN tagging and supports only 802.1Q VLAN single-tag frames for all traffic on the interface.

You configure the hybrid interface encapsulation as flexible Ethernet services. Note that for interfaces with this type of encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511 are no longer reserved for normal Ethernet VLANs. In this example, VLANs

100 through 200 are used for OpenFlow traffic, and VLANs 700 and 800 are used for non-OpenFlow traffic.

All logical interfaces except xe-2/1/0.2 are configured as Layer 2 trunk interfaces by using family **ethernet-switching** and interface mode **trunk**. Logical interfaces xe-2/1/0.0 and xe-2/1/1.0 participate in OpenFlow and receive and forward traffic with OpenFlow VLAN IDs 100 through 200. Logical interfaces xe-2/1/0.1 and xe-2/1/2.0 do not participate in OpenFlow and receive and forward traffic with non-OpenFlow VLAN ID 700.

Logical interface xe-2/1/0.2 is a subinterface with the IP address 198.51.100.10/24 and performs Layer 3 routing. This interface does not participate in OpenFlow and routes traffic with VLAN ID 800.

[Table 187 on page 2142](#) summarizes the logical interfaces, traffic types, and associated VLAN IDs.

**Table 187: Summary of Logical Interfaces**

| Logical Interface | Traffic Type | VLANs           |
|-------------------|--------------|-----------------|
| xe-2/1/0.0        | OpenFlow     | 100 through 200 |
| xe-2/1/0.1        | Non-OpenFlow | 700             |
| xe-2/1/0.2        | Non-OpenFlow | 800             |
| xe-2/1/1.0        | OpenFlow     | 200             |
| xe-2/1/2.0        | Non-OpenFlow | 700             |

You configure the OpenFlow virtual switch and OpenFlow protocol statements at the **[edit protocols openflow]** hierarchy level. The virtual switch 100 connects to the OpenFlow controller over a TCP connection at the IP address 198.51.100.174. The virtual switch configuration must include all of the logical interfaces participating in OpenFlow, which includes xe-2/1/0.0 and xe-2/1/1.0.

An EX9200 switch requires a separate routing instance for a virtual switch. This routing instance isolates the OpenFlow traffic from the non-OpenFlow traffic. When using hybrid interfaces, you configure a virtual switch routing instance for the OpenFlow traffic and another virtual switch routing instance for non-OpenFlow traffic. In this example, you configure routing instance **OF** for the OpenFlow traffic and routing instance **NON-OF** for the non-OpenFlow traffic.

Routing instance **OF** includes the interfaces participating in OpenFlow—xe-2/1/0.0 and xe-2/1/1.0. Within this routing instance, you configure a VLAN to include OpenFlow VLANs 100 through 200. Routing instance **NON-OF** includes the Layer 2 interfaces that do not participate in OpenFlow—xe-2/1/0.1 and xe-2/1/2.0. Within this routing instance, you configure a VLAN to include the non-OpenFlow VLAN 700.





**NOTE:** To direct OpenFlow traffic, the OpenFlow controller must install flow entries that select the appropriate traffic and forward it to the correct OpenFlow interface.

## Configuration

- [Configuring the Interfaces on page 2143](#)
- [Configuring OpenFlow on page 2144](#)
- [Configuring the Virtual Switch Routing Instances on page 2145](#)
- [Results on page 2145](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces xe-2/1/0 flexible-vlan-tagging
set interfaces xe-2/1/0 encapsulation flexible-ethernet-services
set interfaces xe-2/1/0 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-2/1/0 unit 0 family ethernet-switching vlan members 100-200
set interfaces xe-2/1/0 unit 1 family ethernet-switching interface-mode trunk
set interfaces xe-2/1/0 unit 1 family ethernet-switching vlan members 700
set interfaces xe-2/1/0 unit 2 vlan-id 800
set interfaces xe-2/1/0 unit 2 family inet address 198.51.100.10/24
set interfaces xe-2/1/1 vlan-tagging
set interfaces xe-2/1/1 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-2/1/1 unit 0 family ethernet-switching vlan members 200
set interfaces xe-2/1/2 vlan-tagging
set interfaces xe-2/1/2 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-2/1/2 unit 0 family ethernet-switching vlan members 700
set protocols openflow switch 100 controller address 198.51.100.174
set protocols openflow switch 100 controller protocol tcp port 6633
set protocols openflow switch 100 interfaces xe-2/1/0.0
set protocols openflow switch 100 interfaces xe-2/1/1.0
set routing-instances OF instance-type virtual-switch
set routing-instances OF interface xe-2/1/0.0
set routing-instances OF interface xe-2/1/1.0
set routing-instances OF vlans OF-vlan vlan-id-list 100-200
set routing-instances NON-OF instance-type virtual-switch
set routing-instances NON-OF interface xe-2/1/0.1
set routing-instances NON-OF interface xe-2/1/2.0
set routing-instances NON-OF vlans OF-vlan vlan-id-list 700
```

### Configuring the Interfaces

#### Step-by-Step Procedure

To configure the interfaces:

1. On the hybrid physical interface, enable VLAN tagging and configure the encapsulation.

```
[edit interfaces xe-2/1/0]
user@switch# set flexible-vlan-tagging
```

```
user@switch# set encapsulation flexible-ethernet-services
```

2. Configure the OpenFlow logical interface xe-2/1/0.0 as a Layer 2 trunk that supports VLANs 100 through 200.

```
[edit interfaces xe-2/1/0]
user@switch# set unit 0 family ethernet-switching interface-mode trunk
user@switch# set unit 0 family ethernet-switching vlan members 100-200
```

3. Configure the non-OpenFlow logical interface xe-2/1/0.1 as a Layer 2 trunk that supports VLAN 700.

```
[edit interfaces xe-2/1/0]
user@switch# set unit 1 family ethernet-switching interface-mode trunk
user@switch# set unit 1 family ethernet-switching vlan members 700
```

4. Configure the non-OpenFlow logical interface xe-2/1/0.2 as a Layer 3 subinterface.

```
[edit interfaces xe-2/1/0]
user@switch# set unit 2 vlan-id 800
user@switch# set unit 2 family inet address 198.51.100.10/24
```

5. On xe-2/1/1, enable VLAN tagging, and configure the logical interface as a Layer 2 trunk that supports VLAN 200.

```
[edit interfaces xe-2/1/1]
user@switch# set vlan-tagging
user@switch# set unit 0 family ethernet-switching interface-mode trunk
user@switch# set unit 0 family ethernet-switching vlan members 200
```

6. On xe-2/1/2, enable VLAN tagging, and configure the logical interface as a Layer 2 trunk that supports VLAN 700.

```
[edit interfaces xe-2/1/2]
user@switch# set vlan-tagging
user@switch# set unit 0 family ethernet-switching interface-mode trunk
user@switch# set unit 0 family ethernet-switching vlan members 700
```

---

### Configuring OpenFlow

#### Step-by-Step Procedure

To configure OpenFlow:

1. Configure the OpenFlow controller IP address and the connection protocol.

```
[edit protocols openflow switch 100]
user@switch# set controller address 198.51.100.174
user@switch# set controller protocol tcp port 6633
```

2. Specify the logical interfaces participating in OpenFlow under virtual switch 100.

```
[edit protocols openflow switch 100]
user@switch# set interfaces xe-2/1/0.0
user@switch# set interfaces xe-2/1/1.0
```

## Configuring the Virtual Switch Routing Instances

### Step-by-Step Procedure

To configure the routing instances:

1. Configure the routing instance for the OpenFlow traffic.  
  

```
[edit]
user@switch# set routing-instances OF instance-type virtual-switch
user@switch# set routing-instances OF interface xe-2/1/0.0
user@switch# set routing-instances OF interface xe-2/1/1.0
user@switch# set routing-instances OF vlans OF-vlan vlan-id-list 100-200
```
2. Configure the routing instance for the non-OpenFlow traffic on Layer 2 interfaces.  
  

```
[edit]
user@switch# set routing-instances NON-OF instance-type virtual-switch
user@switch# set routing-instances NON-OF interface xe-2/1/0.1
user@switch# set routing-instances NON-OF interface xe-2/1/2.0
user@switch# set routing-instances NON-OF vlans NOF-vlan vlan-id-list 700
```
3. Commit the configuration.  
  

```
[edit]
user@switch# commit
```

## Results

From operational mode, confirm your configuration by entering the **show configuration interfaces**, **show configuration protocols openflow**, and **show configuration routing-instances** commands. If the output does not display the specified configuration, repeat the configuration instructions in this example to correct the configuration.

```
user@switch> show configuration interfaces
xe-2/1/0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members 100-200;
      }
    }
  }
  unit 1 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members 700;
      }
    }
  }
  unit 2 {
    vlan-id 800;
    family inet {
```

```
        address 198.51.100.10/24;
    }
}
xe-2/1/1 {
    vlan-tagging;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members 200;
            }
        }
    }
}
xe-2/1/2 {
    vlan-tagging;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members 700;
            }
        }
    }
}

user@switch> show configuration protocols openflow
switch 100 {
    interfaces {
        xe-2/1/0.0;
        xe-2/1/1.0;
    }
    controller {
        protocol tcp {
            port 6633;
        }
        address 198.51.100.174;
    }
}

user@switch> show configuration routing-instances
OF {
    instance-type virtual-switch;
    interface xe-2/1/0.0;
    interface xe-2/1/1.0;
    vlans {
        OF-vlan {
            vlan-id-list 100-200;
        }
    }
}
NON-OF {
    instance-type virtual-switch;
    interface xe-2/1/0.1;
    interface xe-2/1/2.0;
    vlans {
        NOF-vlan {
            vlan-id 700;
        }
    }
}
```

```

    }
  }
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the OpenFlow Controller Connection on page 2147](#)
- [Verifying the OpenFlow Interfaces on page 2147](#)

### Verifying the OpenFlow Controller Connection

**Purpose** Verify that the OpenFlow controller connection is up.

**Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**. Because the virtual switch configuration has only a single controller, the virtual switch automatically initiates a connection to the controller after you commit the configuration.

```

user@switch> show openflow controller
Openflowd controller information:
Controller socket: 11
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 5
Controller role: equal

```

**Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying the OpenFlow Interfaces

**Purpose** Verify that the OpenFlow interfaces are up.

**Action** Issue the **show openflow interfaces** operational mode command, and verify that the state of each OpenFlow interface is **Up**.

```

user@switch> show openflow interfaces
Switch name: 100
Interface Name: xe-2/1/0.0
Interface port number: 41500
Interface Hardware Address: 00:00:5E:00:53:cf
Interface speed: 10Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up

Switch name: 100
Interface Name: xe-2/1/1.0
Interface port number: 41501
Interface Hardware Address: 00:00:5E:00:53:d0
Interface speed: 10Gb Full-duplex

```

Interface Auto-Negotiation: Disabled  
Interface media type: Fiber  
**Interface state: Up**

**Meaning** The output shows that the state of each OpenFlow interface is **Up**, in addition to other information about the interfaces.

**Related Documentation**

- [Understanding OpenFlow Hybrid Interfaces on Devices Running Junos OS on page 2127](#)
- [Configuring OpenFlow Hybrid Interfaces on EX9200 Switches on page 2138](#)
- *OpenFlow Feature Guide*

# Configuring OpenFlow Traffic Steering Across MPLS Networks

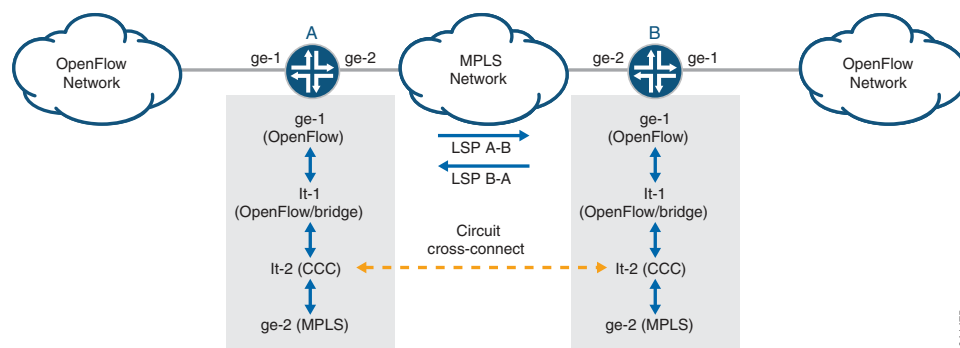
- Understanding OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects on page 2149
- Example: OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects on page 2150

## Understanding OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects

On MX Series devices that support OpenFlow, you can direct traffic from OpenFlow networks over MPLS networks by using logical tunnel interfaces and MPLS LSP tunnel cross-connects. Using logical tunnel interfaces, you can stitch an OpenFlow interface to an MPLS label-switched path (LSP), which enables you to direct traffic from the OpenFlow network onto the MPLS network. MPLS LSP tunnel cross-connects between interfaces and LSPs permit you to connect the OpenFlow network to a remote network by creating MPLS tunnels that use LSPs as the conduit.

The topology in [Figure 35 on page 2149](#) illustrates an MPLS LSP tunnel cross-connect that connects two remote OpenFlow networks through an MPLS network. Circuit cross-connect (CCC) enables you to establish an LSP tunnel between the two domains, through which you can tunnel the traffic from one OpenFlow network across the MPLS network to the second OpenFlow network.

**Figure 35: Connecting OpenFlow Networks Using MPLS LSP Tunnel Cross-Connects**



Router A and Router B are OpenFlow-enabled routers that have MPLS LSPs configured to route traffic across the MPLS network. LSP A-B routes traffic from Router A to Router B, and LSP B-A routes traffic from Router B to Router A.

Each router has an OpenFlow interface, ge-1, and an MPLS interface, ge-2. You can stitch the OpenFlow interface to the MPLS LSP by using two logical tunnel interfaces. You configure the first logical tunnel interface, lt-1, as a Layer 2 interface that participates in OpenFlow. The second logical tunnel interface, lt-2, uses CCC encapsulation. You configure lt-1 and lt-2 interfaces as peers, so that traffic entering one logical interface is automatically directed to the second logical interface.

On each router, MPLS LSP tunnel cross-connects are configured at the **[edit protocols connections remote-interface-switch]** hierarchy level. The cross-connects make an association between the CCC interface, lt-2, and the two LSPs, one for transmitting MPLS packets from the local device to the remote device and the other for receiving MPLS packets on the local device from the remote device.

For traffic flowing from Router A to Router B, the OpenFlow controller must install flow entries on Router A that direct the desired OpenFlow traffic from ge-1 as the OpenFlow ingress port to lt-1 as the output port. On Router B, the OpenFlow controller must install flow entries that direct the OpenFlow traffic from lt-1 as the OpenFlow ingress port to ge-1 as the output port. Similarly for traffic flowing from Router B to Router A, the OpenFlow controller must install flow entries on Router B that direct the desired OpenFlow traffic from ge-1 as the OpenFlow ingress port to lt-1 as the output port. On Router A, the OpenFlow controller must install flow entries that direct the OpenFlow traffic from lt-1 as the OpenFlow ingress port to ge-1 as the output port.

**Related  
Documentation**

- [Example: OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects on page 2150](#)
- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)

---

## Example: OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects

---

On MX series routers that support OpenFlow, you can direct traffic from OpenFlow networks over MPLS networks by using logical tunnel interfaces and MPLS LSP tunnel cross-connects. This example shows how to configure MX Series routers to send traffic between two remote OpenFlow networks over an MPLS-based network using MPLS LSP tunnel cross-connects.

- [Requirements on page 2151](#)
- [Overview on page 2151](#)
- [Configuration on page 2152](#)
- [Verification on page 2163](#)
- [Troubleshooting on page 2166](#)



## Requirements

This example uses the following hardware and software components for the OpenFlow-enabled routers:

- MX240 routers running Junos OS Release 13.3 or a later release.
- OpenFlow software package with a software package release that matches the Junos OS release of the device on which it is installed
- TCP connection between the router and an OpenFlow controller
- Connection between the fxp0 management interface of the router and the management network, which is reachable from the controller IP address

## Overview

In this example, you configure MPLS LSP tunnel cross-connects to connect two remote OpenFlow networks that are separated by an MPLS network. [Figure 36 on page 2152](#) shows the topology used in this example.

This example has three routers: a provider router (P) and two provider edge routers (PE1 and PE2). Router P resides within an MPLS network. Routers PE1 and PE2 are OpenFlow-enabled routers, each with the ge-1/0/0.0 interface configured to accept and forward OpenFlow traffic and two MPLS interfaces that connect to Router P. The network uses OSPF as the IGP, and it has two LSPs: LSP 1-3 routes traffic from PE1 to PE2, and LSP 3-1 routes traffic from PE2 to PE1.

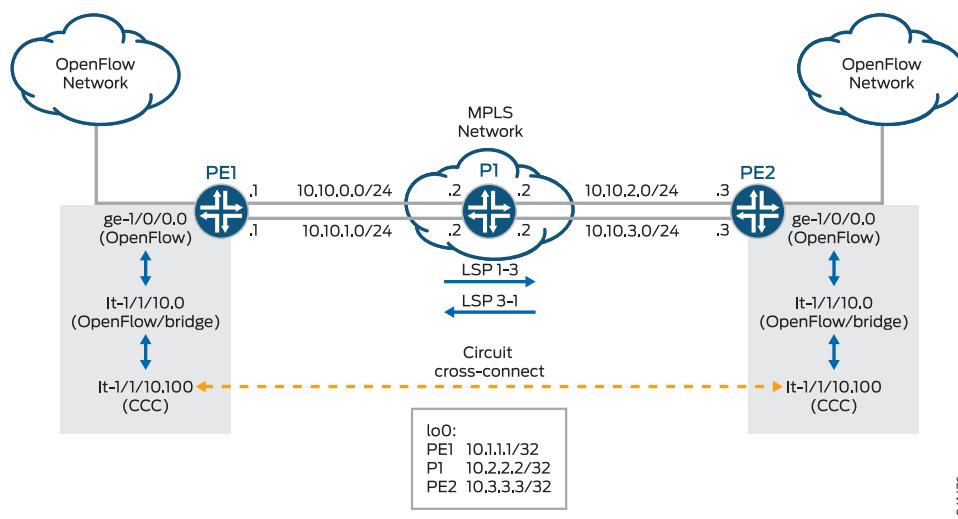
You stitch the OpenFlow interface to the MPLS LSP using two logical tunnel interfaces, lt-1/1/10.0 and lt-1/1/10.100. You configure the first logical tunnel interface, lt-1/1/10.0, as a Layer 2 interface with encapsulation **ethernet-bridge** and family **bridge**. This interface participates in OpenFlow. The second logical tunnel interface, lt-1/1/10.100, uses circuit cross-connect (CCC) encapsulation. You configure lt-1 and lt-2 interfaces as peers, so that traffic entering one logical interface is automatically directed to the second logical interface.

On the PE1 and PE2 routers, you configure an MPLS LSP tunnel cross-connect at the **[edit protocols connections remote-interface-switch]** hierarchy level using the logical tunnel interface with CCC encapsulation. This configuration makes an association between the CCC interface and two LSPs, one for transmitting MPLS packets from the local device to the remote device and the other for receiving MPLS packets on the local device from the remote device.

For traffic flowing from PE1 to PE2, the OpenFlow controller must install flow entries on PE1 that direct the desired OpenFlow traffic from ge-1/0/0.0 as the OpenFlow ingress port to lt-1/1/10.0 as the output port. On PE2, the OpenFlow controller must install flow entries that direct the OpenFlow traffic from lt-1/1/10.0 as the OpenFlow ingress port to ge-1/0/0.0 as the output port. Similarly, for traffic flowing from PE2 to PE1, the OpenFlow controller must install flow entries on PE2 that direct the desired OpenFlow traffic from ge-1/0/0.0 as the OpenFlow ingress port to lt-1/1/10.0 as the output port. On PE1, the OpenFlow controller must install flow entries that direct the OpenFlow traffic from lt-1/1/10.0 as the OpenFlow ingress port to ge-1/1/0.0 as the output port.

## Topology

Figure 36: Connecting OpenFlow Networks Using MPLS Tunnel Cross-Connects



## Configuration

- [Configuring the Ingress Provider Edge Router \(PE1\) on page 2154](#)
- [Configuring the Provider Router \(P\) on page 2158](#)
- [Configuring the Egress Provider Edge Router \(PE2\) on page 2160](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

#### Device PE1

```
set chassis fpc 1 pic 1 tunnel-services bandwidth 1g
set interfaces ge-1/0/0 encapsulation ethernet-bridge
set interfaces ge-1/0/0 unit 0 family bridge
set interfaces ge-1/1/0 unit 0 family inet address 10.10.0.1/24
set interfaces ge-1/1/0 unit 0 family mpls
set interfaces ge-1/1/1 unit 0 family inet address 10.10.1.1/24
set interfaces ge-1/1/1 unit 0 family mpls
set interfaces lt-1/1/10 unit 0 encapsulation ethernet-bridge
set interfaces lt-1/1/10 unit 0 peer-unit 100
set interfaces lt-1/1/10 unit 0 family bridge
set interfaces lt-1/1/10 unit 100 encapsulation ethernet-ccc
set interfaces lt-1/1/10 unit 100 peer-unit 0
set interfaces lt-1/1/10 unit 100 family ccc
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set protocols rsvp interface ge-1/1/0.0
set protocols rsvp interface ge-1/1/1.0
set protocols mpls label-switched-path 1-3 from 10.1.1.1
set protocols mpls label-switched-path 1-3 to 10.3.3.3
set protocols mpls interface ge-1/1/0.0
set protocols mpls interface ge-1/1/1.0
set protocols ospf traffic-engineering
```

```

set protocols ospf area 0.0.0.0 interface ge-1/1/0.0
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0
set protocols connections remote-interface-switch 1-3-ccc interface lt-1/1/10.100
set protocols connections remote-interface-switch 1-3-ccc transmit-lsp 1-3
set protocols connections remote-interface-switch 1-3-ccc receive-lsp 3-1
set protocols openflow switch s1 interfaces ge-1/0/0.0 port-id 1
set protocols openflow switch s1 interfaces lt-1/1/10.0 port-id 2
set protocols openflow switch s1 controller protocol tcp port 6633
set protocols openflow switch s1 controller address 10.94.175.213
set routing-instances r1 instance-type virtual-switch
set routing-instances r1 bridge-domains bd1 interface ge-1/0/0.0
set routing-instances r1 bridge-domains bd1 interface lt-1/1/10.0
set routing-options router-id 10.1.1.1

```

**Device P**

```

set interfaces ge-1/1/0 unit 0 family inet address 10.10.0.2/24
set interfaces ge-1/1/0 unit 0 family mpls
set interfaces ge-1/1/1 unit 0 family inet address 10.10.1.2/24
set interfaces ge-1/1/1 unit 0 family mpls
set interfaces ge-1/1/2 unit 0 family inet address 10.10.2.2/24
set interfaces ge-1/1/2 unit 0 family mpls
set interfaces ge-1/1/3 unit 0 family inet address 10.10.3.2/24
set interfaces ge-1/1/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.2.2.2/32
set protocols rsvp interface ge-1/1/0.0
set protocols rsvp interface ge-1/1/1.0
set protocols rsvp interface ge-1/1/2.0
set protocols rsvp interface ge-1/1/3.0
set protocols mpls interface ge-1/1/0.0
set protocols mpls interface ge-1/1/1.0
set protocols mpls interface ge-1/1/2.0
set protocols mpls interface ge-1/1/3.0
set protocols mpls interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-1/1/0.0
set protocols ospf area 0.0.0.0 interface ge-1/1/1.0
set protocols ospf area 0.0.0.0 interface ge-1/1/2.0
set protocols ospf area 0.0.0.0 interface ge-1/1/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set routing-options router-id 10.2.2.2

```

**Device PE2**

```

set chassis fpc 1 pic 1 tunnel-services bandwidth 1g
set interfaces ge-1/0/0 encapsulation ethernet-bridge
set interfaces ge-1/0/0 unit 0 family bridge
set interfaces ge-1/1/2 unit 0 family inet address 10.10.2.3/24
set interfaces ge-1/1/2 unit 0 family mpls
set interfaces ge-1/1/3 unit 0 family inet address 10.10.3.3/24
set interfaces ge-1/1/3 unit 0 family mpls
set interfaces lt-1/1/10 unit 0 encapsulation ethernet-bridge
set interfaces lt-1/1/10 unit 0 peer-unit 100
set interfaces lt-1/1/10 unit 0 family bridge
set interfaces lt-1/1/10 unit 100 encapsulation ethernet-ccc
set interfaces lt-1/1/10 unit 100 peer-unit 0
set interfaces lt-1/1/10 unit 100 family ccc
set interfaces lo0 unit 0 family inet address 10.3.3.3/32

```

```
set protocols rsvp interface ge-1/1/2.0
set protocols rsvp interface ge-1/1/3.0
set protocols mpls label-switched-path 3-1 from 10.3.3.3
set protocols mpls label-switched-path 3-1 to 10.1.1.1
set protocols mpls interface ge-1/1/2.0
set protocols mpls interface ge-1/1/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-1/1/2.0
set protocols ospf area 0.0.0.0 interface ge-1/1/3.0
set protocols connections remote-interface-switch 3-1-ccc interface lt-1/1/10.100
set protocols connections remote-interface-switch 3-1-ccc transmit-lsp 3-1
set protocols connections remote-interface-switch 3-1-ccc receive-lsp 1-3
set protocols openflow switch s1 interfaces ge-1/0/0.0 port-id 1
set protocols openflow switch s1 interfaces lt-1/1/10.0 port-id 2
set protocols openflow switch s1 controller protocol tcp port 6633
set protocols openflow switch s1 controller address 10.94.175.213
set routing-instances r1 instance-type virtual-switch
set routing-instances r1 bridge-domains bd1 interface ge-1/0/0.0
set routing-instances r1 bridge-domains bd1 interface lt-1/1/10.0
set routing-options router-id 10.3.3.3
```

---

### Configuring the Ingress Provider Edge Router (PE1)

---

#### Step-by-Step Procedure

To configure Router PE1:

1. Create tunnel interfaces by configuring the DPC and its corresponding PIC to use tunneling services.

[edit]

```
user@PE1# set chassis fpc 1 pic 1 tunnel-services bandwidth 1g
```

2. Configure the OpenFlow interface as a Layer 2 interface.

[edit interfaces]

```
user@PE1# set ge-1/0/0 encapsulation ethernet-bridge
```

```
user@PE1# set ge-1/0/0 unit 0 family bridge
```

3. Configure the OpenFlow virtual switch routing instance.

[edit]

```
user@PE1# set routing-instances r1 instance-type virtual-switch
```

```
user@PE1# set routing-instances r1 bridge-domains bd1 interface ge-1/0/0.0
```

```
user@PE1# set routing-instances r1 bridge-domains bd1 interface lt-1/1/10.0
```

4. Configure the OpenFlow controller.

[edit protocols openflow]

```
user@PE1# set switch s1 controller address 10.94.175.213
```

```
user@PE1# set switch s1 controller protocol tcp port 6633
```

5. Configure the interfaces participating in OpenFlow.

[edit protocols openflow]

```
user@PE1# set switch s1 interfaces ge-1/0/0.0 port-id 1
```

```
user@PE1# set switch s1 interfaces lt-1/1/10.0 port-id 2
```

6. Configure the loopback interface and router ID.

[edit]

```

user@PE1# set interfaces lo0 unit 0 family inet address 10.1.1.1/32
user@PE1# set routing-options router-id 10.1.1.1

```

7. Configure the MPLS interfaces.

```

[edit interfaces]
user@PE1# set ge-1/1/0 unit 0 family inet address 10.10.0.1/24
user@PE1# set ge-1/1/0 unit 0 family mpls
user@PE1# set ge-1/1/1 unit 0 family inet address 10.10.1.1/24
user@PE1# set ge-1/1/1 unit 0 family mpls

```

8. Configure the logical tunnel interface.

```

[edit interfaces]
user@PE1# set lt-1/1/10 unit 0 family bridge
user@PE1# set lt-1/1/10 unit 0 encapsulation ethernet-bridge
user@PE1# set lt-1/1/10 unit 0 peer-unit 100
user@PE1# set lt-1/1/10 unit 100 family ccc
user@PE1# set lt-1/1/10 unit 100 encapsulation ethernet-ccc
user@PE1# set lt-1/1/10 unit 100 peer-unit 0

```

9. Enable RSVP, MPLS, and OSPF on the interfaces connected to Router P.

```

[edit protocols]
user@PE1# set rsvp interface ge-1/1/0.0
user@PE1# set rsvp interface ge-1/1/1.0
user@PE1# set mpls interface ge-1/1/0.0
user@PE1# set mpls interface ge-1/1/1.0
user@PE1# set ospf area 0.0.0.0 interface ge-1/1/0.0
user@PE1# set ospf area 0.0.0.0 interface ge-1/1/1.0

```

10. Enable traffic engineering for OSPF.

```

[edit protocols]
user@PE1# set ospf traffic-engineering

```

11. Configure the MPLS LSP from PE1 to PE2.

```

[edit protocols]
user@PE1# set mpls label-switched-path 1-3 from 10.1.1.1
user@PE1# set mpls label-switched-path 1-3 to 10.3.3.3

```

12. Configure the MPLS LSP tunnel cross-connects.

```

[edit protocols]
user@PE1# set connections remote-interface-switch 1-3-ccc interface lt-1/1/10.100
user@PE1# set connections remote-interface-switch 1-3-ccc transmit-lsp 1-3
user@PE1# set connections remote-interface-switch 1-3-ccc receive-lsp 3-1

```

13. Commit the configuration.

```

[edit]
user@PE1# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it. For brevity, this **show** command output includes

only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
chassis {
  fpc 1 {
    pic 1 {
      tunnel-services {
        bandwidth 1g;
      }
    }
  }
}

interfaces {
  ge-1/0/0 {
    encapsulation ethernet-bridge;
    unit 0 {
      family bridge;
    }
  }
  ge-1/1/0 {
    unit 0 {
      family inet {
        address 10.10.0.1/24;
      }
      family mpls;
    }
  }
  ge-1/1/1 {
    unit 0 {
      family inet {
        address 10.10.1.1/24;
      }
      family mpls;
    }
  }
  lt-1/1/10 {
    unit 0 {
      encapsulation ethernet-bridge;
      peer-unit 100;
      family bridge;
    }
    unit 100 {
      encapsulation ethernet-ccc;
      peer-unit 0;
      family ccc;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.1.1.1/32;
      }
    }
  }
}

protocols {
  rsvp {
    interface ge-1/1/0.0;
    interface ge-1/1/1.0;
```

```

}
mpls {
    label-switched-path 1-3 {
        from 10.1.1.1;
        to 10.3.3.3;
    }
    interface ge-1/1/0.0;
    interface ge-1/1/1.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-1/1/0.0;
        interface ge-1/1/1.0;
    }
}
connections {
    remote-interface-switch 1-3-ccc {
        interface lt-1/1/10.100;
        transmit-lsp 1-3;
        receive-lsp 3-1;
    }
}
openflow {
    switch s1 {
        interfaces {
            ge-1/0/0.0 port-id 1;
            lt-1/1/10.0 port-id 2;
        }
        controller {
            protocol {
                tcp {
                    port 6633;
                }
            }
            address 10.94.175.213;
        }
    }
}
}

routing-instances {
    r1 {
        instance-type virtual-switch;
        bridge-domains {
            bd1 {
                interface ge-1/0/0.0;
                interface lt-1/1/10.0;
            }
        }
    }
}

routing-options {
    router-id 10.1.1.1;
}
...

```

### Configuring the Provider Router (P)

---

#### Step-by-Step Procedure

To configure Router P:

1. Configure the loopback interface and router ID.

```
[edit]
user@P# set interfaces lo0 unit 0 family inet address 10.2.2.2/32
user@P# set routing-options router-id 10.2.2.2
```

2. Configure the MPLS interfaces.

```
[edit interfaces]
user@P# set ge-1/1/0 unit 0 family inet address 10.10.0.2/24
user@P# set ge-1/1/0 unit 0 family mpls
user@P# set ge-1/1/1 unit 0 family inet address 10.10.1.2/24
user@P# set ge-1/1/1 unit 0 family mpls
user@P# set ge-1/1/2 unit 0 family inet address 10.10.2.2/24
user@P# set ge-1/1/2 unit 0 family mpls
user@P# set ge-1/1/3 unit 0 family inet address 10.10.3.2/24
user@P# set ge-1/1/3 unit 0 family mpls
```

3. Enable RSVP, MPLS, and OSPF on the interfaces connected to PE1 and PE2.

```
[edit protocols]
user@P# set rsvp interface ge-1/1/0.0
user@P# set rsvp interface ge-1/1/1.0
user@P# set rsvp interface ge-1/1/2.0
user@P# set rsvp interface ge-1/1/3.0
user@P# set mpls interface lo0.0
user@P# set mpls interface ge-1/1/0.0
user@P# set mpls interface ge-1/1/1.0
user@P# set mpls interface ge-1/1/2.0
user@P# set mpls interface ge-1/1/3.0
user@P# set ospf area 0.0.0.0 interface fxp0.0 disable
user@P# set ospf area 0.0.0.0 interface ge-1/1/0.0
user@P# set ospf area 0.0.0.0 interface ge-1/1/1.0
user@P# set ospf area 0.0.0.0 interface ge-1/1/2.0
user@P# set ospf area 0.0.0.0 interface ge-1/1/3.0
user@P# set ospf area 0.0.0.0 interface lo0.0
```

4. Enable traffic engineering for OSPF.

```
[edit protocols]
user@P# set ospf traffic-engineering
```

5. Commit the configuration.

```
[edit]
user@P# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it. For brevity, this **show** command output includes



only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```

interfaces {
  ge-1/1/0 {
    unit 0 {
      family inet {
        address 10.10.0.2/24;
      }
      family mpls;
    }
  }
  ge-1/1/1 {
    unit 0 {
      family inet {
        address 10.10.1.2/24;
      }
      family mpls;
    }
  }
  ge-1/1/2 {
    unit 0 {
      family inet {
        address 10.10.2.2/24;
      }
      family mpls;
    }
  }
  ge-1/1/3 {
    unit 0 {
      family inet {
        address 10.10.3.2/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.2.2.2/32;
      }
    }
  }
}

protocols {
  rsvp {
    interface ge-1/1/0.0;
    interface ge-1/1/1.0;
    interface ge-1/1/2.0;
    interface ge-1/1/3.0;
  }
  mpls {
    interface ge-1/1/0.0;
    interface ge-1/1/1.0;
    interface ge-1/1/2.0;
    interface ge-1/1/3.0;
    interface lo0.0;
  }
  ospf {
    traffic-engineering;
  }
}

```

```
        area 0.0.0.0 {
            interface fxp0.0 {
                disable;
            }
            interface ge-1/1/0.0;
            interface ge-1/1/1.0;
            interface ge-1/1/2.0;
            interface ge-1/1/3.0;
            interface lo0.0;
        }
    }

routing-options {
    router-id 10.2.2.2;
}

...

```

---

### Configuring the Egress Provider Edge Router (PE2)

---

#### Step-by-Step Procedure

To configure Router PE2:

1. Create tunnel interfaces by configuring the DPC and its corresponding PIC to use tunneling services.

```
[edit]
user@PE2# set chassis fpc 1 pic 1 tunnel-services bandwidth 1g

```

2. Configure the OpenFlow interface as a Layer 2 interface.

```
[edit interfaces]
user@PE2# set ge-1/0/0 encapsulation ethernet-bridge
user@PE2# set ge-1/0/0 unit 0 family bridge

```

3. Configure the OpenFlow virtual switch routing instance.

```
[edit]
user@PE2# set routing-instances r1 instance-type virtual-switch
user@PE2# set routing-instances r1 bridge-domains bd1 interface ge-1/0/0.0
user@PE2# set routing-instances r1 bridge-domains bd1 interface lt-1/1/10.0

```

4. Configure the OpenFlow controller.

```
[edit protocols openflow]
user@PE2# set switch s1 controller protocol tcp port 6633
user@PE2# set switch s1 controller address 10.94.175.213

```

5. Configure the interfaces participating in OpenFlow.

```
[edit protocols openflow]
user@PE2# set switch s1 interfaces ge-1/0/0.0 port-id 1
user@PE2# set switch s1 interfaces lt-1/1/10.0 port-id 2

```

6. Configure the loopback interface and router ID.

```
[edit]
user@PE2# set interfaces lo0 unit 0 family inet address 10.3.3.3/32
user@PE2# set routing-options router-id 10.3.3.3

```

7. Configure the MPLS interfaces.

```
[edit interfaces]
user@PE2# set ge-1/1/2 unit 0 family inet address 10.10.2.3/24
user@PE2# set ge-1/1/2 unit 0 family mpls
user@PE2# set ge-1/1/3 unit 0 family inet address 10.10.3.3/24
user@PE2# set ge-1/1/3 unit 0 family mpls
```

8. Configure the logical tunnel interface.

```
[edit interfaces]
user@PE2# set lt-1/1/10 unit 0 family bridge
user@PE2# set lt-1/1/10 unit 0 encapsulation ethernet-bridge
user@PE2# set lt-1/1/10 unit 0 peer-unit 100
user@PE2# set lt-1/1/10 unit 100 family ccc
user@PE2# set lt-1/1/10 unit 100 encapsulation ethernet-ccc
user@PE2# set lt-1/1/10 unit 100 peer-unit 0
```

9. Enable RSVP, MPLS, and OSPF on the interfaces connected to Router P.

```
[edit protocols]
user@PE2# set rsvp interface ge-1/1/2.0
user@PE2# set rsvp interface ge-1/1/3.0
user@PE2# set mpls interface ge-1/1/2.0
user@PE2# set mpls interface ge-1/1/3.0
user@PE2# set ospf area 0.0.0.0 interface ge-1/1/2.0
user@PE2# set ospf area 0.0.0.0 interface ge-1/1/3.0
```

10. Enable traffic engineering for OSPF.

```
[edit protocols]
user@PE2# set ospf traffic-engineering
```

11. Configure the MPLS LSP from PE2 to PE1.

```
[edit protocols]
user@PE2# set mpls label-switched-path 3-1 from 10.3.3.3
user@PE2# set mpls label-switched-path 3-1 to 10.1.1.1
```

12. Configure the MPLS LSP tunnel cross-connects.

```
[edit protocols]
user@PE2# set connections remote-interface-switch 3-1-ccc interface lt-1/1/10.100
user@PE2# set connections remote-interface-switch 3-1-ccc transmit-lsp 3-1
user@PE2# set connections remote-interface-switch 3-1-ccc receive-lsp 1-3
```

13. Commit the configuration.

```
[edit]
user@PE2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it. For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
chassis {
  fpc 1 {
```

```
        pic 1 {
            tunnel-services {
                bandwidth 1g;
            }
        }
    }

    interfaces {
        ge-1/0/0 {
            encapsulation ethernet-bridge;
            unit 0 {
                family bridge;
            }
        }
        ge-1/1/2 {
            unit 0 {
                family inet {
                    address 10.10.2.3/24;
                }
                family mpls;
            }
        }
        ge-1/1/3 {
            unit 0 {
                family inet {
                    address 10.10.3.3/24;
                }
                family mpls;
            }
        }
        lt-1/1/10 {
            unit 0 {
                encapsulation ethernet-bridge;
                peer-unit 100;
                family bridge;
            }
            unit 100 {
                encapsulation ethernet-ccc;
                peer-unit 0;
                family ccc;
            }
        }
        lo0 {
            unit 0 {
                family inet {
                    address 10.3.3.3/32;
                }
            }
        }
    }

    protocols {
        rsvp {
            interface ge-1/1/2.0;
            interface ge-1/1/3.0;
        }
        mpls {
            label-switched-path 3-1 {
                from 10.3.3.3;
                to 10.1.1.1;
            }
        }
    }
}
```

```

    }
    interface ge-1/1/2.0;
    interface ge-1/1/3.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface ge-1/1/2.0;
      interface ge-1/1/3.0;
    }
  }
  connections {
    remote-interface-switch 3-1-ccc {
      interface lt-1/1/10.100;
      transmit-lsp 3-1;
      receive-lsp 1-3;
    }
  }
  openflow {
    switch s1 {
      interfaces {
        ge-1/0/0.0 port-id 1;
        lt-1/1/10.0 port-id 2;
      }
      controller {
        protocol {
          tcp {
            port 6633;
          }
        }
        address 10.94.175.213;
      }
    }
  }
}

routing-instances {
  r1 {
    instance-type virtual-switch;
    bridge-domains {
      bd1 {
        interface ge-1/0/0.0;
        interface lt-1/1/10.0;
      }
    }
  }
}

routing-options {
  router-id 10.3.3.3;
}
...

```

## Verification

Confirm that the configuration is working properly.

- [Verifying that the OpenFlow Controller Connection is Up on page 2164](#)
- [Verifying that the OpenFlow Interfaces Are Up on page 2164](#)

- [Verifying that the MPLS LSP is Operational on page 2165](#)
- [Verifying that the MPLS LSP Cross-Connect is Operational on page 2165](#)
- [Verifying the Routes on page 2166](#)

### Verifying that the OpenFlow Controller Connection is Up

- Purpose** On each of the OpenFlow-enabled routers, verify that the connection state for the OpenFlow controller is **up**.
- Action** Issue the **show openflow controller** operational mode command, and verify that the controller connection state is **up**.
- ```
user@PE1> show openflow controller
Openflowd controller information:
Controller socket: 11
Controller IP address: 10.94.175.213
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 1
Controller role: equal
```
- Meaning** The output shows that the connection state of the OpenFlow controller is **up**, in addition to other information about the controller.

### Verifying that the OpenFlow Interfaces Are Up

- Purpose** On each of the OpenFlow-enabled routers, verify that the OpenFlow interfaces are up.
- Action** Issue the **show openflow interfaces** operational mode command, and verify that the state of each interface is **Up**. For example, on PE1:
- ```
user@PE1> show openflow interfaces
Switch name: s1
Interface Name: ge-1/0/0.0
Interface port number: 1
Interface Hardware Address: 00:00:5e:00:53:b1
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: s1
Interface Name: lt-1/1/10.0
Interface port number: 2
Interface Hardware Address: 00:00:5e:00:53:be
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Disabled
Interface media type: Fiber
Interface state: Up
```
- Meaning** The output shows that the state of each OpenFlow interface is **Up**, in addition to other information about the interfaces.

### Verifying that the MPLS LSP is Operational

**Purpose** On each edge router, verify that the MPLS LSP state is **Up**.

**Action** Issue the **show mpls lsp** operational mode command, and verify that each LSP is operational.

```
user@PE1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt P    ActivePath    LSPName
10.3.3.3     10.1.1.1    Up    0 *
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPName
10.1.1.1     10.3.3.3    Up    0  1 FF  299776      - 3-1
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
user@PE2> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt P    ActivePath    LSPName
10.1.1.1     10.3.3.3    Up    0 *
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPName
10.3.3.3     10.1.1.1    Up    0  1 FF  299840      - 1-3
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** The output shows that each LSP is operational.

### Verifying that the MPLS LSP Cross-Connect is Operational

**Purpose** Verify that the MPLS LSP circuit cross-connect is operational.

**Action** Issue the **show connections remote-interface-switch** operational mode command, and verify that the circuit cross-connect state is **Up**.

```
user@PE1> show connections remote-interface-switch
CCC and TCC connections [Link Monitoring On]
[...Output truncated...]

Connection/Circuit      Type      St      Time last up    # Up trans
1-3-ccc                 rmt-if    Up      Apr 18 22:30:54
1
  1t-1/1/10.100         intf      Up
  1-3                    tlsp      Up
  3-1                    rlsp      Up

user@PE2> show connections remote-interface-switch
```

CCC and TCC connections [Link Monitoring On]  
[...Output truncated...]

| Connection/Circuit | Type   | St | Time last up    | # Up trans |
|--------------------|--------|----|-----------------|------------|
| 3-1-ccc            | rmt-if | Up | Apr 18 15:07:04 |            |
| 1                  |        |    |                 |            |
| lt-1/1/10.100      | intf   | Up |                 |            |
| 3-1                | tlsp   | Up |                 |            |
| 1-3                | rlsp   | Up |                 |            |

**Meaning** The output from both routers indicates that the circuit cross-connect is operational.

### Verifying the Routes

**Purpose** Ensure that the routes from the CCC interface over the LSP are active.

**Action** Issue the **show route ccc lt-1/1/10.100** command.

```
user@PE1> show route ccc lt-1/1/10.100
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
lt-1/1/10.100      *[CCC/7/1] 00:34:54, metric 2
                   > to 10.10.1.2 via ge-1/1/1.0, label-switched-path 1-3
```

```
user@PE2> show route ccc lt-1/1/10.100
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
lt-1/1/10.100      *[CCC/7/1] 00:35:48, metric 2
                   > to 10.10.2.2 via ge-1/1/2.0, label-switched-path 3-1
```

**Meaning** The sample output shows that the circuit cross-connect uses the configured LSPs with the MPLS interface as the exit interface.

## Troubleshooting

### Troubleshooting the Circuit Cross-Connect

**Problem** The OpenFlow-enabled router does not route OpenFlow traffic to the remote OpenFlow network.

**Solution** In order to direct traffic from the local OpenFlow network to the remote OpenFlow network, the OpenFlow controller must install flow entries that select the appropriate traffic and forward it to the correct OpenFlow interface. For traffic flowing from PE1 to PE2, the OpenFlow controller must install flow entries on PE1 that direct OpenFlow traffic from ge-1/0/0.0 to lt-1/1/10.0, and it must install flow entries on PE2 that direct the OpenFlow traffic from lt-1/1/10.0 to ge-1/0/0.0. Similarly, for traffic flowing from PE2 to PE1, the OpenFlow controller must install flow entries on PE2 that direct the desired OpenFlow traffic from ge-1/0/0.0 to lt-1/1/10.0, and it must install flow entries on PE1 that direct the OpenFlow traffic from lt-1/1/10.0 to ge-1/1/0.0.



- Related Documentation**
- [Understanding OpenFlow Traffic Steering Across MPLS Networks Using MPLS LSP Tunnel Cross-Connects on page 2149](#)
  - [Configuring Support for OpenFlow on MX Series Routers on page 2099](#)



## CHAPTER 27

# Configuration Statements

- [\[edit protocols openflow\] Hierarchy Level](#) on page 2169
- [address \(Protocols OpenFlow\)](#) on page 2170
- [controller \(Protocols OpenFlow\)](#) on page 2171
- [default-action \(Protocols OpenFlow\)](#) on page 2172
- [id \(Protocols OpenFlow\)](#) on page 2173
- [interfaces \(Protocols OpenFlow\)](#) on page 2174
- [openflow \(Protocols OpenFlow\)](#) on page 2175
- [port \(Protocols OpenFlow\)](#) on page 2176
- [protocol \(Protocols OpenFlow\)](#) on page 2177
- [purge-flow-timer \(Protocols OpenFlow\)](#) on page 2178
- [role \(Protocols OpenFlow\)](#) on page 2179
- [switch \(Protocols OpenFlow\)](#) on page 2180
- [traceoptions \(Protocols OpenFlow\)](#) on page 2181

### [\[edit protocols openflow\] Hierarchy Level](#)

---

```
protocols {
  openflow {
    switch switch-name {
      controller {
        address address;
        id id;
        protocol tcp {
          port port;
        }
        role equal;
      }
      default-action (drop | packet-in);
      interfaces {
        interface-name port-id port;
      }
      purge-flow-timer seconds;
    }
    traceoptions {
```

```
        file <filename> <files number> <match regular-expression> <size size>
          <world-readable | no-world-readable>;
        flag flag;
      }
    }
  }
```

- Related Documentation**
- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
  - [openflow \(Protocols OpenFlow\) on page 2175](#)
  - [OpenFlow Operational Mode Commands on page 2183](#)

---

## address (Protocols OpenFlow)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>address address;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow switch switch-name controller</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Specify the IPv4 address of the OpenFlow controller that will manage OpenFlow on that virtual switch. The switch establishes a connection to the controller using this address.  |
| <b>Options</b>                  | <b>address</b> —IPv4 address of the OpenFlow controller.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">controller (Protocols OpenFlow) on page 2171</a></li><li>• <a href="#">protocol (Protocols OpenFlow) on page 2177</a></li><li>• <a href="#">switch (Protocols OpenFlow) on page 2180</a></li></ul> |

## controller (Protocols OpenFlow)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> controller {   address address;   id id;   protocol tcp {     port port;   }   role equal; } </pre>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow switch switch-name</a> ]   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>   |
| <b>Description</b>              | Configure the OpenFlow controller connection information for a virtual switch on an OpenFlow-enabled device running Junos OS. If you configure a virtual switch with a single controller, by default, the controller is in active mode, and the switch automatically initiates a connection to the controller.  |
| <b>Options</b>                  | The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li> <li>• <a href="#">Understanding the Virtual Switch Connection to the OpenFlow Controller on Devices Running Junos OS on page 2056</a></li> <li>• <a href="#">address (Protocols OpenFlow) on page 2170</a></li> <li>• <a href="#">protocol (Protocols OpenFlow) on page 2177</a></li> <li>• <a href="#">role (Protocols OpenFlow) on page 2179</a></li> <li>• <a href="#">switch (Protocols OpenFlow) on page 2180</a></li> </ul> |

## default-action (Protocols OpenFlow)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | default-action (drop   packet-in);  |
| <b>Hierarchy Level</b>          | [edit protocols <b>openflow</b> <b>switch</b> <i>switch-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Specify the default action that is executed when an OpenFlow packet does not match an existing flow entry. The default action is specific to the OpenFlow virtual switch and is the same across all filters associated with that virtual switch.  |
| <b>Default</b>                  | If you do not include the <b>default-action</b> statement, the default action is <b>packet-in</b> .   |
| <b>Options</b>                  | <b>drop</b> —Drop packets that do not match an existing flow entry.<br><br><b>packet-in</b> —Accept packets that do not match an existing flow entry, and forward the packet to the controller.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">Understanding OpenFlow Flows and Filters on Devices Running Junos OS on page 2058</a></li><li>• <a href="#">openflow (Protocols OpenFlow) on page 2175</a></li><li>• <a href="#">switch (Protocols OpenFlow) on page 2180</a></li></ul> |

---

## id (Protocols OpenFlow)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>id id;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow switch switch-name controller</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.             |
| <b>Description</b>              | Specify an optional numeric identifier for the OpenFlow controller.  |
| <b>Options</b>                  | <i>id</i> —Numeric identifier for the OpenFlow controller.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">controller (Protocols OpenFlow) on page 2171</a></li></ul> |

## interfaces (Protocols OpenFlow)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>interfaces {<br/>    <i>interface-name</i> port-id <i>port</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow</a> <a href="#">switch</a> <i>switch-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Configure a Layer 2 interface as an OpenFlow-enabled interface.  |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface, including the logical unit number—for example, ge-1/1/0.0.</p> <p><b><i>port-id port</i></b>—(Optional) Unique numeric value specifying the port ID associated with the OpenFlow interface. You can manually configure a port ID in the range 1 through 32640. If you do not specify a port, the system generates a value in the range from 32641 through 65280.</p> <p><b>Range:</b> 1 through 32640</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">openflow (Protocols OpenFlow) on page 2175</a></li><li>• <a href="#">switch (Protocols OpenFlow) on page 2180</a></li></ul>  |



## openflow (Protocols OpenFlow)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> openflow {     switch <i>switch-name</i> {         controller {             address <i>address</i>;             id <i>id</i>;             protocol tcp {                 port <i>port</i>;             }             role equal;         }         default-action (drop   packet-in);         interfaces {             interface-name port-id <i>port</i>;         }         purge-flow-timer <i>seconds</i>;     }     traceoptions {         file &lt;<i>filename</i>&gt; &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt;         &lt;world-readable   no-world-readable&gt;;         flag <i>flag</i>;     } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>   |
| <b>Description</b>              | Configure support for OpenFlow on a device running Junos OS. To configure OpenFlow, the device must be running a release that supports OpenFlow and have the OpenFlow software package installed. The OpenFlow software package release must match the Junos OS release of the device on which the software is installed.   |
| <b>Default</b>                  | OpenFlow is disabled on the device.   |
| <b>Options</b>                  | The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Support on Devices Running Junos OS</a></li> <li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> </ul>  |

## port (Protocols OpenFlow)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>port port;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow switch switch-name controller protocol protocol</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Specify the OpenFlow controller port to which the OpenFlow virtual switch connects.   |
| <b>Options</b>                  | <b>port</b> —Numeric value specifying the OpenFlow controller port to which the device should connect.<br><b>Range:</b> 1024 through 65,535<br><b>Default:</b> 6633   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">address (Protocols OpenFlow) on page 2170</a></li><li>• <a href="#">controller (Protocols OpenFlow) on page 2171</a></li><li>• <a href="#">protocol (Protocols OpenFlow) on page 2177</a></li></ul> |


## protocol (Protocols OpenFlow)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | protocol tcp {<br>port port;<br>}   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow switch switch-name controller</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Specify the connection protocol that the OpenFlow virtual switch uses to connect to the OpenFlow controller.  |
| <b>Options</b>                  | tcp—Establish a TCP connection to the controller.<br><br>The remaining statement is explained separately.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li> <li>• <a href="#">controller (Protocols OpenFlow) on page 2171</a></li> <li>• <a href="#">port (Protocols OpenFlow) on page 2176</a></li> </ul> |

## purge-flow-timer (Protocols OpenFlow)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>purge-flow-timer seconds;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <b>openflow</b> <b>switch</b> <i>switch-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | <p>For an OpenFlow virtual switch, specify the number of seconds after which an invalid OpenFlow flow entry is deleted from the flow table.</p> <p>If you do not configure the <b>purge-flow-timer</b> statement, the device purges invalid flows from hardware, but indefinitely retains the corresponding flow entries in the flow table. If you configure the <b>purge-flow-timer</b> statement, the device purges invalid flows from hardware, and after the specified number of seconds, the device deletes the invalid flow entries from the flow table. Configuring a value of 0 causes the device to immediately delete invalid flow entries from the flow table.</p> <div> <b>NOTE:</b> By default, if you remove an active OpenFlow interface from an existing OpenFlow configuration, flow entries that match on this interface as the ingress interface and flow entries that include this interface in their action list (for OpenFlow v1.0) or flow instructions (for OpenFlow v1.3.1) are invalid and are automatically purged from the flow table and from the hardware regardless of whether you configure the <b>purge-flow-timer</b> statement.</div> |
| <b>Options</b>                  | <p><b>seconds</b>—Number of seconds after which an invalid flow entry is deleted from the flow table.</p> <p><b>Range:</b> 0 through 300</p>   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">Understanding OpenFlow Flow Entry Timers on Devices Running Junos OS on page 2062</a></li></ul>  |

## role (Protocols OpenFlow)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | role equal;   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">openflow switch switch-name controller</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.3.<br>Statement introduced in Junos OS Release 13.3 for EX Series switches.<br>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Specify the role of each OpenFlow controller when configuring more than one controller for a virtual switch. A single controller configuration automatically puts the controller in active mode. In active mode, the virtual switch automatically initiates a connection to the controller. |
| <b>Options</b>                  | <b>equal</b> —Configure the controller as the active controller in a single controller configuration.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li> <li>• <a href="#">controller (Protocols OpenFlow) on page 2171</a></li> </ul>   |

## switch (Protocols OpenFlow)

---

**Syntax**    `switch switch-name {  
              controller {  
                  address address;  
                  id id;  
                  protocol tcp {  
                      port port;  
                  }  
                  role equal;  
              }  
              default-action (drop | packet-in);  
              interfaces {  
                  interface-name port-id port;  
              }  
              purge-flow-timer seconds;  
          }`

**Hierarchy Level**    [edit protocols [openflow](#)]

**Release Information**    Statement introduced in Junos OS Release 13.3.  
Statement introduced in Junos OS Release 13.3 for EX Series switches.  
Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.

**Description**    Configure an OpenFlow virtual switch.

**Options**    *switch-name*—User-configured identifier for the OpenFlow virtual switch. The identifier must be 60 characters or less.

The remaining statements are explained separately.

**Required Privilege Level**    admin—To view this statement in the configuration.  
admin-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Support for OpenFlow on Devices Running Junos OS on page 2049](#)
- [controller \(Protocols OpenFlow\) on page 2171](#)
- [default-action \(Protocols OpenFlow\) on page 2172](#)
- [interfaces \(Protocols OpenFlow\) on page 2174](#)
- [openflow \(Protocols OpenFlow\) on page 2175](#)

## traceoptions (Protocols OpenFlow)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> traceoptions {     file &lt;filename&gt; &lt;files number&gt; &lt;match regular-expression&gt; &lt;size size&gt;     &lt;world-readable   no-world-readable&gt;;     flag flag;     no-remote-trace; } </pre>  |
| <b>Hierarchy Level</b>     | [edit protocols openflow]  |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>  |
| <b>Description</b>         | Define tracing operations for OpenFlow.  |
| <b>Default</b>             | If you do not include this statement, no OpenFlow-specific tracing operations are performed.   |
| <b>Options</b>             | <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. All files are placed in the <code>/var/log</code> directory.</p> <p><b>Default:</b> <code>/var/log/ofd</code></p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed and compressed to <i>trace-file.0.gz</i>. When <i>trace-file</i> again reaches its maximum size, <i>trace-file.0.gz</i> is renamed <i>trace-file.1.gz</i>, and <i>trace-file</i> is renamed and compressed to <i>trace-file.0.gz</i>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size by using the <b>size</b> option and also a filename.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All OpenFlow events.</li> <li>• <b>barrier</b>—OpenFlow barrier events.</li> <li>• <b>configuration</b>—OpenFlow configuration events.</li> <li>• <b>filter</b>—OpenFlow filter events.</li> <li>• <b>flow</b>—OpenFlow flow events.</li> <li>• <b>function</b>—OpenFlow entry and exit events.</li> <li>• <b>group</b>—(Appears only for Juniper Networks devices running OpenFlow v1.3.1 or later) OpenFlow group events.</li> </ul> |

- **interface**—OpenFlow interface events.
- **nh**—OpenFlow next-hop events.
- **packet-io**—OpenFlow packet in and packet out events.
- **packets**—OpenFlow packet events.
- **statistics**—OpenFlow statistics request and reply events.
- **switch**—OpenFlow switch events including controller connection messages and keepalives, and packets sent to and received from the controller.

**match *regular-expression***—(Optional) Only log lines that match the regular expression.

**no-remote-trace**—(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or have passed through the Juniper Networks device.

**no-world-readable**—(Optional) Disable unrestricted file access, which restricts file access to the owner. This is the default.

**size *size***—(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

**Syntax:** *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB

**Range:** 10,240 through 1,073,741,824 bytes

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Understanding Support for OpenFlow on Devices Running Junos OS on page 2049</a></li><li>• <a href="#">openflow (Protocols OpenFlow) on page 2175</a></li></ul> |
|------------------------------|--|



# Operational Commands

- [OpenFlow Operational Mode Commands on page 2183](#)
- [show openflow capability](#)
- [show openflow controller](#)
- [show openflow filters](#)
- [show openflow flows](#)
- [show openflow groups](#)
- [show openflow interfaces](#)
- [show openflow statistics flows](#)
- [show openflow statistics groups](#)
- [show openflow statistics interfaces](#)
- [show openflow statistics packet](#)
- [show openflow statistics queue](#)
- [show openflow statistics summary](#)
- [show openflow statistics tables](#)
- [show openflow summary](#)
- [show openflow switch](#)

## OpenFlow Operational Mode Commands

[Table 188 on page 2183](#) summarizes the operational mode commands that you can use to monitor and troubleshoot OpenFlow operations on an OpenFlow-enabled device running Junos OS. Commands are listed in alphabetical order.

**Table 188: OpenFlow Operational Mode Commands**

| Command                                  | Task  |
|--|---|
| <a href="#">show openflow capability</a> | Display support information for OpenFlow features, actions, and match conditions on the device. |
| <a href="#">show openflow controller</a> | Display OpenFlow controller information and status.   |
| <a href="#">show openflow filters</a>    | Display information for filters bound to OpenFlow interfaces.                                   |

Table 188: OpenFlow Operational Mode Commands (*continued*)

| Command  | Task  |
|--|---|
| <code>show openflow flows</code>                 | Display OpenFlow flow information.  |
| <code>show openflow groups</code>                | Display information about OpenFlow groups.<br><br><i>NOTE:</i> This command is supported only on Juniper Networks devices running OpenFlow v1.3.1 or later. |
| <code>show openflow interfaces</code>            | Display physical characteristics and status information for interfaces participating in OpenFlow.   |
| <code>show openflow statistics flows</code>      | Display statistics for OpenFlow flow entries.   |
| <code>show openflow statistics groups</code>     | Display statistics for OpenFlow groups.<br><br><i>NOTE:</i> This command is supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.    |
| <code>show openflow statistics interfaces</code> | Display statistics for interfaces participating in OpenFlow.  |
| <code>show openflow statistics packet</code>     | Display statistics for packet-in and packet-out actions.  |
| <code>show openflow statistics queue</code>      | Display statistics for OpenFlow queues in hardware.   |
| <code>show openflow statistics summary</code>    | Display summary statistics for all OpenFlow flows.  |
| <code>show openflow statistics tables</code>     | Display statistics for OpenFlow flow tables.  |
| <code>show openflow summary</code>               | Display summary information for OpenFlow flows.   |
| <code>show openflow switch</code>                | Display OpenFlow message statistics for OpenFlow virtual switches.  |

## show openflow capability

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow capability</b><br><action   feature   match-condition>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Display support information for OpenFlow features, actions, and match conditions on the device.  |
| <b>Options</b>                  | <b>none</b> —Display support information for all OpenFlow capabilities.<br><br><b>action</b> —(Optional) Display support information for OpenFlow actions.<br><br><b>feature</b> —(Optional) Display support information for OpenFlow features.<br><br><b>match-condition</b> —(Optional) Display support information for OpenFlow match conditions. |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">OpenFlow Support on Devices Running Junos OS</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show openflow capability on page 2188</a><br><a href="#">show openflow capability (OpenFlow 1.3.1) on page 2188</a><br><a href="#">show openflow capability action on page 2189</a><br><a href="#">show openflow capability feature on page 2189</a><br><a href="#">show openflow capability match-condition on page 2190</a>            |
| <b>Output Fields</b>            | Table 189 on page 2185 lists the output fields for the <b>show openflow capability</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 189: show openflow capability Output Fields**

| Field Name  | Field Description   |
|---|---|
| Supported Features—Indicates Support for the Following OpenFlow Features                |   |
| Flow statistics   | Indicates whether the switch supports OpenFlow flow statistics.       |
| Table statistics  | Indicates whether the switch supports OpenFlow flow table statistics. |
| Port statistics   | Indicates whether the switch supports OpenFlow port statistics.       |
| Group statistics  | Indicates whether the switch supports OpenFlow group statistics.      |
| NOTE: This field appears only if the Junos OS device supports OpenFlow v1.3.1 or later. |   |

Table 189: show openflow capability Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| 802.1d spanning tree  | Indicates whether the switch supports the 802.1D Spanning Tree Protocol.   |
| Reassemble IP fragments   | Indicates whether the switch supports reassembling IP fragments.   |
| Queue statistics  | Indicates whether the switch supports OpenFlow queue statistics.   |
| Match IP addresses in ARP pkts  | Indicates whether the switch supports matching on IP addresses in ARP packets.   |
| <b>Supported Match Conditions—Indicates Support for the Following OpenFlow Match Conditions</b> |  |
| Switch input port   | Displays support for matching against the ingress switch port.   |
| VLAN vid  | Displays support for matching against the VLAN identifier in the outermost VLAN tag.   |
| Ethernet source address   | Displays support for matching against the Ethernet source address.   |
| Ethernet destination address  | Displays support for matching against the Ethernet destination address.  |
| Ethernet frame type   | Displays support for matching against the Ethernet frame type.   |
| IP protocol   | Displays support for matching against the IP protocol or lower 8 bits of the ARP opcode.   |
| TCP/UDP source port   | Displays support for matching against the TCP or UDP source port.  |
| TCP/UDP destination port  | Displays support for matching against the TCP or UDP destination port.   |
| IP(v4) source address   | Displays support for matching against the IPv4 source address.<br><br><b>NOTE:</b> <b>IPv4</b> appears only if the Junos OS device supports OpenFlow v1.3.1 or later.      |
| IP(v4) destination address  | Displays support for matching against the IPv4 destination address.<br><br><b>NOTE:</b> <b>IPv4</b> appears only if the Junos OS device supports OpenFlow v1.3.1 or later. |
| IPv6 source address   | Displays support for matching against the IPv6 source address.<br><br><b>NOTE:</b> This field appears only if the Junos OS device supports OpenFlow v1.3.1 or later.       |

Table 189: show openflow capability Output Fields (*continued*)

| Field Name  | Field Description   |
|---|---|
| IPv6 destination address  | Displays support for matching against the IPv6 destination address.<br><br><b>NOTE:</b> This field appears only if the Junos OS device supports OpenFlow v1.3.1 or later.                                     |
| VLAN priority   | Displays support for matching against the VLAN priority in the outermost VLAN tag.  |
| IP ToS (DSCP field)   | Displays support for matching against the IPv4 ToS bits.  |
| <b>Supported Actions—Indicates Support for the Following OpenFlow Actions</b> |   |
| Output to switch port   | Displays support for forwarding the packet to a specified port.   |
| Set the 802.1q VLAN id  | Displays support for the optional Modify-Field action to modify the existing 802.1Q VLAN ID of the outermost VLAN tag in the frame header or to add a new header with the VLAN ID if none exists.             |
| Set the 802.1q priority   | Displays support for the optional Modify-Field action to modify the existing 802.1Q VLAN priority of the outermost VLAN tag in the frame header or to add a new header with the VLAN priority if none exists. |
| Strip the 802.1q header   | Displays support for the optional Modify-Field action to remove the outermost VLAN header in the frame.   |
| Ethernet source address   | Displays support for the optional Modify-Field action to modify the Ethernet source address field in the frame header.  |
| Ethernet destination address  | Displays support for the optional Modify-Field action to modify the Ethernet destination address field in the frame header.   |
| IP source address   | Displays support for the optional Modify-Field action to modify the IP source address field and update the checksum in the packet header.   |
| IP destination address  | Displays support for the optional Modify-Field action to modify the IP destination address field and update the checksum in the packet header.  |
| IP ToS (DSCP)   | Displays support for the optional Modify-Field action to modify the IPv4 ToS field in the packet header.  |
| TCP/UDP source port   | Displays support for the optional Modify-Field action to modify the TCP or UDP source port field and update the checksum in the packet header.  |
| TCP/UDP destination port  | Displays support for the optional Modify-Field action to modify the TCP or UDP destination port field and update the checksum in the packet header.   |

Table 189: show openflow capability Output Fields (*continued*)

| Field Name      | Field Description   |
|-----------------|---|
| Output to queue | Displays support for the optional Enqueue action to set the queue ID for the packet.  |
| Execute group   | Displays support for a group action to be executed.<br><br><b>NOTE:</b> This field appears only if the Junos OS device supports OpenFlow v1.3.1 or later. |

## Sample Output

### show openflow capability

```

user@host> show openflow capability
Openflowd platform feature support:
Flow statistics:      Yes
Table statistics:     Yes
Port statistics:      Yes
802.1d spanning tree: No
Reassemble IP fragments: No
Queue statistics:     Yes
Match IP addresses in ARP pkts: No

Openflowd platform match condition support:
Switch input port:    Yes
VLAN vid:             Yes
Ethernet source address: Yes
Ethernet destination address: Yes
Ethernet frame type: Yes
IP protocol:          Yes
TCP/UDP source port:  Yes
TCP/UDP destination port: Yes
IP source address:     Yes
IP destination address: Yes
VLAN priority:         Yes
IP ToS (DSCP field):  Yes

Openflowd platform action support:
Output to switch port: Yes
Set the 802.1q VLAN id: Yes
Set the 802.1q priority: No
Strip the 802.1q header: Yes
Ethernet source address: No
Ethernet destination address: No
IP source address:     No
IP destination address: No
IP ToS (DSCP):         No
TCP/UDP source port:   No
TCP/UDP destination port: No
Output to queue:       No

```

### show openflow capability (OpenFlow 1.3.1)

```

user@host> show openflow capability
Openflowd platform feature support:
Flow statistics:      Yes

```

```

Table statistics:    Yes
Port statistics:    Yes
Group statistics:    Yes
802.1d spanning tree:  No
Reassemble IP fragments:  No
Queue statistics:    Yes
Match IP addresses in ARP pkts:  No

Openflowd platform match condition support:
Switch input port:    Yes
VLAN vid:            Yes
Ethernet source address:  Yes
Ethernet destination address:  Yes
Ethernet frame type:  Yes
IP protocol:          Yes
TCP/UDP source port:    Yes
TCP/UDP destination port:  Yes
IPv4 source address:    Yes
IPv4 destination address:  Yes
IPv6 source address:    No
IPv6 destination address:  No
VLAN priority:         Yes
IP ToS (DSCP field):    Yes

Openflowd platform action support:
Output to switch port:    Yes
Set the 802.1q VLAN id    Yes
Set the 802.1q priority:  No
Strip the 802.1q header:  Yes
Ethernet source address:  No
Ethernet destination address:  No
IP source address:        No
IP destination address:    No
IP ToS (DSCP):            No
TCP/UDP source port:      No
TCP/UDP destination port:  No
Output to queue:          No
Execute Group:            Yes

```

#### show openflow capability action

```

user@host> show openflow capability action
Openflowd platform action support:
Output to switch port:    Yes
Set the 802.1q VLAN id    Yes
Set the 802.1q priority:  No
Strip the 802.1q header:  Yes
Ethernet source address:  No
Ethernet destination address:  No
IP source address:        No
IP destination address:    No
IP ToS (DSCP):            No
TCP/UDP source port:      No
TCP/UDP destination port:  No
Output to queue:          No

```

#### show openflow capability feature

```

user@host> show openflow capability feature
Openflowd platform feature support:
Flow statistics:          Yes

```

```
Table statistics:   Yes
Port statistics:   Yes
802.1d spanning tree: No
Reassemble IP fragments: No
Queue statistics:   Yes
Match IP addresses in ARP pkts: No
```

#### **show openflow capability match-condition**

```
user@host> show openflow capability match-condition
Openflowd platform match condition support:
Switch input port:   Yes
VLAN vid:            Yes
Ethernet source address: Yes
Ethernet destination address: Yes
Ethernet frame type: Yes
IP protocol:         Yes
TCP/UDP source port:  Yes
TCP/UDP destination port: Yes
IP source address:    Yes
IP destination address: Yes
VLAN priority:        Yes
IP ToS (DSCP field):  Yes
```



## show openflow controller

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow controller</b><br><b>&lt;address address&gt;</b><br><b>&lt;switch switch-name&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Display OpenFlow controller information and connection status. OpenFlow controllers are configured at the <b>[edit protocols openflow switch switch-name]</b> hierarchy level.   |
| <b>Options</b>                  | <b>none</b> —Display information about all configured controllers.<br><br><b>address address</b> —(Optional) Display information about the controller at the specified IP address.<br><br><b>switch switch-name</b> —(Optional) Display information about controllers associated with the specified virtual switch.  |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">controller (Protocols OpenFlow) on page 2171</a></li> <li>• <a href="#">OpenFlow Support on Devices Running Junos OS</a></li> <li>• <a href="#">Understanding the OpenFlow Version Negotiation Between the Controller and Junos OS Devices on page 2057</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow controller (OpenFlow 1.3.1) on page 2192</a><br><a href="#">show openflow controller address (OpenFlow 1.3.1) on page 2192</a><br><a href="#">show openflow controller switch (OpenFlow 1.3.1) on page 2192</a>  |
| <b>Output Fields</b>            | <a href="#">Table 190 on page 2191</a> lists the output fields for the <b>show openflow controller</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 190: show openflow controller Output Fields**

| Field Name                   | Field Description   |
|------------------------------|---|
| <b>Controller socket</b>     | Socket on the controller to which the OpenFlow virtual switch connects.   |
| <b>Controller IP address</b> | IP address of the OpenFlow controller.                                    |
| <b>Controller protocol</b>   | Protocol used by the switch to initiate a connection with the controller. |

Table 190: show openflow controller Output Fields (*continued*)

| Field Name                          | Field Description  |
|-------------------------------------|--|
| <b>Controller port</b>              | Port on the controller to which the OpenFlow virtual switch connects.  |
| <b>Controller connection state</b>  | Status of the connection between the OpenFlow virtual switch and the controller.   |
| <b>Number of connection attempt</b> | Number of connection attempts made by the virtual switch to the controller.  |
| <b>Controller role</b>              | User-configured role of the controller.  |
| <b>Negotiated version</b>           | <p>A numerical value that represents the OpenFlow version that is negotiated between the Junos OS device and the OpenFlow controller during the initial connection.</p> <p><b>NOTE:</b> This field appears only if the Junos OS device supports OpenFlow version 1.3.1 or later.</p> |

## Sample Output

### show openflow controller (OpenFlow 1.3.1)

```
user@host> show openflow controller
Openflowd controller information:
Controller socket: 15
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 5
Controller role: equal
Negotiated version: 4
```

### show openflow controller address (OpenFlow 1.3.1)

```
user@host> show openflow controller address 198.51.100.174
Openflowd controller information:
Controller socket: 15
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
Controller connection state: up
Number of connection attempt: 5
Controller role: equal
Negotiated version: 4
```

### show openflow controller switch (OpenFlow 1.3.1)

```
user@host> show openflow controller switch OFswitch1
Openflowd controller information:
Controller socket: 15
Controller IP address: 198.51.100.174
Controller protocol: tcp
Controller port: 6633
```

```
Controller connection state: up  
Number of connection attempt: 5  
Controller role: equal  
Negotiated version: 4
```

## show openflow filters

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow filters</b><br><b>&lt;interface <i>interface-name</i>&gt;</b><br><b>&lt;switch <i>switch-name</i>&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Display information for filters bound to OpenFlow interfaces.  |
| <b>Options</b>                  | <p><b>none</b>—Display information for all filters that are bound to OpenFlow interfaces.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display information for the filter bound to the specified OpenFlow interface. The interface name must include the logical unit number.</p> <p><b>switch <i>switch-name</i></b>—(Optional) Display information for filters bound to the interfaces configured under the specified OpenFlow virtual switch.</p> |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">Understanding OpenFlow Flows and Filters on Devices Running Junos OS on page 2058</a></li> </ul>   |
| <b>List of Sample Output</b>    | <a href="#">show openflow filters on page 2195</a><br><a href="#">show openflow filters interface on page 2195</a><br><a href="#">show openflow filters switch on page 2195</a>  |
| <b>Output Fields</b>            | <a href="#">Table 191 on page 2194</a> lists the output fields for the <b>show openflow filters</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 191: show openflow filters Output Fields**

| Field Name        | Field Description   |
|-------------------|---|
| Switch Name       | User-configured identifier for the OpenFlow virtual switch associated with the interface to which the filter is bound.                            |
| Number of filters | Number of filters bound to OpenFlow interfaces on the virtual switch.   |
| Default action    | Default action executed for packets that do not match any existing flow entries. Values are <b>PACKET IN</b> or <b>DROP</b> .                     |
| Filter name       | Filter identifier consisting of the concatenation of the interface name (including the logical unit number) and an internally assigned switch ID. |
| Filter index      | Auto-generated string that identifies the filter.   |

Table 191: show openflow filters Output Fields (*continued*)

| Field Name           | Field Description   |
|----------------------|---|
| Number of terms      | Number of terms in the filter. Each term consists of match conditions and actions.  |
| Number of priorities | Number of unique active flow priorities in the filter.  |
| Term name            | Filter term identifier, which consists of the filter name (interface name and switch ID), the flow priority, and a sequence number. |
| Priority ID          | Flow entry priority. Higher priority terms are installed above lower priority terms.  |
| Flow ID              | Flow identifier associated with that flow entry.  |
| Number of packets    | Number of packets that have matched a filter term. A filter term is equivalent to a flow entry.                                     |
| Number of bytes      | Number of bytes that have matched a filter term. A filter term is equivalent to a flow entry.                                       |

## Sample Output

### show openflow filters

```
user@host> show openflow filters
```

| Switch Name | Filter Index | Number of terms | Number of priorities | Number of packets |
|-------------|--------------|-----------------|----------------------|-------------------|
| OFswitch1   | 96468992     | 0               | 0                    | 0                 |
|             | 96468993     | 0               | 0                    | 0                 |
|             | 96468994     | 0               | 0                    | 0                 |
|             | 96468995     | 0               | 0                    | 0                 |
|             | 96468996     | 1               | 1                    | 7928017621        |

### show openflow filters interface

```
user@host> show openflow filters interface ge-1/1/7.0
```

```
Switch Name: OFswitch1
```

```
Filter name: ge-1/1/7.0_0
```

```
Filter index: 96468996
```

```
Number of terms: 1
```

```
Number of priorities: 1
```

```
Term name: ge-1/1/7.0_0:32766^OF:1
```

```
Priority ID: 32766
```

```
Flow ID: 16842752
```

```
Number of packets: 7941332819
```

```
Number of bytes: 476479969140
```

### show openflow filters switch

```
user@host> show openflow filters switch OFswitch1
```

```
Switch Name: OFswitch1
```

```
Number of filters: 5
```

```
Default action: PACKET IN
```

```
Filter name: ge-1/1/0.0_0
```

```
Filter index: 96468992
```

|                           |                         |            |              |
|---------------------------|-------------------------|------------|--------------|
| Number of terms: 0        | Number of priorities: 0 |            |              |
| Filter name: ge-1/1/1.0_0 |                         |            |              |
| Filter index: 96468993    |                         |            |              |
| Number of terms: 0        | Number of priorities: 0 |            |              |
| Filter name: ge-1/1/2.0_0 |                         |            |              |
| Filter index: 96468994    |                         |            |              |
| Number of terms: 0        | Number of priorities: 0 |            |              |
| Filter name: ge-1/1/3.0_0 |                         |            |              |
| Filter index: 96468995    |                         |            |              |
| Number of terms: 0        | Number of priorities: 0 |            |              |
| Filter name: ge-1/1/7.0_0 |                         |            |              |
| Filter index: 96468996    |                         |            |              |
| Number of terms: 1        | Number of priorities: 1 |            |              |
| Priority                  | Flow                    | Number of  | Number of    |
| ID                        | ID                      | packets    | bytes        |
| 32768                     | 16842752                | 7941332819 | 476479969140 |

## show openflow flows

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow flows</b><br><brief   detail   summary><br>< <i>flow-id</i> ><br><switch <i>switch-name</i> >  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display information for OpenFlow flows.   |
| <b>Options</b>                  | <b>none</b> —Display information for all flows.<br><br><b>brief   detail   summary</b> —(Optional) Display the specified level of output.<br><br><b><i>flow-id</i></b> —(Optional) Display information only for the specified flow.<br><br><b>switch <i>switch-name</i></b> —(Optional) Display information only for the flows on the specified OpenFlow virtual switch.  |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">OpenFlow Support on Devices Running Junos OS</a></li> </ul>   |
| <b>List of Sample Output</b>    | <a href="#">show openflow flows switch brief on page 2198</a><br><a href="#">show openflow flows switch detail on page 2198</a><br><a href="#">show openflow flows switch detail (OpenFlow 1.3.1) on page 2199</a><br><a href="#">show openflow flows switch summary on page 2199</a><br><a href="#">show openflow flows brief (Specific Flow) on page 2199</a><br><a href="#">show openflow flows detail (Specific Flow) on page 2199</a><br><a href="#">show openflow flows detail (Specific Flow, OpenFlow 1.3.1) on page 2199</a><br><a href="#">show openflow flows summary (Specific Flow) on page 2200</a> |
| <b>Output Fields</b>            | Table 192 on page 2197 lists the output fields for the <b>show openflow flows</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 192: show openflow flows Output Fields**

| Field Name             | Field Description   |
|------------------------|---|
| <b>Switch Name</b>     | User-configured identifier for the OpenFlow virtual switch on which the flow resides. |
| <b>Number of flows</b> | Number of active flow entries associated with that OpenFlow virtual switch.           |
| <b>Flow name</b>       | Flow descriptor.  |

Table 192: show openflow flows Output Fields (*continued*)

| Field Name               | Field Description   |
|--------------------------|---|
| <b>Table ID</b>          | Identifier for the flow table from which the flow originated.   |
| <b>Flow ID</b>           | Flow identifier associated with that flow entry.  |
| <b>Number of packets</b> | Number of packets that have matched the flow entry.   |
| <b>Priority</b>          | Flow entry priority. Packets match higher priority entries before matching lower priority entries.  |
| <b>Idle timeout</b>      | Number of seconds after which the flow entry is removed from the flow table provided there are no matching packets.   |
| <b>Hard timeout</b>      | Number of seconds after which the flow entry is removed from the flow table regardless of the number of matching packets.   |
| <b>Cookie</b>            | An identifier, which is specified by the OpenFlow controller when the flow is installed in the flow table. Cookies are used to filter flows for flow modification and delete operations.<br><br><b>NOTE:</b> This field appears only if the Juniper Networks device supports OpenFlow version 1.3.1 or later. |
| <b>Match</b>             | Configured match conditions against which the incoming packet is compared.  |
| <b>Action</b>            | Set of actions (for OpenFlow v1.0) or flow instructions (for OpenFlow v1.3.1) applied to a packet when it matches the flow entry.   |
| <b>Number of match</b>   | Number of match conditions against which the incoming packet is compared.   |
| <b>Number of action</b>  | Number of actions (for OpenFlow v1.0) or flow instructions (for OpenFlow v1.3.1) that are applied to a packet when it matches the flow entry.   |

## Sample Output

### show openflow flows switch brief

```
user@host> show openflow flows switch OFswitch1 brief
```

| Switch Name | Flow ID  | Number of packets | Priority | Number of match | Number of action |
|-------------|----------|-------------------|----------|-----------------|------------------|
| OFswitch1   | 16842752 | 8075372509        | 32768    | 1               | 1                |

### show openflow flows switch detail

```
user@host> show openflow flows switch OFswitch1 detail
```

```
Flow name: flow-16842752
Table ID: 1      Flow ID: 16842752
Priority: 32768  Idle timeout(in sec):0      Hard timeout(in sec): 0
Match: Input port: 45549
      Ethernet src addr: wildcard
      Ethernet dst addr: wildcard
```



```

        Input vlan id: wildcard          Input VLAN priority: wildcard
        Ether type: wildcard
        IP ToS: wildcard                 IP protocol: wildcard
        IP src addr: wildcard            IP dst addr: wildcard
        Source port: wildcard            Destination port: wildcard
    Action: Output port 41350,

```

### show openflow flows switch detail (OpenFlow 1.3.1)

```

user@host> show openflow flows switch OFswitch1 detail
Flow name: flow-16842752
Table ID: 1      Flow ID: 16842752
Priority: 32768   Idle timeout(in sec):0      Hard timeout(in sec): 0
Cookie: 0
Match: Input port: 45549
      Ethernet src addr: wildcard
      Ethernet dst addr: wildcard
      Input vlan id: wildcard          Input VLAN priority: wildcard
      Ether type: wildcard
      IP ToS: wildcard                 IP protocol: wildcard
      IPv4 src addr: wildcard
      IPv4 dst addr: wildcard
      IPv6 src addr: none
      IPv6 dst addr: none
      IP src addr: wildcard            IP dst addr: wildcard
      Source port: wildcard            Destination port: wildcard
    Action: Group 20,

```

### show openflow flows switch summary

```
user@host> show openflow flows switch OFswitch1 summary
```

| Switch Name | Number of Flows |
|-------------|-----------------|
| OFswitch1   | 1               |

### show openflow flows brief (Specific Flow)

```
user@host> show openflow flows 16842752 brief
```

| Switch Name | Flow ID  | Number of packets | Priority | Number of match | Number of action |
|-------------|----------|-------------------|----------|-----------------|------------------|
| OFswitch1   | 16842752 | 8056139439        | 32768    | 1               | 1                |

### show openflow flows detail (Specific Flow)

```

user@host> show openflow flows 16842752 detail
Flow name: flow-16842752
Table ID: 1      Flow ID: 16842752
Priority: 32768   Idle timeout(in sec):0      Hard timeout(in sec): 0
Match: Input port: 45549
      Ethernet src addr: wildcard
      Ethernet dst addr: wildcard
      Input vlan id: wildcard          Input VLAN priority: wildcard
      Ether type: wildcard
      IP ToS: wildcard                 IP protocol: wildcard
      IP src addr: wildcard            IP dst addr: wildcard
      Source port: wildcard            Destination port: wildcard
    Action: Output port 41350,

```

### show openflow flows detail (Specific Flow, OpenFlow 1.3.1)

```
user@host> show openflow flows 16842752 detail
```

```
Flow name: flow-16842752
Table ID: 1      Flow ID: 16842752
Priority: 32768  Idle timeout(in sec):0      Hard timeout(in sec): 0
Cookie: 0
Match: Input port: 45549
      Ethernet src addr: wildcard
      Ethernet dst addr: wildcard
      Input vlan id: wildcard      Input VLAN priority: wildcard
      Ether type: wildcard
      IP ToS: wildcard      IP protocol: wildcard
      IPv4 src addr: wildcard
      IPv4 dst addr: wildcard
      IPv6 src addr: none
      IPv6 dst addr: none
      Source port: wildcard      Destination port: wildcard
Action: Group 20,
```

#### show openflow flows summary (Specific Flow)

```
user@host> show openflow flows 16842752 summary
Flow name: flow-16842752
Number of packets: 8066495711
```

## show openflow groups

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | show openflow groups<br><brief   details   summary><br><group-id>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display information about OpenFlow groups. Groups are supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.   |
| <b>Options</b>                  | <p><b>none</b>—Display information for all groups.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>group-id</b>—(Optional) Display information about the specified group only.</p>  |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding How the OpenFlow Group Action Works on page 2061</a></li> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow statistics groups on page 2211</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show openflow groups on page 2202</a><br><a href="#">show openflow groups brief on page 2202</a><br><a href="#">show openflow groups detail on page 2202</a><br><a href="#">show openflow groups summary on page 2203</a><br><a href="#">show openflow groups (Specific Group) on page 2203</a><br><a href="#">show openflow groups brief (Specific Group) on page 2203</a><br><a href="#">show openflow groups detail (Specific Group) on page 2203</a><br><a href="#">show openflow groups summary (Specific Group) on page 2203</a> |
| <b>Output Fields</b>            | <p><a href="#">Table 193 on page 2201</a> describes the output fields for the <b>show openflow groups</b> command.</p> <p><a href="#">Table 193 on page 2201</a> lists the output fields in the approximate order in which they are displayed in the sample output.</p>  |

**Table 193: show openflow groups Output Fields**

| Field Name               | Field Description   |
|--------------------------|---|
| Group ID                 | Unique identifier assigned to a group by the OpenFlow controller.           |
| Type                     | Group type, which can be either All or Indirect.                            |
| Number of Buckets        | Number of buckets for a particular group. A group can have 0 to 32 buckets. |
| Number of Flow Reference | Number of flow entries that point to a particular group.                    |

Table 193: show openflow groups Output Fields (*continued*)

| Field Name              | Field Description   |
|-------------------------|---|
| <b>Bucket</b>           | Information about each bucket for a particular group.                                 |
| <b>Actions</b>          | Set of action(s) applied to a packet when it matches the flow entry.                  |
| <b>Switch Name</b>      | User-configured identifier for the OpenFlow virtual switch on which the flow resides. |
| <b>Number of Groups</b> | Number of groups that currently exist in the OpenFlow virtual switch.                 |
| <b>Flow</b>             | Identifier associated with a particular flow entry.                                   |

## Sample Output

### show openflow groups

```
user@host> show openflow groups
```

| Group ID | Type     | Number of Buckets | Number of Flow Reference |
|----------|----------|-------------------|--------------------------|
| 50       | All      | 2                 | 1                        |
| 51       | All      | 2                 | 1                        |
| 60       | Indirect | 1                 | 0                        |

### show openflow groups brief

```
user@host> show openflow groups brief
```

| Group ID | Type     | Number of Buckets | Number of Flow Reference |
|----------|----------|-------------------|--------------------------|
| 50       | All      | 2                 | 1                        |
| 51       | All      | 2                 | 1                        |
| 60       | Indirect | 1                 | 0                        |

### show openflow groups detail

```
user@host> show openflow groups detail
```

```

Group Id: 50                                Type: All
Bucket Bucket 1
  Actions: VLAN ID 2022, Output port 2,
Bucket Bucket 2
  Actions: VLAN ID 3022, Output port 4,

Group Id: 51                                Type: All
Bucket Bucket 3
  Actions: VLAN ID 2001, Output port 1,
Bucket Bucket 4
  Actions: VLAN ID 3001, Output port 3,

Group Id: 60                                Type: Indirect
Bucket Bucket 5
  Actions: VLAN ID 2060, Output port 3,
```

**show openflow groups summary**

```
user@host> show openflow groups summary
```

|             |                  |
|-------------|------------------|
| Switch Name | Number of Groups |
| OF-ex92k    | 3                |

**show openflow groups (Specific Group)**

```
user@host> show openflow groups 50
```

|          |      |                   |                          |
|----------|------|-------------------|--------------------------|
| Group ID | Type | Number of Buckets | Number of Flow Reference |
| 50       | All  | 2                 | 1                        |

**show openflow groups brief (Specific Group)**

```
user@host> show openflow groups 50 brief
```

|          |      |                   |                          |
|----------|------|-------------------|--------------------------|
| Group ID | Type | Number of Buckets | Number of Flow Reference |
| 50       | All  | 2                 | 1                        |

**show openflow groups detail (Specific Group)**

```
user@host> show openflow groups 50 detail
```

```
Group Id: 50                                Type: All
Bucket 1
  Actions: VLAN ID 2022, Output port 2,
Bucket 2
  Actions: VLAN ID 3022, Output port 4,
Flow 570710622208
```

**show openflow groups summary (Specific Group)**

```
user@host> show openflow groups 50 summary
```

|          |      |                   |                          |
|----------|------|-------------------|--------------------------|
| Group ID | Type | Number of Buckets | Number of Flow Reference |
| 50       | All  | 2                 | 1                        |

## show openflow interfaces

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow interfaces</b><br><b>&lt;interface-name&gt;</b><br><b>&lt;switch switch-name&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.   |
| <b>Description</b>              | Display physical characteristics and status information for interfaces participating in OpenFlow.  |
| <b>Options</b>                  | <b>none</b> —Display information for all interfaces participating in OpenFlow.<br><br><b>interface-name</b> —(Optional) Display information only for the specified interface. Specify the interface name including the logical unit number—for example, ge-1/1/0.0.<br><br><b>switch switch-name</b> —(Optional) Display information only for those interfaces configured under the specified OpenFlow virtual switch. |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow filters on page 2194</a></li> <li>• <a href="#">show openflow flows on page 2197</a></li> <li>• <a href="#">show openflow statistics interfaces on page 2213</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show openflow interfaces on page 2205</a><br><a href="#">show openflow interfaces (Specific Interface) on page 2206</a><br><a href="#">show openflow interfaces switch on page 2206</a>  |
| <b>Output Fields</b>            | Table 194 on page 2204 lists the output fields for the <b>show openflow interfaces</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 194: show openflow interfaces Output Fields**

| Field Name                 | Field Description   |
|----------------------------|---|
| Switch name                | User-configured identifier for the OpenFlow virtual switch to which the interface is bound. |
| Interface Name             | Name of the logical interface.  |
| Interface port number      | Port identifier associated with the OpenFlow interface.                                     |
| Interface Hardware Address | Media access control (MAC) address of the interface.  |

Table 194: show openflow interfaces Output Fields (*continued*)

| Field Name                 | Field Description   |
|----------------------------|---|
| Interface speed            | Speed and duplex mode of the interface.                     |
| Interface Auto-Negotiation | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> . |
| Interface media type       | Media type of the interface. For example, copper or fiber.  |
| Interface state            | Current state of the interface.                             |

## Sample Output

### show openflow interfaces

```

user@host> show openflow interfaces
Switch name: OFswitch1
Interface Name: ge-1/1/2.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:b4
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/3.0
Interface port number: 44383
Interface Hardware Address: 00:00:5e:00:53:b5
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/1.0
Interface port number: 41350
Interface Hardware Address: 00:00:5e:00:53:b7
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/7.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:b6
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up

Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2

```

```
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

#### show openflow interfaces (Specific Interface)

```
user@host> show openflow interfaces ge-1/1/0.0
Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

#### show openflow interfaces switch

```
user@host> show openflow interfaces switch OFswitch1
Switch name: OFswitch1
Interface Name: ge-1/1/2.0
Interface port number: 41507
Interface Hardware Address: 00:00:5e:00:53:b4
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: ge-1/1/3.0
Interface port number: 44383
Interface Hardware Address: 00:00:5e:00:53:b5
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: ge-1/1/1.0
Interface port number: 41350
Interface Hardware Address: 00:00:5e:00:53:b7
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: ge-1/1/7.0
Interface port number: 45549
Interface Hardware Address: 00:00:5e:00:53:b6
Interface speed: 1Gb Full-duplex
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

```
Switch name: OFswitch1
Interface Name: ge-1/1/0.0
Interface port number: 44538
Interface Hardware Address: 00:00:5e:00:53:b2
Interface speed: 1Gb Full-duplex
```



```
Interface Auto-Negotiation: Enabled
Interface media type: Fiber
Interface state: Up
```

## show openflow statistics flows

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow statistics flows</b><br><b>&lt;flow-id&gt;</b><br><b>&lt;switch switch-name&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display statistics for OpenFlow flows.  |
| <b>Options</b>                  | <b>none</b> —Display flow statistics for all flows for all OpenFlow virtual switches.<br><br><b>flow-id</b> —(Optional) Display flow statistics only for the specified flow.<br><br><b>switch switch-name</b> —(Optional) Display flow statistics only for the specified OpenFlow virtual switch.   |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow flows on page 2197</a></li> <li>• <a href="#">show openflow statistics interfaces on page 2213</a></li> <li>• <a href="#">show openflow statistics packet on page 2216</a></li> <li>• <a href="#">show openflow statistics tables on page 2223</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics flows on page 2209</a><br><a href="#">show openflow statistics flows (Specific Flow) on page 2209</a><br><a href="#">show openflow statistics flows switch on page 2209</a>  |
| <b>Output Fields</b>            | Table 195 on page 2208 lists the output fields for the <b>show openflow statistics flows</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 195: show openflow statistics flows Output Fields**

| Field Name        | Field Description   |
|-------------------|---|
| Switch Name       | User-configured identifier for the OpenFlow virtual switch on which the flow resides. |
| Table ID          | Identifier for the flow table from which the flow originated.                         |
| Flow ID           | OpenFlow flow entry identifier.   |
| Duration(in sec)  | Number of seconds the flow has been active.   |
| Duration(in nsec) | Number of nanoseconds the flow has been active beyond the <b>Duration(in sec)</b> .   |

Table 195: show openflow statistics flows Output Fields (*continued*)

| Field Name        | Field Description   |
|-------------------|---|
| Priority          | Flow entry priority. Packets match higher priority entries before matching lower priority entries.                        |
| Idle timeout      | Number of seconds after which the flow entry is removed from the flow table provided there are no matching packets.       |
| Hard timeout      | Number of seconds after which the flow entry is removed from the flow table regardless of the number of matching packets. |
| Number of packets | Number of packets that have matched the flow entry.   |
| Number of bytes   | Number of bytes that have matched the flow entry.   |
| Match             | Fields against which the incoming packet is compared.   |
| Action            | Set of actions applied to a packet when it matches the flow entry.  |

## Sample Output

### show openflow statistics flows

```

user@host> show openflow statistics flows
Switch Name: OFswitch1
Table ID: 1      Flow ID: 16842752
Duration(in sec): 58772      Duration(in nsec): 215702000
Priority: 32768  Idle timeout(in sec):0      Hard timeout(in sec): 0
Number of packets: 8745275026
Number of bytes:  524716501560
Match: IN_PORT,
Action: OUTPUT,

```

### show openflow statistics flows (Specific Flow)

```

user@host> show openflow statistics flows 16842752
Switch Name: OFswitch1
Table ID: 1      Flow ID: 16842752
Duration(in sec): 58803      Duration(in nsec): 4127548296
Priority: 32768  Idle timeout(in sec):0      Hard timeout(in sec): 0
Number of packets: 8749713419
Number of bytes:  524982805140
Match: IN_PORT,
Action: OUTPUT,

```

### show openflow statistics flows switch

```

user@host> show openflow statistics flows switch OFswitch1
Switch Name: OFswitch1
Table ID: 1      Flow ID: 16842752
Duration(in sec): 58829      Duration(in nsec): 4124448296
Priority: 32768  Idle timeout(in sec):0      Hard timeout(in sec): 0
Number of packets: 8752672358
Number of bytes:  525160341480

```

Match: IN\_PORT,  
Action: OUTPUT,

## show openflow statistics groups

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow statistics groups</b><br><b>&lt;group-id&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display statistics for OpenFlow groups. Groups are supported only on Juniper Networks devices running OpenFlow v1.3.1 or later.  |
| <b>Options</b>                  | <b>none</b> —Display statistics for all groups defined on all OpenFlow virtual switches.<br><br><b>group-id</b> —(Optional) Display statistics only for the specified group.   |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding How the OpenFlow Group Action Works on page 2061</a></li> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow groups on page 2201</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics group on page 2212</a><br><a href="#">show openflow statistics group (Specific Group) on page 2212</a>  |
| <b>Output Fields</b>            | <a href="#">Table 196 on page 2211</a> lists the output fields for the <b>show openflow statistics groups</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 196: show openflow statistics groups Output Fields**

| Field                      | Description   |
|----------------------------|---|
| Switch Name                | User-configured identifier for the OpenFlow virtual switch in which the groups resides. |
| Group ID                   | Unique identifier assigned to the group by the OpenFlow controller.                     |
| Ref Count                  | Number of flow entries that reference the group.  |
| Number of packets (group)  | Number of packets handled by the group.   |
| Number of bytes (group)    | Number of bytes handled by the group.   |
| Duration(in sec)           | Number of seconds the group has been active.  |
| Duration(in nsec)          | Number of nanoseconds the group has been active beyond the <b>Duration(in sec)</b>      |
| Bucket <i>number</i>       | Statistics for a particular bucket.   |
| Number of packets (bucket) | Number of packets handled by a bucket in the group.                                     |

Table 196: show openflow statistics groups Output Fields (*continued*)

| Field                           | Description                                       |
|---------------------------------|---|
| <b>Number of bytes</b> (bucket) | Number of bytes handled by a bucket in the group. |

For a group with the group type of all, the values specified in the **Number of packets** (group) and **Number of bytes** (group) fields are usually the same as those specified in the **Number of packets** (bucket) and **Number of bytes** (bucket) fields because all buckets in the group are executed.

## Sample Output

### show openflow statistics group

```

user@host> show openflow statistics groups

Switch Name: OF-ex92k
Group ID: 50                      Ref Count: 1
Number of packets: 62161
Number of bytes: 7956608
Duration(in sec): 22687           Duration(in nsec): 4255381296
  Bucket 0
    Number of packets: 62161
    Number of bytes: 7956608
  Bucket 1
    Number of packets: 62161
    Number of bytes: 7956608

Switch Name: OF-ex92k
Group ID: 51                      Ref Count: 1
Number of packets: 0
Number of bytes: 0
Duration(in sec): 22673           Duration(in nsec): 8549000
  Bucket 0
    Number of packets: 0
    Number of bytes: 0
  Bucket 1
    Number of packets: 0
    Number of bytes: 0
...

```

### show openflow statistics group (Specific Group)

```

user@host> show openflow statistics groups 50

Switch Name: OF-ex92k
Group ID: 50                      Ref Count: 1
Number of packets: 64886
Number of bytes: 8305408
Duration(in sec): 22789           Duration(in nsec): 586200000
  Bucket 0
    Number of packets: 64886
    Number of bytes: 8305408
  Bucket 1
    Number of packets: 64886
    Number of bytes: 8305408

```

## show openflow statistics interfaces

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow statistics interfaces</b><br><b>&lt;switch <i>switch-name</i>&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display statistics for interfaces participating in OpenFlow.  |
| <b>Options</b>                  | <b>none</b> —Display statistics for all interfaces participating in OpenFlow for all configured OpenFlow virtual switches.<br><br><b>switch <i>switch-name</i></b> —(Optional) Display statistics only for those interfaces on the specified OpenFlow virtual switch.   |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow interfaces on page 2204</a></li> <li>• <a href="#">show openflow statistics flows on page 2208</a></li> <li>• <a href="#">show openflow statistics tables on page 2223</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics interfaces on page 2214</a>  |

**Output Fields** Table 197 on page 2213 lists the output fields for the **show openflow statistics interfaces** command. Output fields are listed in the approximate order in which they appear.

**Table 197: show openflow statistics interfaces Output Fields**

| Field Name      | Field Description   |
|-----------------|---|
| Switch Name     | User-configured identifier for the OpenFlow virtual switch to which the interface is bound. |
| Interface Name  | Name of the logical interface.  |
| Port Number     | Port identifier associated with the OpenFlow interface.                                     |
| Num of rx pkts  | Number of packets received on the OpenFlow interface.                                       |
| Num of tx pkts  | Number of packets transmitted on the OpenFlow interface.                                    |
| Num of rx bytes | Number of bytes received on the OpenFlow interface.   |
| Num of tx bytes | Number of bytes transmitted on the OpenFlow interface.                                      |

Table 197: show openflow statistics interfaces Output Fields (*continued*)

| Field Name                      | Field Description                                   |
|---------------------------------|---|
| Num of rx error                 | Number of receive errors.                           |
| Num of tx error                 | Number of transmit errors.                          |
| Number of packets dropped by RX | Number of packets dropped by the ingress interface. |
| Number of packets dropped by TX | Number of packets dropped by the egress interface.  |
| Number of rx frame error        | Number of packets with frame alignment errors.      |
| Number of rx overrun error      | Number of packets with RX overrun.                  |
| Number of CRC error             | Number of CRC errors.                               |
| Number of collisions            | Number of Ethernet collisions.                      |

## Sample Output

### show openflow statistics interfaces

```

user@host> show openflow statistics interfaces
Switch Name: OFswitch1
Interface Name: ge-1/1/2.0      Port Number: 41507
Num of rx pkts: 0                Num of tx pkts: 1372301
Num of rx bytes: 0              Num of tx bytes: 88665532
Num of rx error: 0              Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:      0
Number of rx overrun error:    0
Number of CRC error:           0
Number of collisions:          0

Switch Name: OFswitch1
Interface Name: ge-1/1/3.0      Port Number: 44383
Num of rx pkts: 0                Num of tx pkts: 1372285
Num of rx bytes: 0              Num of tx bytes: 88664476
Num of rx error: 0              Num of tx error:0
Number of packets dropped by RX: 0
Number of packets dropped by TX: 0
Number of rx frame error:      0
Number of rx overrun error:    0
Number of CRC error:           0
Number of collisions:          0

Switch Name: OFswitch1
Interface Name: ge-1/1/1.0      Port Number: 41350

```



|                                    |                               |
|------------------------------------|-------------------------------|
| Num of rx pkts: 0                  | Num of tx pkts: 8776241344    |
| Num of rx bytes: 0                 | Num of tx bytes: 526580807026 |
| Num of rx error: 0                 | Num of tx error:0             |
| Number of packets dropped by RX: 0 |                               |
| Number of packets dropped by TX: 0 |                               |
| Number of rx frame error: 0        |                               |
| Number of rx overrun error: 0      |                               |
| Number of CRC error: 0             |                               |
| Number of collisions: 0            |                               |

Switch Name: OFswitch1

Interface Name: ge-1/1/7.0

Port Number: 45549

Num of rx pkts: 8840952127

Num of tx pkts: 1047701

Num of rx bytes: 530457127620

Num of tx bytes: 69187816

Num of rx error: 0

Num of tx error:0

Number of packets dropped by RX: 0

Number of packets dropped by TX: 0

Number of rx frame error: 0

Number of rx overrun error: 0

Number of CRC error: 0

Number of collisions: 0

Switch Name: OFswitch1

Interface Name: ge-1/1/0.0

Port Number: 44538

Num of rx pkts: 0

Num of tx pkts: 1372031

Num of rx bytes: 0

Num of tx bytes: 88647712

Num of rx error: 0

Num of tx error:0

Number of packets dropped by RX: 0

Number of packets dropped by TX: 0

Number of rx frame error: 0

Number of rx overrun error: 0

Number of CRC error: 0

Number of collisions: 0

## show openflow statistics packet

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow statistics packet (in   out)</b><br><b>&lt;switch switch-name&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.             |
| <b>Description</b>              | Display statistics for packet-in and packet-out (send-packet) actions.   |
| <b>Options</b>                  | <b>none</b> —Display statistics for all OpenFlow virtual switches.<br><br><b>switch switch-name</b> —(Optional) Display statistics only for the specified OpenFlow virtual switch.                             |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics packet in on page 2216</a><br><a href="#">show openflow statistics packet out on page 2217</a><br><a href="#">show openflow statistics packet out switch on page 2217</a> |
| <b>Output Fields</b>            | <a href="#">Table 198 on page 2216</a> lists the output fields for the <b>show openflow statistics packet</b> command. Output fields are listed in the approximate order in which they appear.                 |

**Table 198: show openflow statistics packet Output Fields**

| Field Name   | Field Description  |
|--------------|--|
| Switch Name  | User-configured identifier for the OpenFlow virtual switch.  |
| Rx packets   | Number of packets received by the OpenFlow virtual switch that have been sent to the OpenFlow controller. The switch includes the packet in the data portion of an OFPT_PACKET_IN message. |
| Tx packets   | Number of packets sent by the OpenFlow controller to an egress interface. The controller includes the packet in the data portion of an OFPT_PACKET_OUT message.                            |
| Drop packets | Number of dropped packets.   |

## Sample Output

### show openflow statistics packet in

```
user@host> show openflow statistics packet in
```

Openflow packet-in statistics information:

| Switch Name | Rx packets | Drop packets |
|-------------|------------|--------------|
| OFswitch1   | 1044137    | 0            |

#### show openflow statistics packet out

user@host> show openflow statistics packet out

Openflow packet-out statistics information:

| Switch Name | Tx packets | Drop packets |
|-------------|------------|--------------|
| OFswitch1   | 5260759    | 0            |

#### show openflow statistics packet out switch

user@host> show openflow statistics packet out switch OFswitch1

Openflow packet-out statistics information:

| Switch Name | Tx packets | Drop packets |
|-------------|------------|--------------|
| OFswitch1   | 5260759    | 0            |

## show openflow statistics queue

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow statistics queue</b><br><b>&lt;interface <i>interface-name</i>&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display statistics for hardware queues for interfaces participating in OpenFlow.  |
| <b>Options</b>                  | <b>none</b> —Display queue statistics for all interfaces participating in OpenFlow.<br><br><b>interface <i>interface-name</i></b> —(Optional) Display queue statistics only for the specified interface. Specify the interface name including the logical unit number—for example, ge-1/1/0.0 |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow statistics flows on page 2208</a></li> <li>• <a href="#">show openflow statistics tables on page 2223</a></li> </ul>                            |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics queue on page 2219</a><br><a href="#">show openflow statistics queue interface on page 2219</a>  |
| <b>Output Fields</b>            | Table 199 on page 2218 lists the output fields for the <b>show openflow statistics queue</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 199: show openflow statistics queue Output Fields**

| Field Name  | Field Description   |
|-------------|---|
| Switch Name | User-configured identifier for the OpenFlow virtual switch. |
| Port No     | Port identifier associated with the OpenFlow interface.     |
| Queue Id    | Priority queue identifier.                                  |
| TX bytes    | Number of bytes transmitted through the queue.              |
| TX packets  | Number of packets transmitted through the queue.            |
| Tx errors   | Number of packets dropped by the queue due to overrun.      |

## Sample Output

### show openflow statistics queue

```

user@host> show openflow statistics queue
Openflow queue statistics information:
Switch Name      Port No Queue Id  TX bytes    TX packets  Tx errors
OFswitch1        41507  0      115327076   1372459     0
OFswitch1        41507  1        0          0           0
OFswitch1        41507  2        0          0           0
OFswitch1        41507  3        0          0           0
OFswitch1        41507  4        0          0           0
OFswitch1        41507  5        0          0           0
OFswitch1        41507  6        0          0           0
OFswitch1        41507  7        0          0           0
OFswitch1        44383  0      115325732   1372443     0
OFswitch1        44383  1        0          0           0
OFswitch1        44383  2        0          0           0
OFswitch1        44383  3        0          0           0
OFswitch1        44383  4        0          0           0
OFswitch1        44383  5        0          0           0
OFswitch1        44383  6        0          0           0
OFswitch1        44383  7        0          0           0
OFswitch1        41350  0      752072717540 8953246155  0
OFswitch1        41350  1        0          0           0
OFswitch1        41350  2        0          0           0
OFswitch1        41350  3        0          0           0
OFswitch1        41350  4        0          0           0
OFswitch1        41350  5        0          0           0
OFswitch1        41350  6        0          0           0
OFswitch1        41350  7        0          0           0
OFswitch1        45549  0      88060496    1047859     0
OFswitch1        45549  1        0          0           0
OFswitch1        45549  2        0          0           0
OFswitch1        45549  3        0          0           0
OFswitch1        45549  4        0          0           0
OFswitch1        45549  5        0          0           0
OFswitch1        45549  6        0          0           0
OFswitch1        45549  7        0          0           0
OFswitch1        44538  0      115304396   1372189     0
OFswitch1        44538  1        0          0           0
OFswitch1        44538  2        0          0           0
OFswitch1        44538  3        0          0           0
OFswitch1        44538  4        0          0           0
OFswitch1        44538  5        0          0           0
OFswitch1        44538  6        0          0           0
OFswitch1        44538  7        0          0           0

```

### show openflow statistics queue interface

```

user@host> show openflow statistics queue interface ge-1/1/2.0
Openflow queue statistics information:
Switch Name      Port No Queue Id  TX bytes    TX packets  Tx errors
OFswitch1        41507  0      115327076   1372459     0
OFswitch1        41507  1        0          0           0
OFswitch1        41507  2        0          0           0
OFswitch1        41507  3        0          0           0
OFswitch1        41507  4        0          0           0
OFswitch1        41507  5        0          0           0
OFswitch1        41507  6        0          0           0
OFswitch1        41507  7        0          0           0

```



## show openflow statistics summary

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show openflow statistics summary</b>  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.3.</p> <p>Command introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>  |
| <b>Description</b>              | Display summary statistics for all installed OpenFlow flow entries for all OpenFlow virtual switches.  |
| <b>Options</b>                  | This command has no options.   |
| <b>Required Privilege Level</b> | admin  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow statistics flows on page 2208</a></li> <li>• <a href="#">show openflow statistics tables on page 2223</a></li> <li>• <a href="#">show openflow summary on page 2225</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics summary on page 2222</a>  |
| <b>Output Fields</b>            | Table 200 on page 2221 lists the output fields for the <b>show openflow statistics summary</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 200: show openflow statistics summary Output Fields**

| Field Name                   | Field Description   |
|------------------------------|---|
| Switch Name                  | User-configured identifier for the OpenFlow virtual switch. |
| Port Number                  | Port identifier associated with the OpenFlow interface.     |
| Number of RX packets         | Number of packets received on the OpenFlow interface.       |
| Number of TX packets         | Number of packets transmitted on the OpenFlow interface.    |
| Num of packets dropped by RX | Number of packets dropped by the ingress interface.         |
| Flow ID                      | Flow identifier associated with that flow entry.            |
| Number of packets            | Number of packets that have matched the flow entry.         |
| Duration (in sec)            | Number of seconds the flow has been active.                 |

Table 200: show openflow statistics summary Output Fields (*continued*)

| Field Name   | Field Description   |
|--------------|---|
| Priority     | Flow entry priority. Packets match higher priority entries before matching lower priority entries.                        |
| Idle Timeout | Number of seconds after which the flow entry is removed from the flow table provided there are no matching packets.       |
| Hard Timeout | Number of seconds after which the flow entry is removed from the flow table regardless of the number of matching packets. |

## Sample Output

### show openflow statistics summary

```
user@host> show openflow statistics summary
```

| Switch Name | Port Number | Number of RX packets | Number of TX packets | Num of packets dropped by RX |
|-------------|-------------|----------------------|----------------------|------------------------------|
| OFswitch1   | 41507       | 0                    | 1372609              | 0                            |
| OFswitch1   | 44383       | 0                    | 1372593              | 0                            |
| OFswitch1   | 41350       | 0                    | 9119477900           | 0                            |
| OFswitch1   | 45549       | 9184188377           | 1048009              | 0                            |
| OFswitch1   | 44538       | 0                    | 1372339              | 0                            |

| Switch Name | Flow ID  | Number of packets | Duration (in sec) | Priority | Idle Timeout | Hard Timeout |
|-------------|----------|-------------------|-------------------|----------|--------------|--------------|
| OFswitch1   | 16842752 | 9117212928        | 61278             | 32768    | 0            | 0            |



## show openflow statistics tables

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow statistics tables</b><br><b>&lt;switch switch-name&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series routers.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series routers.  |
| <b>Description</b>              | Display statistics for OpenFlow flow tables.  |
| <b>Options</b>                  | <b>none</b> —Display statistics for flow tables on all OpenFlow virtual switches.<br><br><b>switch switch-name</b> —(Optional) Display statistics only for flow tables on the specified OpenFlow virtual switch.  |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow statistics flows on page 2208</a></li> <li>• <a href="#">show openflow statistics interfaces on page 2213</a></li> <li>• <a href="#">show openflow statistics summary on page 2221</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow statistics tables on page 2224</a><br><a href="#">show openflow statistics tables switch on page 2224</a>   |
| <b>Output Fields</b>            | <a href="#">Table 201 on page 2223</a> lists the output fields for the <b>show openflow statistics tables</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 201: show openflow statistics tables Output Fields**

| Field Name                     | Field Description   |
|--------------------------------|---|
| Table Name                     | String identifier for the OpenFlow flow table.  |
| Table id                       | Numeric identifier for the OpenFlow flow table.   |
| Supported wildcards            | Wildcards supported by the flow table.  |
| Max number of entries          | Maximum number of entries supported in the flow table.  |
| Number of active entries       | Number of active entries in the flow table.   |
| Number of idle timeout entries | Number of entries in the flow table that have been removed because the idle timeout expired and no packets matched those entries. |

Table 201: show openflow statistics tables Output Fields (*continued*)

| Field Name                     | Field Description   |
|--------------------------------|---|
| Number of hard timeout entries | Number of entries in the flow table that have been removed because the hard timeout expired.    |
| Number of flow delete entries  | Number of entries in the flow table that have been removed in response to controller requests.  |
| Number of flow add entries     | Number of entries in the flow table that have been added in response to controller requests.    |
| Number of flow modify entries  | Number of entries in the flow table that have been modified in response to controller requests. |
| Number of total delete entries | Number of entries in the flow table that have been removed for any reason.                      |

## Sample Output

### show openflow statistics tables

```

user@host> show openflow statistics tables
Table name: Default flow table           Table id:1
Supported wildcards: IN_PORT, DL_VLAN, DL_SRC, DL_DST, DL_TYPE, NW_PROTO, TP_SRC,
TP_DST, NW_SRC, NW_DST, DL_VLAN_PCP, NW_TOS,
Max number of entries: 65535             Number of active entries: 1
Number of idle timeout entries: 0
Number of hard timeout entries: 0
Number of flow delete entries: 0
Number of flow add entries: 1
Number of flow modify entries: 0
Number of total delete entries: 0

```

### show openflow statistics tables switch

```

user@host> show openflow statistics tables switch OFswitch1
Table name: Default flow table           Table id:1
Supported wildcards: IN_PORT, DL_VLAN, DL_SRC, DL_DST, DL_TYPE, NW_PROTO, TP_SRC,
TP_DST, NW_SRC, NW_DST, DL_VLAN_PCP, NW_TOS,
Max number of entries: 65535             Number of active entries: 1
Number of idle timeout entries: 0
Number of hard timeout entries: 0
Number of flow delete entries: 0
Number of flow add entries: 1
Number of flow modify entries: 0
Number of total delete entries: 0

```

## show openflow summary

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow summary</b>  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.3.</p> <p>Command introduced in Junos OS Release 13.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.</p>   |
| <b>Description</b>              | Display summary information for OpenFlow including the number of configured virtual switches, controllers, interfaces, and flows.   |
| <b>Options</b>                  | This command has no options.  |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow statistics summary on page 2221</a></li> <li>• <a href="#">show openflow switch on page 2226</a></li> </ul> |
| <b>Output Fields</b>            | <p><a href="#">Table 202 on page 2225</a> lists the output fields for the <b>show openflow summary</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

**Table 202: show openflow summary Output Fields**

| Field Name                    | Field Description                                       |
|-------------------------------|---|
| Number of switches            | Total number of configured OpenFlow virtual switches.   |
| Number of controllers         | Total number of configured OpenFlow controllers.        |
| Number of interfaces          | Number of logical interfaces participating in OpenFlow. |
| Number of active flow entries | Number of active entries in the flow table.             |

## Sample Output

### show openflow summary

```
user@host> show openflow summary
Number of switches:      1
Number of controllers:   1
Number of interfaces:    5
Number of active flow entries: 1
```

## show openflow switch

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show openflow switch</b><br><b>&lt;switch-name&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 13.3.<br>Command introduced in Junos OS Release 13.3 for EX Series switches.<br>Command introduced in Junos OS Release 14.1X53-D10 for QFX Series switches.  |
| <b>Description</b>              | Display OpenFlow message statistics for OpenFlow virtual switches.  |
| <b>Options</b>                  | <b>none</b> —Display information for all OpenFlow virtual switches.<br><br><b>switch switch-name</b> —(Optional) Display information only for the specified OpenFlow virtual switch.  |
| <b>Required Privilege Level</b> | admin   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">OpenFlow Operational Mode Commands on page 2183</a></li> <li>• <a href="#">show openflow statistics tables on page 2223</a></li> <li>• <a href="#">show openflow summary on page 2225</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show openflow switch on page 2227</a><br><a href="#">show openflow switch (Specific OpenFlow Virtual Switch) on page 2227</a>   |
| <b>Output Fields</b>            | Table 203 on page 2226 lists the output fields for the <b>show openflow switch</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 203: show openflow switch Output Fields**

| Field Name        | Field Description   |
|-------------------|---|
| Switch Name       | User-configured identifier for the OpenFlow virtual switch.   |
| Switch ID         | Device identifier for the OpenFlow virtual switch.  |
| Switch DPID       | Data path ID uniquely identifying the OpenFlow instance. This value is a concatenation of the switch ID for the virtual switch and the management port MAC address. |
| Flow mod received | Number of Modify Flow Entry messages (OFPT_FLOW_MOD) received from the controller.  |
| Vendor received   | Number of messages with vendor-specific extensions.   |
| Packets sent      | Number of packets sent to the controller.   |
| Packets received  | Number of packets received from the controller.   |

Table 203: show openflow switch Output Fields (*continued*)

| Field Name          | Field Description   |
|---------------------|---|
| Echo req sent       | Number of Echo Request messages (OFPT_ECHO_REQUEST) sent to the controller.             |
| Echo req received   | Number of Echo Request messages (OFPT_ECHO_REQUEST) received from the controller.       |
| Echo reply sent     | Number of Echo Reply messages (OFPT_ECHO_REPLY) sent to the controller.                 |
| Echo reply received | Number of Echo Reply messages (OFPT_ECHO_REPLY) received from the controller.           |
| Port Status sent    | Number of Port Status messages (OFPT_PORT_STATUS) sent to the controller.               |
| Port mod received   | Number of Port Modification messages (OFPT_PORT_MOD) received from the controller.      |
| Barrier request     | Number of Barrier Request messages (OFPT_BARRIER_REQUEST) received from the controller. |
| Barrier reply       | Number of Barrier Reply messages (OFPT_BARRIER_REPLY) sent to the controller.           |
| Error msg sent      | Number of error messages (OFPT_ERROR) sent to the controller.                           |
| Error msg received  | Number of error messages (OFPT_ERROR) received from the controller.                     |

## Sample Output

### show openflow switch

```

user@host> show openflow switch
Switch Name:      OFswitch1
Switch ID:        0
Flow mod received: 4
Packets sent:     1048258
Echo req sent:    4115
Echo reply sent:  0
Port Status sent: 1
Barrier request:  0
Error msg sent:   1
Switch DPID:      00:00:00:00:5e:00:53:d0
Vendor received:  0
Packets received: 1089664
Echo req received: 0
Echo reply received: 4115
Port mod received: 0
Barrier reply:    0
Error msg received: 0

```

### show openflow switch (Specific OpenFlow Virtual Switch)

```

user@host> show openflow switch OFswitch1
Switch Name:      OFswitch1
Switch ID:        0
Flow mod received: 4
Packets sent:     1048259
Echo req sent:    4116
Echo reply sent:  0
Port Status sent: 1
Switch DPID:      00:00:00:00:5e:00:53:d0
Vendor received:  0
Packets received: 1089675
Echo req received: 0
Echo reply received: 4116
Port mod received: 0

```

|                  |   |                     |   |
|------------------|---|---------------------|---|
| Barrier request: | 0 | Barrier reply:      | 0 |
| Error msg sent:  | 1 | Error msg received: | 0 |

## PART 8

# High Availability

- [Overview on page 2231](#)
- [Configuration on page 2261](#)
- [Administration on page 2341](#)
- [Troubleshooting on page 2385](#)





## CHAPTER 29

# Overview

- [Software Feature Overview on page 2231](#)

### Software Feature Overview

---

- [Understanding Graceful Routing Engine Switchover on page 2231](#)
- [Graceful Routing Engine Switchover System Requirements on page 2237](#)
- [Nonstop Active Routing Concepts on page 2240](#)
- [Nonstop Active Routing System Requirements on page 2243](#)
- [Nonstop Bridging Concepts on page 2254](#)
- [Nonstop Bridging System Requirements on page 2256](#)
- [Graceful Restart Concepts on page 2257](#)
- [Understanding VRRP on page 2258](#)

### Understanding Graceful Routing Engine Switchover

This topic contains the following sections:

- [Graceful Routing Engine Switchover Concepts on page 2231](#)
- [Effects of a Routing Engine Switchover on page 2235](#)

#### Graceful Routing Engine Switchover Concepts

---

The graceful Routing Engine switchover (GRES) feature in Junos OS enables a routing platform with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails. GRES preserves interface and kernel information. Traffic is not interrupted. However, GRES does not preserve the control plane.



**NOTE:** On T Series routers, TX Matrix routers, and TX Matrix Plus routers, the control plane is preserved in case of GRES with nonstop active routing (NSR), and nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES.

Neighboring routers detect that the router has experienced a restart and react to the event in a manner prescribed by individual routing protocol specifications.

To preserve routing during a switchover, GRES must be combined with either:

- Graceful restart protocol extensions
- Nonstop active routing

Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur.

Mastership switches to the backup Routing Engine if:

- The master Routing Engine kernel stops operating.
- The master Routing Engine experiences a hardware failure.
- The administrator initiates a manual switchover,



**NOTE:** To quickly restore or to preserve routing protocol state information during a switchover, GRES must be combined with either graceful restart or nonstop active routing, respectively. For more information about graceful restart, see [“Graceful Restart Concepts” on page 2257](#). For more information about nonstop active routing, see [“Nonstop Active Routing Concepts” on page 2240](#).

If the backup Routing Engine does not receive a keepalive from the master Routing Engine after 2 seconds (4 seconds on M20 routers), it determines that the master Routing Engine has failed and: takes mastership.

The Packet Forwarding Engine:

- Seamlessly disconnects from the old master Routing Engine
- Reconnects to the new master Routing Engine
- Does not reboot
- Does not interrupt traffic

The new master Routing Engine and the Packet Forwarding Engine then become synchronized. If the new master Routing Engine detects that the Packet Forwarding Engine state is not up to date, it resends state update messages.



**NOTE:** If adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the [hold-time](#) for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.



**NOTE:** Successive Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to **Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset**. Do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the FPCs should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.



**NOTE:** In a routing matrix with TX Matrix Plus router with 3D SIBs, for successive Routing Engine switchover, events must be a minimum of 900 seconds (15 minutes) apart after both Routing Engines have come up.

GRES must be performed on one line-card chassis (LCC) (of a TX Matrix router with 3D SIBs) at a time to avoid synchronization issues.

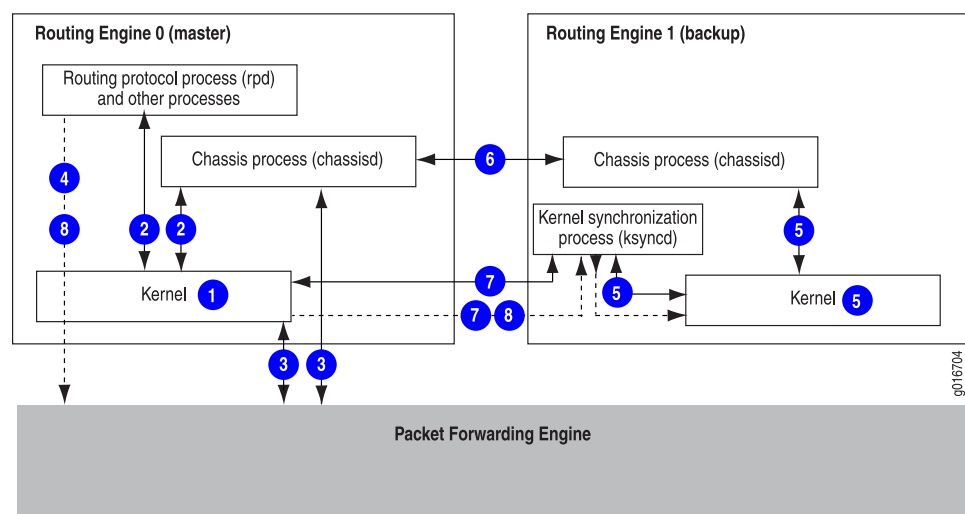


**NOTE:**

- We do *not* recommend performing a commit operation on the backup Routing Engine when GRES is enabled on the router or switch.
- We do *not* recommend enabling GRES on the backup Routing Engine in *any* scenario.

Figure 37 on page 2233 shows the system architecture of graceful Routing Engine switchover and the process a routing platform follows to prepare for a switchover.

**Figure 37: Preparing for a Graceful Routing Engine Switchover**





**NOTE:** Check GRES readiness by executing both:

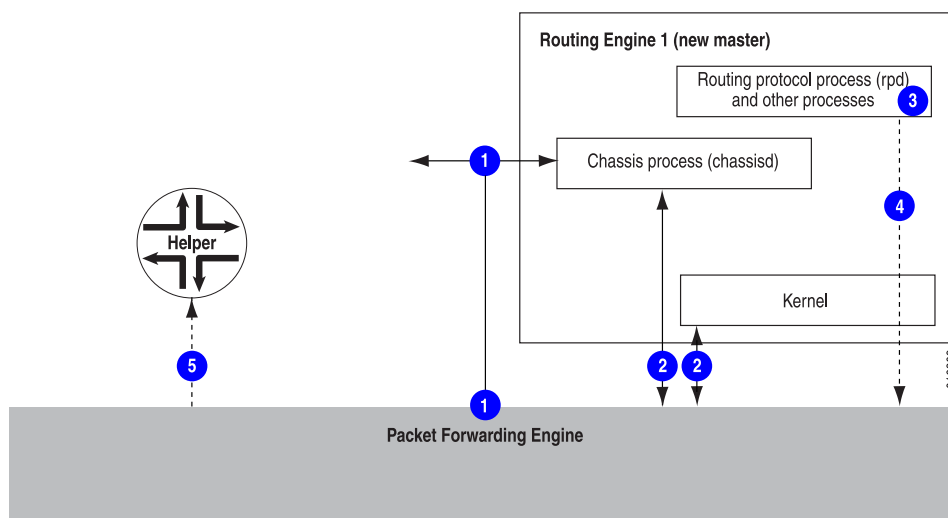
- The **request chassis routing-engine master switch check** command from the master Routing Engine
- The **show system switchover** command from the Backup Routing Engine

The switchover preparation process for GRES is as follows:

1. The master Routing Engine starts.
2. The routing platform processes (such as the chassis process [chassisd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts.
6. The system determines whether GRES has been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. After ksyncd completes the synchronization, all state information and the forwarding table are updated.

Figure 38 on page 2234 shows the effects of a switchover on the routing (or switching) platform.

**Figure 38: Graceful Routing Engine Switchover Process**



When a switchover occurs, the switchover process is as follows:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master.
3. Routing platform processes that are not part of GRES (such as the routing protocol process [rpd]) restart.
4. State information learned from the point of the switchover is updated in the system.
5. If configured, graceful restart protocol extensions collect and restore routing information from neighboring peer *helper* routers.



**NOTE:** On T Series and M320 routers during GRES, the Switch Interface Boards (SIBs) are taken offline and restarted one by one. This is done to provide the Switch Processor Mezzanine Board (SPMB) that manages the SIB enough time to populate state information for its associated SIB. However, on a fully populated chassis where all FPCs are sending traffic at full line rate, there might be momentary packet loss during the switchover.



**NOTE:** When GRES is configured and the `restart chassis-control` command is executed on a TX Matrix Plus router with 3D SIBs, we cannot ascertain which Routing Engine becomes the master. This is because the `chassisd` process restarts with the execution of the `restart chassis-control` command. The `chassisd` process is responsible for maintaining and retaining mastership and when it is restarted, the new `chassisd` is processed based on the router or switch load. As a result, any one of the Routing Engines is made the master.

### Effects of a Routing Engine Switchover

Table 204 on page 2236 describes the effects of a Routing Engine switchover when different features are enabled:

- No high availability features
- Graceful Routing Engine switchover
- Graceful restart
- Nonstop active routing

Table 204: Effects of a Routing Engine Switchover

| Feature   | Benefits   | Considerations   |
|---|--|--|
| Dual Routing Engines only (no features enabled) | <ul style="list-style-type: none"> <li>When the switchover to the new master Routing Engine is complete, routing convergence takes place and traffic is resumed.</li> </ul>  | <ul style="list-style-type: none"> <li>All physical interfaces are taken offline.</li> <li>Packet Forwarding Engines restart.</li> <li>The standby Routing Engine restarts the routing protocol process (rpd).</li> <li>All hardware and interfaces are discovered by the new master Routing Engine.</li> <li>The switchover takes several minutes.</li> <li>All of the router's adjacencies are aware of the physical (interface alarms) and routing (topology) changes.</li> </ul> |
| GRES enabled                                    | <ul style="list-style-type: none"> <li>During the switchover, interface and kernel information is preserved.</li> <li>The switchover is faster because the Packet Forwarding Engines are not restarted.</li> </ul>   | <ul style="list-style-type: none"> <li>The new master Routing Engine restarts the routing protocol process (rpd).</li> <li>All hardware and interfaces are acquired by a process that is similar to a warm restart.</li> <li>All adjacencies are aware of the router's change in state.</li> </ul>   |
| GRES <i>and</i> nonstop active routing enabled  | <ul style="list-style-type: none"> <li>Traffic is not interrupted during the switchover.</li> <li>Interface and kernel information are preserved.</li> </ul>   | <ul style="list-style-type: none"> <li>Unsupported protocols must be refreshed using the normal recovery mechanisms inherent in each protocol.</li> </ul>  |
| GRES <i>and</i> graceful restart enabled        | <ul style="list-style-type: none"> <li>Traffic is not interrupted during the switchover.</li> <li>Interface and kernel information are preserved.</li> <li>Graceful restart protocol extensions quickly collect and restore routing information from the neighboring routers.</li> </ul> | <ul style="list-style-type: none"> <li>Neighbors are required to support graceful restart, and a wait interval is required.</li> <li>The routing protocol process (rpd) restarts.</li> <li>For certain protocols, a significant change in the network can cause graceful restart to stop.</li> <li>If adjacencies between the restarting router and the neighboring peer 'helper' routers time out, graceful restart can stop and cause interruptions in traffic.</li> </ul>         |

**Related Documentation**

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Graceful Routing Engine Switchover System Requirements on page 2237](#)
- [Configuring Graceful Routing Engine Switchover on page 2268](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2270](#)
- [Requirements for Routers with a Backup Router Configuration](#)

- *Example: Configuring IS-IS for GRES with Graceful Restart*
- [hold-time on page 3931](#)

## Graceful Routing Engine Switchover System Requirements

Graceful Routing Engine switchover is supported on all routing (or switching) platforms that contain dual Routing Engines. All Routing Engines configured for graceful Routing Engine switchover must run the same Junos OS release. Hardware and software support for graceful Routing Engine switchover is described in the following sections:

- [Graceful Routing Engine Switchover Platform Support on page 2237](#)
- [Graceful Routing Engine Switchover Feature Support on page 2238](#)
- [Graceful Routing Engine Switchover DPC Support on page 2239](#)
- [Graceful Routing Engine Switchover and Subscriber Access on page 2239](#)
- [Graceful Routing Engine Switchover PIC Support on page 2239](#)

### Graceful Routing Engine Switchover Platform Support

To enable graceful Routing Engine switchover, your system must meet these minimum requirements:

- M20 and M40e routers—Junos OS Release 5.7 or later
- M10i router—Junos OS Release 6.1 or later
- M320 router—Junos OS Release 6.2 or later
- T320 router, T640 router, and TX Matrix router—Junos OS Release 7.0 or later
- M120 router—Junos OS Release 8.2 or later
- MX960 router—Junos OS Release 8.3 or later
- MX480 router—Junos OS Release 8.4 or later (8.4R2 recommended)
- MX240 router—Junos OS Release 9.0 or later
- PTX5000 router—Junos OS Release 12.1X48 or later
- Standalone T1600 router—Junos OS Release 8.5 or later
- Standalone T4000 router—Junos OS Release 12.1R2 or later
- TX Matrix Plus router—Junos OS Release 9.6 or later
- TX Matrix Plus router with 3D SIBs—Junos Release 13.1 or later
- EX Series switches with dual Routing Engines or in a Virtual Chassis — Junos OS Release 9.2 or later for EX Series switches
- QFX Series switches in a Virtual Chassis —Junos OS Release 13.2 or later for the QFX Series
- EX Series or QFX Series switches in a Virtual Chassis Fabric —Junos OS Release 13.2X51-D20 or later for the EX Series and QFX Series switches

For more information about support for graceful Routing Engine switchover, see the sections that follow.

### Graceful Routing Engine Switchover Feature Support

Graceful Routing Engine switchover supports most Junos OS features in Release 5.7 and later. Particular Junos OS features require specific versions of Junos OS. See [Table 205 on page 2238](#).

**Table 205: Graceful Routing Engine Switchover Feature Support**

| Application  | Junos OS Release               |
|--|--------------------------------|
| Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP) and aggregated SONET interfaces   | 6.2                            |
| Asynchronous Transfer Mode (ATM) virtual circuits (VCs)  | 6.2                            |
| Logical systems  | 6.3                            |
| <b>NOTE:</b> In Junos OS Release 9.3 and later, the logical router feature is renamed to logical system.   |                                |
| Multicast  | 6.4 (7.0 for TX Matrix router) |
| Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR)   | 7.0                            |
| Automatic Protection Switching (APS)—The current active interface (either the designated working or the designated protect interface) remains the active interface during a Routing Engine switchover. | 7.4                            |
| Point-to-multipoint Multiprotocol Label Switching MPLS LSPs (transit only)   | 7.4                            |
| Compressed Real-Time Transport Protocol (CRTP)   | 7.6                            |
| Virtual private LAN service (VPLS)   | 8.2                            |
| Ethernet Operation, Administration, and Management (OAM) as defined by IEEE 802.3ah  | 8.5                            |
| Extended DHCP relay agent  | 8.5                            |
| Ethernet OAM as defined by IEEE 802.1ag  | 9.0                            |
| Packet Gateway Control Protocol (PGCP) process (pgcpd) on Multiservices 500 PICs on T640 routers.  | 9.0                            |
| Subscriber access  | 9.4                            |
| Layer 2 Circuit and LDP-based VPLS pseudowire redundant configuration  | 9.6                            |



The following constraints apply to graceful Routing Engine switchover feature support:

- When graceful Routing Engine switchover and aggregated Ethernet interfaces are configured in the same system, the aggregated Ethernet interfaces must not be configured for fast-polling LACP. When fast polling is configured, the LACP polls time out at the remote end during the Routing Engine mastership switchover. When LACP polling times out, the aggregated link and interface are disabled. The Routing Engine mastership change is fast enough that standard and slow LACP polling do not time out during the procedure. However, note that this restriction does not apply to MX Series Routers that are running Junos OS Release 9.4 or later and have distributed periodic packet management (PPM) enabled—which is the default configuration—on them. In such cases, you can configure graceful Routing Engine switchover and have aggregated Ethernet interfaces configured for fast-polling LACP on the same device.
- When a graceful Routing Engine switchover occurs, the VRRP state does not change. VRRP is supported by graceful Routing Engine switchover only in the case that PPM delegation is enabled (which is the default) and that VRRP is not running on logical interfaces like aggregate-interface (ae) or integrated-routing-bridging-interfaces (irb).

### Graceful Routing Engine Switchover DPC Support

Graceful Routing Engine switchover supports all Dense Port Concentrators (DPCs) on the MX Series 3D Universal Edge Routers running the appropriate version of Junos OS as shown in “[Graceful Routing Engine Switchover Platform Support](#)” on page 2237. For more information about DPCs, see the *MX Series DPC Guide*.

### Graceful Routing Engine Switchover and Subscriber Access

Graceful Routing Engine switchover currently supports most of the features directly associated with dynamic DHCP and dynamic PPPoE subscriber access. Graceful Routing Engine switchover also supports the unified in-service software upgrade (ISSU) for the DHCP access model and the PPPoE access model used by subscriber access.

### Graceful Routing Engine Switchover PIC Support

Graceful Routing Engine switchover is supported on most PICs, except for the services PICs listed in this section. The PIC must be on a supported routing platform running the appropriate version of Junos OS. For information about FPC types, FPC/PIC compatibility, and the initial Junos OS Release in which an FPC supported a particular PIC, see the PIC guide for your router platform.

The following constraints apply to graceful Routing Engine switchover support for services PICs:

- You can include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level on a router with Adaptive Services, Multiservices, and Tunnel Services PICs configured on it and successfully commit the configuration. However, all services on these PICs—except the Layer 2 service packages and extension-provider and SDK applications on Multiservices PICs—are reset during a switchover.
- Graceful Routing Engine switchover is not supported on any Monitoring Services PICs or Multilink Services PICs. If you include the **graceful-switchover** statement at the **[edit**

**chassis redundancy]** hierarchy level on a router with either of these PIC types configured on it and issue the **commit** command, the commit fails.

- Graceful Routing Engine switchover is not supported on Multiservices 400 PICs configured for monitoring services applications. If you include the **graceful-switchover** statement, the commit fails.



**NOTE:** When an unsupported PIC is online, you cannot enable graceful Routing Engine switchover. If graceful Routing Engine switchover is already enabled, an unsupported PIC cannot come online.

---

**Related  
Documentation**

- *Understanding High Availability Features on Juniper Networks Routers*
- [Understanding Graceful Routing Engine Switchover on page 2231](#)
- [Configuring Graceful Routing Engine Switchover on page 2268](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2270](#)
- *Requirements for Routers with a Backup Router Configuration*

## Nonstop Active Routing Concepts

Nonstop active routing (NSR) uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop active routing also saves routing protocol information by running the routing protocol process (rpd) on the backup Routing Engine. By saving this additional information, nonstop active routing is self-contained and does not rely on helper routers (or switches) to assist the routing platform in restoring routing protocol information. Nonstop active routing is advantageous in networks where neighbor routers (or switches) do not support graceful restart protocol extensions. As a result of this enhanced functionality, nonstop active routing is a natural replacement for graceful restart.

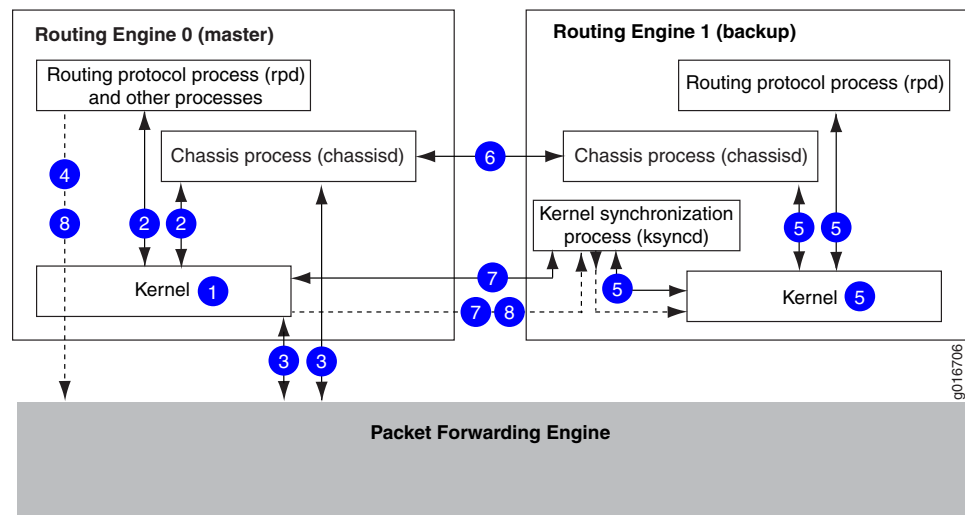


**NOTE:** To use nonstop active routing, you must first enable graceful Routing Engine switchover on your routing (or switching) platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover” on page 2231](#).

---

Figure 39 on page 2241 shows the system architecture of nonstop active routing and the process a routing (or switching) platform follows to prepare for a switchover.

**Figure 39: Nonstop Active Routing Switchover Preparation Process**

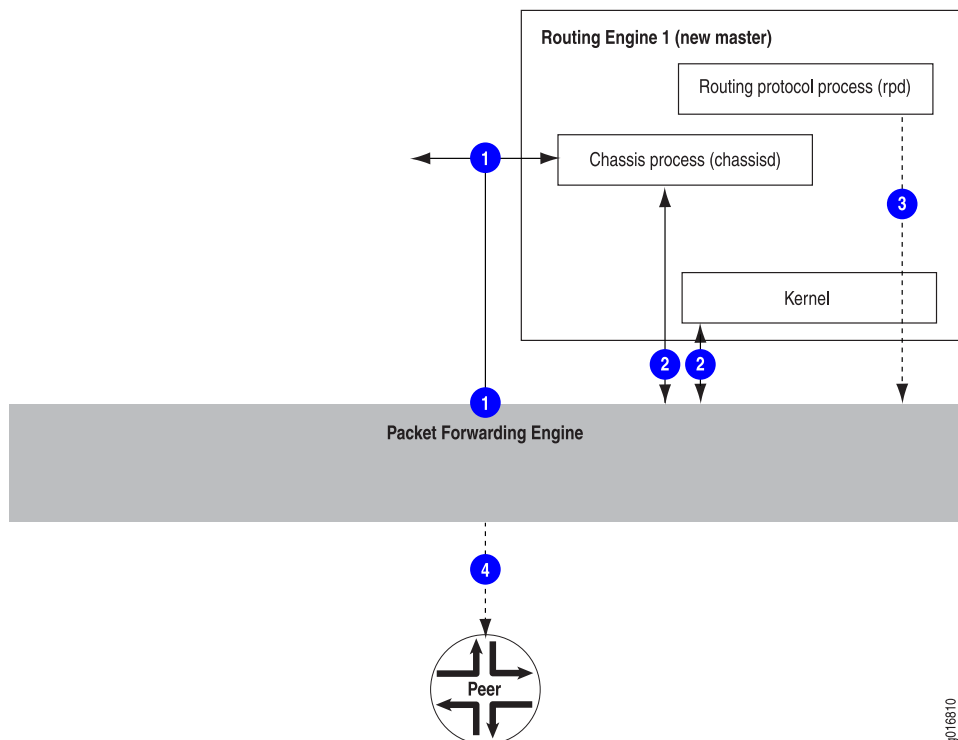


The switchover preparation process for nonstop active routing follows these steps:

1. The master Routing Engine starts.
2. The routing (or switching) platform processes on the master Routing Engine (such as the chassis process [chassisd] and the routing protocol process [rpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the routing protocol process (rpd).
6. The system determines whether graceful Routing Engine switchover and nonstop active routing have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the routing protocol processes on the master and backup Routing Engines.

Figure 40 on page 2242 shows the effects of a switchover on the routing platform.

Figure 40: Nonstop Active Routing During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the routing protocol process (rpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and routing are continued during the switchover, resulting in minimal packet loss.
4. Peer routers (or switches) continue to interact with the routing platform as if no change had occurred. Routing adjacencies and session state relying on underlying routing information are preserved and not reset.



**CAUTION:** We recommend that you do not restart Routing Protocol Process (rpd) on master Routing Engine after enabling nonstop active routing, as it disrupts the protocol adjacency/peering sessions, resulting in traffic loss.

#### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Active Routing System Requirements on page 2243](#)
- [Configuring Nonstop Active Routing](#)

- [Configuring Nonstop Active Routing on Switches on page 2277](#)

## Nonstop Active Routing System Requirements

This section contains the following topics:

- [Nonstop Active Routing Platform and Switching Platform Support on page 2243](#)
- [Nonstop Active Routing Protocol and Feature Support on page 2244](#)
- [Nonstop Active Routing BFD Support on page 2247](#)
- [Nonstop Active Routing BGP Support on page 2248](#)
- [Nonstop Active Routing Layer 2 Circuit and VPLS Support on page 2249](#)
- [Nonstop Active Routing PIM Support on page 2249](#)
- [Nonstop Active Routing MSDP Support on page 2252](#)
- [Nonstop Active Routing Support for RSVP-TE LSPs on page 2252](#)

### Nonstop Active Routing Platform and Switching Platform Support

[Table 206 on page 2243](#) lists the platforms that support nonstop active routing (NSR).

**Table 206: Nonstop Active Routing Platform Support**

| Platform                            | Junos OS Release |
|-------------------------------------|------------------|
| M10i router                         | 8.4 or later     |
| M20 router                          | 8.4 or later     |
| M40e router                         | 8.4 or later     |
| M120 router                         | 9.0 or later     |
| M320 router                         | 8.4 or later     |
| MX Series routers                   | 9.0 or later     |
| PTX Series Packet Transport Routers | 12.1R4 or later  |

#### NOTE:

Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:

- Labeled BGP
- Layer 2 VPNs excluding Layer 2 interworking (Layer 2 switching)
- Layer 3 VPNs
- LDP
- RSVP

Table 206: Nonstop Active Routing Platform Support (*continued*)

| Platform  | Junos OS Release   |
|---|--|
| PTX Series Packet Transport Routers   | 12.1R4 or later  |
| <p><b>NOTE:</b></p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> <li>• Labeled BGP</li> <li>• Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching)</li> <li>• Layer 3 VPNs</li> <li>• LDP</li> <li>• RSVP</li> </ul> |  |
| PTX Series Packet Transport Routers   | 12.1R4 or later  |
| <p><b>NOTE:</b></p> <p>Nonstop active routing (NSR) switchover on PTX series is supported only for the following MPLS and VPN protocols and applications using chained composite next hops:</p> <ul style="list-style-type: none"> <li>• Labeled BGP</li> <li>• Layer 2 VPNs excluding Layer 2 interworking (Layer 2 stitching)</li> <li>• Layer 3 VPNs</li> <li>• LDP</li> <li>• RSVP</li> </ul> |  |
| T320 router, T640 router, and TX Matrix router  | 8.4 or later   |
| Standalone T1600 router   | 8.5 or later   |
| Standalone T4000 router   | 12.1R2 or later  |
| TX Plus Matrix router   | 10.0 or later  |
| TX Plus Matrix router with 3D SIBs  | 13.1 or later  |
| EX Series switch with dual Routing Engines or in a Virtual Chassis  | 10.4 or later for EX Series switches                           |
| EX Series or QFX Series switches in a Virtual Chassis Fabric  | 13.2X51-D20 or later for the EX Series and QFX Series switches |



**NOTE:** All Routing Engines configured for nonstop active routing must be running the same Junos OS release.

### Nonstop Active Routing Protocol and Feature Support

Table 207 on page 2245 lists the protocols that are supported by nonstop active routing.

Table 207: Nonstop Active Routing Protocol and Feature Support

| Protocol  | Junos OS Release   |
|---|--|
| Aggregated Ethernet interfaces with Link Aggregation Control Protocol (LACP)  | 9.4 or later   |
| Bidirectional Forwarding Detection (BFD)<br><br>For more information, see <a href="#">“Nonstop Active Routing BFD Support” on page 2247</a> .   | 8.5 or later   |
| BGP<br><br>For more information, see <a href="#">“Nonstop Active Routing BGP Support” on page 2248</a> .  | 8.4 or later   |
| Labeled BGP (PTX Series Packet Transport Routers only)  | 12.1R4 or later  |
| IS-IS   | 8.4 or later   |
| LDP   | 8.4 or later   |
| LDP-based virtual private LAN service (VPLS)  | 9.3 or later   |
| LDP OAM (operation, administration, and management) features  | 9.6 or later   |
| LDP (PTX Series Packet Transport Routers only)<br><br>Nonstop active routing support for LDP includes: <ul style="list-style-type: none"> <li>• LDP unicast transit LSPs</li> <li>• LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP)</li> <li>• LDP over RSVP transit LSPs</li> <li>• LDP transit LSPs with indexed next hops</li> <li>• LDP transit LSPs with unequal cost load balancing</li> </ul> NOTE: Nonstop active routing is not supported for LDP Point-to-Multipoint LSPs and LDP ingress LSPs. | 12.3R4 or later  |
| Layer 2 circuits  | (on LDP-based VPLS) 9.2 or later<br><br>(on RSVP-TE LSP) 11.1 or later |
| Layer 2 VPNs  | 9.1 or later   |
| Layer 2 VPNs (PTX Series Packet Transport Routers only)<br><br>NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).  | 12.1R4 or later  |
| Layer 3 VPNs (see the first Note after this table for restrictions)   | 9.2 or later   |
| Layer 3 VPNs (PTX Series Packet Transport Routers only)   | 12.1R4 or later  |

Table 207: Nonstop Active Routing Protocol and Feature Support (*continued*)

| Protocol   | Junos OS Release              |
|--|-------------------------------|
| Multicast Source Discovery Protocol (MSDP)   | 12.1 or later                 |
| For more information, see <a href="#">“Nonstop Active Routing MSDP Support” on page 2252</a> .   |                               |
| OSPF/OSPFv3  | 8.4 or later                  |
| Protocol Independent Multicast (PIM)   | (for IPv4) 9.3 or later       |
| For more information, see <a href="#">“Nonstop Active Routing PIM Support” on page 2249</a> .  | (for IPv6) 10.4 or later      |
| RIP and RIP next generation (RIPng)  | 9.0 or later                  |
| RSVP (PTX Series Packet Transport Routers only)  | 12.1R4 or later               |
| Nonstop active routing support for RSVP includes:  |                               |
| <ul style="list-style-type: none"> <li>Point-to-Multipoint LSPs <ul style="list-style-type: none"> <li>RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop.</li> <li>RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes.</li> </ul> </li> <li>Point-to-Point LSPs <ul style="list-style-type: none"> <li>RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops.</li> <li>RSVP Point-to-Point transit LSPs using chained composite next hops.</li> </ul> </li> </ul> |                               |
| RSVP-TE LSP  | 9.5 or later                  |
| For more information, see <a href="#">“Nonstop Active Routing Support for RSVP-TE LSPs” on page 2252</a> .   |                               |
| VPLS   | (LDP-based) 9.1 or later      |
|  | (RSVP-TE-based) 11.2 or later |
| VRRP   | 13.2 or later                 |
| VRRP   | 13.2 or later                 |



**NOTE:** Layer 3 VPN support does not include dynamic GRE tunnels, multicast VPNs, or BGP flow routes.





**NOTE:** If you configure a protocol that is not supported by nonstop active routing, the protocol operates as usual. When a switchover occurs, the state information for the unsupported protocol is not preserved and must be refreshed using the normal recovery mechanisms inherent in the protocol.



**NOTE:** On routers that have logical systems configured on them, only the master logical system supports nonstop active routing.

### Nonstop Active Routing BFD Support

Nonstop active routing supports the Bidirectional Forwarding Detection (BFD) protocol, which uses the topology discovered by routing protocols to monitor neighbors. The BFD protocol is a simple hello mechanism that detects failures in a network. Because BFD is streamlined to be efficient at fast liveness detection, when it is used in conjunction with routing protocols, routing recovery times are improved. With nonstop active routing enabled, BFD session states are not restarted when a Routing Engine switchover occurs.



**NOTE:** BFD session states are saved only for clients using aggregate or static routes or for BGP, IS-IS, OSPF/OSPFv3, or PIM.

When a BFD session is distributed to the Packet Forwarding Engine, BFD packets continue to be sent during a Routing Engine switchover. If nondistributed BFD sessions are to be kept alive during a switchover, you must ensure that the session failure detection time is greater than the Routing Engine switchover time. The following BFD sessions are not distributed to the Packet Forwarding Engine: multihop sessions, tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping. The minimum-interval configuration statement is a BFD liveness detection parameter.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions, and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing is configured, specify a minimum interval of 10 seconds for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

---

### Nonstop Active Routing BGP Support

---

Nonstop active routing BGP support is subject to the following conditions:

- You must include the **path-selection external-router-ID** statement at the **[edit protocols bgp]** hierarchy level to ensure consistent path selection between the master and backup Routing Engines during and after the nonstop active routing switchover.
- You must include the **advertise-from-main-vpn-tables** statement at the **[edit protocols bgp]** hierarchy level to prevent BGP sessions from going down when route reflector (RR) or autonomous system border router (ASBR) functionality is enabled or disabled on a routing device that has VPN address families configured.
- BGP session uptime and downtime statistics are not synchronized between the primary and backup Routing Engines during Nonstop Active Routing and ISSU. The backup Routing Engine maintains its own session uptime based on the time when the backup first becomes aware of the established sessions. For example, if the backup Routing Engine is rebooted (or if you run **restart routing** on the backup Routing Engine), the backup's uptime is a short duration, because the backup has just learned about the established sessions. If the backup is operating when the BGP sessions first come up on the primary, the uptime on the primary and the uptime on the backup are almost the same duration. After a Routing Engine switchover, the new master continues from the time left on the standby Routing Engine.
- If the BGP peer in the master Routing Engine has negotiated address-family capabilities that are not supported for nonstop active routing, then the corresponding BGP neighbor state on the backup Routing Engine shows as idle. On switchover, the BGP session is reestablished from the new master Routing Engine.

Only the following address families are supported for nonstop active routing.



**NOTE:** Address families are supported only on the main instance of BGP. Only unicast is supported on VRF instances.

- inet unicast
- inet labeled-unicast
- inet multicast
- inet6 labeled-unicast
- inet6 multicast
- inet6 unicast
- route-target
- l2vpn signaling
- inet6-vpn unicast
- inet-vpn unicast
- inet-mdt
- iso-vpn
- BGP route dampening does not work on the backup Routing Engine when nonstop active routing is enabled.

### Nonstop Active Routing Layer 2 Circuit and VPLS Support

Nonstop active routing supports Layer 2 circuit and VPLS on both LDP-based and RSVP-TE-based networks. Nonstop active routing support enables the backup Routing Engine to track the label advertised by Layer 2 circuit and VPLS on the primary Routing Engine, and to use the same label after the Routing Engine switchover.

in Junos OS Release 9.6 and later, nonstop active routing support is extended to the Layer 2 circuit and LDP-based VPLS pseudowire redundant configurations.

### Nonstop Active Routing PIM Support

Nonstop active routing supports Protocol Independent Multicast (PIM) with stateful replication on backup Routing Engines. State information replicated on the backup Routing Engine includes information about neighbor relationships, join and prune events, rendezvous point (RP) sets, synchronization between routes and next hops, multicast session states, and the forwarding state between the two Routing Engines.



**NOTE:** Nonstop active routing for PIM is supported for IPv4 on Junos OS Release 9.3 and later, and for IPv6 on Junos OS Release 10.4 and later. Starting with Release 11.1, Junos OS also supports nonstop active routing for PIM on devices that have both IPv4 and IPv6 configured on them.

To configure nonstop active routing for PIM, include the same statements in the configuration as for other protocols: the **nonstop-routing** statement at the **[edit routing-options]** hierarchy level and the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level. To trace PIM nonstop active routing events, include the **flag nsr-synchronization** statement at the **[edit protocols pim traceoptions]** hierarchy level.



**NOTE:** The **clear pim join**, **clear pim register**, and **clear pim statistics** operational mode commands are not supported on the backup Routing Engine when nonstop active routing is enabled.

Nonstop active routing support varies for different PIM features. The features fall into the following three categories: supported features, unsupported features, and incompatible features.

**Supported features:**

- Auto-RP



**NOTE:** Nonstop active routing PIM support on IPv6 does not support auto-RP because IPv6 does not support auto-RP.

- Bootstrap router (BSR)
- Static RPs
- Embedded RP on non-RP IPv6 routers
- Local RP



**NOTE:** RP set information synchronization is supported for local RP and BSR (on IPv4 and IPv6), autoRP (on IPv4), and embedded RP (on IPv6).

- BFD
- Dense mode
- Sparse mode
- Source-specific multicast (SSM)
- Draft Rosen multicast VPNs (MVPNs)
- Anycast RP (anycast RP set information synchronization and anycast RP register state synchronization on IPv4 and IPv6 configurations)
- Flow maps
- Unified ISSU
- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies

- Upstream assert synchronization
- PIM join load balancing

Starting with Release 12.2, Junos OS extends the nonstop active routing PIM support to draft Rosen MVPNs. Nonstop active routing PIM support for draft Rosen MVPNs enables nonstop active routing-enabled devices to preserve draft Rosen MPVN-related information—such as default and data multicast distribution tree (MDT) states—across switchovers. In releases earlier than 12.2, nonstop active routing PIM configuration was incompatible with draft Rosen MPVN configuration.

The backup Routing Engine sets up the default MDT based on the configuration and the information it receives from the master Routing Engine, and keeps updating the default MDT state information.

However, for data MDTs, the backup Routing Engine relies on the master Routing Engine to provide updates when data MDTs are created, updated, or deleted. The backup Routing Engine neither monitors data MDT flow rates nor triggers a data MDT switchover based on variations in flow rates. Similarly, the backup Routing Engine does not maintain the data MDT delay timer or timeout timer. It does not send MDT join TLV packets for the data MDTs until it takes over as the master Routing Engine. After the switchover, the new master Routing Engine starts sending MDT join TLV packets for each data MDT, and also resets the data MDT timers. Note that the expiration time for the timers might vary from the original values on the previous master Routing Engine.

Starting with Release 12.3, Junos OS extends the Protocol Independent Multicast (PIM) nonstop active routing support to IGMP-only interfaces.

In Junos OS releases earlier than 12.3, the PIM joins created on IGMP-only interfaces were not replicated on the backup Routing Engine. Thus, the corresponding multicast routes were marked as pruned (meaning discarded) on the backup Routing Engine. Because of this limitation, after a switchover, the new master Routing Engine had to wait for the IGMP module to come up and start receiving reports to create PIM joins and to install multicast routes. This caused traffic loss until the multicast joins and routes were reinstated.

However, in Junos OS Release 12.3 and later, the multicast joins on the IGMP-only interfaces are mapped to PIM states, and these states are replicated on the backup Routing Engine. If the corresponding PIM states are available on the backup, the multicast routes are marked as forwarding on the backup Routing Engine. This enables uninterrupted traffic flow after a switchover. This enhancement covers IGMPv2, IGMPv3, MLDv1, and MLDv2 reports and leaves.

**Unsupported features:** You can configure the following PIM features on a router along with nonstop active routing, but they function as if nonstop active routing is not enabled. In other words, during Routing Engine switchover and other outages, their state information is not preserved, and traffic loss is to be expected.

- Internet Group Management Protocol (IGMP) exclude mode
- IGMP snooping

Nonstop active routing is not supported for next-generation MVPNs with PIM provider tunnels. The commit operation fails if the configuration includes both nonstop active routing and next-generation MVPNs with PIM provider tunnels.

Junos OS provides a configuration statement that disables nonstop active routing for PIM only, so that you can activate incompatible PIM features and continue to use nonstop active routing for the other protocols on the router. Before activating an incompatible PIM feature, include the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. Note that in this case, nonstop active routing is disabled for all PIM features, not just incompatible features.

### Nonstop Active Routing MSDP Support

---

Starting with Release 12.1, Junos OS extends nonstop active routing support to the Multicast Source Discovery Protocol (MSDP).

Nonstop active routing support for MSDP preserves the following MSDP-related information across the switchover:

- MSDP configuration and peer information
- MSDP peer socket information
- Source-active and related information

However, note that the following restrictions or limitations apply to nonstop active routing MSDP support:

- Because the backup Routing Engine learns the active source information by processing the source-active messages from the network, synchronizing of source active information between the master and backup Routing Engines might take up to 60 seconds. So, no planned switchover is allowed within 60 seconds of the initial replication of the sockets.
- Similarly, Junos OS does not support two planned switchovers within 240 seconds of each other.

Junos OS enables you to trace MSDP nonstop active routing events by including the **flag nsr-synchronization** statement at the **[edit protocols msdp traceoptions]** hierarchy level.

### Nonstop Active Routing Support for RSVP-TE LSPs

---

Junos OS extends nonstop active routing support to label-switching routers (LSRs) and Layer 2 Circuits that are part of an RSVP-TE LSP. Nonstop active routing support on LSRs ensures that the master to backup Routing Engine switchover on an LSR remains transparent to the network neighbors and that the LSP information remains unaltered during and after the switchover.

You can use the **show rsvp version** command to view the nonstop active routing mode and state on an LSR. Similarly, you can use the **show mpls lsp** and **show rsvp session** commands on the standby Routing Engine to view the state recreated on the standby Routing Engine.

The Junos OS nonstop active routing feature is also supported on RSVP point-to-multipoint LSPs. Nonstop active routing support for RSVP point-to-multipoint egress and transit LSPs was added in Junos OS Release 11.4, and for ingress LSPs in Release 12.1. During the switchover, the LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint transit and egress LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover.

However, Junos OS nonstop active routing support for RSVP point-to-multipoint LSPs does not include support for dynamically created point-to-multipoint LSPs, such as VPLS.

Starting with Release 14.1, Junos OS extends nonstop active routing support to the next-generation multicast VPNs.

The **show rsvp session detail** command enables you to check the point-to-multipoint LSP remerge state information (**P2MP LSP re-merge**; possible values are **head**, **member**, and **none**).

However, Junos OS does not support nonstop active routing for the following features:

- Generalized Multiprotocol Label Switching (GMPLS) and LSP hierarchy
- Interdomain or loose-hop expansion LSPs
- BFD liveness detection

Nonstop active routing support for RSVP-TE LSPs is subject to the following limitations and restrictions:

- Detour LSPs are not maintained across a switchover and so, detour LSPs might fail to come back online after the switchover.
- Control plane statistics corresponding to the **show rsvp statistics** and **show rsvp interface detail | extensive** commands are not maintained across Routing Engine switchovers.
- Statistics from the backup Routing Engine are not reported for **show mpls lsp statistics** and **monitor mpls label-switched-path** commands. However, if a switchover occurs, the backup Routing Engine, after taking over as the master, starts reporting statistics. Note that the **clear statistics** command issued on the old master Routing Engine does not have any effect on the new master Routing Engine, which reports statistics, including any uncleared statistics.
- State timeouts might take additional time during nonstop active routing switchover. For example, if a switchover occurs after a neighbor has missed sending two hello messages to the master, the new master Routing Engine waits for another three hello periods before timing out the neighbor.
- On the RSVP ingress router, if you configure auto-bandwidth functionality, the bandwidth adjustment timers are set in the new master after the switchover. This causes a one-time increase in the length of time required for the bandwidth adjustment after the switchover occurs.

- RSVP ingress LSPs that have BFD liveness detection enabled on them do not come up on the backup Routing Engine during the switchover. Such BFD-enabled LSPs have to be reestablished after the switchover.
- Backup LSPs —LSPs that are established between the point of local repair (PLR) and the merge point after a node or link failure—are not preserved during a Routing Engine switchover.
- When nonstop active routing is enabled, graceful restart is not supported. However, graceful restart helper mode is supported.

**Related Documentation**

- [Nonstop Active Routing Concepts on page 2240](#)
- [Configuring Nonstop Active Routing](#)
- [Configuring Nonstop Active Routing on Switches on page 2277](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2274](#)

## Nonstop Bridging Concepts

Nonstop bridging uses the same infrastructure as graceful Routing Engine switchover (GRES) to preserve interface and kernel information. However, nonstop bridging also saves Layer 2 Control Protocol (L2CP) information by running the Layer 2 Control Protocol process (l2cpd) on the backup Routing Engine.




**NOTE:** To use nonstop bridging, you must first enable graceful Routing Engine switchover on your routing (or switching) platform. For more information about graceful Routing Engine switchover, see [“Understanding Graceful Routing Engine Switchover” on page 2231](#).



Figure 41 on page 2255 shows the system architecture of nonstop bridging and the process a routing (or switching) platform follows to prepare for a switchover.

**Figure 41: Nonstop Bridging Switchover Preparation Process**

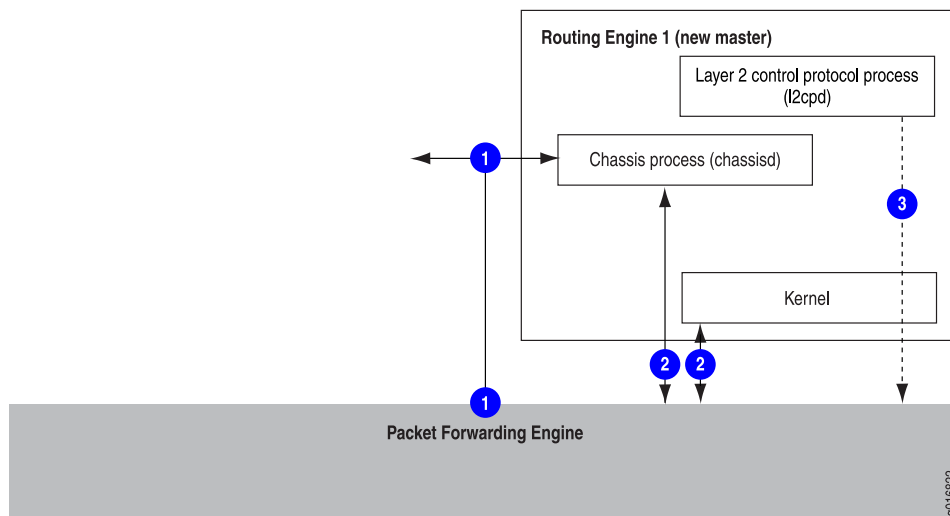


The switchover preparation process for nonstop bridging follows these steps:

1. The master Routing Engine starts.
2. The routing platform processes on the master Routing Engine (such as the chassis process [chassisd] and the Layer 2 Control Protocol process [l2cpd]) start.
3. The Packet Forwarding Engine starts and connects to the master Routing Engine.
4. All state information is updated in the system.
5. The backup Routing Engine starts, including the chassis process (chassisd) and the Layer 2 Control Protocol process (l2cpd).
6. The system determines whether graceful Routing Engine switchover and nonstop bridging have been enabled.
7. The kernel synchronization process (ksyncd) synchronizes the backup Routing Engine with the master Routing Engine.
8. For supported protocols, state information is updated directly between the l2cpds on the master and backup Routing Engines.

Figure 42 on page 2256 shows the effects of a switchover on the routing platform.

Figure 42: Nonstop Bridging During a Switchover



The switchover process follows these steps:

1. When keepalives from the master Routing Engine are lost, the system switches over gracefully to the backup Routing Engine.
2. The Packet Forwarding Engine connects to the backup Routing Engine, which becomes the new master. Because the Layer 2 Control Protocol process (l2cpd) and chassis process (chassisd) are already running, these processes do not need to restart.
3. State information learned from the point of the switchover is updated in the system. Forwarding and bridging are continued during the switchover, resulting in minimal packet loss.

#### Related Documentation

- [Understanding High Availability Features on Juniper Networks Routers](#)
- [Nonstop Bridging System Requirements on page 2256](#)
- [Configuring Nonstop Bridging](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2272](#)

## Nonstop Bridging System Requirements

This topic contains the following sections:

- [Platform Support on page 2256](#)
- [Protocol Support on page 2257](#)

### Platform Support

Nonstop bridging is supported on MX Series 3D Universal Edge Routers. Your system must be running Junos OS Release 8.4 or later.

Nonstop bridging is supported on EX Series switches with redundant Routing Engines in a Virtual Chassis or in a Virtual Chassis Fabric.

Nonstop bridging is supported on QFX Series switches in a Virtual Chassis, EX4600 switches in a Virtual Chassis, or in a Virtual Chassis Fabric. Limited support for NSB is also provided on a QFX5100 and EX4600 standalone switches, but NSB is enabled *only* during an ISSU.



**NOTE:** If you are using a QFX5100 or EX4600 standalone switch and you want to use ISSU, configure Graceful Routing Engine switchover (GRES), NSB and nonstop active routing (NSR). You must configure NSB, GRES, and NSR in order to run ISSU. However, GRES, NSB and NSR are enabled *only* during the upgrade. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.

For a list of the EX Series switches and Layer 2 protocols that support nonstop bridging, see *EX Series Switch Software Features Overview*.



**NOTE:** All Routing Engines configured for nonstop bridging must be running the same Junos OS release.

## Protocol Support

Nonstop bridging is supported for the following Layer 2 control protocols:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)

### Related Documentation

- [Nonstop Bridging Concepts on page 2254](#)
- [Configuring Nonstop Bridging](#)
- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2272](#)

## Graceful Restart Concepts

With routing protocols, any service interruption requires that an affected router recalculate adjacencies with neighboring routers, restore routing table entries, and update other protocol-specific information. An unprotected restart of a router can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Three main types of graceful restart are available on Juniper Networks routing platforms:

- Graceful restart for aggregate and static routes and for routing protocols—Provides protection for aggregate and static routes and for Border Gateway Protocol (BGP), End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), next-generation RIP (RIPng), and Protocol Independent Multicast (PIM) sparse mode routing protocols.
- Graceful restart for MPLS-related protocols—Provides protection for Label Distribution Protocol (LDP), Resource Reservation Protocol (RSVP), circuit cross-connect (CCC), and translational cross-connect (TCC).
- Graceful restart for virtual private networks (VPNs)—Provides protection for Layer 2 and Layer 3 VPNs.

Graceful restart works similarly for routing protocols and MPLS protocols and combines components of these protocol types to enable graceful restart in VPNs. The main benefits of graceful restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful restart thus enables a router to pass through intermediate convergence states that are hidden from the rest of the network.

Most graceful restart implementations define two types of routers—the restarting router and the helper router. The restarting router requires rapid restoration of forwarding state information so it can resume the forwarding of network traffic. The helper router assists the restarting router in this process. Graceful restart configuration statements typically affect either the restarting router or the helper router.

#### **Related Documentation**

- *Understanding High Availability Features on Juniper Networks Routers*
- *Graceful Restart System Requirements*
- *Graceful Restart for Aggregate and Static Routes*
- *Graceful Restart and Routing Protocols*
- *Graceful Restart and MPLS-Related Protocols*
- *Graceful Restart and Layer 2 and Layer 3 VPNs*
- *Graceful Restart on Logical Systems*
- *Configuring Graceful Restart*
- *Configuring Graceful Restart for QFabric Systems*

## **Understanding VRRP**

Juniper Networks QFX Series and EX4600 switches support the Virtual Router Redundancy Protocol (VRRP) and VRRPv3 (for IPv6). This topic covers:

- [Overview of VRRP on page 2259](#)
- [Sample VRRP Topology on page 2259](#)

## Overview of VRRP

Configuring end hosts on your network with static default routes minimizes configuration effort and complexity and reduces processing overhead on the end hosts. When hosts are configured with static routes, the failure of the default gateway normally results in a catastrophic event, isolating all hosts that are unable to detect available alternate paths to their gateway. Using Virtual Router Redundancy Protocol (VRRP) enables you to dynamically provide alternative gateways for end hosts if the primary gateway fails.

VRRP (defined in RFC 3768) provides dynamic failover of IP addresses from one router to another in the event of failure. You can implement VRRP to provide a highly available default path to a gateway without needing to configure dynamic routing or router discovery protocols on end hosts.

Switches configured with VRRP share a virtual IP address, which is the address you configure as the default route on the hosts. At any time, one of the switches is the VRRP master, meaning that it owns the virtual IP address and is the active default gateway. The other devices are backups. The switches dynamically assign master and backup roles based on priorities that you configure (1 **through 255**). If the master fails, the backup switch with the highest priority becomes the master within a few seconds. This is done without any interaction with the hosts.

In VRRP operation, the master sends advertisements to the backup switches at regular intervals. The default interval is 1 second. If the backup switches do not receive an advertisement for a set period, the backup with the highest priority takes over as master within a few seconds and begins forwarding packets. This is done without any interaction with the hosts.

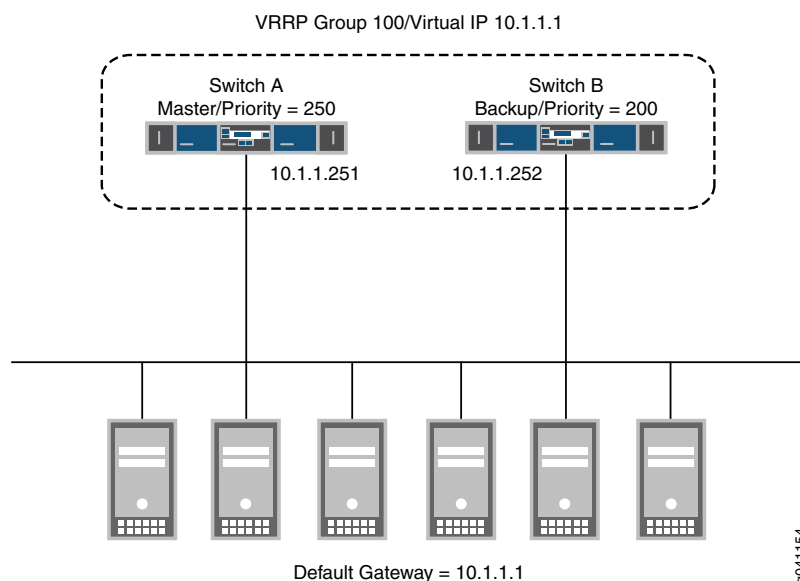


**NOTE:** Priority 255 cannot be set for routed VLAN interfaces (RVIs).

You can configure two QFabric systems to participate in a VRRP configuration as if they were two standalone switches. One benefit of this configuration is if you use VMware's vMotion, virtual machines can transition between hosts connected to the QFabric systems without updating their default gateway information. For example, a virtual machine running on a host connected to a QFabric system in data center A can transition to a host connected to a QFabric system in data center B without needing to resolve a new gateway IP address and MAC address.

## Sample VRRP Topology

Figure 43 on page 2260 illustrates a basic VRRP topology. In this example, switches A and B are running VRRP and share the virtual IP address 10.1.1.1. The default gateway for each of the clients is 10.1.1.1.

**Figure 43: Basic VRRP Topology**

The following illustrates basic VRRP behavior using [Figure 43 on page 2260](#) for reference:

1. When any of the servers wants to send traffic out of the LAN, it sends the traffic to the default gateway address of 10.1.1.1. This is a virtual IP address (VIP) owned by VRRP group 100. Because switch A is the master of the group, the VIP is associated with the “real” address 10.1.1.251 on switch A, and traffic from the servers is actually sent to this address. (Switch A is the master because it has been configured with a higher priority value.)
2. If there is a failure on switch A that prevents it from forwarding traffic to or from the servers—for example, if the interface connected to the LAN fails—switch B becomes the master and assumes ownership of the VIP. The servers continue to send traffic to the VIP, but because the VIP is now associated with the “real” address 10.1.1.252 on switch B (because of change of master), the traffic is sent to switch B instead of switch A.
3. If the problem that caused the failure on switch A is corrected, switch A becomes the master again and reasserts ownership of the VIP. In this case, the servers resume sending traffic to switch A.

Notice that no configuration changes are required on the servers for them to switch between sending traffic to switch A and switch B. When the VIP moves between 10.1.1.251 and 10.1.1.252, the change is detected by normal TCP-IP behavior and no configuration or intervention is required on the servers.

#### Related Documentation

- [Configuring Basic VRRP Support on page 2285](#)
- [Example: Configuring VRRP for Load Sharing on page 2279](#)
- [Understanding VRRP Between QFabric Systems](#)

## CHAPTER 30

# Configuration

- [Configuration Tasks for Graceful Restart on page 2261](#)
- [Configuration Tasks for Graceful Switchover on page 2268](#)
- [Configuration Tasks for Nonstop Bridging on page 2271](#)
- [Configuration Example for Nonstop Active Routing on page 2274](#)
- [Configuration Tasks for Nonstop Active Routing on page 2277](#)
- [Configuration Example for VRRP on page 2279](#)
- [Configuration Tasks for VRRP on page 2284](#)
- [Configuration Statements for Graceful Restart on page 2294](#)
- [Configuration Statement for Graceful Switchover on page 2308](#)
- [Configuration Statement for Nonstop Bridging on page 2310](#)
- [Configuration Statements for Nonstop Routing on page 2311](#)
- [Configuration Statements for VRRP on page 2317](#)

### Configuration Tasks for Graceful Restart

---

- [Configuring Routing Protocols Graceful Restart on page 2261](#)

### Configuring Routing Protocols Graceful Restart

This topic includes the following sections:

- [Enabling Graceful Restart on page 2262](#)
- [Configuring Graceful Restart Options for BGP on page 2262](#)
- [Configuring Graceful Restart Options for ES-IS on page 2263](#)
- [Configuring Graceful Restart Options for IS-IS on page 2263](#)
- [Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2264](#)
- [Configuring Graceful Restart Options for RIP and RIPng on page 2266](#)
- [Configuring Graceful Restart Options for PIM Sparse Mode on page 2266](#)
- [Tracking Graceful Restart Events on page 2267](#)

## Enabling Graceful Restart

---

By default, graceful restart is disabled. To enable graceful restart, include the **graceful-restart** statement at the **[edit routing-instance *instance-name* routing-options]** or **[edit routing-options]** hierarchy level.

For example:

```
routing-options {  
  graceful-restart;  
}
```

To configure the duration of the graceful restart period, include the **restart-duration** at the **[edit routing-options graceful-restart]** hierarchy level.



**NOTE:** Helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode on a per-protocol basis.

```
[edit]  
routing-options {  
  graceful-restart {  
    disable;  
    restart-duration seconds;  
  }  
}
```

To disable graceful restart globally, include the **disable** statement at the **[edit routing-options graceful-restart]** hierarchy level.

When graceful restart is enabled for all routing protocols at the **[edit routing-options graceful-restart]** hierarchy level, you can disable graceful restart on a per-protocol basis.



**NOTE:** If you configure graceful restart after a BGP or LDP session has been established, the BGP or LDP session restarts and the peers negotiate graceful restart capabilities. Also, the BGP peer routing statistics are reset to zero.

## Configuring Graceful Restart Options for BGP

---

To configure the duration of the BGP graceful restart period, include the **restart-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the **stale-routes-time** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.

```
[edit]  
protocols {  
  bgp {  
    graceful-restart {  
      disable;  
    }  
  }  
}
```



```

        restart-time seconds;
        stale-routes-time seconds;
    }
}
routing-options {
    graceful-restart;
}

```

To disable BGP graceful restart capability for all BGP sessions, include the **disable** statement at the **[edit protocols bgp graceful-restart]** hierarchy level.



**NOTE:** To set BGP graceful restart properties or disable them for a group, include the desired statements at the **[edit protocols bgp group group-name graceful-restart]** hierarchy level.

To set BGP graceful restart properties or disable them for a specific neighbor in a group, include the desired statements at the **[edit protocols bgp group group-name neighbor ip-address graceful-restart]** hierarchy level.



**NOTE:** Configuring graceful restart for BGP resets the BGP peer routing statistics to zero. Also, existing BGP sessions restart, and the peers negotiate graceful restart capabilities.

### Configuring Graceful Restart Options for ES-IS

On J Series Services Routers, to configure the duration of the ES-IS graceful restart period, include the **restart-duration** statement at the **[edit protocols esis graceful-restart]** hierarchy level.

```

[edit]
protocols {
    esis {
        graceful-restart {
            disable;
            restart-duration seconds;
        }
    }
}
routing-options {
    graceful-restart;
}

```

To disable ES-IS graceful restart capability, include the **disable** statement at the **[edit protocols esis graceful-restart]** hierarchy level.

### Configuring Graceful Restart Options for IS-IS

To configure the duration of the IS-IS graceful restart period, include the **restart-duration** statement at the **[edit protocols isis graceful-restart]** hierarchy level.

```

[edit]

```

```

protocols {
  isis {
    graceful-restart {
      disable;
      helper-disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}

```

To disable IS-IS graceful restart helper capability, include the **helper-disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level. To disable IS-IS graceful restart capability, include the **disable** statement at the **[edit protocols isis graceful-restart]** hierarchy level.



**NOTE:** If you configure Bidirectional Forwarding Detection (BFD) and graceful restart for IS-IS, graceful restart might not work as expected.



**NOTE:** If adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the hold-time for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds. See *Example: Configuring IS-IS for GRES with Graceful Restart* for more information.



**NOTE:** You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols isis]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 2267](#).

### Configuring Graceful Restart Options for OSPF and OSPFv3

To configure the duration of the OSPF/OSPFv3 graceful restart period, include the **restart-duration** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. To specify the length of time for which the router notifies helper routers that it has completed graceful restart, include the **notify-duration** at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level. Strict OSPF link-state advertisement (LSA) checking results in the termination of graceful restart by a helping router. To disable strict LSA checking, include the **no-strict-lsa-checking** statement at the **[edit protocols (ospf | ospf3) graceful-restart]** hierarchy level.

```

[edit]
protocols {

```

```
ospf | ospfv3{
  graceful-restart {
    disable;
    helper-disable
    no-strict-lsa-checking;
    notify-duration seconds;
    restart-duration seconds;
  }
}
routing-options {
  graceful-restart;
}
```

To disable OSPF/OSPFv3 graceful restart, include the **disable** statement at the **[edit protocols (ospf | ospfv3) graceful-restart]** hierarchy level.

Starting with Release 11.3, the Junos OS supports both the standard (based on RFC 3623, *Graceful OSPF Restart*) and the restart signaling-based (as specified in RFC 4811, RFC 4812, and RFC 4813) helper modes for OSPF version 2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default. To disable the helper mode for OSPF version 2 graceful restart configurations, include the **helper-disable <both | restart-signaling | standard>** statement at the **[edit protocols ospf graceful-restart]** hierarchy level. Note that the last committed statement always takes precedence over the previous one.

```
[edit protocols ospf]
  graceful-restart {
    helper-disable <both | restart-signaling | standard>
  }
```

To reenabling the helper mode, delete the **helper-disable** statement from the configuration by using the **delete protocols ospf graceful-restart helper-disable <restart-signaling | standard | both>** command. In this case also, the last executed command takes precedence over the previous ones.



#### NOTE:

Restart signaling-based helper mode is not supported for OSPFv3 configurations. To disable helper mode for OSPFv3 configurations, include the **helper-disable** statement at the **[edit protocols ospfv3 graceful-restart]** hierarchy level.



**TIP:** You can also track graceful restart events with the **traceoptions** statement at the **[edit protocols (ospf | ospfv3)]** hierarchy level. For more information, see [“Tracking Graceful Restart Events” on page 2267](#).



**NOTE:** You cannot enable OSPFv3 graceful restart between a routing platform running Junos OS Release 7.5 and earlier and a routing platform running Junos OS Release 7.6 or later. As a workaround, make sure both routing platforms use the same Junos OS version.



**NOTE:** If you configure BFD and graceful restart for OSPF, graceful restart might not work as expected.

---

### Configuring Graceful Restart Options for RIP and RIPng

---

To configure the duration of the RIP or RIPng graceful restart period, include the **restart-time** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

```
[edit]
protocols {
  (rip | ripng) {
    graceful-restart {
      disable;
      restart-time seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable RIP or RIPng graceful restart capability, include the **disable** statement at the **[edit protocols (rip | ripng) graceful-restart]** hierarchy level.

---

### Configuring Graceful Restart Options for PIM Sparse Mode

---

PIM sparse mode continues to forward existing multicast packet streams during a graceful restart, but does not forward new streams until after the restart is complete. After a restart, the routing platform updates the forwarding state with any updates that were received from neighbors and occurred during the restart period. For example, the routing platform relearns the join and prune states of neighbors during the restart, but does not apply the changes to the forwarding table until after the restart.

PIM sparse mode-enabled routing platforms generate a unique 32-bit random number called a generation identifier. Generation identifiers are included by default in PIM hello messages, as specified in the IETF Internet draft *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*. When a routing platform receives PIM hellos containing generation identifiers on a point-to-point interface, Junos OS activates an algorithm that optimizes graceful restart.

Before PIM sparse mode graceful restart occurs, each routing platform creates a generation identifier and sends it to its multicast neighbors. If a PIM sparse mode-enabled routing platform restarts, it creates a new generation identifier and sends it to its neighbors. When a neighbor receives the new identifier, it resends multicast updates to the restarting

router to allow it to exit graceful restart efficiently. The restart phase completes when either the PIM state becomes stable or when the restart interval timer expires.

If a routing platform does not support generation identifiers or if PIM is enabled on multipoint interfaces, the PIM sparse mode graceful restart algorithm does not activate, and a default restart timer is used as the restart mechanism.

To configure the duration of the PIM graceful restart period, include the **restart-duration** statement at the **[edit protocols pim graceful-restart]** hierarchy level:

```
[edit]
protocols {
  pim {
    graceful-restart {
      disable;
      restart-duration seconds;
    }
  }
}
routing-options {
  graceful-restart;
}
```

To disable PIM sparse mode graceful restart capability, include the **disable** statement at the **[edit protocols pim graceful-restart]** hierarchy level.



**NOTE:** Multicast forwarding can be interrupted in two ways. First, if the underlying routing protocol is unstable, multicast reverse-path-forwarding (RPF) checks can fail and cause an interruption. Second, because the forwarding table is not updated during the graceful restart period, new multicast streams are not forwarded until graceful restart is complete.

### Tracking Graceful Restart Events

To track the progress of a graceful restart event, you can configure graceful restart trace options flags for IS-IS and OSPF/OSPFv3. To configure graceful restart trace options, include the **graceful-restart** statement at the **[edit protocols protocol traceoptions flag]** hierarchy level:

```
[edit protocols]
isis {
  traceoptions {
    flag graceful-restart;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag graceful-restart;
  }
}
```

- Related Documentation**
- [Graceful Restart Concepts on page 2257](#)
  - [Graceful Restart System Requirements](#)

- *Graceful Restart and Routing Protocols*
- [Verifying Graceful Restart Operation on page 2341](#)
- *Configuring Graceful Restart*
- *Example: Configuring IS-IS for GRES with Graceful Restart*

## Configuration Tasks for Graceful Switchover

---

- [Configuring Graceful Routing Engine Switchover on page 2268](#)
- [Configuring Graceful Routing Engine Switchover in a Virtual Chassis \(CLI Procedure\) on page 2270](#)
- [Resetting Local Statistics on page 2270](#)

### Configuring Graceful Routing Engine Switchover

This section contains the following topics:

- [Enabling Graceful Routing Engine Switchover on page 2268](#)
- [Configuring Graceful Routing Engine Switchover with Graceful Restart on page 2268](#)
- [Synchronizing the Routing Engine Configuration on page 2269](#)
- [Verifying Graceful Routing Engine Switchover Operation on page 2269](#)

#### Enabling Graceful Routing Engine Switchover

---

By default, graceful Routing Engine switchover (GRES) is disabled. To configure GRES, include the **graceful-switchover** statement at the **[edit chassis redundancy]** hierarchy level.

```
[edit chassis redundancy]
graceful-switchover;
```

When you enable GRES, the command-line interface (CLI) indicates which Routing Engine you are using. For example:

```
{master} [edit]
user@host#
```

To disable GRES, delete the **graceful-switchover** statement from the **[edit chassis redundancy]** hierarchy level.

#### Configuring Graceful Routing Engine Switchover with Graceful Restart

---

When using GRES with Graceful Restart, if adjacencies between the Routing Engine and the neighboring peer 'helper' routers time out, graceful restart protocol extensions are unable to notify the peer 'helper' routers about the impending restart. Graceful restart can then stop and cause interruptions in traffic.

To ensure that these adjacencies are kept, change the **hold-time** for IS-IS protocols from the default of 27 seconds to a value higher than 40 seconds.

## Synchronizing the Routing Engine Configuration



**NOTE:** A newly inserted backup Routing Engine automatically synchronizes its configuration with the master Routing Engine configuration.

When you configure GRES, you can bring the backup Routing Engine online after the master Routing Engine is already running. There is no requirement to start the two Routing Engines simultaneously.

## Verifying Graceful Routing Engine Switchover Operation

To verify whether GRES is enabled on the backup Routing Engine, issue the **show system switchover** command. When the output of the command indicates that the **Graceful switchover** field is set to , GRES is operational. The status of the kernel database and configuration database synchronization between Routing Engines is also provided. For example:

```
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady state
```



**NOTE:** You must issue the **show system switchover** command on the backup Routing Engine. This command is not supported on the master Routing Engine.

For more information about the **show system switchover** command, see the [CLI Explorer](#).

### Related Documentation

- [Understanding Graceful Routing Engine Switchover on page 2231](#)
- [Graceful Routing Engine Switchover System Requirements on page 2237](#)
- [Requirements for Routers with a Backup Router Configuration](#)
- [Resetting Local Statistics on page 2270](#)
- [graceful-switchover](#)
- [graceful-switchover on page 2309](#)
- [Example: Configuring IS-IS for GRES with Graceful Restart](#)
- [hold-time on page 3931](#)

## Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)

In a Virtual Chassis, one member switch is assigned the master role and has the master Routing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master to backup without interruption to packet forwarding. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and the forwarding state.

To set up the Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two switches in a Virtual Chassis configuration with mastership priority of 255:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255
[edit]
user@switch# set virtual-chassis member 1 mastership-priority 255
```

2. Set up graceful Routing Engine switchover:

```
[edit]
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.



**NOTE:** We recommend that you use the `commit synchronize` command to save any configuration changes that you make to a multimember Virtual Chassis.

---

### Related Documentation

- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- *High Availability Features for EX Series Switches Overview*
- *Understanding EX Series Virtual Chassis Configuration*
- [Understanding QFX Series Virtual Chassis on page 6907](#)

## Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine



switchover, we recommend that you issue the **clear interface statistics** (*interface-name* | **all**) command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the [CLI Explorer](#).



**NOTE:** The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

**Related  
Documentation**

- [Understanding Graceful Routing Engine Switchover on page 2231](#)
- [Configuring Graceful Routing Engine Switchover on page 2268](#)

## Configuration Tasks for Nonstop Bridging

- [Configuring Nonstop Bridging on Switches \(CLI Procedure\) on page 2272](#)
- [Resetting Local Statistics on page 2273](#)

## Configuring Nonstop Bridging on Switches (CLI Procedure)



**NOTE:** This task uses switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 43.

You can configure nonstop bridging (NSB) to provide resilience for Layer 2 protocol sessions on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Limited support for NSB is also provided on QFX5100 and EX4600 standalone switches, but NSB is enabled *only* during an ISSU.

NSB operates by synchronizing all protocol information for NSB-supported Layer 2 protocols between the master and backup Routing Engines. If the switch has a Routing Engine switchover, the NSB-supported Layer 2 protocol sessions remain active because they are already synchronized on the backup Routing Engine. The Routing Engine switchover is transparent to neighbor devices, which do not detect any changes related to the Layer 2 protocol sessions. The neighboring devices and other devices on the network do not, therefore, have to resynchronize their Layer 2 protocol states to respond to the downtime on the switch—a process that adds network overhead and risks disrupting network performance—when a Routing Engine switchover occurs when NSB is enabled.



**NOTE:** If you are using a QFX5100 or EX4600 standalone switch and you want to use ISSU, configure Graceful Routing Engine switchover (GRES), NSB and nonstop active routing (NSR). You must configure NSB, GRES, and NSR in order to run ISSU. However, GRES, NSB and NSR are enabled *only* during the upgrade. During an ISSU, the Junos OS runs in two separate virtual machines (VMs)—one VM is in the master role acting as the master Routing Engine, and the other VM is in the backup role acting as the backup Routing Engine. The Junos OS is upgraded on the backup VM. After a successful software upgrade, the backup VM then becomes the master VM, and the original master VM is no longer needed and is shut down.

To configure NSB:

1. Enable graceful Routing Engine switchover (GRES):

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```

2. Enable NSB:

```
[edit protocols layer2-control]
user@switch# set nonstop-bridging
```

3. Synchronize configuration changes between the Routing Engines:

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit a configuration that includes NSB without including the **commit synchronize** statement, the commit fails.



**NOTE:** There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you use the **commit synchronize** statement, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes online, its configuration is automatically synchronized with that of the master.



**BEST PRACTICE:** After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics (*interface-name* | all)** command to reset the cumulative values for local statistics on the new master Routing Engine.

#### Related Documentation

- [Performing an In-Service Software Upgrade \(ISSU\) on page 119](#)
- [Understanding Nonstop Bridging on EX Series Switches](#)
- [Nonstop Bridging Concepts on page 2254](#)
- [Understanding In-Service Software Upgrade \(ISSU\) on page 25](#)

## Resetting Local Statistics

When you enable graceful Routing Engine switchover, the master Routing Engine configuration is copied and loaded to the backup Routing Engine. User files, accounting information, and trace options information are not replicated to the backup Routing Engine.

When a graceful Routing Engine switchover occurs, local statistics such as process statistics and networking statistics are displayed as a cumulative value from the time the process first came online. Because processes on the master Routing Engine can start at different times from the processes on the backup Routing Engine, the statistics on the two Routing Engines for the same process might differ. After a graceful Routing Engine switchover, we recommend that you issue the **clear interface statistics (*interface-name* | all)** command to reset the cumulative values for local statistics. Forwarding statistics are not affected by graceful Routing Engine switchover.

For information about how to use the **clear** command to clear statistics and protocol database information, see the [CLI Explorer](#).



**NOTE:** The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover.

- Related Documentation**
- [Understanding Graceful Routing Engine Switchover on page 2231](#)
  - [Configuring Graceful Routing Engine Switchover on page 2268](#)

## Configuration Example for Nonstop Active Routing

---

- [Example: Configuring Nonstop Active Routing on Switches on page 2274](#)

### Example: Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides high availability for Routing Engines by enabling transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

This example describes how to configure nonstop active routing on switches with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

- [Requirements on page 2274](#)
- [Overview and Topology on page 2274](#)
- [Configuration on page 2275](#)
- [Verification on page 2276](#)
- [Troubleshooting on page 2276](#)

#### Requirements

---

This example uses the following hardware and software components:

- An EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration
- Junos OS Release 10.4 or later for EX Series switches
- Junos OS Release 13.2X51-D20 or later for QFX Series switches

#### Overview and Topology

---

Configure nonstop active routing on any EX Series with multiple Routing Engines or on an EX Series or a QFX series switch in a Virtual Chassis or Virtual Chassis Fabric configuration. Nonstop active routing is advantageous in networks where neighbor routing devices do not support graceful restart protocol extensions.

The topology used in this example consists of an EX8200 switch with redundant Routing Engines connected to neighbor routing devices that are not configured to support graceful restart of protocols.

## Configuration

- CLI Quick Configuration** To quickly configure nonstop active routing, copy the following commands and paste them into the switch terminal window:
- ```
[edit]
set chassis redundancy graceful-switchover
set routing-options nonstop-routing
set system commit synchronize
```
- Step-by-Step Procedure** To configure nonstop active routing on a switch:
1. Enable graceful Routing Engine switchover (GRES):
 

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```
  2. Enable nonstop active routing (by default, nonstop active routing is disabled):
 

```
[edit routing-options]
user@switch# set nonstop-routing
```
  3. Synchronize configuration changes between the Routing Engines:
 

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.



**NOTE:** If the backup Routing Engine is down when you issue the commit, a warning is displayed and the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine comes up, its configuration is automatically synchronized with that of the master. If you subsequently insert or bring up a backup Routing Engine, it automatically synchronizes its configuration with the master Routing Engine configuration.

## Results

Check the results of the configuration:

```
[edit]
user@switch# show
chassis {
  redundancy {
    graceful-switchover;
  }
}
routing-options {
  nonstop-routing;
}
system {
  commit synchronize;
}
```

## Verification

---

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Nonstop Active Routing Is Working Correctly on the Switch on page 2276](#)

### *Verifying That Nonstop Active Routing Is Working Correctly on the Switch*

**Purpose** Verify that nonstop active routing is enabled.

**Action** Issue the [show task replication](#) command:

```
user@switch# show task replication
Stateful Replication: Enabled
RE mode: Master
```

Protocol	Synchronization Status
OSPF	Complete
RIP	Complete
PIM	Complete
RSVP	Complete

**Meaning** This output shows that nonstop active routing (Stateful Replication) is enabled on master routing engine. If nonstop routing is not enabled, instead of the output shown above:

- On the backup routing engine the following error message is displayed: **“error: the routing subsystem is not running.”**
- On the master routing engine, the following output is displayed if nonstop routing is not enabled:

```
Stateful Replication: Disabled
RE mode: Master
```

## Troubleshooting

---

To troubleshoot nonstop active routing, perform these tasks:

- [Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled on page 2276](#)

### *Investigating Problems with Synchronization of Routing Engines When NSR Is Enabled*

**Problem** A protocol loses connectivity with neighbors after a graceful Routing Engine switchover (GRES) occurs with nonstop active routing (NSR) enabled.

**Solution** Use trace options to help isolate the problem and gather troubleshooting information. Using the information gathered from trace options, you can confirm or eliminate the synchronization of the Routing Engines as the cause of the loss of connectivity for the protocol. See [“Tracing Nonstop Active Routing Synchronization Events” on page 2278](#).

**Related Documentation**

- [Configuring Nonstop Active Routing on Switches on page 2277](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2278](#)

- [Understanding Nonstop Active Routing on EX Series Switches](#)
- [Nonstop Active Routing Concepts on page 2240](#)

## Configuration Tasks for Nonstop Active Routing

- [Configuring Nonstop Active Routing on Switches on page 2277](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2278](#)

### Configuring Nonstop Active Routing on Switches

Nonstop active routing (NSR) provides a mechanism for transparent switchover of the Routing Engines without necessitating restart of supported routing protocols. Both Routing Engines are fully active in processing protocol sessions, and so each can take over for the other. The switchover is transparent to neighbors.

You can configure NSR on an on a Juniper Networks EX Series switch with multiple Routing Engines or an EX Series or QFX Series switch in a Virtual Chassis or Virtual Chassis Fabric configuration.

To configure nonstop active routing:

1. Enable graceful Routing Engine switchover (GRES):  

```
[edit chassis redundancy]
user@switch# set graceful-switchover
```
2. Enable nonstop active routing (by default, nonstop active routing is disabled):  

```
[edit routing-options]
user@switch# set nonstop-routing
```
3. Synchronize configuration changes between the Routing Engines:  

```
[edit system]
user@switch# set commit synchronize
```

If you try to commit the nonstop active routing configuration without including the **commit synchronize** statement, the commit fails.



**NOTE:** There is no requirement to start the two Routing Engines simultaneously. If the backup Routing Engine is not up when you issue the **commit synchronize** command, the candidate configuration is committed in the master Routing Engine. When the backup Routing Engine is inserted or comes online, its configuration is automatically synchronized with that of the master.



**BEST PRACTICE:** After a graceful Routing Engine switchover, we recommend that you issue the clear interface statistics (*interface-name* | all) command to reset the cumulative values for local statistics on the new master Routing Engine.

To disable nonstop active routing:

```
[edit routing-options]
user@switch# delete nonstop-routing
```

**Related  
Documentation**

- [Example: Configuring Nonstop Active Routing on Switches on page 2274](#)
- [Tracing Nonstop Active Routing Synchronization Events on page 2278](#)
- [Understanding Nonstop Active Routing on EX Series Switches](#)
- [Nonstop Active Routing Concepts on page 2240](#)

## Tracing Nonstop Active Routing Synchronization Events

To track the progress of nonstop active routing synchronization between Routing Engines, you can configure nonstop active routing trace options flags for each supported protocol and for BFD sessions and record these operations to a log file.

To configure nonstop active routing trace options for supported routing protocols, include the **nsr-synchronization** statement at the **[edit protocols protocol-name traceoptions flag]** hierarchy level and optionally specify one or more of the **detail**, **disable**, **receive**, and **send** options:

```
[edit protocols]
bgp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
isis {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ldp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
mpls {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-synchronization-detail;
  }
}
msdp {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
(ospf | ospf3) {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
```



```

rip {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
ripng {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}
pim {
  traceoptions {
    flag nsr-synchronization <detail> <disable> <receive> <send>;
  }
}

```

To configure nonstop active routing trace options for BFD sessions, include the **nsr-synchronization** and **nsr-packet** statements at the **[edit protocols bfd traceoptions flag]** hierarchy level.

```

[edit protocols]
bfd {
  traceoptions {
    flag nsr-synchronization;
    flag nsr-packet;
  }
}

```

To trace the Layer 2 VPN signaling state replicated from routes advertised by BGP, include the **nsr-synchronization** statement at the **[edit routing-options traceoptions flag]** hierarchy level. This flag also traces the label and logical interface association that VPLS receives from the kernel replication state.

```

[edit routing-options]
traceoptions {
  flag nsr-synchronization;
}

```

#### Related Documentation

- [Configuring Nonstop Active Routing](#)
- [Configuring Nonstop Active Routing on Switches on page 2277](#)
- [Example: Configuring Nonstop Active Routing on Switches on page 2274](#)
- [Configuring Nonstop Active Routing](#)

## Configuration Example for VRRP

- [Example: Configuring VRRP for Load Sharing on page 2279](#)

### Example: Configuring VRRP for Load Sharing

If you do not want to dedicate a switch to be a VRRP backup (and therefore leave it idle unless the master fails), you can create a load-sharing configuration in which each participating switch simultaneously acts as a master and a backup.

One reason to use a load-sharing (active-active) configuration is that you are more likely to actively monitor and maintain both switches and notice if a problem occurs on either of them. If you use a configuration in which one switch is only a backup (an active-backup configuration), you might be less likely to pay attention to the backup switch while it is idle. In the worst case, this could lead to the backup switch developing an undetected problem and not being able to perform adequately when a failover occurs.

- [Requirements on page 2280](#)
- [Overview and Topology on page 2280](#)
- [Configuring VRRP on Both Switches on page 2281](#)
- [Verification on page 2283](#)

## Requirements

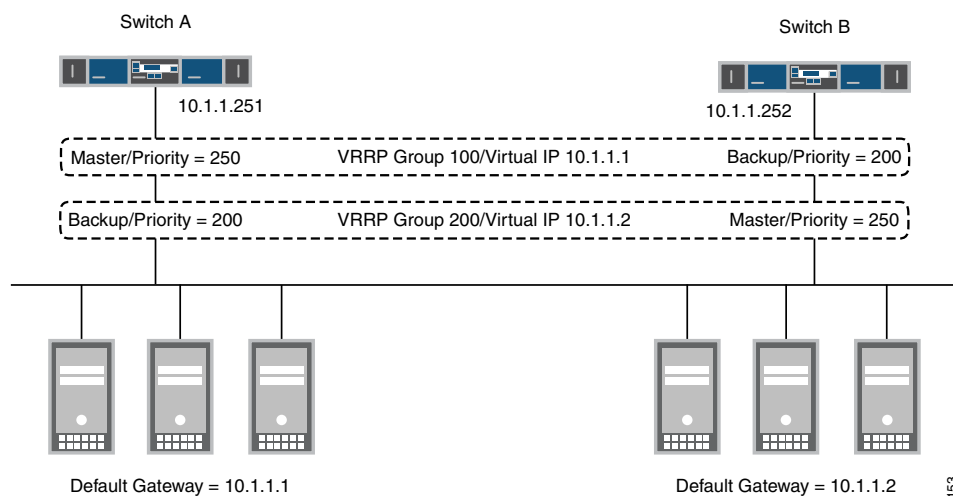
This example uses the following hardware and software components:

- Two QFX3500 switches
- Junos OS Release 11.3 or later
- Static routing or a dynamic routing protocol enabled on both switches.

## Overview and Topology

This example uses two VRRP groups, each of which has its own virtual IP address. Devices on the LAN use one of these virtual IP addresses as their default gateway. If one of the switches fails, the other switch takes over for it. In the topology shown in [Figure 44 on page 2280](#), for example, Switch A is the master for VRRP group 100. If Switch A fails, Switch B takes over and forwards traffic that the end devices send to the default gateway address 10.1.1.1.

**Figure 44: VRRP Load-Sharing Configuration**



9041153

This example shows a simple configuration to illustrate the basic steps for configuring two switches running VRRP to back each other up. [Table 208 on page 2281](#) lists VRRP settings for each switch.

**Table 208: Settings for VRRP Load-Sharing Example**

Switch A	Switch B
VRRP Group 100: <ul style="list-style-type: none"> <li>Interface address: 10.1.1.251</li> <li>VIP: 10.1.1.1</li> <li>Priority: 250</li> </ul>	VRRP Group 100: <ul style="list-style-type: none"> <li>Interface address: 10.1.1.252</li> <li>VIP: 10.1.1.1</li> <li>Priority: 200</li> </ul>
VRRP Group 200: <ul style="list-style-type: none"> <li>Interface address: 10.1.1.251</li> <li>VIP: 10.1.1.2</li> <li>Priority: 200</li> </ul>	VRRP Group 200: <ul style="list-style-type: none"> <li>Interface address: 10.1.1.252</li> <li>VIP: 10.1.1.2</li> <li>Priority: 250</li> </ul>

In addition to configuring the two switches as shown, you must configure your end devices so that some of them use one of the virtual IP addresses as their default gateway and the remaining end devices use the other virtual IP address as their default gateway.

Note that if a failover occurs, the remaining switch might be unable to handle all of the traffic, depending on the demand.

### Configuring VRRP on Both Switches

#### CLI Quick Configuration

Enter the following on Switch A:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 100 priority 250
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group 200 priority 200
```

Enter the following on Switch B:

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 priority 200
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 virtual-address 10.1.1.2
set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 200 priority 250
```

#### Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch A:

1. Create VRRP group 100 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group 100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 250
```

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 200
```

#### Step-by-Step Procedure

Configure the VRRP groups and priorities on Switch B:

1. Create VRRP group 100 on Switch B and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
100 virtual-address 10.1.1.1
```

2. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
100 priority 200
```

Switch A remains the master for group 100 because it has the highest priority for this group.

3. Create VRRP group 200 on Switch A and configure the virtual IP address for the group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.252/24 vrrp-group
200 virtual-address 10.1.1.2
```

4. Assign the VRRP priority for this interface in this group:

```
[edit]
user@switch# set interfaces xe-0/0/0 unit 0 family inet address 10.1.1.251/24 vrrp-group
100 priority 250
```

Switch B becomes the master for group 200 because it has the highest priority for this group.

**Results** Display the results of the configuration on Switch A:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.251 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 250
          }
        }
      }
    }
  }
}
```

```

        vrrp-group 200 {
            virtual address 10.1.1.2
            priority 200
        }
    }
}

```

Display the results of the configuration on Switch B:

```

user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 10.1.1.252 {
          vrrp-group 100 {
            virtual address 10.1.1.1
            priority 200
          }
          vrrp-group 200 {
            virtual address 10.1.1.2
            priority 250
          }
        }
      }
    }
  }
}

```

## Verification

- [Verifying that VRRP is Working on Switch A on page 2283](#)
- [Verifying that VRRP is Working on Switch B on page 2284](#)

### Verifying that VRRP is Working on Switch A

**Purpose** Verify that VRRP is active on Switch A and that the master and backup roles are correct.

**Action** Use the following command to verify that VRRP is active on Switch A and that the switch is master for group 100 and backup for group 200.

```

user@switch> show vrrp

```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	master	A .0327 1c1	10.1.1.251 vip 10.1.1.1
xe-0/0/0.0	up	200	backup	A .0327 1c1	10.1.1.251 vip 10.1.1.2

**Meaning** The `show vrrp` command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The `lcl` address is the physical address of the

interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 200 does not arrive before the timer expires, Switch A asserts itself as the master for this group.

#### ***Verifying that VRRP is Working on Switch B***

**Purpose** Verify that VRRP is active on Switch B and that the master and backup roles are correct.

**Action** Use the following command to verify that VRRP is active on Switch B and that the switch is backup for group 100 and master for group 200.

```
user@switch> show vrrp
```

Interface Address	State	Group	VR state	Timer	Type
xe-0/0/0.0	up	100	backup	A .0327 1c1 10.1.1.252 vip 10.1.1.1	
xe-0/0/0.0	up	200	master	A .0327 1c1 10.1.1.252 vip 10.1.1.2	

**Meaning** The **show vrrp** command displays fundamental information about the VRRP configuration. This output shows that both VRRP groups are active and that this switch has assumed the correct master and backup roles. The **lcl** address is the physical address of the interface and the **vip** address is the virtual address shared by both switches. The **Timer** value (**A .0327**) indicates the remaining time (in seconds) in which this switch expects to receive a VRRP advertisement from the other switch. If an advertisement for group 100 does not arrive before the timer expires, Switch B asserts itself as the master for this group.

**Related Documentation**

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)

---

## Configuration Tasks for VRRP

- [Configuring Basic VRRP Support on page 2285](#)
- [Configuring VRRP Authentication \(IPv4 Only\) on page 2286](#)
- [Configuring the Startup Period for VRRP Operations on page 2287](#)
- [Configuring the Advertisement Interval for the VRRP Master on page 2287](#)
- [Configuring VRRP Preemption and Hold Time on page 2288](#)
- [Configuring a Route to Be Tracked on page 2289](#)
- [Configuring a Logical Interface to Be Tracked on page 2290](#)
- [Configuring a Backup to Accept Packets Destined for the Virtual IP Address on page 2292](#)
- [Configuring Passive ARP Learning for VRRP Backups on page 2292](#)
- [Configuring the Silent Period on page 2293](#)
- [Configuring Inheritance for a VRRP Group on page 2293](#)

## Configuring Basic VRRP Support

To configure basic VRRP support, configure VRRP groups on interfaces by including the **vrrp-group** statement:

```
vrrp-group group-id {
  priority number;
  virtual-address [ addresses ];
}
```

An interface can be a member of multiple VRRP groups.

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]**

For each interface, you must configure the following:

- Group identifier—Assign a value from 0 through 255. You must use the same identifier for each switch in the VRRP group.
- Priority—Assign a value from 1 through 255. The switch with the highest priority becomes the VRRP master. Assign different priorities to each switch in the VRRP group. If there are two or more switches with the same priority, the switch with the VRRP interface that has the highest IP address becomes the master.
- Virtual IP address—Normally, you configure only one address per group, but you can configure as many as eight addresses. Do not include a prefix length in a virtual IP address. The following considerations apply to configuring a virtual IP address:
  - You must configure the same address on all the switches in the VRRP group.
  - If you configure a virtual IP address to be the same as a physical interface address, the switch with that interface becomes the master for the group. You must configure the priority to be 255, and you must configure preemption by including the **preempt** statement.
  - If the virtual IP address is not the same as the physical interface address, you must ensure that the address does not appear anywhere else in the switch configuration. For example, verify that you do not use this address for another interface (including an aggregated Ethernet interface) or for a static ARP entry.



**NOTE:** If you enable MAC source address filtering on an interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement at the **[edit interfaces *interface-name*]** hierarchy. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

- Related Documentation**
- [Understanding VRRP on page 2258](#)
  - [Configuring the Startup Period for VRRP Operations on page 2287](#)
  - [Configuring VRRP Authentication \(IPv4 Only\) on page 2286](#)

## Configuring VRRP Authentication (IPv4 Only)

VRRP (IPv4 only) protocol exchanges can be authenticated to guarantee that only trusted switches participate in a VRRP group. By default, VRRP authentication is disabled. You can configure one of the following authentication methods for a group, and each switch in the same group must use the same method:

- Simple authentication—Uses a text password included in the transmitted packet. The receiving switch uses an authentication key (password) to verify the packet.
- Message Digest 5 (MD5) algorithm—Adds an authentication header (AH) to the IP packet that encapsulates the VRRP packet. You create an authentication key that is used to create a hash of the packet, and the hash is stored in the AH. A receiving switch recalculates the hash on the incoming packet and compares the hashes. If they are identical, the packet is valid and is accepted. Otherwise the switch drops the incoming packet.

To enable authentication and specify an authentication method, include the **authentication-type** statement.

**authentication-type** *authentication*;

**authentication** can be **simple** or **md5**. The authentication type must be the same for all the switches in the VRRP group.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

If you include the **authentication-type** statement, you can configure a key (password) on each interface by including the **authentication-key** statement:

**authentication-key** *key*;

**key** (the password) is an ASCII string. For simple authentication, it can be from 1 through 8 characters long. For MD5 authentication, it can be from 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").



**NOTE:** The key must be the same for all switches in the VRRP group.

---



You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

**Related  
Documentation**

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)

## Configuring the Startup Period for VRRP Operations

Configure the startup-silent period interval to avoid alarms caused by the delay or interruption of the incoming VRRP advertisement packets while an interface is coming online. The period starts when the state of a VRRP interface is changed from down to up. During this period, Master Down Events are ignored.

To configure the startup period for VRRP operations, include the **startup-silent-period** statement at the [edit protocols vrrp] hierarchy level:

```
[edit protocols vrrp]
  startup-silent-period seconds;
```

**Related  
Documentation**

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)
- [Example: Configuring VRRP for Load Sharing on page 2279](#)

## Configuring the Advertisement Interval for the VRRP Master

By default, the master switch sends VRRP advertisement packets every second to all members of the VRRP group. These packets indicate that the master switch is still operational. If the master switch fails or becomes unreachable, the backup switch with the highest priority value becomes the new master switch.

You can modify the advertisement interval in seconds or in milliseconds; the interval must be the same for all routing platforms in the VRRP group.

This topic contains the following sections:

- [Modifying the Advertisement Interval in Seconds on page 2287](#)
- [Modifying the Advertisement Interval in Milliseconds on page 2288](#)

### Modifying the Advertisement Interval in Seconds

To modify the time, in seconds, between the sending of VRRP advertisement packets, include the **advertise-interval** statement:

```
advertise-interval seconds;
```

The interval can be from 1 through 255 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

### Modifying the Advertisement Interval in Milliseconds

---

To modify the time, in milliseconds, between the sending of VRRP advertisement packets, include the **fast-interval** statement:

**fast-interval** *milliseconds*;

The interval can be from 100 through 999 milliseconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-id*]



**NOTE:** Junos OS sets the advertisement interval to 0 in VRRP packets. When you configure VRRP with other vendors' equipment, the **fast-interval** statement works correctly only when the other equipment also has the advertisement interval set to 0 in the VRRP packet. Otherwise, Junos OS interprets other routers' settings as advertisement timer errors.

---

#### Related Documentation

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)
- [Example: Configuring VRRP for Load Sharing on page 2279](#)

## Configuring VRRP Preemption and Hold Time

- [Configuring VRRP Preemption on page 2288](#)
- [Configuring the Preemption Hold Time on page 2289](#)
- [Overriding the Hold Time on page 2289](#)

### Configuring VRRP Preemption

---

By default, a higher-priority VRRP backup switch preempts a lower-priority master switch. To explicitly enable this behavior, include the following statement:

**preempt**;

To prohibit a higher-priority VRRP backup switch from preempting a lower-priority master switch, include the following statement on the lower-priority switch:

**no-preempt**;

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

### Configuring the Preemption Hold Time

You can also configure a preemption hold time, which is the number of seconds a higher-priority backup router that has just started up waits before preempting the master router. You might want to configure a hold time so that routing protocols or other Junos OS components converge before preemption occurs.

The hold time is applied only on startup. By default, the hold-time value is 0 seconds, meaning that preemption can occur immediately after the backup router starts up.

To modify the preemption hold-time value, configure the following statement:

**hold-time** *seconds*;

The hold time can be from 0 through 3600 seconds.

You can include this statement at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address vrrp-group *group-id*] preempt

### Overriding the Hold Time

You can use the **asymmetric-hold-time** statement to configure a VRRP master to fail over to the backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.

When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.

You can include this statement at the following hierarchy level:

- [edit protocols vrrp]

#### Related Documentation

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)
- [Example: Configuring VRRP for Load Sharing on page 2279](#)

## Configuring a Route to Be Tracked

A VRRP master can track a route and dynamically trigger a new master router election if the route becomes unreachable. To enable this behavior, you must configure a cost that will be subtracted from the priority of the master if the tracked route becomes unreachable. The new priority must be less than the priority of one of the backups so that the backup becomes the new master.

To configure a route to be tracked, include the following statements:

```
track {  
  priority-hold-time seconds;  
  route prefix/prefix-length routing-instance default priority-cost priority;  
}
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]

The **prefix** and **prefix-length** values specify the route to be tracked. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority of the master changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **priority-cost** option is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).

#### Related Documentation

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)
- [Example: Configuring VRRP for Load Sharing on page 2279](#)
- [Configuring a Logical Interface to Be Tracked on page 2290](#)

## Configuring a Logical Interface to Be Tracked

VRRP can track whether a logical interface is up, down, or not present, and can change the priority of the switch based on the state of the interface, which might trigger a new master election. VRRP can also track the operational speed of a logical interface and update the priority of the switch when the speed crosses a configured threshold. For each VRRP group, you can track as many as 10 logical interfaces.

When interface tracking is enabled on a switch, you cannot assign the switch a priority of 255 to make it the master for the group.

To configure a logical interface to be tracked, include the following statements:

```
track {  
  interface interface-name {  
    bandwidth-threshold bits-per-second priority-cost priority;  
    priority-cost priority;  
  }  
  priority-hold-time seconds;  
}
```

You can include these statements at the following hierarchy level:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* *vrrp-group group-id*]

The interface specified is the interface to be tracked for the VRRP group. The **priority-hold-time** statement is the minimum length of time that must elapse between priority changes. If the priority changes because of a tracking event, the priority hold timer begins. If another tracking event or manual configuration change occurs while the timer is running, the new priority update is postponed until the timer expires. You might configure the **priority-hold-time** statement to prevent problems that could occur if there were multiple VRRP transitions in a short period of time.

The **bandwidth-threshold** statement specifies a threshold for the tracked interface. If the bandwidth of the tracked interface drops below the threshold value, the system subtracts the bandwidth threshold **priority-cost** value from the VRRP priority for the switch. You can create as many as five **bandwidth-threshold** statements for each tracked interface.

The interface **priority-cost** statement is the value to be subtracted from the VRRP priority when the tracked route goes down. The value can be 1 through 254. The sum of the costs for all tracked interfaces and routes must be less than or equal to the configured priority (so that subtracting all the costs results in a priority equal to or greater than 0).



**WARNING:** On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

If you configure tracking for more than one interface, Junos OS subtracts the sum of the priority costs for the tracked interfaces from the VRRP priority if all the tracked interfaces fail. However, if you configure the interface **priority-cost** statement and the bandwidth threshold **priority-cost** statement, they are not added together. The switch uses only one priority cost for a tracked interface, as indicated in [Table 209 on page 2291](#):

**Table 209: Interface State and Priority Cost Usage**

Tracked Interface State	Priority Cost Usage
Down	<b>priority cost</b> <i>priority</i>
Not down; media speed below one or more bandwidth thresholds	Priority-cost of the lowest applicable bandwidth threshold

You must configure an interface priority cost only if you do not configure any bandwidth thresholds. If you do not configure an interface **priority-cost** value and the interface fails, Junos OS subtracts the bandwidth threshold **priority-cost** value of the lowest bandwidth threshold from the priority of the switch.

- Related Documentation**
- [Understanding VRRP on page 2258](#)
  - [Configuring Basic VRRP Support on page 2285](#)
  - [Example: Configuring VRRP for Load Sharing on page 2279](#)
  - [Configuring a Route to Be Tracked on page 2289](#)

## Configuring a Backup to Accept Packets Destined for the Virtual IP Address

By default, a switch configured to be a VRRP backup but acting as the master does not process packets sent to the virtual IP address—that is, packets in which the destination address is the virtual IP address. To configure a backup switch to process packets sent to the virtual IP address while it is acting as the master, include the **accept-data** statement on the backup:

```
accept-data;
```

You can include this statement at the following hierarchy level:

- `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group] group-id`

To explicitly prohibit the backup from accepting packets destined for the virtual IP address while acting as master, include the **no-accept-data** statement:

```
no-accept-data;
```

If you include the **accept-data** statement, configure the connected hosts so that they:

- Process gratuitous ARP requests.
- Do not use packets other than ARP replies to update their ARP cache.

This statement is disabled by default. If you enable it, your configuration does not comply with RFC 3768.

To restrict incoming IP packets to ICMP only, you must configure firewall filters to accept only ICMP packets.

- Related Documentation**
- [Understanding VRRP on page 2258](#)
  - [Configuring Basic VRRP Support on page 2285](#)
  - [Example: Configuring VRRP for Load Sharing on page 2279](#)

## Configuring Passive ARP Learning for VRRP Backups

By default, VRRP backup switches drop ARP requests for the MAC address of the VRRP IP. This means that backups do not learn the ARP mappings (IP address to MAC address mappings) for the hosts sending the requests. If it becomes the master, the configured backup must learn all the entries that were present in the ARP cache of the original master. In environments with many directly attached hosts, the number of ARP entries to learn can be very large. This can cause a significant delay while the backup transitions

to the master state, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup to learn approximately the same contents as the ARP cache in the master, thus preventing the problem of needing to learn many ARP entries quickly. To enable passive ARP learning, include the **passive-learning** statement at the **[edit system arp]** hierarchy level:

```
[edit system arp]
passive-learning;
```

We recommend setting passive learning on both the backup and master VRRP switches. Doing so prevents the need to manually configure a master that fails and becomes a backup. While a switch operates as the master, the passive learning configuration has no impact. The configuration takes effect only when a switch operates as a backup.

- Related Documentation**
- [Understanding VRRP on page 2258](#)
  - [Configuring Basic VRRP Support on page 2285](#)
  - [Example: Configuring VRRP for Load Sharing on page 2279](#)

## Configuring the Silent Period

When the state of a VRRP interface changes from down to up, a silent period begins. During this period, any master down events are ignored. Configure the silent period interval to avoid problems that can be caused if incoming VRRP advertisement packets are delayed or interrupted while an interface starts up.

To configure the silent period, include the **startup-silent-period** statement at the **[edit protocols vrrp]** hierarchy level:

```
[edit protocols vrrp]
startup-silent-period seconds;
```

- Related Documentation**
- [Understanding VRRP on page 2258](#)
  - [Configuring Basic VRRP Support on page 2285](#)
  - [Example: Configuring VRRP for Load Sharing on page 2279](#)

## Configuring Inheritance for a VRRP Group

Junos OS enables you to configure VRRP groups on the various subnets of a VLAN to inherit the state and configuration of one of the groups, which is known as the *active VRRP group*. By configuring inheritance, you can prevent VRRP groups other than the active group from sending out VRRP advertisements. When the **vrrp-inherit-from** configuration statement is included in the configuration, only the active VRRP group from which the other VRRP groups are inheriting the state sends out VRRP advertisements; the groups inheriting the state do not send any VRRP advertisements, because the state is maintained only on the group from which the state is inherited.

If the **vrrp-inherit-from** statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

To configure inheritance for a VRRP group, include the **vrrp-inherit-from** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*]**:

```
[edit interfaces interface-name unit logical-unit-number family inet address address
  vrrp-group group-id]
  vrrp-inherit-from vrrp-group;
```

When you configure a group to inherit a state from another group, note the following conditions:

- Both inheriting groups and active groups must be on the same physical interface and logical system. However, the groups need not necessarily be on the same VLAN or logical interface.
- Both inheriting groups and active groups must be on the same routing instances; however, this limitation does not apply for groups on the integrated routing and bridging (IRB) interfaces.

When you include the **vrrp-inherit-from** statement for a VRRP group, the VRRP group inherits the following parameters from the active group:

- **advertise-interval**
- **authentication-key**
- **authentication-type**
- **fast-interval**
- **preempt | no-preempt**
- **priority**
- **track interfaces**
- **track route**

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

#### Related Documentation

- [Understanding VRRP on page 2258](#)
- [Configuring Basic VRRP Support on page 2285](#)
- [Example: Configuring VRRP for Load Sharing on page 2279](#)

---

## Configuration Statements for Graceful Restart

- [disable on page 2295](#)
- [disable \(BGP Graceful Restart\) on page 2296](#)



- [graceful-restart \(Enabling Globally\) on page 2297](#)
- [graceful-restart \(Protocols BGP\) on page 2299](#)
- [graceful-restart \(Protocols OSPF\) on page 2300](#)
- [helper-disable \(OSPF\) on page 2302](#)
- [no-strict-lsa-checking on page 2303](#)
- [notify-duration on page 2304](#)
- [redundancy \(Graceful Switchover\) on page 2305](#)
- [restart-duration on page 2306](#)
- [restart-time \(BGP Graceful Restart\) on page 2307](#)
- [stale-routes-time on page 2308](#)

## disable

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (bgp   isis   ldp   ospf   ospf3   pim   rip   ripng   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (bgp   ldp   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit protocols (bgp   isis   isis   ospf   ospf3   ldp   pim   rip   ripng   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit protocols bgp group <i>group-name</i> <a href="#">graceful-restart</a>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>ip-address</i> <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (bgp   ldp   ospf   ospf3   pim) <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">graceful-restart</a>],</p> <p>[edit routing-options <a href="#">graceful-restart</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
<b>Description</b>	Disable graceful restart.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling Graceful Restart</a></li> <li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2261</a></li> <li>• <a href="#">Configuring Graceful Restart for MPLS-Related Protocols</a></li> <li>• <a href="#">Configuring VPN Graceful Restart</a></li> <li>• <a href="#">Configuring Logical System Graceful Restart</a></li> <li>• <a href="#">Graceful Restart Configuration Statements</a></li> <li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li> </ul>

## disable (BGP Graceful Restart)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> graceful-restart],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> graceful-restart],</p> <p>[edit protocols bgp graceful-restart],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> graceful-restart],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> graceful-restart]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.



**NOTE:** When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp **group** *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp **group** *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp **group** *group-name* **neighbor** *address*] hierarchy level.

<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li> <li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li> <li>• <a href="#">graceful-restart on page 2299</a></li> </ul>

## graceful-restart (Enabling Globally)

<b>Syntax</b>	<pre> graceful-restart {   disable;   helper-disable;   maximum-helper-recovery-time <i>seconds</i>;   maximum-helper-restart-time <i>seconds</i>;   notify-duration <i>seconds</i>;   recovery-time <i>seconds</i>;   restart-duration <i>seconds</i>;   stale-routes-time <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.




### NOTE:

- For VPNs, the `graceful-restart` statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
- For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
- LDP sessions flap when `graceful-restart` configurations change.

<b>Default</b>	Graceful restart is disabled by default.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling Graceful Restart</a></li> <li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2261</a></li> </ul>

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

## graceful-restart (Protocols BGP)

<b>Syntax</b>	<pre> graceful-restart {   disable;   restart-time seconds;   stale-routes-time seconds; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp],  [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],  [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],  [edit protocols bgp],  [edit protocols bgp <b>group</b> <i>group-name</i>],  [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.  Statement introduced in Junos OS Release 9.0 for EX Series switches.  Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the <b>restart-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the <b>stale-routes-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <hr/> <div>  <p><b>NOTE:</b> If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <hr/> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li> <li>• <i>Configuring Graceful Restart for QFabric Systems</i></li> <li>• <i>Junos OS High Availability Library for Routing Devices</i></li> </ul>

## graceful-restart (Protocols OSPF)

---

<b>Syntax</b>	<pre>graceful-restart {   disable;   helper-disable (standard   restart-signaling   both);   no-strict-lsa-checking;   notify-duration <i>seconds</i>;   restart-duration <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   (<b>ospf</b>   ospf3)], [edit protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the <b>no-strict-lsa-checking</b> statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the <b>[edit routing-options]</b> hierarchy level.</p>
<b>Options</b>	<p><b>disable</b>—Disable graceful restart for OSPF.</p> <p><b>helper-disable (standard   restart-signaling   both)</b>—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options are only supported for OSPFv2. Specify <b>standard</b> to disable helper mode for standard graceful restart (based on RFC 3623). Specify <b>restart-signaling</b> to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify <b>both</b> to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p><b>Default:</b> Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p><b>no-strict-lsa-checking</b>—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

**notify-duration seconds**—Estimated time needed to send out purged grace LSAs over all the interfaces.

**Range:** 1 through 3600 seconds

**Default:** 30 seconds

**restart-duration seconds**—Estimated time needed to reacquire a full OSPF neighbor from each area.

**Range:** 1 through 3600 seconds


**Default:** 180 seconds

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Graceful Restart for OSPF on page 4136</a></li> <li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 4140</a></li> <li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4144</a></li> <li>• <a href="#">Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 4147</a></li> <li>• <i>Junos OS High Availability Library for Routing Devices</i></li> </ul>
------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## helper-disable (OSPF)

---

<b>Syntax</b>	helper-disable < both   restart-signaling   standard >;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ospf <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">graceful-restart</a> ], [edit protocols ospf <a href="#">graceful-restart</a> ], [edit routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Options <b>both</b> , <b>restart-signaling</b> , and <b>standard</b> introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Disable helper mode for graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart. The last committed statement takes precedence over the previously configured statement.
<b>Default</b>	Helper mode is enabled by default for OSPF.
<b>Options</b>	<b>both</b> —(Optional) Disable helper mode for both standard and restart signaling-based graceful restart.  <b>restart-signaling</b> —(Optional) Disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813).  ..... <div> <b>NOTE:</b> Restart signaling-based helper mode is not supported for OSPFv3 configurations.</div> ..... <b>standard</b> —(Optional) Disable helper mode for standard graceful restart (based on RFC 3623).
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2261</a></li><li>• <i>Configuring Graceful Restart for MPLS-Related Protocols</i></li></ul>



---

## no-strict-lsa-checking

---

<b>Syntax</b>	no-strict-lsa-checking;
<b>Hierarchy Level</b>	[edit protocols (ospf   ospf3) <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router or switch.
<b>Default</b>	By default, LSA checking is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2264</a></li><li>• <i>Configuring Graceful Restart for QFabric Systems</i></li><li>• <a href="#">maximum-neighbor-recovery-time on page 4550</a></li><li>• <i>recovery-time</i></li></ul>

## notify-duration

---

<b>Syntax</b>	<code>notify-duration <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols (ospf   ospf3) <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a> ], [edit routing-instances <i>instance-name</i> protocols (ospf   ospf3) <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the length of time the router or switch notifies helper OSPF routers that it has completed graceful restart.
<b>Options</b>	<b><i>seconds</i></b> —Length of time in the router notifies helper OSPF routers that it has completed graceful restart. <b>Range:</b> 1 through 3600 <b>Default:</b> 30
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for OSPF and OSPFv3 on page 2264</a></li><li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li><li>• <a href="#">restart-duration on page 2306</a></li></ul>

## redundancy (Graceful Switchover)

<b>Syntax</b>	<pre> redundancy {   failover {     on-disk-failure;     on-loss-of-keepalives;   }   graceful-switchover; } </pre>
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	Redundancy is enabled for the Routing Engines.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">graceful-switchover on page 2309</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 2270</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2268</a></li> <li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li> <li>• <a href="#">High Availability Features for EX Series Switches Overview</a></li> </ul>

## restart-duration

---

<b>Syntax</b>	<code>restart-duration <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols (isis   ospf   ospf3   pim) graceful-restart],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3   pim) graceful-restart],</code> <code>[edit protocols (esis   isis   ospf   ospf3   pim) graceful-restart],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3   pim) graceful-restart],</code> <code>[edit routing-options graceful-restart]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Configure the grace period for graceful restart globally.</p> <p>Additionally, you can individually configure the duration of the graceful restart period for the End System-to-Intermediate System (ES-IS), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), and OSPFv3 protocols and for Protocol Independent Multicast (PIM) sparse mode.</p>
<b>Options</b>	<p><b><i>seconds</i></b>—Time for the graceful restart period.</p> <p><b>Range:</b></p> <p>The range of values varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"><li>• <b>[edit routing-options graceful-restart]</b> (global setting)—120 through 900</li><li>• ES-IS—30 through 300</li><li>• IS-IS—30 through 300</li><li>• OSPF/OSPFv3—1 through 3600</li><li>• PIM—30 through 300</li></ul> <p><b>Default:</b></p> <p>The default value varies according to whether the graceful restart period is being set globally or for a particular protocol:</p> <ul style="list-style-type: none"><li>• <b>[edit routing-options graceful-restart]</b> (global setting)—300</li><li>• ES-IS—180</li><li>• IS-IS—210</li><li>• OSPF/OSPFv3—180</li><li>• PIM—60</li></ul>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Enabling Graceful Restart*
  - [Configuring Routing Protocols Graceful Restart on page 2261](#)
  - *Configuring Graceful Restart for MPLS-Related Protocols*
  - *Configuring VPN Graceful Restart*
  - *Configuring Graceful Restart for VPNs*
  - *Configuring Logical System Graceful Restart*
  - *Graceful Restart Configuration Statements*

## restart-time (BGP Graceful Restart)

<b>Syntax</b>	<code>restart-time seconds;</code>
<b>Hierarchy Level</b>	<p>[edit protocols (bgp   rip   ripng) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (bgp   rip   ripng) <a href="#">graceful-restart (Enabling Globally)</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.
<b>Options</b>	<p><b>seconds</b>—Length of time for the graceful restart period.</p> <p><b>Range:</b> 1 through 600 seconds</p> <p><b>Default:</b> Varies by protocol:</p> <ul style="list-style-type: none"> <li>• BGP—120 seconds</li> <li>• RIP and RIPng—60 seconds</li> </ul>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li> <li>• <a href="#">Configuring Graceful Restart Options for RIP and RIPng on page 2266</a></li> <li>• <i>Configuring Graceful Restart for QFabric Systems</i></li> <li>• <a href="#">stale-routes-time on page 2308</a></li> </ul>

## stale-routes-time

---

<b>Syntax</b>	<code>stale-routes-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-routing-name</i> protocols bgp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ], [edit protocols bgp <a href="#">graceful-restart</a> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the maximum time that stale routes are kept during a restart. The <b>stale-routes-time</b> statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.
<b>Options</b>	<b>seconds</b> —Time the router device waits to receive messages from restarting neighbors before declaring them down. <b>Range:</b> 1 through 600 seconds <b>Default:</b> 300 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li><li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li><li>• <a href="#">restart-time (BGP Graceful Restart) on page 2307</a></li></ul>

## Configuration Statement for Graceful Switchover

---

- [graceful-switchover on page 2309](#)
- [redundancy \(Graceful Switchover\) on page 2310](#)

## graceful-switchover

---

<b>Syntax</b>	graceful-switchover;
<b>Hierarchy Level</b>	[edit chassis <a href="#">redundancy</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	For switches with more than one Routing Engine, including those in a Virtual Chassis or a Virtual Chassis Fabric, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.
<b>Default</b>	Graceful Routing Engine switchover (GRES) is disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2274</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2268</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 2270</a></li> <li>• <a href="#">Configuring Nonstop Active Routing on Switches on page 2277</a></li> <li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li> </ul>

## redundancy (Graceful Switchover)

---

<b>Syntax</b>	<pre>redundancy {     failover {         on-disk-failure;         on-loss-of-keepalives;     }     graceful-switchover; }</pre>
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	Redundancy is enabled for the Routing Engines.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">graceful-switchover on page 2309</a></li><li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 2270</a></li><li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2268</a></li><li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li><li>• <a href="#">High Availability Features for EX Series Switches Overview</a></li></ul>

## Configuration Statement for Nonstop Bridging

---

- [nonstop-bridging on page 2311](#)



## nonstop-bridging

---

<b>Syntax</b>	nonstop-bridging;
<b>Hierarchy Level</b>	[edit protocols layer2-control]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4.
<b>Description</b>	For platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and preserve Layer 2 Control Protocol (L2CP) information.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Synchronizing the Routing Engine Configuration</i></li> <li>• <i>Configuring Nonstop Bridging</i></li> <li>• For information about configuring NSB on EX Series switches that do not support the Enhanced Layer 2 Software (ELS) CLI style, see <i>Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)</i></li> <li>• For information about configuring NSB on switches that support ELS, see <a href="#">Configuring Nonstop Bridging on Switches (CLI Procedure) on page 2272</a></li> </ul>

## Configuration Statements for Nonstop Routing

---

- [nonstop-routing on page 2312](#)
- [synchronize on page 2313](#)
- [traceoptions \(Routing Options\) on page 2315](#)

## nonstop-routing

---

<b>Syntax</b>	nonstop-routing;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 10.4 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series switches
<b>Description</b>	For routing platforms with two Routing Engines, configure a master Routing Engine to switch over gracefully to a backup Routing Engine and to preserve routing protocol information.
<b>Default</b>	disabled
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Nonstop Active Routing</a></li><li>• <a href="#">Configuring Nonstop Active Routing on Switches on page 2277</a></li><li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2274</a></li><li>• <a href="#">Example: Configuring Nonstop Active Routing on Switches on page 2274</a></li></ul>

## synchronize

<b>Syntax</b>	synchronize;
<b>Hierarchy Level</b>	[edit system commit]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4 for EX Series switches.
<b>Description</b>	For devices with multiple Routing Engines only. Configure the <b>commit</b> command to automatically perform a <b>commit synchronize</b> action between dual Routing Engines within the same chassis. The Routing Engine on which you execute the <b>commit</b> command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engine. Each Routing Engine then performs a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on both Routing Engines.



**NOTE:** When you configure nonstop active routing (NSR), you must configure the **commit synchronize** statement. Otherwise, the commit operation fails.

On the TX Matrix router, synchronization only occurs between the Routing Engines within the same chassis. When synchronization is complete, the new configuration is then distributed to the Routing Engines on the T640 routers. That is, the master Routing Engine on the TX Matrix router distributes the configuration to the master Routing Engine on each T640 router. Likewise, the backup Routing Engine on the TX Matrix router distributes the configuration to the backup Routing Engine on each T640 router.

On the TX Matrix Plus router, synchronization only occurs between the Routing Engines within the switch-fabric chassis and when synchronization is complete, the new configuration is then distributed to the Routing Engines on the line-card chassis (LCC). That is, the master Routing Engine on the TX Matrix Plus router distributes the configuration to the master Routing Engine on each LCC. Likewise, the backup Routing Engine on the TX Matrix Plus router distributes the configuration to the backup Routing Engine on each LCC.

In EX Series Virtual Chassis configurations:

- On EX4200 switches in Virtual Chassis, synchronization occurs between the switch in the master role and the switch in the backup role.
- On EX8200 switches in a Virtual Chassis, synchronization occurs only between the master and backup XRE200 External Routing Engines.

<b>Options</b>	<b>and-quit</b> —(Optional) Quit configuration mode if the commit synchronization succeeds.
	<b>at</b> —(Optional) Time at which to activate configuration changes.
	<b>comment</b> —(Optional) Write a message to the commit log.

**force**—(Optional) Force a commit synchronization on the other Routing Engine (ignore warnings).

**scripts**—(Optional) Push scripts to the other Routing Engine.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Synchronizing the Routing Engine Configuration</i></li><li>• <i>Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically</i></li></ul>
------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## traceoptions (Routing Options)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>nsr-synchronization</b> flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p><b>nsr-synchronization</b> and <b>nsr-packet</b> flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>nsr-synchronization</b> flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p><b>nsr-synchronization</b> flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p><b>nsr-synchronization</b> flag for PIM added in Junos OS Release 9.3.</p> <p><b>nsr-synchronization</b> flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>nsr-synchronization</b> flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
<b>Description</b>	<p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>
<b>Default</b>	If you do not include this statement, no global tracing operations are performed.
<b>Options</b>	<p><b>Values:</b></p> <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and</p>

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.


<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Tracing Global Routing Protocol Operations</i></li> <li>• <a href="#">Tracing Nonstop Active Routing Synchronization Events on page 2278</a></li> </ul>

## Configuration Statements for VRRP

---


- [accept-data on page 2318](#)
- [advertise-interval on page 2319](#)
- [asymmetric-hold-time on page 2320](#)
- [authentication-key on page 2321](#)
- [authentication-type on page 2322](#)
- [bandwidth-threshold on page 2323](#)
- [failover-delay on page 2324](#)
- [fast-interval on page 2325](#)
- [hold-time \(VRRP\) on page 2326](#)
- [interface \(VRRP Group\) on page 2327](#)
- [preempt \(VRRP\) on page 2328](#)
- [priority \(Protocols VRRP\) on page 2329](#)
- [priority-cost \(VRRP\) on page 2330](#)
- [priority-hold-time on page 2331](#)
- [route \(Interfaces\) on page 2332](#)
- [startup-silent-period on page 2333](#)
- [traceoptions on page 2334](#)
- [track \(VRRP\) on page 2336](#)
- [virtual-address on page 2337](#)
- [vrrp-group on page 2338](#)

## accept-data

<b>Syntax</b>	(accept-data   no-accept-data);
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	<p>In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a router that is acting as the master router accepts all packets destined for the virtual IP address.</p> <ul style="list-style-type: none"> <li>• <b>accept-data</b>—Enable the master router to accept all packets destined for the virtual IP address.</li> <li>• <b>no-accept-data</b>—Prevent the master router from accepting packets other than the ARP packets destined for the virtual IP address.</li> </ul>
<b>Default</b>	<p>If the router acting as the master router is the IP address owner or has its priority set to 255, the master router, by default, responds to all packets sent to the virtual IP address. However, if the router acting as the master router does not own the IP address or has its priority set to a value less than 255, the master router responds only to ARP requests.</p>
<div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>• If you want to restrict the incoming IP packets to ICMP packets only, you must configure firewall filters to accept only ICMP packets.</li> <li>• If you include the <b>accept-data</b> statement, your routing platform configuration does not comply with RFC 3768 (see section 6.4.3 of RFC 3768, <i>Virtual Router Redundancy Protocol (VRRP)</i>).</li> </ul> </div>	
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring an Interface to Accept All Packets Destined for the Virtual IP Address of a VRRP Group</i></li> </ul>



## advertise-interval


<b>Syntax</b>	<code>advertise-interval seconds;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv4 advertisement packets.  All routers in the VRRP group must use the same advertisement interval.
<div>  <p><b>NOTE:</b> When VRRPv3 is enabled, the <code>advertise-interval</code> statement cannot be used to configure advertisement intervals. Instead, use the <code>fast-interval</code> statement to configure advertisement intervals.</p> </div>	
<b>Options</b>	<i>seconds</i> —Interval between advertisement packets. <b>Range:</b> 1 through 255 seconds <b>Default:</b> 1 second
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Advertisement Interval for the VRRP Master Router</i></li> <li>• <a href="#">fast-interval on page 2325</a></li> <li>• <i>inet6-advertise-interval</i></li> <li>• <i>version-3</i></li> </ul>

## asymmetric-hold-time

---

<b>Syntax</b>	asymmetric-hold-time;
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	<p>Configure a VRRP master to fail over to a backup immediately—without waiting for the preemption hold time to expire—when a tracked route goes down. Otherwise, the master waits for the hold time to expire before it initiates a failover when a tracked route goes down.</p> <p>When the tracked route comes up again, the new backup (original master) router waits for the preemption hold time to expire before it reasserts mastership.</p>
<b>Default</b>	asymmetric-hold-time is disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP Preemption and Hold Time on page 2288</a></li></ul>

## authentication-key

<b>Syntax</b>	<code>authentication-key key;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 authentication key. You also must specify a VRRP authentication scheme by including the <b>authentication-type</b> statement.  All routers in the VRRP group must use the same authentication scheme and password.
<div>  <b>NOTE:</b> When VRRPv3 is enabled, the <b>authentication-type</b> and <b>authentication-key</b> statements cannot be configured for any VRRP groups. </div>	
<b>Options</b>	<b>key</b> —Authentication password. For simple authentication, it can be 1 through 8 characters long. For Message Digest 5 (MD5) authentication, it can be 1 through 16 characters long. If you include spaces, enclose all characters in quotation marks (" ").
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring VRRP Authentication (IPv4 Only)</i></li> <li>• <a href="#">Configuring VRRP Authentication (IPv4 Only) on page 2286</a></li> <li>• <a href="#">authentication-type on page 2322</a></li> <li>• <i>version-3</i></li> </ul>

## authentication-type

---

<b>Syntax</b>	<code>authentication-type <i>authentication</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	Enable Virtual Router Redundancy Protocol (VRRP) IPv4 authentication and specify the authentication scheme for the VRRP group. If you enable authentication, you must specify a password by including the <b>authentication-key</b> statement.  All routers in the VRRP group must use the same authentication scheme and password.



**NOTE:** When VRRPv3 is enabled, the **authentication-type** and **authentication-key** statements cannot be configured for any VRRP groups.

<b>Options</b>	<b><i>authentication</i></b> —Authentication scheme: <ul style="list-style-type: none"><li>• <b>simple</b>—Use a simple password. The password is included in the transmitted packet, so this method of authentication is relatively insecure.</li><li>• <b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing platform uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li></ul> <b>Default:</b> none (no authentication is performed).
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring VRRP Authentication (IPv4 Only)</a></li><li>• <a href="#">Configuring VRRP Authentication (IPv4 Only) on page 2286</a></li><li>• <a href="#">authentication-key on page 2321</a></li><li>• <a href="#">version-3</a></li></ul>

## bandwidth-threshold

<b>Syntax</b>	<code>bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> <b>track interface</b> <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> <b>track interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> <b>track interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> <b>track interface</b> <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	Specify the bandwidth threshold for Virtual Router Redundancy Protocol (VRRP) logical interface tracking.
<b>Options</b>	<p><b><i>bits-per-second</i></b>—Bandwidth threshold for the tracked interface. When the bandwidth of the tracked interface drops below the specified value, the VRRP group uses the bandwidth threshold priority cost value. You can include up to five bandwidth threshold statements for each interface you track.</p> <p><b>Range:</b> 1 through 10000000000000 bits per second</p> <p><b><i>priority-cost priority</i></b>—The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li> <li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2290</a></li> </ul>

## failover-delay

---

<b>Syntax</b>	<code>failover-delay <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).</p> <p>If you configure a failover delay, the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.</p>
<b>Options</b>	<p><i>milliseconds</i>—Specify the failover delay time, in milliseconds.</p> <p><b>Range:</b> 50 through 2000</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Troubleshooting VRRP on page 2385</a></li><li>• <a href="#">show vrrp on page 2374</a></li></ul>

## fast-interval

<b>Syntax</b>	<code>fast-interval milliseconds;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure the interval, in milliseconds, between Virtual Router Redundancy Protocol (VRRP) advertisement packets.</p> <p>All routers in the VRRP group must use the same advertisement interval.</p>
<b>Options</b>	<p><i>milliseconds</i>—Interval between advertisement packets.</p> <p><b>Range:</b> 10 through 40,950 milliseconds (range extended from 100–999 to 10–40,950 in Junos OS Release 12.2).</p>



**NOTE:** When configuring VRRP for IPv4, if you have chosen not to enable VRRPv3, you cannot set a value less than 100 for *fast-interval*. Commit check fails if a value less than 100 is configured.

**Default:** 1 second

<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Advertisement Interval for the VRRP Master Router</i></li> <li>• <a href="#">Configuring the Advertisement Interval for the VRRP Master on page 2287</a></li> <li>• <a href="#">advertise-interval on page 2319</a></li> <li>• <a href="#">advertise-interval on page 2319</a></li> <li>• <i>inet6-advertise-interval</i></li> <li>• <i>version-3</i></li> </ul>

## hold-time (VRRP)

---

<b>Syntax</b>	<code>hold-time seconds;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id preempt</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id preempt</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	In a Virtual Router Redundancy Protocol (VRRP) configuration, set the hold time before a higher-priority backup router preempts the master router.
<b>Default</b>	VRRP preemption is not timed.
<b>Options</b>	<b><i>seconds</i></b> —Hold-time period. <b>Range:</b> 0 through 3600 seconds <b>Default:</b> 0 seconds (VRRP preemption is not timed.)
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Backup Router to Preempt the VRRP Master Router</a></li><li>• <a href="#">Configuring VRRP Preemption and Hold Time on page 2288</a></li></ul>



## interface (VRRP Group)


<b>Syntax</b>	<pre>interface <i>interface-name</i> {     <b>bandwidth-threshold</b> <i>bits-per-second</i> <i>priority-cost</i> <i>priority</i>;     <b>priority-cost</b> <i>priority</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <b>vrrp-group</b> <i>group-id</i> <b>track</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <b>vrrp-inet6-group</b> <i>group-id</i> <b>track</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>bandwidth-threshold</b> statement added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	Enable logical interface tracking for a Virtual Router Redundancy Protocol (VRRP) group.



**WARNING:** On a QFabric system, do not apply interface tracking to a multichassis link aggregation group (MC-LAG) that includes an interface belonging to a network Node group device and an interface belonging to a server Node group device. If you do apply interface tracking to an MC-LAG configured in this way, a priority update will not occur if the state of the MC-LAG interface changes.

<b>Options</b>	<p><b>interface-name</b>—Interface to be tracked for this VRRP group.</p> <p><b>Range:</b> 1 through 10 interfaces</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring a Logical Interface to Be Tracked for a VRRP Group</i></li> <li>• <i>Configuring a Logical Interface to Be Tracked on page 2290</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> </ul>

## preempt (VRRP)

<b>Syntax</b>	(preempt   no-preempt) { hold-time seconds; }
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	In a Virtual Router Redundancy Protocol (VRRP) configuration, determine whether or not a backup router can preempt a master router: <ul style="list-style-type: none"> <li>• <b>preempt</b>—Allow the master router to be preempted.</li> </ul> <p>.....</p> <div>  <p><b>NOTE:</b> By default, a higher-priority backup router can preempt a lower-priority master router.</p> <p>.....</p> </div> <ul style="list-style-type: none"> <li>• <b>no-preempt</b>—Prohibit the preemption of the master router. When <b>no-preempt</b> is configured, the backup router cannot preempt the master router even if the backup router has a higher priority.</li> </ul> <p>The remaining statement is explained separately.</p>
<b>Default</b>	By default the <b>preempt</b> statement is enabled, and a higher-priority backup router preempts a lower-priority master router even if the <b>preempt</b> statement is not explicitly configured.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Backup Router to Preempt the VRRP Master Router</a></li> <li>• <a href="#">Configuring VRRP Preemption and Hold Time on page 2288</a></li> </ul>

## priority (Protocols VRRP)


<b>Syntax</b>	<code>priority <i>priority</i>;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
<b>Options</b>	<p><b>priority</b>—Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.</p> <p><b>Range:</b> 1 through 255</p> <p><b>Default:</b> 100 (for backup routers)</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Basic VRRP Support</a></li> <li>• <a href="#">Configuring Basic VRRP Support on page 2285</a></li> </ul>

## priority-cost (VRRP)

---

<b>Syntax</b>	<code>priority-cost priority;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id track interface interface-name],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet6 address address vrrp-inet6-group group-id track interface interface-name]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX2000 Universal Access Routers.
<b>Description</b>	Configure a Virtual Router Redundancy Protocol (VRRP) router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
<b>Options</b>	<b>priority</b> —The value subtracted from the configured VRRP priority when the tracked interface or route is down to force a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group. <b>Range:</b> 1 through 254
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2290</a></li></ul>

## priority-hold-time

<b>Syntax</b>	<code>priority-hold-time seconds;</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <a href="#">vrrp-group group-id track</a>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <a href="#">vrrp-inet6-group group-id track</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <a href="#">vrrp-group group-id track</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <a href="#">vrrp-inet6-group group-id track</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure a Virtual Router Redundancy Protocol (VRRP) router's priority hold time to define the minimum length of time that must elapse between dynamic priority changes. If the dynamic priority changes because of a tracking event, the priority hold timer begins running. If another tracking event or manual configuration change occurs while the timer is running, the new dynamic priority update is postponed until the timer expires.</p>
<div>  <p><b>NOTE:</b> When the track feature is configured, and if VRRP should pre-empt due to the tracking interface or route transition, any configured pre-empt hold time will be ignored. VRRP master will pre-empt according to the configuration of the priority-hold time.</p> </div>	
<b>Options</b>	<p><b>seconds</b>—Minimum length of time that must elapse between dynamic priority changes.</p> <p><b>Range:</b> 0through 3600 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li> <li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2290</a></li> </ul>

## route (Interfaces)

---

<b>Syntax</b>	<code>route <i>prefix</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id track</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id track</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS 11.3 for QFX Series. Statement introduced in Junos OS 12.1 for EX Series switches.
<b>Description</b>	Enable route tracking for a Virtual Router Redundancy Protocol (VRRP) group.
<b>Options</b>	<p><b><i>prefix</i></b>—Route to be tracked for this VRRP group.</p> <p><b><i>priority-cost priority</i></b>—The value subtracted from the configured VRRP priority when the tracked interface or route is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the VRRP group.</p> <p><b><i>routing-instance instance-name</i></b>—Routing instance in which the route is to be tracked. If the route is in the default, or global, routing instance, the value for <b><i>instance-name</i></b> must be <b>default</b>.</p>
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Route to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Route to Be Tracked on page 2289</a></li></ul>

---

## startup-silent-period

---

<b>Syntax</b>	startup-silent-period <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols vrrp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	Instruct the system to ignore the Master Down Event when an interface transitions from the down state to the up state. This statement is used to avoid incorrect error alarms caused by the delay or interruption of incoming Virtual Router Redundancy Protocol (VRRP) advertisement packets during the interface startup phase.
<b>Options</b>	<b>seconds</b> —Number of seconds for the startup period. <b>Default:</b> 4 seconds <b>Range:</b> 1 through 2000 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Startup Period for VRRP Operations</i></li><li>• <a href="#">Configuring the Startup Period for VRRP Operations on page 2287</a></li></ul>

## traceoptions

---

**Syntax**    traceoptions {  
              file <filename> <files number> <match regular-expression> <microsecond-stamp>  
                  <size size> <world-readable | no-world-readable>;  
              flag flag;  
              no-remote-trace;  
              }

**Hierarchy Level**    [edit protocols vrrp]

**Release Information**    Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Define tracing operations for the Virtual Router Redundancy Protocol (VRRP) process.

To specify more than one tracing operation, include multiple **flag** statements.

By default, VRRP logs the error, dcd configuration, and routing socket events in a file in the directory **/var/log**.



**NOTE:** The traceoptions statement is not supported on a QFabric system.

---

**Default**    If you do not include this statement, no VRRP-specific tracing operations are performed.

**Options**    **filename filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, VRRP tracing output is placed in the file **vrrpd**.

**files number**—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. When the maximum number is reached, the oldest trace file is overwritten.

**Range:** 0 through 4,294,967,296 files

**Default:** 3 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the VRRP-specific tracing options:

- **all**—All VRRP tracing operations
- **database**—Database changes
- **general**—General events
- **interfaces**—Interface changes



- **normal**—Normal events
- **packets**—Packets sent and received
- **state**—State transitions
- **timer**—Timer events

**match *regex***—(Optional) Refine the output to include only those lines that match the given regular expression.

**microsecond-stamp**—(Optional) Provide a timestamp with microsecond granularity.

**no-world-readable**—Restrict users from reading the log file.

**size *size***—(Optional) Maximum size of each trace file, in kilobytes, megabytes, or gigabytes. When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your routing platform

**Default:** 1 MB

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**world-readable**—Allow users to read the log file.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Tracing VRRP Operations</i></li> </ul>

## track (VRRP)

---

<b>Syntax</b>	<pre>track {   interface <i>interface-name</i> {     bandwidth-threshold <i>bits-per-second</i> priority-cost <i>priority</i>;     priority-cost <i>priority</i>;   }   priority-hold-time <i>seconds</i>;   route <i>prefix/prefix-length</i> routing-instance <i>instance-name</i> priority-cost <i>priority</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6 address <i>address</i> <i>vrrp-inet6-group group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>priority-hold-time</b> statement added in Junos OS Release 8.1.</p> <p><b>route</b> statement added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	Enable logical interface tracking, route tracking, or both, for a Virtual Router Redundancy Protocol (VRRP) group.
<b>Options</b>	The remaining statements are described separately.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Logical Interface to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Route to Be Tracked for a VRRP Group</a></li><li>• <a href="#">Configuring a Logical Interface to Be Tracked on page 2290</a></li><li>• <a href="#">Configuring a Route to Be Tracked on page 2289</a></li></ul>

## virtual-address

---

<b>Syntax</b>	<code>virtual-address [ <i>addresses</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> <i>vrrp-group group-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv4 or IPv6 group. You can configure up to eight addresses.
<b>Options</b>	<b><i>addresses</i></b> —Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Basic VRRP Support</a></li> <li>• <a href="#">Configuring Basic VRRP Support on page 2285</a></li> </ul>

## vrrp-group

<b>Syntax</b>	<pre> vrrp-group <i>group-id</i> {   (<b>accept-data</b>   <b>no-accept-data</b>);   <b>advertise-interval</b> <i>seconds</i>;   <b>advertisements-threshold</b> <i>number</i>;   <b>authentication-key</b> <i>key</i>;   <b>authentication-type</b> <i>authentication</i>;   <b>fast-interval</b> <i>milliseconds</i>;   (<b>preempt</b>   <b>no-preempt</b>) {     <b>hold-time</b> <i>seconds</i>;   }   <b>priority</b> <i>number</i>;   <b>track</b> {     <b>interface</b> <i>interface-name</i> {       <b>bandwidth-threshold</b> <i>bits-per-second</i> <b>priority-cost</b> <i>priority</i>;       <b>priority-cost</b> <i>priority</i>;     }     <b>priority-hold-time</b> <i>seconds</i>;     <b>route</b> <i>prefix/prefix-length</i> <b>routing-instance</b> <i>instance-name</i> <b>priority-cost</b> <i>priority</i>;   }   <b>virtual-address</b> [ <i>addresses</i> ];   <b>vrrp-inherit-from</b> <i>vrrp-group</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],  [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.  Statement introduced in Junos OS 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure a Virtual Router Redundancy Protocol (VRRP) IPv4 group. As of Junos OS Release 13.2, VRRP nonstop active routing (NSR) is enabled only when you configure the <b>nonstop-routing</b> statement at the [edit routing-options] or [edit logical system logical-system-name routing-options] hierarchy level.</p>
<b>Options</b>	<p><b><i>group-id</i></b>—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the <b>source-address-filter</b> statement. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p><b>Range:</b> 0 through 255</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.  interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring Basic VRRP Support*
  - *Configuring VRRP*
  - [Configuring Basic VRRP Support on page 2285](#)
  - [Example: Configuring VRRP for Load Sharing on page 2279](#)
  - *vrrp-inet6-group*



## CHAPTER 31

# Administration

- [Operational Mode Commands for Graceful Restart on page 2341](#)
- [Operational Mode Command for Graceful Switchover on page 2365](#)
- [Operational Mode Command for Nonstop Routing on page 2371](#)
- [Operational Mode Commands for VRRP on page 2373](#)

### Operational Mode Commands for Graceful Restart

---

- [Verifying Graceful Restart Operation on page 2341](#)
- [show bgp neighbor](#)
- [show log](#)
- [show \(ospf | ospf3\) overview](#)

### Verifying Graceful Restart Operation

This topic contains the following sections:

- [Graceful Restart Operational Mode Commands on page 2341](#)
- [Verifying BGP Graceful Restart on page 2342](#)
- [Verifying IS-IS and OSPF Graceful Restart on page 2342](#)
- [Verifying CCC and TCC Graceful Restart on page 2343](#)

### Graceful Restart Operational Mode Commands

---

To verify proper operation of graceful restart, use the following commands:

- **show bgp neighbor** (for BGP graceful restart)
- **show log** (for IS-IS and OSPF/OSPFv3 graceful restart)
- **show (ospf | ospfv3) overview** (for OSPF/OSPFv3 graceful restart)
- **show rsvp neighbor detail** (for RSVP graceful restart—helper router)
- **show rsvp version** (for RSVP graceful restart—restarting router)
- **show ldp session detail** (for LDP graceful restart)
- **show connections** (for CCC and TCC graceful restart)

- **show route instance detail** (for Layer 3 VPN graceful restart and for any protocols using graceful restart in a routing instance)
- **show route protocol l2vpn** (for Layer 2 VPN graceful restart)

For more information about these commands and a description of their output fields, see the [CLI Explorer](#).

### Verifying BGP Graceful Restart

To view graceful restart information for BGP sessions, use the **show bgp neighbor** command:

```
user@PE1> show bgp neighbor 192.255.10.1
Peer: 192.255.10.1+179 AS 64595 Local: 192.255.5.1+1106 AS 64595
  Type: Internal   State: Established   Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ static ]
  Options:<Preference LocalAddress HoldTime GracefulRestart Damping PeerAS Refresh>

  Local Address: 192.255.5.1 Holdtime: 90 Preference: 170
  IPsec SA Name: hope
  Number of flaps: 0
  Peer ID: 192.255.10.1      Local ID: 192.255.5.1      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 180
  Stale routes from peer are kept for: 180
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast
  NLRI that peer saved forwarding for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 0
  Suppressed due to damping: 0
  Last traffic (seconds): Received 19   Sent 19   Checked 19
  Input messages: Total 2      Updates 1      Refreshes 0      Octets 42
  Output messages: Total 3      Updates 0      Refreshes 0      Octets 116
  Output Queue[0]: 0
```

### Verifying IS-IS and OSPF Graceful Restart

To view graceful restart information for IS-IS and OSPF, configure traceoptions (see [“Tracking Graceful Restart Events” on page 2267](#)).

Here is the output of a traceoptions log from an OSPF restarting router:

```
Oct  8 05:20:12 Restart mode - sending grace lsas
Oct  8 05:20:12 Restart mode - estimated restart duration timer triggered
Oct  8 05:20:13 Restart mode - Sending more grace lsas
```



Here is the output of a traceoptions log from an OSPF helper router:

```
Oct  8 05:20:14 Helper mode for neighbor 192.255.5.1
Oct  8 05:20:14 Received multiple grace lsa from 192.255.5.1
```

### Verifying CCC and TCC Graceful Restart

To view graceful restart information for CCC and TCC connections, use the **show connections** command. The following example assumes four remote interface CCC connections between CE1 and CE2:

```
user@PE1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP
```

#### CCC Graceful restart : Restarting

Connection/Circuit	Type	St	Time last up	# Up trans
CE1-CE2-0	rmt-if	Restart	-----	0
fe-1/1/0.0	intf	Up		
PE1-PE2-0	tlsp	Up		
PE2-PE1-0	rlsp	Up		
CE1-CE2-1	rmt-if	Restart	-----	0
fe-1/1/0.1	intf	Up		
PE1-PE2-1	tlsp	Up		
PE2-PE1-1	rlsp	Up		
CE1-CE2-2	rmt-if	Restart	-----	0
fe-1/1/0.2	intf	Up		
PE1-PE2-2	tlsp	Up		
PE2-PE1-2	rlsp	Up		
CE1-CE2-3	rmt-if	Restart	-----	0
fe-1/1/0.3	intf	Up		
PE1-PE2-3	tlsp	Up		
PE2-PE1-3	rlsp	Up		

- Related Documentation**
- [Graceful Restart Concepts on page 2257](#)
  - [Configuring Graceful Restart for QFabric Systems](#)

## show bgp neighbor

---

<b>List of Syntax</b>	<a href="#">Syntax on page 2344</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 2344</a>
<b>Syntax</b>	<pre>show bgp neighbor &lt;exact-instance <i>instance-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>neighbor-address</i>&gt; &lt;orf (detail   <i>neighbor-address</i>)&gt;</pre>
<b>Syntax (EX Series Switch and QFX Series)</b>	<pre>show bgp neighbor &lt;instance <i>instance-name</i>&gt; &lt;exact-instance <i>instance-name</i>&gt; &lt;<i>neighbor-address</i>&gt; &lt;orf (<i>neighbor-address</i>   detail)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>orf</b> option introduced in Junos OS Release 9.2.</p> <p><b>exact-instance</b> option introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Display information about BGP peers.
<b>Options</b>	<p><b>none</b>—Display information about all BGP peers.</p> <p><b>exact-instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show bgp neighbor instance cust1</b> command).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>neighbor-address</i></b>—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p><b>orf (detail   <i>neighbor-address</i>)</b>—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the <b>detail</b> option to display detailed output.</p>
<b>Additional Information</b>	For information about the <b>local-address</b> , <b>nlri</b> , <b>hold-time</b> , and <b>preference</b> statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
<b>Required Privilege Level</b>	view

**Related Documentation** • [clear bgp neighbor on page 3751](#)

**List of Sample Output** [show bgp neighbor on page 2351](#)  
[show bgp neighbor \(CLNS\) on page 2352](#)  
[show bgp neighbor \(Layer 2 VPN\) on page 2353](#)  
[show bgp neighbor \(Layer 3 VPN\) on page 2355](#)  
[show bgp neighbor neighbor-address on page 2355](#)  
[show bgp neighbor neighbor-address on page 2356](#)  
[show bgp neighbor orf neighbor-address detail on page 2357](#)

**Output Fields** [Table 210 on page 2345](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

**Table 210: show bgp neighbor Output Fields**

Field Name	Field Description
<b>Peer</b>	Address of the BGP neighbor. The address is followed by the neighbor port number.
<b>AS</b>	AS number of the peer.
<b>Local</b>	Address of the local routing device. The address is followed by the peer port number.
<b>Type</b>	Type of peer: <b>Internal</b> or <b>External</b> .
<b>State</b>	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul>

Table 210: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
<b>Flags</b>	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> <li>• <b>Aggregate Label</b>—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label.</li> <li>• <b>CleanUp</b>—The peer session is being shut down.</li> <li>• <b>Delete</b>—This peer has been deleted.</li> <li>• <b>Idled</b>—This peer has been permanently idled.</li> <li>• <b>ImportEval</b>—At the last commit operation, this peer was identified as needing to reevaluate all received routes.</li> <li>• <b>Initializing</b>—The peer session is initializing.</li> <li>• <b>SendRtn</b>—Messages are being sent to the peer.</li> <li>• <b>Sync</b>—This peer is synchronized with the rest of the peer group.</li> <li>• <b>TryConnect</b>—Another attempt is being made to connect to the peer.</li> <li>• <b>Unconfigured</b>—This peer is not configured.</li> <li>• <b>WriteFailed</b>—An attempt to write to this peer failed.</li> </ul>
<b>Last state</b>	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul>
<b>Last event</b>	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Closed</b>—The BGP session closed.</li> <li>• <b>ConnectRetry</b>—The transport protocol connection failed, and BGP is trying again to connect.</li> <li>• <b>HoldTime</b>—The session ended because the hold timer expired.</li> <li>• <b>KeepAlive</b>—The local routing device sent a BGP keepalive message to the peer.</li> <li>• <b>Open</b>—The local routing device sent a BGP open message to the peer.</li> <li>• <b>OpenFail</b>—The local routing device did not receive an acknowledgment of a BGP open message from the peer.</li> <li>• <b>RecvKeepAlive</b>—The local routing device received a BGP keepalive message from the peer.</li> <li>• <b>RecvNotify</b>—The local routing device received a BGP notification message from the peer.</li> <li>• <b>RecvOpen</b>—The local routing device received a BGP open message from the peer.</li> <li>• <b>RecvUpdate</b>—The local routing device received a BGP update message from the peer.</li> <li>• <b>Start</b>—The peering session started.</li> <li>• <b>Stop</b>—The peering session stopped.</li> <li>• <b>TransportError</b>—A TCP error occurred.</li> </ul>

Table 210: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Cease</b>—An error occurred, such as a version mismatch, that caused the session to close.</li> <li>• <b>Finite State Machine Error</b>—In setting up the session, BGP received a message that it did not understand.</li> <li>• <b>Hold Time Expired</b>—The session's hold time expired.</li> <li>• <b>Message Header Error</b>—The header of a BGP message was malformed.</li> <li>• <b>Open Message Error</b>—A BGP open message contained an error.</li> <li>• <b>None</b>—No errors occurred in the BGP session.</li> <li>• <b>Update Message Error</b>—A BGP update message contained an error.</li> </ul>
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.
Options	<p>Configured BGP options:</p> <ul style="list-style-type: none"> <li>• <b>AddressFamily</b>—Configured address family: <b>inet</b> or <b>inet-vpn</b>.</li> <li>• <b>AuthKeyChain</b>—Authentication key change is enabled.</li> <li>• <b>DropPathAttributes</b>—Certain path attributes are configured to be dropped from neighbor updates during inbound processing.</li> <li>• <b>GracefulRestart</b>—Graceful restart is configured.</li> <li>• <b>HoldTime</b>—Hold time configured with the <b>hold-time</b> statement. The hold time is three times the interval at which keepalive messages are sent.</li> <li>• <b>IgnorePathAttributes</b>—Certain path attributes are configured to be ignored in neighbor updates during inbound processing.</li> <li>• <b>Local Address</b>—Address configured with the <b>local-address</b> statement.</li> <li>• <b>Multihop</b>—Allow BGP connections to external peers that are not on a directly connected network.</li> <li>• <b>NLRI</b>—Configured MBGP state for the BGP group: <b>multicast</b>, <b>unicast</b>, or both if you have configured <b>nlri any</b>.</li> <li>• <b>Peer AS</b>—Configured peer autonomous system (AS).</li> <li>• <b>Preference</b>—Preference value configured with the <b>preference</b> statement.</li> <li>• <b>Refresh</b>—Configured to refresh automatically when the policy changes.</li> <li>• <b>Rib-group</b>—Configured routing table group.</li> </ul>
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the <b>authentication-keychain</b> statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the <b>authentication-algorithm</b> statement has been configured) Type of authentication algorithm enabled: <b>hmac</b> or <b>md5</b> .
Address families configured	Names of configured address families for the VPN.

Table 210: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Local Address	Address of the local routing device.
Remove-private options	Options associated with the <code>remove-private</code> statement.
Holdtime	Hold time configured with the <code>hold-time</code> statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> <li>• <b>TrafficStatistics</b>—Collection of statistics for labeled-unicast traffic is enabled.</li> </ul>
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> <li>• <b>Options</b>—Options configured for collecting statistics about labeled-unicast traffic.</li> <li>• <b>File</b>—Name and location of statistics log files.</li> <li>• <b>size</b>—Size of all the log files, in bytes.</li> <li>• <b>files</b>—Number of log files.</li> </ul>
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the <code>preference</code> statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the <code>out-delay</code> parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.

Table 210: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRI advertised by peer	Address families supported by the peer: <b>unicast</b> or <b>multicast</b> .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the <b>end-of-rib</b> marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route.  Possible value is <b>inet-unicast</b> .

Table 210: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route.  Possible value is <b>inet-unicast</b> .
Table inet.number	Information about the routing table: <ul style="list-style-type: none"> <li>• <b>RIB State</b>—BGP is in the graceful restart process for this routing table: <b>restart is complete</b> or <b>restart in progress</b>.</li> <li>• <b>Bit</b>—Number that represents the entry in the routing table for this peer.</li> <li>• <b>Send state</b>—State of the BGP group: <b>in sync</b>, <b>not in sync</b>, or <b>not advertising</b>.</li> <li>• <b>Active prefixes</b>—Number of prefixes received from the peer that are active in the routing table.</li> <li>• <b>Received prefixes</b>—Total number of prefixes from the peer, both active and inactive, that are in the routing table.</li> <li>• <b>Accepted prefixes</b>—Total number of prefixes from the peer that have been accepted by a routing policy.</li> <li>• <b>Suppressed due to damping</b>—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.</li> </ul>
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	Information about dropped path attributes: <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>
Input ignored path attributes	Information about ignored path attributes: <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.  It also specifies the routing table name and the NLRI they represent in the format ( <b>routing table name, NLRI</b> ).  <b>NOTE:</b> The output queues of routing tables that are not advertised, will only show up at <b>extensive</b> output level.
Trace options	Configured tracing of BGP protocol packets and operations.



Table 210: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates rcv	(orf option only) Number of outbound-route filters received for each configured address family.  <b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.  <b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: <b>prefix-based</b> or <b>extended-community</b> .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to <b>permit</b> or <b>deny</b> route updates.

## Sample Output

### show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast

```

```

NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages: Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214 Local ID: 10.255.167.205 Active Holdtime: 90
  Keepalive Interval: 30 Group index: 0 Peer index: 1

```

### show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1 Local ID: 10.245.245.3 Active Holdtime: 90
  Keepalive Interval: 30 Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          3
    Received prefixes:        3
    Suppressed due to damping: 0
    Advertised prefixes:      3
  Table aaaa.iso.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:          3
    Received prefixes:        3
    Suppressed due to damping: 0
  Last traffic (seconds): Received 6    Sent 5    Checked 5
  Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
  Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305

```

```
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)
```

### show bgp neighbor (Layer 2 VPN)

```
user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established  Flags: <ImportEval>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast l2vpn
NLRI advertised by peer: inet-vpn-unicast l2vpn
NLRI for this session: inet-vpn-unicast l2vpn
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast l2vpn
NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:      10
  Received prefixes:    10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:      1
```

```

Received prefixes:          1
Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           2
Received prefixes:         2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           2
Received prefixes:         2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           1
Received prefixes:         1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           2
Received prefixes:         2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           2
Received prefixes:         2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           1
Received prefixes:         1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:           1
Received prefixes:         1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)

```

```
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)
```

### show bgp neighbor (Layer 3 VPN)

```
user@host> show bgp neighbor
Peer: 4.4.4.4+179 AS 10045 Local: 5.5.5.5+1214 AS 10045
Type: Internal State: Established Flags: <ImportEval>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Export: [ match-all ] Import: [ match-all ]
Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
Address families configured: inet-vpn-unicast
Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
Flags for NLRI inet-labeled-unicast: TrafficStatistics
Traffic Statistics: Options: all File: /var/log/bstat.log
size 131072 files 10

Traffic Statistics Interval: 60
Number of flaps: 0
Peer ID: 192.168.1.110 Local ID: 192.168.1.111 Active Holdtime: 90
Keepalive Interval: 30
NLRI for restart configured on peer: inet-vpn-unicast
NLRI advertised by peer: inet-vpn-unicast
NLRI for this session: inet-vpn-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-vpn-unicast
NLRI peer can save forwarding state: inet-vpn-unicast
NLRI that peer saved forwarding for: inet-vpn-unicast
NLRI that restart is negotiated for: inet-vpn-unicast
NLRI of received end-of-rib markers: inet-vpn-unicast
NLRI of all end-of-rib markers sent: inet-vpn-unicast
Table bgp.l3vpn.0 Bit: 10000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes: 2
Received prefixes: 2
Suppressed due to damping: 0
Table vpn-green.inet.0 Bit: 20001
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes: 2
Received prefixes: 2
Suppressed due to damping: 0
Last traffic (seconds): Received 15 Sent 20 Checked 20
Input messages: Total 40 Updates 2 Refreshes 0 Octets 856
Output messages: Total 44 Updates 2 Refreshes 0 Octets 1066
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
Trace options: detail packets
Trace file: /var/log/bgpgr.log size 131072 files 10
```

### show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.1.111
```

```

Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
  Refresh>
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3 Sent 3 Checked 3
  Input messages: Total 9 Updates 6 Refreshes 0 Octets 403
  Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
  Output Queue[0]: 0 (inet.0, inet-unicast)
  Output Queue[1]: 0 (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpr size 131072 files 10

```

### show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External State: Established Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6 Local ID: 10.255.245.5 Active Holdtime: 60000
  Keepalive Interval: 20000 Peer index: 0
  BFD: disabled, down

```

```

Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:       3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           0
  Received prefixes:         0
  Accepted prefixes:         0
  Suppressed due to damping: 0
  Advertised prefixes:       0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

#### show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:           1 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:           0 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

## show log

---

<b>List of Syntax</b>	<a href="#">Syntax on page 2358</a> <a href="#">Syntax (QFabric System) on page 2358</a> <a href="#">Syntax (TX Matrix Routers) on page 2358</a>
<b>Syntax</b>	<code>show log</code> <code>&lt;filename   user &lt;username&gt;&gt;</code>
<b>Syntax (QFabric System)</b>	<code>show log filename</code> <code>&lt;device-type (device-id   device-alias)&gt;</code>
<b>Syntax (TX Matrix Routers)</b>	<code>show log</code> <code>&lt;all-lcc   lcc number   scc&gt;</code> <code>&lt;filename   user &lt;username&gt;&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option <i>device-type (device-id   device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.
<b>Description</b>	List log files, display log file contents, or display information about users who have logged in to the router or switch.
<b>Options</b>	<b>none</b> —List all log files.  <b>&lt;all-lcc   lcc number   scc&gt;</b> —(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).  <b>device-type</b> —(QFabric system only) (Optional) Display log messages for only one of the following device types: <ul style="list-style-type: none"><li>• <b>director-device</b>—Display logs for Director devices.</li><li>• <b>infrastructure-device</b>—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).</li><li>• <b>interconnect-device</b>—Display logs for Interconnect devices.</li><li>• <b>node-device</b>—Display logs for Node devices.</li></ul>



**NOTE:** If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

---



**(device-id | device-alias)**—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

**filename**—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



**NOTE:** The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

**user <username>**—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

**Required Privilege Level** trace

**List of Sample Output** [show log on page 2359](#)  
[show log filename on page 2359](#)  
[show log filename \(QFabric System\) on page 2360](#)  
[show log user on page 2360](#)

## Sample Output

### show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin        19656 Oct  1 19:37 wtmp
```

### show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT rcv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT rcv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT rcv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

### show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

### show log user

```

user@host> show log user
darius  mg2546          Thu Oct  1 19:37   still logged in
darius  mg2529          Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518          Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575          Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```

## show (ospf | ospf3) overview

<b>List of Syntax</b>	<a href="#">Syntax on page 2361</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 2361</a>
<b>Syntax</b>	<pre>show (ospf   ospf3) overview &lt;brief   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>
<b>Syntax (EX Series Switch and QFX Series)</b>	<pre>show (ospf   ospf3) overview &lt;brief   extensive&gt; &lt;instance <i>instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Database protection introduced in Junos 10.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display Open Shortest Path First (OSPF) overview information.
<b>Options</b>	<p><b>none</b>—Display standard information about all OSPF neighbors for all routing instances.</p> <p><b>brief   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ospf overview on page 2363</a> <a href="#">show ospf overview (With Database Protection) on page 2364</a> <a href="#">show ospf3 overview (With Database Protection) on page 2364</a> <a href="#">show ospf overview extensive on page 2364</a>
<b>Output Fields</b>	<p><a href="#">Table 211 on page 2362</a> lists the output fields for the <b>show ospf overview</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 211: show ospf overview Output Fields

Field name	Field Description	Level of Output
<b>Instance</b>	OSPF routing instance.	All levels
<b>Router ID</b>	Router ID of the routing device.	All levels
<b>Route table index</b>	Route table index.	All levels
<b>Configured overload</b>	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
<b>Topology</b>	Topology identifier.	All levels
<b>Prefix export count</b>	Number of prefixes exported into OSPF.	All levels
<b>Full SPF runs</b>	Number of complete Shortest Path First calculations.	All levels
<b>SPF delay</b>	Delay before performing consecutive Shortest Path First calculations.	All levels
<b>SPF holddown</b>	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
<b>SPF rapid runs</b>	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
<b>LSA refresh time</b>	Refresh period for link-state advertisement (in minutes).	All levels
<b>Database protection state</b>	Current state of database protection.	All levels
<b>Warning threshold</b>	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
<b>Non self-generated LSAs</b>	Number of LSAs whose router ID is not equal to the local router ID: <b>Current</b> , <b>Warning</b> (threshold), and <b>Allowed</b> .	All levels
<b>Ignore time</b>	How long the database has been in the ignore state.	All levels
<b>Reset time</b>	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
<b>Ignore count</b>	Number of times the database has been in the ignore state: <b>Current</b> and <b>Allowed</b> .	All levels
<b>Restart</b>	Graceful restart capability: <b>enabled</b> or <b>disabled</b> .	All levels
<b>Restart duration</b>	Time period for complete reacquisition of OSPF neighbors.	All levels
<b>Restart grace period</b>	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 211: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): <b>enabled</b> or <b>disabled</b> .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): <b>enabled</b> or <b>disabled</b> .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .	All levels
Trace options	OSPF-specific trace options.	<b>extensive</b>
Trace file	Name of the file to receive the output of the tracing operation.	<b>extensive</b>
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: <b>Normal Stub</b> , <b>Not Stub</b> , or <b>Not so Stubby Stub</b> .	All levels
Authentication Type	Type of authentication: <b>None</b> , <b>Password</b> , or <b>MD5</b> .  <b>NOTE:</b> The <b>Authentication Type</b> field refers to the authentication configured at the <code>[edit protocols ospf area area-id]</code> level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

## Sample Output

### show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

### show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

### show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

### show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

---

## Operational Mode Command for Graceful Switchover

- `show system switchover`
- `show task replication`

## show system switchover

---

<b>List of Syntax</b>	<a href="#">Syntax on page 2366</a> <a href="#">Syntax (TX Matrix Router) on page 2366</a> <a href="#">Syntax (TX Matrix Plus Router) on page 2366</a> <a href="#">Syntax (MX Series Router) on page 2366</a>
<b>Syntax</b>	show system switchover
<b>Syntax (TX Matrix Router)</b>	show system switchover <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system switchover <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show system switchover <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. sfc option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.
<b>Description</b>	Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.



**NOTE:** Issue the `show system switchover` command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine, because the kernel-replication process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Beginning Junos OS Release 9.6, the `show system switchover` command has been deprecated on the master Routing Engine on all routers other than a TX Matrix (switch-card chassis) or a TX Matrix Plus (switch-fabric chassis) router.

However, in a routing matrix, if you issue the `show system switchover` command on the master Routing Engine of the TX Matrix router (or switch-card chassis), the CLI displays graceful switchover information for the master Routing Engine of the T640 routers (or line-card chassis) in the routing matrix. Likewise, if you issue the `show system switchover` command on the master Routing Engine of a TX Matrix Plus router (or switch-fabric chassis), the CLI displays output for the master Routing Engine of T1600 or T4000 routers in the routing matrix.

---



**Options** **all-chassis**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix router and the T640 routers configured in the routing matrix. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all Routing Engines on the TX Matrix Plus router and the T1600 or T4000 routers configured in the routing matrix.

**all-lcc**—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for all connected T1600 or T4000 LCCs.

Note that in this instance, packets get dropped. The LCCs perform GRES on their own chassis (GRES cannot be handled by one particular chassis for the entire router) and synchronization is not possible as the LCC plane bringup time varies for each LCC. Therefore, when there is traffic on these planes, there may be a traffic drop.

**all-members**—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on all members of the Virtual Chassis configuration.

**lcc *number***—(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display graceful Routing Engine switchover information for a specific T640 router connected to the TX Matrix router. On a TX Matrix Plus router, display graceful Routing Engine switchover information for a specific router connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(MX Series routers only) (Optional) Display graceful Routing Engines switchover information for all Routing Engines on the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display graceful Routing Engine switchover information for all Routing Engines on the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix router only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display graceful Routing Engine switchover information for the TX Matrix Plus router.

<b>Additional Information</b>	<p>If you issue the <b>show system switchover</b> command on a TX Matrix backup Routing Engine, the command is broadcast to all the T640 backup Routing Engines that are connected to it.</p> <p>Likewise, if you issue the <b>show system switchover</b> command on a TX Matrix Plus backup Routing Engine, the command is broadcast to all the T1600 or T4000 backup Routing Engines that are connected to it.</p> <p>If you issue the <b>show system switchover</b> command on the active Routing Engine in the master router of an MX Series Virtual Chassis, the router displays an error message that graceful Routing Engine switchover (GRES) is not enabled on this member.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Routing Matrix with a TX Matrix Plus Router Solutions Page</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show system switchover (Backup Routing Engine) on page 2369</a></p> <p><a href="#">show system switchover all-lcc (Routing Matrix) on page 2369</a></p>
<b>Output Fields</b>	<p>Table 212 on page 2368 describes the output fields for the <b>show system switchover</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 212: show system switchover Output Fields

Field Name	Field Description
<b>Graceful switchover</b>	<p>Display graceful Routing Engine switchover status:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—Indicates <b>graceful-switchover</b> is specified for the <b>routing-options</b> configuration command.</li> <li>• <b>Off</b>—Indicates <b>graceful-switchover</b> is not specified for the <b>routing-options</b> configuration command.</li> </ul>
<b>Configuration database</b>	<p>State of the configuration database:</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Configuration database has synchronized.</li> <li>• <b>Synchronizing</b>—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds.</li> <li>• <b>Synchronize failed</b>—Configuration database synchronize process failed.</li> </ul>
<b>Kernel database</b>	<p>State of the kernel database:</p> <ul style="list-style-type: none"> <li>• <b>Ready</b>—Kernel database has synchronized.</li> <li>• <b>Synchronizing</b>—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds.</li> <li>• <b>Version incompatible</b>—The primary and standby Routing Engines are running incompatible kernel database versions.</li> <li>• <b>Replication error</b>—An error occurred when the state was replicated from the primary Routing Engine. Inspect <b>Steady State</b> for possible causes, or notify Juniper Networks customer support.</li> </ul>
<b>Peer state</b>	<p>Routing Engine peer state:</p> <ul style="list-style-type: none"> <li>• <b>Steady State</b>—Peer completed switchover transition.</li> <li>• <b>Peer Connected</b>—Peer in switchover transition.</li> </ul>

## Sample Output

### show system switchover (Backup Routing Engine)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

### show system switchover all-lcc (Routing Matrix)

```
user@host> show system switchover all-lcc
```

```
lcc0-re0:
```

```
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

```
lcc2-re0:
```

```
-----
Multichassis replication: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

## show task replication

<b>Syntax</b>	<b>show task replication</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p>
<b>Description</b>	Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show task replication (Issued on the Master Routing Engine) on page 2370</a> <a href="#">show task replication (Issued on the Backup Routing Engine) on page 2371</a>
<b>Output Fields</b>	<a href="#">Table 213 on page 2370</a> lists the output fields for the <b>show task replication</b> command. Output fields are listed in the approximate order in which they appear.

**Table 213: show task replication Output Fields**

Field Name	Field Description
<b>Stateful replication</b>	Displays whether or not graceful Routing Engine switchover is configured. The status can be <b>Enabled</b> or <b>Disabled</b> .
<b>RE mode</b>	Displays the Routing Engine on which the command is issued: <b>Master</b> , <b>Backup</b> , or <b>Not applicable</b> (when the router has only one Routing Engine).
<b>Protocol</b>	Protocols that are supported by nonstop active routing.
<b>Synchronization Status</b>	Nonstop active routing synchronization status for the supported protocols. States are <b>NotStarted</b> , <b>InProgress</b> , and <b>Complete</b> .

## Sample Output

### show task replication (Issued on the Master Routing Engine)

```

user@host> show task replication
  Stateful Replication: Enabled
      RE mode: Master

  Protocol              Synchronization Status
  ---
  OSPF                  NotStarted
  BGP                   Complete
  IS-IS                 NotStarted

```

LDP	Complete
PIM	Complete

#### show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```

### Operational Mode Command for Nonstop Routing

---

- `show task replication`

## show task replication

<b>Syntax</b>	<b>show task replication</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series switches.</p> <p>Support for logical systems introduced in Junos OS Release 13.3</p>
<b>Description</b>	Displays nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show task replication (Issued on the Master Routing Engine) on page 2372</a> <a href="#">show task replication (Issued on the Backup Routing Engine) on page 2373</a>
<b>Output Fields</b>	<a href="#">Table 213 on page 2370</a> lists the output fields for the <b>show task replication</b> command. Output fields are listed in the approximate order in which they appear.

**Table 214: show task replication Output Fields**

Field Name	Field Description
<b>Stateful replication</b>	Displays whether or not graceful Routing Engine switchover is configured. The status can be <b>Enabled</b> or <b>Disabled</b> .
<b>RE mode</b>	Displays the Routing Engine on which the command is issued: <b>Master</b> , <b>Backup</b> , or <b>Not applicable</b> (when the router has only one Routing Engine).
<b>Protocol</b>	Protocols that are supported by nonstop active routing.
<b>Synchronization Status</b>	Nonstop active routing synchronization status for the supported protocols. States are <b>NotStarted</b> , <b>InProgress</b> , and <b>Complete</b> .

## Sample Output

### show task replication (Issued on the Master Routing Engine)

```

user@host> show task replication
  Stateful Replication: Enabled
      RE mode: Master

  Protocol              Synchronization Status
  ---
  OSPF                  NotStarted
  BGP                   Complete
  IS-IS                 NotStarted

```

LDP	Complete
PIM	Complete

#### show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
Stateful Replication: Enabled
RE mode: Backup
```

## Operational Mode Commands for VRRP

---

- `show vrrp`

## show vrrp

<b>Syntax</b>	<pre>show vrrp &lt;brief   detail   extensive   summary&gt; &lt;interface <i>interface-name</i>&gt; &lt;track interfaces&gt;</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display information and status about VRRP groups.
<b>Options</b>	<p><b>none</b>—(Same as brief) Display brief status information about all VRRP interfaces.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display information and status about the specified VRRP interface.</p> <p><b>track interfaces</b>—(Optional) Display information and status about VRRP track interfaces.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring VRRP for IPv6 (CLI Procedure)</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show vrrp on page 2379</a> <a href="#">show vrrp brief on page 2379</a> <a href="#">show vrrp detail (IPv6) on page 2379</a> <a href="#">show vrrp detail (Route Track) on page 2380</a> <a href="#">show vrrp extensive on page 2380</a> <a href="#">show vrrp interface on page 2381</a> <a href="#">show vrrp summary on page 2382</a> <a href="#">show vrrp track detail on page 2382</a> <a href="#">show vrrp track summary on page 2383</a>
<b>Output Fields</b>	Table 215 on page 2374 lists the output fields for the <b>show vrrp</b> command. Output fields are listed in the approximate order in which they appear.

**Table 215: show vrrp Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the logical interface.	<b>none, brief, extensive, summary</b>
<b>Interface index</b>	Physical interface index number, which reflects its initialization sequence.	<b>extensive</b>
<b>Groups</b>	Total number of VRRP groups configured on the interface.	<b>extensive</b>



Table 215: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Active</b>	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	<b>extensive</b>
<b>Interface VRRP PDU statistics</b>	Nonerrored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Advertisement sent</b>—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted.</li> <li>• <b>Advertisement received</b>—Number of VRRP advertisement PDUs received by the interface.</li> <li>• <b>Packets received</b>—Number of VRRP packets received for VRRP groups on the interface.</li> <li>• <b>No group match received</b>—Number of VRRP packets received for VRRP groups that do not exist on the interface.</li> </ul>	<b>extensive</b>
<b>Interface VRRP PDU error statistics</b>	Errored statistics for the logical interface: <ul style="list-style-type: none"> <li>• <b>Invalid IPAH next type received</b>—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets.</li> <li>• <b>Invalid VRRP ttl value received</b>—Number of packets received whose IP time-to-live (TTL) value is not 255.</li> <li>• <b>Invalid VRRP version received</b>—Number of packets received whose VRRP version is not 2.</li> <li>• <b>Invalid VRRP pdu type received</b>—Number of packets received whose VRRP PDU type is not 1.</li> <li>• <b>Invalid VRRP authentication type received</b>—Number of packets received whose VRRP authentication is not none, simple, or md5.</li> <li>• <b>Invalid VRRP IP count received</b>—Number of packets received whose VRRP IP count exceeds 8.</li> <li>• <b>Invalid VRRP checksum received</b>—Number of packets received whose VRRP checksum does not match the calculated value.</li> </ul>	<b>extensive</b>
<b>Physical interface</b>	Name of the physical interface.	<b>detail, extensive</b>
<b>Unit</b>	Logical unit number.	All levels
<b>Address</b>	Address of the physical interface.	<b>none, brief, detail, extensive</b>
<b>Index</b>	Physical interface index number, which reflects its initialization sequence.	<b>detail, extensive</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail, extensive</b>
<b>VRRP-Traps</b>	Status of VRRP traps: <b>Enabled</b> or <b>Disabled</b> .	<b>detail, extensive</b>

Table 215: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Type and Address</b>	Identifier for the address and the address itself: <ul style="list-style-type: none"> <li>• <b>lcl</b>—Configured local interface address.</li> <li>• <b>mas</b>—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router.</li> <li>• <b>vip</b>—Configured virtual IP addresses.</li> </ul>	none, brief, summary
<b>Interface state or Int state</b>	State of the physical interface: <ul style="list-style-type: none"> <li>• <b>down</b>—The device is present and the link is unavailable.</li> <li>• <b>not present</b>—The interface is configured, but no physical device is present.</li> <li>• <b>unknown</b>—The VRRP process has not had time to query the kernel about the state of the interface.</li> <li>• <b>up</b>—The device is present and the link is established.</li> </ul>	none, brief, extensive, summary
<b>Group</b>	VRRP group number.	none, brief, extensive, summary
<b>State</b>	VRRP state: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>initializing</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul>	extensive
<b>Priority</b>	Configured VRRP priority for the interface.	detail, extensive
<b>Advertisement interval</b>	Configured VRRP advertisement interval.	detail, extensive
<b>Authentication type</b>	Configured VRRP authentication type: <b>none</b> , <b>simple</b> , or <b>md5</b> .	detail, extensive
<b>Preempt</b>	Whether preemption is allowed on the interface: <b>yes</b> or <b>no</b> .	detail, extensive
<b>Accept-data mode</b>	Whether the interface is configured to accept packets destined for the virtual IP address: <b>yes</b> or <b>no</b> .	detail, extensive
<b>VIP count</b>	Number of virtual IP addresses that have been configured on the interface.	detail, extensive
<b>VIP</b>	List of virtual IP addresses configured on the interface.	detail, extensive
<b>Advertisement timer</b>	Time until the advertisement timer expires.	detail, extensive

Table 215: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is <b>N/A</b> .	detail, extensive
Virtual router uptime	Time that the virtual router has been up.	detail, extensive
Master router uptime	Time that the master router has been up.	detail, extensive
Virtual MAC	MAC address associated with the virtual IP address.	detail, extensive
Tracking	Whether tracking is <b>enabled</b> or <b>disabled</b> .	detail, extensive
Current priority	Current operational priority for being the VRRP master.	detail, extensive
Configured priority	Configured base priority for being the VRRP master.	detail, extensive
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. <b>Disabled</b> indicates no minimum interval.	detail, extensive
Remaining-time	( <b>track</b> option only) Displays the time remaining in the priority hold-time interval.	detail
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	detail extensive
Interface/Tracked interface	Name of the tracked interface.	detail extensive
Int state/Interface state	Current operational state of the tracked interface: <b>up</b> or <b>down</b> .	detail, extensive
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	detail, extensive
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	detail, extensive
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred.  An entry of <b>down</b> means that the corresponding priority cost is incurred when the interface is down.	detail, extensive
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	detail, extensive
Route count	The number of routes being tracked.	detail, extensive

Table 215: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Route</b>	The IP address of the route being tracked.	<b>detail, extensive</b>
<b>VRF name</b>	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	<b>detail, extensive</b>
<b>Route state</b>	The state of the route being tracked: <b>up</b> , <b>down</b> , or <b>unknown</b> .	<b>detail, extensive</b>
<b>Priority cost</b>	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	<b>detail, extensive</b>
<b>Active</b>	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	<b>detail, extensive</b>
<b>Group VRRP PDU statistics</b>	Number of VRRP advertisements sent and received by the group.	<b>extensive</b>
<b>Group VRRP PDU error statistics</b>	Errored statistics for the VRRP group: <ul style="list-style-type: none"> <li>• <b>Bad authentication type received</b>—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be <b>none</b>, <b>simple</b>, or <b>md5</b> and must be the same for all routers in the VRRP group.</li> <li>• <b>Bad password received</b>—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group.</li> <li>• <b>Bad MD5 digest received</b>—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router.</li> <li>• <b>Bad advertisement timer received</b>—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group.</li> <li>• <b>Bad VIP count received</b>—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance.</li> <li>• <b>Bad VIPADDR received</b>—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.</li> </ul>	<b>extensive</b>
<b>Group state transition statistics</b>	State transition statistics for the VRRP group: <ul style="list-style-type: none"> <li>• <b>Idle to master transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the master state.</li> <li>• <b>Idle to backup transitions</b>—Number of times that the VRRP instance transitioned from the idle state to the backup state.</li> <li>• <b>Backup to master transitions</b>—Number of times that the VRRP instance transitioned from the backup state to the master state.</li> <li>• <b>Master to backup transitions</b>—Number of times that the VRRP instance transitioned from the master state to the backup state.</li> </ul>	<b>extensive</b>
<b>Vlan-id</b>	ID of Vlan	<b>detail</b>

Table 215: show vrrp Output Fields (*continued*)

Field Name	Field Description	Level of Output
VR state	VRRP information: <ul style="list-style-type: none"> <li>• <b>backup</b>—The interface is acting as the backup router interface.</li> <li>• <b>bringup</b>—VRRP is just starting, and the physical device is not yet present.</li> <li>• <b>idle</b>—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established.</li> <li>• <b>initializing</b>—VRRP is initializing.</li> <li>• <b>master</b>—The interface is acting as the master router interface.</li> <li>• <b>transition</b>—The interface is changing between being the backup and being the master router.</li> </ul>	none, brief
Timer	VRRP timer information: <ul style="list-style-type: none"> <li>• <b>A</b>—Time, in seconds, until the advertisement timer expires.</li> <li>• <b>D</b>—Time, in seconds, until the Master is Dead timer expires.</li> </ul>	none, brief

## Sample Output

### show vrrp

```

user@host> show vrrp
Interface      State      Group  VR state  Timer  Type  Address
ge-0/0/0.121   up         1      master   A 1.052 1c1  gec0::12:1:1:1
                                     vip  ge80::12:1:1:99
                                     vip  gec0::12:1:1:99
ge-0/0/2.131   up         1      master   A 0.364 1c1  gec0::13:1:1:1
                                     vip  ge80::13:1:1:99
                                     vip  gec0::13:1:1:99

```

### show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see [show vrrp on page 2379](#).

### show vrrp detail (IPv6)

```

user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled

```

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled  
 Interface state: up, Group: 1, State: master  
 Priority: 200, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,  
 gec0::13:1:1:99  
 Advertisement timer: 0.327s, Master router: ge80::13:1:1:1  
 Virtual router uptime: 00:03:47, Master router uptime: 00:03:41  
 Virtual MAC: 00:00:5e:00:02:01  
 Tracking: disabled

### show vrrp detail (Route Track)

user@host> show vrrp detail

Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24  
 Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled  
 Interface state: up, Group: 100, State: master  
 Priority: 150, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100  
 Advertisement timer: 1.218s, Master router: 30.30.30.30  
 Virtual router uptime: 00:04:28, Master router uptime: 00:00:13  
 Virtual MAC: 00:00:5e:00:01:64  
 Tracking: enabled  
 Current priority: 150, Configured priority: 150  
 Priority hold-time: disabled  
 Interface tracking: disabled  
 Route tracking: enabled, Route count: 1

Route	VRF name	Route state	Priority cost
192.168.40.0/22	default	up	30

### show vrrp extensive

user@host> show vrrp extensive

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

#### Interface VRRP PDU statistics

Advertisement sent	:	188
Advertisement received	:	0
Packets received	:	0
No group match received	:	0

#### Interface VRRP PDU error statistics

Invalid IPAH next type received	:	0
Invalid VRRP TTL value received	:	0
Invalid VRRP version received	:	0
Invalid VRRP PDU type received	:	0
Invalid VRRP authentication type received	:	0
Invalid VRRP IP count received	:	0
Invalid VRRP checksum received	:	0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled  
 Interface state: up, Group: 1, State: master  
 Priority: 200, Advertisement interval: 1, Authentication type: none  
 Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,  
 gec0::12:1:1:99  
 Advertisement timer: 1.034s, Master router: ge80::12:1:1:1  
 Virtual router uptime: 00:04:04, Master router uptime: 00:03:58  
 Virtual MAC: 00:00:5e:00:02:01  
 Tracking: disabled  
 Group VRRP PDU statistics

```

    Advertisement sent           :          188
    Advertisement received       :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received           :         0
    Bad MD5 digest received         :         0
    Bad advertisement timer received:         0
    Bad VIP count received          :         0
    Bad VIPADDR received           :         0
Group state transition statistics
    Idle to master transitions       :         0
    Idle to backup transitions       :         1
    Backup to master transitions     :         1
    Master to backup transitions     :         0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
Interface VRRP PDU statistics
    Advertisement sent             :          186
    Advertisement received         :           0
    Packets received               :           0
    No group match received        :           0
Interface VRRP PDU error statistics
    Invalid IPAH next type received :         0
    Invalid VRRP TTL value received :         0
    Invalid VRRP version received   :         0
    Invalid VRRP PDU type received  :         0
    Invalid VRRP authentication type received:         0
    Invalid VRRP IP count received  :         0
    Invalid VRRP checksum received  :         0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
Interface state: up, Group: 1, State: master
Priority: 200, Advertisement interval: 1, Authentication type: none
Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
Virtual MAC: 00:00:5e:00:02:01
Tracking: disabled
Group VRRP PDU statistics
    Advertisement sent             :          186
    Advertisement received         :           0
Group VRRP PDU error statistics
    Bad authentication type received:         0
    Bad password received           :         0
    Bad MD5 digest received         :         0
    Bad advertisement timer received:         0
    Bad VIP count received          :         0
    Bad VIPADDR received           :         0
Group state transition statistics
    Idle to master transitions       :         0
    Idle to backup transitions       :         1
    Backup to master transitions     :         1
    Master to backup transitions     :         0

```

### show vrrp interface

```
user@host> show vrrp interface
```

Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1

Interface VRRP PDU statistics

```

Advertisement sent           :          205
Advertisement received       :           0
Packets received            :           0
No group match received      :           0

```

Interface VRRP PDU error statistics

```

Invalid IPAH next type received :          0
Invalid VRRP TTL value received :          0
Invalid VRRP version received  :          0
Invalid VRRP PDU type received :          0
Invalid VRRP authentication type received:          0
Invalid VRRP IP count received :          0
Invalid VRRP checksum received :          0

```

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled

Interface state: up, Group: 1, State: master

Priority: 200, Advertisement interval: 1, Authentication type: none

Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,  
gec0::12:1:1:99

Advertisement timer: 0.789s, Master router: ge80::12:1:1:1

Virtual router uptime: 00:04:26, Master router uptime: 00:04:20

Virtual MAC: 00:00:5e:00:02:01

Tracking: disabled

Group VRRP PDU statistics

```

Advertisement sent           :          205
Advertisement received       :           0

```

Group VRRP PDU error statistics

```

Bad authentication type received:          0
Bad password received           :          0
Bad MD5 digest received         :          0
Bad advertisement timer received:          0
Bad VIP count received          :          0
Bad VIPADDR received            :          0

```

Group state transition statistics

```

Idle to master transitions      :          0
Idle to backup transitions      :          1
Backup to master transitions    :          1
Master to backup transitions    :          0

```

## show vrrp summary

user@host> show vrrp summary

Interface	State	Group	VR state	Type	Address
ge-4/1/0.0	up	1	backup	lcl	10.57.0.2
				vip	10.57.0.100

## show vrrp track detail

user@host> show vrrp track detail

Tracked interface: ae1.211

State: up, Speed: 400m

Incurred priority cost: 0

Threshold	Priority cost	Active
400m	10	
300m	60	
200m	110	
100m	160	
down	190	



```
Tracking VRRP interface: ae0.210, Group: 1
VR State: master
Current priority: 200, Configured priority: 200
Priority hold-time: disabled,    Remaining-time: 50.351
```

#### show vrrp track summary

```
user@host> show vrrp track summary
```

Track if	State	Speed	VRRP if	Group	VR State	Current priority
ae1.211	up	400m	ae0.210	1	master	200



## CHAPTER 32

# Troubleshooting

- [Troubleshooting Procedures on page 2385](#)

## Troubleshooting Procedures

---

- [Troubleshooting VRRP on page 2385](#)

## Troubleshooting VRRP

**Problem**    **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

**Solution**    Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

**Related Documentation**    • [failover-delay on page 2324](#)



## PART 9

# Interfaces

- [Overview on page 2389](#)
- [Configuration on page 2457](#)
- [Administration on page 2747](#)
- [Troubleshooting on page 2887](#)



## CHAPTER 33

# Overview

- [Interfaces Overview on page 2389](#)

## Interfaces Overview

---

- [Interfaces Overview on page 2389](#)
- [Overview of Uplink Failure Detection on page 2392](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2393](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396](#)
- [Understanding Interface Naming Conventions on page 2401](#)
- [Understanding Interface Ranges on page 2406](#)
- [Understanding Layer 3 Logical Interfaces on page 2408](#)
- [Understanding Local Link Bias on page 2408](#)
- [Understanding Management Interfaces on page 2410](#)
- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Understanding Port Ranges and System Modes on page 2421](#)
- [Understanding Redundant Trunk Links on page 2447](#)
- [Understanding Generic Routing Encapsulation on page 2449](#)
- [Understanding Ethernet OAM Link Fault Management on page 2452](#)
- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 2454](#)

## Interfaces Overview

Juniper Networks devices have two types of interfaces: network interfaces and special interfaces. This topic provides brief information about these interfaces. For additional information, see the *Junos OS Network Interfaces Library for Routing Devices*.

- [Network Interfaces on page 2390](#)
- [Special Interfaces on page 2391](#)

## Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 216 on page 2390](#) lists the types of network interfaces supported.

**Table 216: Network Interface Types and Purposes**

Type	Purpose
Aggregated Ethernet interfaces	You can group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.
Channelized Interfaces	<p>Depending on the device and software package, 40-Gbps QSFP+ ports can be configured to operate as the following types of interfaces:</p> <ul style="list-style-type: none"> <li>10-Gigabit Ethernet interfaces (<i>xe</i>)</li> <li>40-Gigabit Ethernet interfaces (<i>et</i> and <i>xle</i>)</li> <li>40-Gigabit data plane uplink interfaces (<i>fte</i>)</li> </ul> <p>When an <i>et</i> port is channelized to four <i>xe</i> ports, a colon is used to signify the four separate channels. For example, on a QFX3500 standalone switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are <i>xe-0/1/2:0</i>, <i>xe-0/1/2:1</i>, <i>xe-0/1/2:2</i>, and <i>xe-0/1/2:3</i></p> <p><b>NOTE:</b> You cannot configure channelized interfaces to operate as Virtual Chassis ports.</p>
Ethernet Interfaces	You can configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches. You can configure 40-Gigabit data plane uplink ports to connect a Node device to an Interconnect devices as well as for Virtual Chassis ports (VCPs).
Fibre Channel interfaces	You can use Fibre Channel interfaces to connect the switch to a Fibre Channel over Ethernet (FCoE) forwarder or a Fibre Channel switch in a storage area network (SAN). You can configure Fibre Channel interfaces only on ports 0 through 5 and 42 through 47 on QFX3500 devices. Fibre Channel interfaces do not forward Ethernet traffic.
LAN access interfaces	You can use these interfaces to connect to other servers, storage, and switches. When you power on a QFX Series product and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports.
Multichassis aggregated Ethernet (MC-AE) interfaces	You can group a LAG on one standalone switch with a LAG on another standalone switch to create a MC-AE. The MC-AE provides load balancing and redundancy across the two standalone switches.
Tagged-access mode interfaces	You can use tagged-access interfaces to connect a switch to an access layer device. Tagged-access interfaces can accept VLAN-tagged packets from multiple VLANs.
Trunk interfaces	You can use trunk interfaces to connect to other switches or routers. To use a port for this type of connection, you must explicitly configure the port interface for trunk mode. The interfaces from the switches or routers must also be configured for trunk mode. In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.
Virtual Chassis ports (VCPs)	You can use Virtual Chassis ports to send and receive Virtual Chassis Control Protocol (VCCP) traffic, and to create, monitor, and maintain the Virtual Chassis. On QFX3500, QFX3600, QFX5100, and EX4600 standalone switches, you can configure 40-Gigabit Ethernet QSFP+ uplink ports (non-channelized) or fixed SFP+ 10-Gigabit Ethernet ports as VCPs by issuing the <b>request virtual-chassis-vc-port-set</b> CLI command.



## Special Interfaces

Table 217 on page 2391 lists the types of special interfaces supported on the QFX Series.

**Table 217: Special Interface Types and Purposes**

Type	Purpose
Console port	Each QFX Series product has a serial port, labeled <b>CON</b> or <b>CONSOLE</b> , for connecting tty-type terminals to the switch using standard PC-type tty cables. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch.
Loopback interface	All QFX Series products have this software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	The Juniper Networks Junos OS for the QFX Series includes management Ethernet interfaces. The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.
Routed VLAN interfaces (RVI and IRB interfaces)	<p>QFX Series products use a Layer 3 routed VLAN interface (called RVI in the original CLI, and called IRB in Enhanced Layer 2 Software) <b>vlan</b> to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The RVI or IRB functions as a logical router, eliminating the need for having both a switch and a router. The RVI or IRB must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.</p>

- Related Documentation**
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2393](#)
  - [Understanding Interface Naming Conventions on page 2401](#)
  - [Understanding Layer 3 Logical Interfaces on page 2408](#)
  - [Understanding Management Interfaces on page 2410](#)
  - [Understanding Integrated Routing and Bridging on page 1539](#)
  - [Overview of Fibre Channel on page 5508](#)

## Overview of Uplink Failure Detection

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate this information to the downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all of the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure that the traffic of the failed link is not dropped.

This topic describes:


- [Uplink Failure Detection Configuration on page 2392](#)
- [Failure Detection Pair on page 2393](#)

### Uplink Failure Detection Configuration

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces bound to the uplink interface. A server that is connected to the disabled downlink interface triggers a network adapter failover to a secondary link to avoid any traffic loss.

[Figure 45 on page 2392](#) illustrates a typical setup for uplink failure detection.

**Figure 45: Uplink Failure Detection Configuration on Switches**



g040577

For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.

The switch can monitor both physical interface links and logical interface links for uplink failures, but you must put the two types of interfaces into separate groups.



**NOTE:** For logical interfaces, the server must send keepalives between the switch and the server to detect failure of logical links.

### Failure Detection Pair

Uplink failure detection requires that you create pairs of uplink and downlink interfaces in a group. Each pair includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplinks the switch monitors. You can configure a maximum of eight uplink interfaces as link-to-monitor interfaces for a group.
- A link-to-disable interface—The link-to-disable interfaces specify the downlinks the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlinks to disable in the group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

#### Related Documentation

- [Configuring Interfaces for Uplink Failure Detection on page 2592](#)
- [Example: Configuring Interfaces for Uplink Failure Detection on page 2457](#)

## Understanding Aggregated Ethernet Interfaces and LACP

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface, also known as a *link aggregation group (LAG)* or *bundle*.

Link aggregation is used to aggregate Ethernet interfaces between two devices. You can create a LAG between a Juniper Networks device and a router, switch, aggregation switch, server, or other devices. The aggregated Ethernet interfaces that participate in a LAG are called member links. Because a LAG is composed of multiple member links, even if one member link fails, the LAG continues to carry traffic over the remaining links.



**NOTE:** On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard and is used as a discovery protocol.



**NOTE:** To ensure load balancing across the aggregated Ethernet (AE) interfaces on a redundant server Node group, the members of the AE must be equally distributed across the redundant server Node group.



**NOTE:** During a network Node group switchover, traffic might be dropped for a few seconds.

- [Link Aggregation Group on page 2394](#)
- [Link Aggregation Control Protocol \(LACP\) on page 2395](#)

---

## Link Aggregation Group

---

To create a LAG:

1. Create a logical aggregated Ethernet interface.
2. Define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and Link Aggregation Control Protocol (LACP).
3. Define the member links to be contained within the aggregated Ethernet interface—for example, two 10-Gigabit Ethernet interfaces.
4. Configure LACP for link detection.

Keep in mind these hardware and software guidelines:

- Up to 32 Ethernet interfaces can be grouped to form a LAG on a redundant server Node group, a server Node group, and a network Node group on a QFabric system. Up to 48 LAGs are supported on redundant server Node groups and server Node groups on a QFabric system, and up to 128 LAGs are supported on network Node groups on a QFabric system. You can configure LAGs across Node devices in redundant server Node groups, server Node groups, and network Node groups.



**NOTE:** If you try to commit a configuration containing more than 32 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 32 has been exceeded, and the configuration checkout has failed.

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 448 LAGs are supported on QFX3500, QFX3600, QFX5100, and EX4600 switches.



**NOTE:** If you try to commit a configuration containing more than 64 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 64 has been exceeded, and the configuration checkout has failed.

- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.



**NOTE:** On a QFX5100 and EX4600 standalone switch or QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure mixed rate aggregated Ethernet bundles (LAGs with different link speeds).



**NOTE:** Junos OS assigns a unique ID and port priority to each port. The ID and priority are not configurable.

- QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across a QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG. See *Understanding FCoE LAGs* for more information.

### Link Aggregation Control Protocol (LACP)

LACP is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure Ethernet links to actively transmit protocol data units (PDUs), or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. You can configure both VLAN-tagged

and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the LAG without user intervention.
- Link monitoring to check whether both ends of the bundle are connected to the correct group.

When a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server may not be able to exchange LACP PDUs. In such a situation you can configure an interface to be in the **up** state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When there are no received PDUs, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

**Related  
Documentation**

- [Configuring Link Aggregation on page 2593](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [Verifying the Status of a LAG Interface on page 2750](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

## Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Juniper Networks EX Series and QFX Series use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

This topic contains the following sections:

- [Understanding the Hashing Algorithm on page 2397](#)
- [IP \(IPv4 and IPv6\) on page 2398](#)
- [MPLS on page 2399](#)

- [MAC-in-MAC Packet Hashing on page 2400](#)
- [Layer 2 Header Hashing on page 2400](#)

### Understanding the Hashing Algorithm

---

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. The packet fields used by the hashing algorithm varies by the packet's EtherType and, in some instances, by the configuration on the switch. The hashing algorithm recognizes the following EtherTypes:

- IP (IPv4 and IPv6)
- MPLS
- MAC-in-MAC

Traffic that is not recognized as belonging to any of these EtherTypes is hashed based on the Layer 2 header. IP and MPLS traffic are also hashed based on the Layer 2 header when a user configures the hash mode as Layer 2 header.

You can configure some fields that are used by the hashing algorithm to make traffic forwarding decisions. You cannot, however, configure how certain values within a header are used by the hashing algorithm.

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.
- The same fields are used by the hashing algorithm to hash ECMP and LAG traffic, but the hashing algorithm hashes ECMP and LAG traffic differently. The different hashing ensures that traffic is not polarized when a LAG bundle is part of the ECMP next-hop path.
- The same fields are used for hashing regardless of whether the switch is or is not participating in a mixed or non-mixed Virtual Chassis or Virtual Chassis Fabric (VCF).

The fields used for hashing by each EtherType as well as the fields used by the Layer 2 header are discussed in the following sections.

## IP (IPv4 and IPv6)

Payload fields in IPv4 and IPv6 packets are used by the hashing algorithm when IPv4 or IPv6 packets need to be placed onto a member link in a LAG bundle or sent to the next-hop device when ECMP is enabled.

The hash mode is set to Layer 2 payload field, by default. IPv4 and IPv6 payload fields are used for hashing when the hash mode is set to Layer 2 payload.

If the hash mode is configured to Layer 2 header, IPv4, IPv6, and MPLS packets are hashed using the Layer 2 header fields. If you want incoming IPv4, IPv6, and MPLS packets hashed by the source MAC address, destination MAC address, or EtherType fields, you must set the hash mode to Layer 2 header.

Table 218 on page 2398 displays the IPv4 and IPv6 payload fields that are used by the hashing algorithm, by default.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

**Table 218: IPv4 and IPv6 Hashing Fields**

Fields	EX4300		QFX5100	
	LAG	ECMP	LAG	ECMP
Source MAC	X	X	X	X
Destination MAC	X	X	X	X
EtherType	X	X	X	X
VLAN ID	X (configurable)	X (configurable)	X (configurable)	X (configurable)
Source IP or IPv6	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
Destination IP or IPv6	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
Protocol (IPv4 only)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
Next header (IPv6 only)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)



Table 218: IPv4 and IPv6 Hashing Fields (*continued*)

Fields	EX4300		QFX5100	
Layer 4 Source Port	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
Layer 4 Destination Port	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
IPv6 Flow label (IPv6 only)	X	X	X	X

### MPLS

The hashing algorithm hashes MPLS packets using the source IP, destination IP, MPLS label 0, MPLS label 1, and MPLS label 2 fields. See [Table 219 on page 2399](#).

The fields used by the hashing algorithm for MPLS packet hashing are not user-configurable.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

The source IP and destination IP fields are not always used for hashing. For non-terminated MPLS packets, the payload is checked if the packet has a single MPLS label. If the payload is IPv4 or IPv6, then the IP source address and IP destination address fields are used for hashing along with the MPLS labels. If the packet has more than one MPLS label, only the MPLS labels are used for hashing.

Table 219: MPLS Hashing Fields

Field	EX4300	QFX5100
Source MAC	X	X
Destination MAC	X	X
EtherType	X	X
VLAN ID	X	X
Source IP	✓	✓
Destination IP	✓	✓
Protocol (for IPv4 packets)	X	X
Next header (for IPv6 packets)	X	X

Table 219: MPLS Hashing Fields (*continued*)

Field	EX4300	QFX5100
Layer 4 Source Port	X	X
Layer 4 Destination Port	X	X
IPv6 Flow lab	X	X
MPLS label 0	✓	✓
MPLS label 1	✓	✓
MPLS label 2	✓	✓

### MAC-in-MAC Packet Hashing

Packets using the MAC-in-MAC EtherType are hashed by the hashing algorithm using the Layer 2 payload source MAC, Layer 2 payload destination MAC, and Layer 2 payload EtherType fields. See [Table 220 on page 2400](#).

Hashing using the fields in the MAC-in-MAC EtherType packet is first supported on EX4300 switches in Release 13.2X51-D20. Hashing using the fields in the MAC-in-MAC EtherType is not supported on earlier releases.

The fields used by the hashing algorithm for MAC-in-MAC hashing are not user-configurable.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

Table 220: MAC-in-MAC Hashing Fields

Field	EX4300	QFX5100
Layer 2 Payload Source MAC	✓	✓
Layer 2 Payload Destination MAC	✓	✓
Layer 2 Payload EtherType	✓	✓
Layer 2 Payload Outer VLAN	X	X

### Layer 2 Header Hashing

Layer 2 header fields are used by the hashing algorithm when a packet's EtherType is not recognized as IP (IPv4 or IPv6), MPLS, or MAC-in-MAC. The Layer 2 header fields are also used for hashing IPv4, IPv6, and MPLS traffic instead of the payload fields when the hash mode is set to Layer 2 header.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

**Table 221: Layer 2 Header Hashing Fields**

Field	EX4300	QFX5100
Source MAC	✓ (configurable)	✓ (configurable)
Destination MAC	✓ (configurable)	✓ (configurable)
EtherType	✓ (configurable)	✓ (configurable)
VLAN ID	X (configurable)	X (configurable)

**Related  
Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2590](#)

## Understanding Interface Naming Conventions

The QFX Series and the EX4600 device uses a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the QFX Series and on EX4600 switches.

This topic describes:

- [Physical Part of an Interface Name on page 2401](#)
- [Logical Part of an Interface Name on a Switch Running QFabric Software Package on page 2405](#)
- [Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software on page 2406](#)
- [Wildcard Characters in Interface Names on page 2406](#)

### Physical Part of an Interface Name

Interfaces in Junos OS are specified as follows:

*device-name:type-fpc/pic/port*

The convention is as follows:

- *device-name*—(QFabric systems only) The *device-name* is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.
- *type*—The QFX Series and EX4600 device interfaces use the following media types:
  - **fc**—Fibre Channel interface
  - **ge**—Gigabit Ethernet interface
  - **xe**—10-Gigabit Ethernet interface
  - **xle**—40-Gigabit Ethernet interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
  - **et**—40-Gigabit Ethernet interface (QFX3500, QFX3600, QFX5100, and EX4600 switches running Enhanced Layer 2 Software)
  - **fte**—40-Gigabit data plane uplink interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
  - **me**—Management interface
  - **em**—Management interface on QFX5100 and EX4600 switches.
- *fpc*—Flexible PIC Concentrator. QFX Series interfaces use the following convention for the FPC number in interface names:
  - On QFX3500, QFX3600, and QFX5100 devices running a QFabric software package, the FPC number is always **0**.  
  
The FPC number indicates the slot number of the line card that contains the physical interface.
  - On QFX3500, QFX3600, QFX5100, and EX4600 switches running Enhanced Layer 2 Software, the member ID of a member in a Virtual Chassis determines the FPC number.



**NOTE:** Every member in a Virtual Chassis must have a unique member ID, otherwise the Virtual Chassis will not be created.

- On standalone QFX5100 and EX4600 switches, the FPC number is always **0**.
- *pic*—QFX Series and EX4600 device interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
  - On a QFX3500 switch running a QFabric software package, PIC **0** can support 48 ports, PIC **1** can support 16 10-Gigabit Ethernet ports, and PIC **2** can support 4 40-Gigabit Ethernet ports.
  - On a QFX3500 switch running Enhanced Layer 2 software, PIC **0** can support 48 ports, and PIC **1** can support 16 10-Gigabit Ethernet ports, and 4 40-Gigabit Ethernet ports.

- On a QFX3500 Node device running a QFabric software package, PIC 0 can support 48 ports and PIC 1 can support four 40-Gigabit data plane uplink ports.
- On a QFX3600 switch running a QFabric software package, PIC 0 can support 64 10-Gigabit Ethernet ports, and PIC 1 can support 16 40-Gigabit Ethernet ports.
- On a QFX3600 switch running Enhanced Layer 2 software, PIC 0 can support 64 10-Gigabit Ethernet ports and can also support 16 40-Gigabit Ethernet ports.
- On a QFX3600 Node device running a QFabric software package, PIC 0 can support 56 10-Gigabit Ethernet ports, and PIC 1 can support 8 40-Gigabit data plane uplink ports, and up to 14 40-Gigabit Ethernet ports.
- On a QFX5100-48S switch running Enhanced Layer 2 software, PIC 0 provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.
- On an EX4600 device running Enhanced Layer 2 software, PIC 0 provides 4 40-Gbps QSFP+ ports and 24 10-Gigabit Ethernet interfaces. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provides 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 40-Gbps QSFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each.
- On a QFX5100-48S switch running a QFabric software package, PIC 1 provides six 40-Gbps QSFP+ ports, and PIC 0 provides 48 10-Gigabit Ethernet interfaces.
- On a QFX5100-24Q switch running Enhanced Layer 2 software, PIC 0 provides 24 40-Gbps QSFP+ ports. PIC 1 and PIC 2 can each contain a QFX-EM-4Q expansion module, and each expansion module provides 4 40-Gbps QSFP+ ports.
- On a QFX5100-96S switch running Enhanced Layer 2 software, PIC 0 provides 96 10-Gigabit Ethernet interfaces and 8 40-Gbps QSFP+ ports.
- *port*—Interfaces use the following convention for port numbers:
  - On a QFX3500 switch running a QFabric software package, there are 48 network access ports (10-Gigabit Ethernet) labeled 0 through 47 on PIC 0 and, 16 network access ports labeled 0 through 15 on PIC 1, and four 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 2. You can use the QSFP+ ports to connect the Node device to Interconnect devices.

By default, the 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. Optionally, you can choose to configure the QSFP+ ports as 40-Gigabit Ethernet ports (see *Configuring the QSFP+ Port Type on QFX3500 Standalone Switches*).

- On a QFX3500 switch running Enhanced Layer 2 software, there are 48 network access ports labeled 0 through 47 on PIC 0 and 4 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 1. See [“Channelizing Interfaces” on page 2608](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- On a QFX3600 switch running a QFabric software package, there are 64 network access ports (10-Gigabit Ethernet) labeled Q0 through Q15 on PIC 0, and there are 16 network access ports (40-Gigabit Ethernet) labeled Q0 through Q15 on PIC 1.

By default, all the QSFP+ ports are configured to operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the QSFP+ ports as 10-Gigabit Ethernet ports (see *Configuring the Port Type on QFX3600 Standalone Switches*) and use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.

- On a QFX3600 Node device running a QFabric software package, PIC 0 can support up to 56 10-Gigabit Ethernet ports labeled Q2 through Q15, and PIC 1 can support up to 8 40-Gigabit data plane uplink ports labeled Q0 through Q7, and up to 14 40-Gigabit Ethernet ports labeled Q2 through Q15. See *Configuring the Port Type on QFX3600 Node Devices* for information on how to configure the 40-Gbps QSFP+ ports.

On a QFX3600 Node device, by default, four 40-Gbps QSFP+ ports (labeled Q0 through Q3) are configured for uplink connections between your Node device and your Interconnect devices, and twelve 40-Gbps QSFP+ ports (labeled Q4 through Q15) use QSFP+ to four SFP+ copper breakout cables to support up to 48 10-Gigabit Ethernet ports for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and ports Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see *Configuring the Port Type on QFX3600 Node Devices*).

- On a QFX3600 switch running Enhanced Layer 2 software, PIC 0 can support 64 network access ports (10-Gigabit Ethernet ports) labeled Q0 through Q15 and 16 40-Gigabit Ethernet ports labeled Q0 through Q15. See [“Channelizing Interfaces” on page 2608](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On a QFX5100-48S switch running Enhanced Layer 2 software, PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47 and 6 40-Gbps QSFP+ ports labeled 48 through 53. See [“Channelizing Interfaces” on page 2608](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On an EX4600 switch running Enhanced Layer 2 software, PIC 0 can support 24 network access ports (10-Gigabit Ethernet ports) labeled 0 through 23 and 4 40-Gbps QSFP+ ports labeled 24 through 27. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 40-Gbps QSFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each. See [“Channelizing Interfaces” on page 2608](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- On a QFX5100-48S switch running a QFabric software package, PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47, and PIC 1 can support 6 40-Gbps QSFP+ ports labeled 0 through 5. See *Configuring the QSFP+ Port Type on QFX5100 Switches* for information on how to configure the port mode of 40-Gbps QSFP+ ports.
- On a QFX5100-24Q switch running Enhanced Layer 2 software, PIC 0 can support 24 40-Gbps QSFP+ ports labeled 0 through 24. PIC 1 and PIC 2 each support one 40-Gbps QSFP+ port, for a total of two 40-Gbps QSFP+ ports. See [“Channelizing Interfaces” on page 2608](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.



**NOTE:** You cannot channelize the 40-Gbps QSFP+ ports provided in the two QFX-EM-4Q expansion modules. Also, even though there is a total of 128 physical ports, only 104 logical ports can be channelized.

You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See [“Configuring the System Mode” on page 2610](#) for information on how to configure the system mode.

- On a QFX5100-96S switch running Enhanced Layer 2 software, PIC 0 can support 96 10-Gigabit Ethernet ports labeled 0 through 95, and 8 40-Gbps QSFP+ ports labeled 96 through 103. See [“Channelizing Interfaces” on page 2608](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.



**NOTE:** You can only channelize the 40-Gbps QSFP+ ports provided in ports 96 and 100, because only 104 logical ports can be channelized.

You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See [“Configuring the System Mode” on page 2610](#) for information on how to configure the system mode.

### Logical Part of an Interface Name on a Switch Running QFabric Software Package

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *device-name* (QFabric systems only): *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

Interface	State	VLAN members	Blocking
node-device1:xe-0/0/1.0	down	remote-analyzer	unblocked
node-device1:xe-0/0/2.0	down	default	unblocked
node-device1:xe-0/0/3.0	down	default	unblocked

When you configure aggregated Ethernet interfaces, you configure a logical interface, which is called a *bundle* or a *LAG*. Each LAG can include up to eight Ethernet interfaces, depending on the switch model.

### Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software

---

Channelizing enables you to configure four 10-Gigabit Ethernet interfaces from a 40-Gigabit Ethernet QSFP+ interface. By default, a 40-Gigabit Ethernet QSFP+ interface is named *et-fpc/pic/port*. The resulting 10-Gigabit Ethernet interfaces appear in the following format: *xe-fpc/pic/port:channel*, where channel can be a value of 0 through 3.

For example, if an *et* interface named **et-0/0/3** is channelized to four 10-Gigabit Ethernet interfaces, the resulting 10-Gigabit Ethernet interface names will be **xe-0/0/3:0**, **xe-0/0/3:1**, **xe-0/0/3:2**, and **xe-0/0/3:3**:

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/3:0	up	down			
xe-0/0/3:1	up	down			
xe-0/0/3:2	up	down			
xe-0/0/3:3	up	down			

### Wildcard Characters in Interface Names

---

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the *interface-name* option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (\*) in quotation marks (" ").

#### Related Documentation

- [Interfaces Overview on page 2389](#)
- [Channelizing Interfaces on page 2608](#)
- [Configuring the System Mode on page 2610](#)
- [Understanding Management Interfaces on page 2410](#)
- [Understanding Port Ranges and System Modes on page 2421](#)
- *Rear Panel of a QFX3500 Device*
- *Front Panel of a QFX3600 Device*
- *Junos OS Network Interfaces Library for Routing Devices*

## Understanding Interface Ranges

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements



common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations is also a valid definition.



**NOTE:** The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** statement is used in the following configuration hierarchies:

- `ethernet-switching-options analyzer name input egress interface`
- `ethernet-switching-options analyzer name input ingress interface`
- `ethernet-switching-options analyzer output interface`
- `ethernet-switching-options bpdv-block interface`
- `ethernet-switching-options interfaces`
- `ethernet-switching-options redundant-trunk-group group-name interface`
- `ethernet-switching-options secure-access-port interface`
- `ethernet-switching-options voip interface`
- `protocols igmp-snooping vlan vlan-name interface`
- `protocols isis interface`
- `protocols link-management peer lmp-control-channel interface`
- `protocols link-management te-link name interface`
- `protocols lldp interface`
- `protocols mstp interface`
- `protocols mstp msti-id interface`
- `protocols mstp msti-id vlan vlan-id interface`
- `protocols sflow interfaces`
- `protocols stp interface`
- `protocols vstp vlan vlan-id interface`
- `vllans vlan-name interface`

#### Related Documentation

- [Interfaces Overview on page 2389](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)
- [Configuring Link Aggregation on page 2593](#)
- [Configuring a Layer 3 Logical Interface on page 2593](#)

- *Junos OS Network Interfaces Library for Routing Devices*
- [interface-range on page 2684](#)

## Understanding Layer 3 Logical Interfaces

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks switch to a Layer 2 switch. Only one physical connection is required between the switches.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series and EX4600 switches support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.


### Related Documentation

- [Interfaces Overview on page 2389](#)
- [Configuring a Layer 3 Logical Interface on page 2593](#)
- *Configuring DHCP and BOOTP Relay*
- *Junos OS Network Interfaces Library for Routing Devices*

## Understanding Local Link Bias


Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in [Figure 46 on page 2409](#).

Figure 46: Egress Traffic Flow with Local Link Bias



When local link bias is disabled, egress traffic exiting a Virtual Chassis or VCF on a LAG bundle can be forwarded out of any member link in the LAG bundle. Traffic forwarding decisions are made by an internal algorithm that attempts to load-balance traffic between the member links in the bundle. VCP bandwidth is frequently consumed by egress traffic when local link bias is disabled because the egress traffic traverses the VCPs to reach the destination egress member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is disabled is illustrated in Figure 47 on page 2409.

Figure 47: Egress Traffic Flow without Local Link Bias



Local link bias is configured in a LAG bundle. A Virtual Chassis or VCF that has multiple LAG bundles can contain bundles that have and have not enabled local link bias. Local link bias only impacts the forwarding of unicast traffic exiting a Virtual Chassis or VCF; ingress traffic handling is not impacted by the local link bias setting. Egress multicast, unknown unicast, and broadcast traffic exiting a Virtual Chassis or VCF over a LAG bundle is not impacted by the local link bias setting and is always load-balanced among the member links. Local link bias is disabled, by default.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced across the member links in the LAG bundle as it exits the Virtual Chassis or VCF.

#### Related Documentation

- [Configuring Local Link Bias \(CLI Procedure\) on page 2596](#)

## Understanding Management Interfaces

You use management interfaces to access devices remotely. Typically, a management interface is not connected to the in-band network, but is connected to a device in the internal network. Through a management interface, you can access the device over the network using utilities such as **ssh** and **telnet** and configure it from anywhere, regardless of its physical location. As a security feature, users cannot log in as **root** through a management interface. To access the device as **root**, you must use the console port. You can also use **root** to log in using SSH.



**NOTE:** Before you can use the management interfaces on the QFX3500, QFX3600, QFX5100, and EX4600 devices, you must configure the logical interfaces with valid IP addresses. Juniper Networks does not support configuring two management interfaces in the same subnet.

Management interface port ranges vary based on device type:

- QFX3500 devices:

The valid port range for a management interface (**me**) on a QFX3500 device is between 0 and 6, with a total of seven available ports. On a QFX3500 standalone switch, however, you can only configure **me0** and **me1** as management interfaces. The management interfaces are labeled **C0** and **C1**, and they correspond to **me0** and **me1**. On a QFX3500 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**.

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0s** and **C1s**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. Each pair of management interfaces correspond to one Ethernet interface—for example, both RJ-45 management interfaces (labeled **C0** and **C0s**) can correspond to **me0**, and both SFP management interfaces (labeled **C1** and **C1s**) can correspond to **me1**. By default, both RJ-45 management interfaces are active. If you insert an SFP interface into the SFP management port (**C0s**, for example), the SFP interface would become the active management interface, and the corresponding RJ-45 management interface (**C0**) is disabled.



**NOTE:** On a QFX3600 device, you can use either the RJ-45 or the SFP management interfaces, but not both at the same time.

- On QFX5100 and EX4600 switches, there is one RJ-45 management interface (labeled **C0**) and one SFP management interface (labeled **C1**), and they correspond to em0 and em1. You can use both management interfaces simultaneously.

- QFabric system:

On a QFabric system, there are management interfaces on the Node devices, Interconnect devices, and Director devices. However, you cannot access the management interfaces on the Node devices or Interconnect devices directly. You can only manage and configure these devices using the Director device. You can connect to the management interface over the network using utilities such as SSH.

For information on how to use management interfaces on a QFabric system, see *Performing the QFabric System Initial Setup on a QFX3100 Director Group* and *Gaining Access to the QFabric System Through the Default Partition*.

**Related  
Documentation**

- [Interfaces Overview on page 2389](#)

## Understanding Multichassis Link Aggregation

Layer 2 networks are increasing in scale mainly because of technologies such as virtualization. Protocol and control mechanisms that limit the disastrous effects of a topology loop in the network are necessary. Spanning Tree Protocol (STP) is the primary solution to this problem because it provides a loop-free Layer 2 environment. STP has gone through a number of enhancements and extensions, and although it scales to very large network environments, it still only provides one active path from one device to another, regardless of how many actual connections might exist in the network. Although STP is a robust and scalable solution to redundancy in a Layer 2 network, the single logical link creates two problems: At least half of the available system bandwidth is off-limits to data traffic, and network topology changes occur. The Rapid Spanning Tree Protocol (RSTP) reduces the overhead of the rediscovery process and allows a Layer 2 network to reconverge faster, but the delay is still high.

Link aggregation (IEEE 802.3ad) solves some of these problems by enabling users to use more than one link connection between switches. All physical connections are considered one logical connection. The problem with standard link aggregation is that the connections are point to point.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers. An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running the Spanning Tree Protocol (STP).

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of the MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use Interchassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.



**NOTE:** You must specify a service identifier (service-id) for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG), otherwise multichassis link aggregation will not work.

See [Table 222 on page 2413](#) for information about ICCP failure scenarios.

The following sections provide an overview of the terms and features associated with MC-LAG:

- [Active-Active Mode on page 2412](#)
- [ICCP and ICL-PL on page 2413](#)
- [Failure Handling on page 2413](#)
- [Multichassis Link Protection on page 2414](#)
- [MC-LAG Packet Forwarding on page 2414](#)
- [Layer 3 Routing on page 2414](#)
- [Spanning Tree Protocol \(STP\) Guidelines on page 2414](#)
- [MC-LAG Upgrade Guidelines on page 2415](#)
- [Layer 2 Unicast Features Supported on page 2415](#)
- [Layer 2 Multicast Features Supported on page 2416](#)
- [IGMP Snooping on an Active-Active MC-LAG on page 2416](#)
- [Layer 3 Unicast Features Supported on page 2417](#)
- [VRRP Active-Standby Support on page 2417](#)
- [Routed VLAN Interface \(RVI\) MAC Address Synchronization on page 2417](#)
- [Address Resolution Protocol \(ARP\) on page 2418](#)
- [DHCP Relay with Option 82 on page 2418](#)
- [Private VLAN \(PVLAN\) on page 2419](#)
- [Layer 3 Multicast on page 2419](#)

---

### Active-Active Mode

In active-active mode, all member links are active on the MC-LAG. In this mode, MAC addresses learned on one MC-LAG peer are propagated to the other MC-LAG peer. Active-active mode is the only mode supported at this time.

### ICCP and ICL-PL

ICCP replicates control traffic and forwarding states across the MC-LAG peers and communicates the operational state of the MC-LAG members. Because ICCP uses TCP/IP to communicate between the peers, the two peers must be connected to each other. ICCP messages exchange MC-LAG configuration parameters and ensure that both peers use the correct LACP parameters.

The interchassis link-protection link (ICL-PL) provides redundancy when a link failure (for example, an MC-LAG trunk failure) occurs on one of the active links. The ICL-PL can be either a 10-Gigabit Ethernet interface or an aggregated Ethernet interface. You can configure only one ICL-PL between the two peers, although you can configure multiple MC-LAGs between them.

### Failure Handling

Configuring ICCP adjacency over aggregated links mitigates the possibility of a split-brain state. A split brain state occurs when the ICL-PL configured between the MC-LAG peers goes down. To work around this problem, enable backup liveness detection. With backup liveness enabled, the MC-LAG peers can communicate through the keepalive link.

During a split-brain state, the standby peer brings down local members in the MC-LAG links by changing the LACP system ID. When the ICCP connection is active, both of the MC-LAG peers use the configured LACP system ID. If the LACP system ID is changed during failures, the server that is connected over the MC-LAG removes these links from the aggregated Ethernet bundle.

When the ICL-PL is operationally down and the ICCP connection is active, the LACP state of the links with status control configured as standby is set to the standby state. When the LACP state of the links is changed to standby, the server that is connected over the MC-LAG makes these links inactive and does not use them for sending data.

[Table 222 on page 2413](#) describes the different ICCP failure scenarios. The dash means that the item is not applicable.

**Table 222: ICCP Failure Scenarios**

ICCP Connection Status	ICL-PL Status	Backup Liveness Peer Status	Action on Multichassis Aggregated Ethernet (MC-AE) Interface with Status Set to Standby
Down	Down or Up	Not configured	LACP system ID is changed to default value.
Down	Down or Up	Active	LACP system ID is changed to default value.
Down	Down or Up	Inactive	No change in LACP system ID.
Up	Down	–	LACP state is set to standby. MUX state moves to waiting state.

Split-brain states bring down the MC-LAG link completely if the primary peer members are also down for other reasons. Recovery from the split-brain state occurs automatically when the ICCP adjacency comes up between the MC-LAG peers.

### Multichassis Link Protection

---

Multichassis link protection provides link protection between the two MC-LAG peers hosting an MC-LAG. If the ICCP connection is up and the ICL-PL comes up, the peer configured as standby brings up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer. Multichassis protection must be configured on each MC-LAG peer that is hosting an MC-LAG.

### MC-LAG Packet Forwarding

---

To prevent the server from receiving multiple copies from both of the MC-LAG peers, a block mask is used to prevent forwarding of traffic received on the ICL-PL toward the MC-AE interface. Preventing forwarding of traffic received on the ICL-PL interface toward the MC-AE interface ensures that traffic received on MC-LAG links is not forwarded back to the same link on the other peer. The forwarding block mask for a given MC-LAG link is cleared if all of the local members of the MC-LAG link go down on the peer. To achieve faster convergence, if all local members of the MC-LAG link are down, outbound traffic on the MC-LAG is redirected to the ICL-PL interface on the data plane.

### Layer 3 Routing

---

To provide Layer 3 routing functions to downstream clients, configure the same gateway address on both MC-LAG network peers. To upstream routers, the MC-LAG network peers could be viewed as either equal-cost multipath (ECMP) or two routes with different preference values.

Junos OS supports active-active MC-LAGs by using Virtual Router Redundancy Protocol (VRRP) over routed VLAN interfaces (RVIs). Junos OS also supports active-active MC-LAGs by using RVI MAC address synchronization. You must configure the RVI using the same IP address across MC-LAG peers.

### Spanning Tree Protocol (STP) Guidelines

---

- Enable STP globally.

STP might detect local miswiring loops within the peer or across MC-LAG peers.

STP might not detect network loops introduced by MC-LAG peers.

- Disable STP on ICL-PL links; otherwise, it might block ICL-PL ports and disable protection.
- Do not enable bridge protocol data unit (BPDU) block on interfaces connected to aggregation switches.

For more information about BPDU block, see [“Understanding BPDU Protection for STP, RSTP, and MSTP” on page 1558](#).



## MC-LAG Upgrade Guidelines

Upgrade the MC-LAG peers according to the following guidelines. See “[Upgrading Software](#)” on page 134 for exact details about how to perform a software upgrade.



**NOTE:** After a reboot, the MC-AE interfaces come up immediately and might start receiving packets from the server. If routing protocols are enabled, and the routing adjacencies have not been formed, packets might be dropped.

To prevent this scenario, issue the `set interfaces interface-name aggregated-ether-options mc-ae init-delay-time time` to set a time by which the routing adjacencies are formed.

1. Make sure that both of the MC-LAG peers (node1 and node2) are in the active-active state using the following command on any one of the MC-LAG peers:

```
user@switch> show interfaces mc-ae id 1
Member Link           : ae0
Current State Machine's State: mcae active state
Local Status          : active<<<<<<<<
Local State           : up
Peer Status           : active<<<<<<<<
Peer State            : up
  Logical Interface    : ae0.0
  Topology Type        : bridge
  Local State          : up
  Peer State           : up
  Peer Ip/MCP/State    : 20.1.1.2 ae2.0 up
```

2. Upgrade node1 of the MC-LAG.

When node1 is upgraded it is rebooted, and all traffic is sent across the available LAG interfaces of node2, which is still up. The amount of traffic lost depends on how quickly the neighbor devices detect the link loss and rehash the flows of the LAG.

3. Verify that node1 is running the software you just installed. Issue the `show version` command.
4. Make sure that both nodes of the MC-LAG (node1 and node2) are in the active-active state after the reboot of node1.
5. Upgrade node2 of the MC-LAG.

Repeat step 1 through step 3 to upgrade node2.

## Layer 2 Unicast Features Supported

The following Layer 2 unicast features are supported:

- L2 unicast: learning and aging

- Learned MAC addresses are propagated across MC-LAG peers for all of the VLANs that are spawned across the peers.
- Aging of MAC addresses occurs when the MAC address is not seen on both of the peers.
- MAC learning is disabled on the ICL-PL automatically.
- MAC addresses learned on single-homed links are propagated across all of the VLANs that have MC-LAG links as members.

---

## Layer 2 Multicast Features Supported

The following Layer 2 multicast features are supported:

- L2 multicast: unknown unicast and IGMP snooping
  - Flooding happens on all links across peers if both peers have virtual LAN membership. Only one of the peers forwards traffic on a given MC-LAG link.
  - Known and unknown multicast packets are forwarded across the peers by adding the ICL-PL port as a multicast router port.
  - IGMP membership learned on MC-LAG links is propagated across peers.
  - During an MC-LAG peer reboot, known multicast traffic is flooded until the IGMP snooping state is synced with the peer.

---

## IGMP Snooping on an Active-Active MC-LAG

IGMP snooping controls multicast traffic in a switched network. When IGMP snooping is not enabled, the Layer 2 device broadcasts multicast traffic out of all of its ports, even if the hosts on the network do not want the multicast traffic. With IGMP snooping enabled, a Layer 2 device monitors the IGMP join and leave messages sent from each connected host to a multicast router. This enables the Layer 2 device to keep track of the multicast groups and associated member ports. The Layer 2 device uses this information to make intelligent decisions and to forward multicast traffic to only the intended destination hosts. IGMP uses Protocol Independent Multicast (PIM) to route the multicast traffic. PIM uses distribution trees to determine which traffic is forwarded.

In an active-active MC-LAG configuration, IGMP snooping replicates the Layer 2 multicast routes so that each MC-LAG peer has the same routes. If a device is connected to an MC-LAG peer by way of a single-homed interface, IGMP snooping replicates join message to its IGMP snooping peer. If a multicast source is connected to an MC-LAG by way of a Layer 3 device, the Layer 3 device passes this information to the RVI that is configured on the MC-LAG. The first hop DR is responsible for sending the register and register-stop messages for the multicast group. The last hop DR is responsible for sending PIM join and leave messages toward the rendezvous point and source for the multicast group. The routing device with the smallest preference metric forwards traffic on transit LANs.

Configure the ICL-PL interface as a router-facing interface. For the scenario in which traffic arrives by way of a Layer 3 interface, PIM and IGMP must be enabled on the RVI interface configured on the MC-LAG peers.

---

### Layer 3 Unicast Features Supported

---

The following Layer 3 unicast features are supported:

- VRRP active-standby support enables Layer 3 routing over MC-AE interfaces.
- Routed VLAN interface (RVI) MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on MC-AE interfaces with either its own RVI MAC address or its peer's RVI MAC address.
- Address Resolution Protocol (ARP) synchronization enables ARP resolution on both of the MC-LAG peers.
- DHCP Relay with option 82 enables option 82 on the MC-LAG peers. Option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client.

---

### VRRP Active-Standby Support

---

VRRP in active-standby mode enables Layer 3 routing over the MC-AE interfaces on the MC-LAG peers. In this mode, the MC-LAG peers act as virtual routers. The virtual routers share the virtual IP address that corresponds to the default route configured on the host or server connected to the MC-LAG. This virtual IP address, known as a routed VLAN interface (RVI), maps to either of the VRRP MAC addresses or the logical interfaces of the MC-LAG peers. The host or server uses the VRRP MAC address to send any Layer 3 upstream packets. At any time, one of the VRRP routers is the master (active), and the other is a backup (standby). Both VRRP active and VRRP backup routers forward Layer 3 traffic arriving on the MC-AE interface. If the master router fails, all the traffic shifts to the MC-AE link on the backup router.



**NOTE:** You must configure VRRP on both MC-LAG peers in order for both the active and standby members to accept and route packets. Additionally, configure the VRRP backup router to send and receive ARP requests.

Routing protocols run on the primary IP address of the RVI, and both of the MC-LAG peers run routing protocols independently. The routing protocols use the primary IP address of the RVI and the RVI MAC address to communicate with the MC-LAG peers. The RVI MAC address of each MC-LAG peer is replicated on the other MC-LAG peer and is installed as a MAC address that has been learned on the ICL-PL.

---

### Routed VLAN Interface (RVI) MAC Address Synchronization

---

Routed VLAN interface (RVI) MAC address synchronization enables MC-LAG peers to forward Layer 3 packets arriving on MC-AE interfaces with either its own RVI MAC address or its peer's RVI MAC address. Each MC-LAG peer installs its own RVI MAC address as well as the peer's RVI MAC address in the hardware. Each MC-LAG peer treats the packet as if it were its own packet. If RVI MAC address synchronization is not enabled, the RVI MAC address is installed on the MC-LAG peer as if it was learned on the ICL-PL.



**NOTE:** If you need routing capability, configure both VRRP and routing protocols on each MC-LAG peer.

Control packets destined for a particular MC-LAG peer that arrive on an MC-AE interface of its MC-LAG peer are not forwarded on the ICL-PL interface. Additionally, using the gateway IP address as a source address when you issue either a ping, traceroute, telnet, or FTP request is not supported.

To enable RVI MAC address synchronization, issue the **set vlan *vlan-name* l3\_interface *rvi-name* mcae-mac-synchronize** on each MC-LAG peer. Configure the same IP address on both MC-LAG peers. This IP address is used as the default gateway for the MC-LAG servers or hosts.

### Address Resolution Protocol (ARP)

---

Address Resolution Protocol (ARP) maps IP addresses to MAC addresses. Without synchronization, if one MC-LAG peer sends an ARP request, and the other MC-LAG peer receives the response, ARP resolution is not successful. With synchronization, the MC-LAG peers synchronize the ARP resolutions by sniffing the packet at the MC-LAG peer receiving the ARP response and replicating this to the other MC-LAG peer. This ensures that the entries in ARP tables on the MC-LAG peers are consistent.

When one of the MC-LAG peers restarts, the ARP destinations on its MC-LAG peer are synchronized. Because the ARP destinations are already resolved, its MC-LAG peer can forward Layer 3 packets out of the MC-AE interface.



**NOTE:** For integrated routing and bridging (IRB) interfaces, static ARP configuration is required on the MC-LAG IRB peer for any routing protocol to come up over that IRB interface.

### DHCP Relay with Option 82

---

DHCP relay with option 82 provides information about the network location of DHCP clients. The DHCP server uses this information to implement IP addresses or other parameters for the client. With DHCP relay enabled, DHCP request packets might take the path to the DHCP server through either of the MC-LAG peers. Because the MC-LAG peers have different host names, chassis MAC addresses, and interface names, you need to observe these requirements when you configure DHCP relay with option 82:

- Use the interface description instead of the interface name.
- Do not use the hostname as part of the circuit ID or remote ID strings.
- Do not use the chassis MAC address as part of the remote ID string.
- Do not enable the vendor ID.
- If the ICL-PL interface receives DHCP request packets, the packets are dropped to avoid duplicate packets in the network.

A counter called *Due to received on ICL interface* has been added to the **show helper statistics** command, which tracks the packets that the ICL-PL interface drops.

An example of the CLI output follows:

```
user@switch> show helper statistics
BOOTP:
  Received packets: 6
  Forwarded packets: 0
  Dropped packets: 6
    Due to no interface in DHCP Relay database: 0
    Due to no matching routing instance: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
    Due to invalid server address: 0
    Due to no valid local address: 0
    Due to no route to server/client: 0
    Due to received on ICL interface: 6
```

The output shows that six packets received on the ICL-PL interface have been dropped.

### Private VLAN (PVLAN)

Private VLANs allow you to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside of a VLAN. A PVLAN can span multiple peers on an MC-LAG.

When configuring a PVLAN, you must configure the ICL-PL interface as the PVLAN trunk interface for the PVLAN. This is essential for traffic to be switched to the required primary and secondary ports of the PVLAN across the MC-LAG peers.

### Layer 3 Multicast

- [PIM Operation With Normal Mode DR Election on page 2419](#)
- [PIM Operation with Dual-DR Mode on page 2420](#)
- [Configuration Guidelines and Caveats on page 2420](#)

Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) provide support for Layer 3 multicast. In addition to the standard mode of PIM operation, there is a special mode called PIM dual DR (designated router). PIM dual DR minimizes traffic loss in case of failures.

#### ***PIM Operation With Normal Mode DR Election***

In normal mode DR election, the RVI interfaces on both of the MC-LAG peers are configured with PIM enabled. In this mode, one of the MC-LAG peers becomes the DR through the PIM DR election mechanism. The elected DR maintains the rendezvous-point tree (RPT) and shortest-path tree (SPT) so it can receive data from the source device. The elected DR participates in periodic PIM join and prune activities toward the rendezvous point (RP) or the source.

The trigger for initiating these join and prune activities is the IGMP membership reports that are received from interested receivers. IGMP reports received over MC-AE interfaces (potentially hashing on either of the MC-LAG peers) and single-homed links are synchronized to the MC-LAG peer through ICCP.

Both MC-LAG peers receive traffic on their incoming interface (IIF). The non-DR receives traffic by way of the ICL-PL interface, which acts as a multicast router (mrouter) interface.

If the DR fails, the non-DR has to build the entire forwarding tree (RPT and SPT), which can cause multicast traffic loss.

#### ***PIM Operation with Dual-DR Mode***

In this mode, both of MC-LAG peers act as DRs (active and backup) and send periodic join and prune messages upstream towards the RP, or source, and eventually join the RPT or SPT.

The primary MC-LAG peer forwards the multicast traffic to the receiver devices even if the standby MC-LAG peer has a smaller preference metric.

The standby MC-LAG peer also joins the forwarding tree and receives the multicast data. The standby MC-LAG peer drops the data because it has an empty outgoing interface list (OIL). When the standby MC-LAG peer detects the primary MC-LAG peer failure, it adds the receiver VLAN to the OIL, and starts to forward the multicast traffic

To enable a multicast dual DR, issue the **set protocols pim interface interface-name dual-dr** command on the VLAN interfaces of each MC-LAG peer.

#### ***Configuration Guidelines and Caveats***

- Configure the IP address on the active MC-LAG peer with a high IP address or a high DR priority. To ensure that the active MC-LAG peer retains the DR membership designation if PIM neighborship with the peer goes down.
- Using Bidirectional Forwarding Detection (BFD) and RVI MAC synchronization together is not supported because ARP fails.
- When using RVI MAC synchronization, make sure that you configure the primary IP address on both MC-LAG peers. Doing this ensures that both MC-LAG peers cannot become assert winners.
- The number of BFD sessions on RVIs with PIM enabled is restricted to 100. Also, if you have more than 100 RVIs configured, do not configure BFD, and make sure that the hello interval is 2 seconds.

#### **Related Documentation**

- [Configuring Link Aggregation on page 2593](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2493](#)
- [Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization on page 2530](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\) on page 2551](#)

## Understanding Port Ranges and System Modes

QFX Series devices and EX4600 switches can support different port ranges depending on the device, media type of the interface, the software that is running on the device, and the system mode.

This topic describes:

- [Port Ranges for Different Media Types on page 2421](#)
- [Supported System Modes on page 2444](#)

### Port Ranges for Different Media Types

The following media types support the following port ranges:

- On a QFX3500 device:
  - The valid port range for a Fibre Channel (fc) interface is **0** through **5** and **42** through **47** on PIC **0**, with a total of 12 available Fibre Channel ports.



**NOTE:** Fibre Channel ports are not supported on QFX3500, QFX3600, and QFX5100 switches running Enhanced Layer 2 software.

- The valid port range for a Gigabit Ethernet (ge) interface is **6** through **41** on PIC **0** because the ports between **0** and **5** and **42** and **47** are reserved as Fibre Channel ports. The total number of available Gigabit Ethernet ports is 36, because 12 of the remaining 48 ports are reserved for Fibre Channel and 10-Gigabit Ethernet interfaces. Fibre Channel ports cannot be configured as Gigabit Ethernet ports.
- The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **47** on PIC **0**. The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **15** on PIC **1**. The total number of available 10-Gigabit Ethernet ports is 64.
- The valid port range for a 40-Gigabit data plane uplink interface is **0** through **3** on PIC **1**
- The valid port range for a 40-Gigabit Ethernet interface is **0** through **3** on PIC **2**. There are four available ports.
- On a QFX3600 Node device:
  - The valid port range for a 10-Gigabit Ethernet interface is **8** through **63** on PIC **0**. There are 56 available ports.
  - The valid port range for a 40-Gigabit Ethernet interface is **2** through **15** on PIC **1**. There are 14 available ports.
  - The valid port range for a 40-Gigabit data plane uplink interface is **0** through **7** on PIC **1**. There are eight available ports.

See [Table 225 on page 2430](#) for physical port to logical port mappings.

- On a QFX3600 switch running Enhanced Layer 2 Software:

- The valid port range for a 10-Gigabit Ethernet interface is **0** through **63** on PIC **0**. There are 64 available ports.
- The valid port range for a 40-Gigabit Ethernet interface is **0** through **15** on PIC **0**. There are 16 available ports.

See [Table 226 on page 2433](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running Enhanced Layer 2 Software:
  - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC **0**. There are 48 available ports. When you channelize the 6 40-Gbps QSFP+ ports on **0** through **5** on PIC **1**, there are 72 available ports.



**NOTE:** On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

- The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC **1**. There are six available ports.

See [Table 228 on page 2438](#) for physical port to logical port mappings.

- On EX4600 switches running Enhanced Layer 2 Software:
  - The valid port range for a 10-Gigabit Ethernet interface is **0** through **23** on PIC **0**. There are 24 available ports. When you channelize the 4 40-Gbps QSFP+ ports on **24** through **27** on PIC **0**. There are 40 available ports.

See [Table 228 on page 2438](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running a QFabric software package:
  - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC **0**. There are 48 available ports.
  - The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC **1**. There are six available ports.



**NOTE:** On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

See [Table 229 on page 2441](#) for physical port to logical port mappings.

- For QFX5100-24Q and QFX5100-96S switches running Enhanced Layer 2 Software, see [Table 230 on page 2445](#) for physical port to logical port mappings for different system modes.



Table 223: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
0	fc-0/0/0	Not supported on this port	xe-0/0/0	Not supported on this port	Not supported on this port
1	fc-0/0/1	Not supported on this port	xe-0/0/1	Not supported on this port	Not supported on this port
2	fc-0/0/2	Not supported on this port	xe-0/0/2	Not supported on this port	Not supported on this port
3	fc-0/0/3	Not supported on this port	xe-0/0/3	Not supported on this port	Not supported on this port
4	fc-0/0/4	Not supported on this port	xe-0/0/4	Not supported on this port	Not supported on this port
5	fc-0/0/5	Not supported on this port	xe-0/0/5	Not supported on this port	Not supported on this port
6	Not supported on this port	ge-0/0/6	xe-0/0/6	Not supported on this port	Not supported on this port
7	Not supported on this port	ge-0/0/7	xe-0/0/7	Not supported on this port	Not supported on this port
8	Not supported on this port	ge-0/0/8	xe-0/0/8	Not supported on this port	Not supported on this port
9	Not supported on this port	ge-0/0/9	xe-0/0/9	Not supported on this port	Not supported on this port
10	Not supported on this port	ge-0/0/10	xe-0/0/10	Not supported on this port	Not supported on this port
11	Not supported on this port	ge-0/0/11	xe-0/0/11	Not supported on this port	Not supported on this port
12	Not supported on this port	ge-0/0/12	xe-0/0/12	Not supported on this port	Not supported on this port
13	Not supported on this port	ge-0/0/13	xe-0/0/13	Not supported on this port	Not supported on this port
14	Not supported on this port	ge-0/0/14	xe-0/0/14	Not supported on this port	Not supported on this port

Table 223: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
15	Not supported on this port	ge-0/0/15	xe-0/0/15	Not supported on this port	Not supported on this port
16	Not supported on this port	ge-0/0/16	xe-0/0/16	Not supported on this port	Not supported on this port
17	Not supported on this port	ge-0/0/17	xe-0/0/17	Not supported on this port	Not supported on this port
18	Not supported on this port	ge-0/0/18	xe-0/0/18	Not supported on this port	Not supported on this port
19	Not supported on this port	ge-0/0/19	xe-0/0/19	Not supported on this port	Not supported on this port
20	Not supported on this port	ge-0/0/20	xe-0/0/20	Not supported on this port	Not supported on this port
21	Not supported on this port	ge-0/0/21	xe-0/0/21	Not supported on this port	Not supported on this port
22	Not supported on this port	ge-0/0/22	xe-0/0/22	Not supported on this port	Not supported on this port
23	Not supported on this port	ge-0/0/23	xe-0/0/23	Not supported on this port	Not supported on this port
24	Not supported on this port	ge-0/0/24	xe-0/0/24	Not supported on this port	Not supported on this port
25	Not supported on this port	ge-0/0/25	xe-0/0/25	Not supported on this port	Not supported on this port
26	Not supported on this port	ge-0/0/26	xe-0/0/26	Not supported on this port	Not supported on this port
27	Not supported on this port	ge-0/0/27	xe-0/0/27	Not supported on this port	Not supported on this port
28	Not supported on this port	ge-0/0/28	xe-0/0/28	Not supported on this port	Not supported on this port
29	Not supported on this port	ge-0/0/29	xe-0/0/29	Not supported on this port	Not supported on this port

**Table 223: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)**

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
30	Not supported on this port	ge-0/0/30	xe-0/0/30	Not supported on this port	Not supported on this port
31	Not supported on this port	ge-0/0/31	xe-0/0/31	Not supported on this port	Not supported on this port
32	Not supported on this port	ge-0/0/32	xe-0/0/32	Not supported on this port	Not supported on this port
33	Not supported on this port	ge-0/0/33	xe-0/0/33	Not supported on this port	Not supported on this port
34	Not supported on this port	ge-0/0/34	xe-0/0/34	Not supported on this port	Not supported on this port
35	Not supported on this port	ge-0/0/35	xe-0/0/35	Not supported on this port	Not supported on this port
36	Not supported on this port	ge-0/0/36	xe-0/0/36	Not supported on this port	Not supported on this port
37	Not supported on this port	ge-0/0/37	xe-0/0/37	Not supported on this port	Not supported on this port
38	Not supported on this port	ge-0/0/38	xe-0/0/38	Not supported on this port	Not supported on this port
39	Not supported on this port	ge-0/0/39	xe-0/0/39	Not supported on this port	Not supported on this port
40	Not supported on this port	ge-0/0/40	xe-0/0/40	Not supported on this port	Not supported on this port
41	Not supported on this port	ge-0/0/41	xe-0/0/41	Not supported on this port	Not supported on this port
42	fc-0/0/42	Not supported on this port	xe-0/0/42	Not supported on this port	Not supported on this port
43	fc-0/0/43	Not supported on this port	xe-0/0/43	Not supported on this port	Not supported on this port
44	fc-0/0/44	Not supported on this port	xe-0/0/44	Not supported on this port	Not supported on this port

Table 223: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (*continued*)

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
45	fc-0/0/45	Not supported on this port	xe-0/0/45	Not supported on this port	Not supported on this port
46	fc-0/0/46	Not supported on this port	xe-0/0/46	Not supported on this port	Not supported on this port
47	fc-0/0/47	Not supported on this port	xe-0/0/47	Not supported on this port	Not supported on this port
Q0	Not supported on this port	Not supported on this port	xe-0/1/0 xe-0/1/1 xe-0/1/2 xe-0/1/3  <i>NOTE:</i> Supported on QFX3500 standalone switch only.	fte-0/1/0	xle-0/2/0
Q1	Not supported on this port	Not supported on this port	xe-0/1/4 xe-0/1/5 xe-0/1/6 xe-0/1/7  <i>NOTE:</i> Supported on QFX3500 standalone switch only.	fte-0/1/1	xle-0/2/1
Q2	Not supported on this port	Not supported on this port	xe-0/1/8 xe-0/1/9 xe-0/1/10 xe-0/1/11  <i>NOTE:</i> Supported on QFX3500 standalone switch only.	fte-0/1/2	xle-0/2/2

**Table 223: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)**

Port Number	Fibre Channel Interfaces (On PIC 0)	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 2)
Q3	Not supported on this port	Not supported on this port	xe-0/1/12 xe-0/1/13 xe-0/1/14 xe-0/1/15  NOTE: Supported on QFX3500 standalone switch only.	fte-0/1/3	xle-0/2/3

**Table 224: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software**

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
0	Not supported on this port	xe-0/0/0	Not supported on this port
1	Not supported on this port	xe-0/0/1	Not supported on this port
2	Not supported on this port	xe-0/0/2	Not supported on this port
3	Not supported on this port	xe-0/0/3	Not supported on this port
4	Not supported on this port	xe-0/0/4	Not supported on this port
5	Not supported on this port	xe-0/0/5	Not supported on this port
6	ge-0/0/6	xe-0/0/6	Not supported on this port
7	ge-0/0/7	xe-0/0/7	Not supported on this port
8	ge-0/0/8	xe-0/0/8	Not supported on this port
9	ge-0/0/9	xe-0/0/9	Not supported on this port
10	ge-0/0/10	xe-0/0/10	Not supported on this port
11	ge-0/0/11	xe-0/0/11	Not supported on this port
12	ge-0/0/12	xe-0/0/12	Not supported on this port

Table 224: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
13	ge-0/0/13	xe-0/0/13	Not supported on this port
14	ge-0/0/14	xe-0/0/14	Not supported on this port
15	ge-0/0/15	xe-0/0/15	Not supported on this port
16	ge-0/0/16	xe-0/0/16	Not supported on this port
17	ge-0/0/17	xe-0/0/17	Not supported on this port
18	ge-0/0/18	xe-0/0/18	Not supported on this port
19	ge-0/0/19	xe-0/0/19	Not supported on this port
20	ge-0/0/20	xe-0/0/20	Not supported on this port
21	ge-0/0/21	xe-0/0/21	Not supported on this port
22	ge-0/0/22	xe-0/0/22	Not supported on this port
23	ge-0/0/23	xe-0/0/23	Not supported on this port
24	ge-0/0/24	xe-0/0/24	Not supported on this port
25	ge-0/0/25	xe-0/0/25	Not supported on this port
26	ge-0/0/26	xe-0/0/26	Not supported on this port
27	ge-0/0/27	xe-0/0/27	Not supported on this port
28	ge-0/0/28	xe-0/0/28	Not supported on this port
29	ge-0/0/29	xe-0/0/29	Not supported on this port
30	ge-0/0/30	xe-0/0/30	Not supported on this port
31	ge-0/0/31	xe-0/0/31	Not supported on this port
32	ge-0/0/32	xe-0/0/32	Not supported on this port
33	ge-0/0/33	xe-0/0/33	Not supported on this port
34	ge-0/0/34	xe-0/0/34	Not supported on this port

**Table 224: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)**

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
35	ge-0/0/35	xe-0/0/35	Not supported on this port
36	ge-0/0/36	xe-0/0/36	Not supported on this port
37	ge-0/0/37	xe-0/0/37	Not supported on this port
38	ge-0/0/38	xe-0/0/38	Not supported on this port
39	ge-0/0/39	xe-0/0/39	Not supported on this port
40	ge-0/0/40	xe-0/0/40	Not supported on this port
41	ge-0/0/41	xe-0/0/41	Not supported on this port
42	Not supported on this port	xe-0/0/42	Not supported on this port
43	Not supported on this port	xe-0/0/43	Not supported on this port
44	Not supported on this port	xe-0/0/44	Not supported on this port
45	Not supported on this port	xe-0/0/45	Not supported on this port
46	Not supported on this port	xe-0/0/46	Not supported on this port
47	Not supported on this port	xe-0/0/47	Not supported on this port
Q0	Not supported on this port	xe-0/1/0:0 xe-0/1/0:1 xe-0/1/0:2 xe-0/1/0:3	et-0/1/0
Q1	Not supported on this port	xe-0/1/1:0 xe-0/1/1:1 xe-0/1/1:2 xe-0/1/1:3	et-0/1/1

**Table 224: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (*continued*)**

Port Number	Gigabit Ethernet Interfaces (On PIC 0)	10-Gigabit Ethernet Interfaces (On PIC 0 and 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q2	Not supported on this port	xe-0/1/2:0 xe-0/1/2:1 xe-0/1/2:2 xe-0/1/2:3	et-0/1/2
Q3	Not supported on this port	xe-0/1/3:0 xe-0/1/3:1 xe-0/1/3:2 xe-0/1/3:3	et-0/1/3

**Table 225: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	xe-0/0/0 xe-0/0/1 xe-0/0/2 xe-0/0/3	xle-0/1/0
Q1	xe-0/0/4 xe-0/0/5 xe-0/0/6 xe-0/0/7	xle-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	xle-0/1/2



Table 225: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q3	xe-0/0/12	xle-0/1/3
	xe-0/0/13	
	xe-0/0/14	
	xe-0/0/15	
Q4	xe-0/0/16	xle-0/1/4
	xe-0/0/17	
	xe-0/0/18	
	xe-0/0/19	
Q5	xe-0/0/20	xle-0/1/5
	xe-0/0/21	
	xe-0/0/22	
	xe-0/0/23	
Q6	xe-0/0/24	xle-0/1/6
	xe-0/0/25	
	xe-0/0/26	
	xe-0/0/27	
Q7	xe-0/0/28	xle-0/1/7
	xe-0/0/29	
	xe-0/0/30	
	xe-0/0/31	
Q8	xe-0/0/32	xle-0/1/8
	xe-0/0/33	
	xe-0/0/34	
	xe-0/0/35	

Table 225: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q9	xe-0/0/36	xle-0/1/9
	xe-0/0/37	
	xe-0/0/38	
	xe-0/0/39	
Q10	xe-0/0/40	xle-0/1/10
	xe-0/0/41	
	xe-0/0/42	
	xe-0/0/43	
Q11	xe-0/0/44	xle-0/1/11
	xe-0/0/45	
	xe-0/0/46	
	xe-0/0/47	
Q12	xe-0/0/48	xle-0/1/12
	xe-0/0/49	
	xe-0/0/50	
	xe-0/0/51	
Q13	xe-0/0/52	xle-0/1/13
	xe-0/0/53	
	xe-0/0/54	
	xe-0/0/55	
Q14	xe-0/0/56	xle-0/1/14
	xe-0/0/57	
	xe-0/0/58	
	xe-0/0/59	

**Table 225: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (*continued*)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q15	xe-0/0/60	xle-0/1/15
	xe-0/0/61	
	xe-0/0/62	
	xe-0/0/63	

**Table 226: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q0	xe-0/0/0:0	et-0/0/0
	xe-0/0/0:1	
	xe-0/0/0:2	
	xe-0/0/0:3	
Q1	xe-0/0/1:0	et-0/0/1
	xe-0/0/1:1	
	xe-0/0/1:2	
	xe-0/0/1:3	
Q2	xe-0/0/2:0	et-0/0/2
	xe-0/0/2:1	
	xe-0/0/2:2	
	xe-0/0/2:3	
Q3	xe-0/0/3:0	et-0/0/3
	xe-0/0/3:1	
	xe-0/0/3:2	
	xe-0/0/3:3	

Table 226: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q4	xe-0/0/4:0	et-0/0/4
	xe-0/0/4:1	
	xe-0/0/4:2	
	xe-0/0/4:3	
Q5	xe-0/0/5:0	et-0/0/5
	xe-0/0/5:1	
	xe-0/0/5:2	
	xe-0/0/5:3	
Q6	xe-0/0/6:0	et-0/0/6
	xe-0/0/6:1	
	xe-0/0/6:2	
	xe-0/0/6:3	
Q7	xe-0/0/7:0	et-0/0/7
	xe-0/0/7:1	
	xe-0/0/7:2	
	xe-0/0/7:3	
Q8	xe-0/0/8:0	et-0/0/8
	xe-0/0/8:1	
	xe-0/0/8:2	
	xe-0/0/8:3	
Q9	xe-0/0/9:0	et-0/0/9
	xe-0/0/9:1	
	xe-0/0/9:2	
	xe-0/0/9:3	

**Table 226: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (*continued*)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
Q10	xe-0/0/10:0	et-0/0/10
	xe-0/0/10:1	
	xe-0/0/10:2	
	xe-0/0/10:3	
Q11	xe-0/0/11:0	et-0/0/11
	xe-0/0/11:1	
	xe-0/0/11:2	
	xe-0/0/11:3	
Q12	xe-0/0/12:0	et-0/0/12
	xe-0/0/12:1	
	xe-0/0/12:2	
	xe-0/0/12:3	
Q13	xe-0/0/13:0	et-0/0/13
	xe-0/0/13:1	
	xe-0/0/13:2	
	xe-0/0/13:3	
Q14	xe-0/0/14:0	et-0/0/14
	xe-0/0/14:1	
	xe-0/0/14:2	
	xe-0/0/14:3	
Q15	xe-0/0/15:0	et-0/0/15
	xe-0/0/15:1	
	xe-0/0/15:2	
	xe-0/0/15:3	

Table 227: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q0	Not supported on this port	fte-0/1/0	xle-0/1/0
Q1	Not supported on this port	fte-0/1/1	xle-0/1/1
Q2	xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11	fte-0/1/2	xle-0/1/2
Q3	xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15	fte-0/1/3	xle-0/1/3
Q4	xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19	fte-0/1/4	xle-0/1/4
Q5	xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23	fte-0/1/5	xle-0/1/5
Q6	xe-0/0/24 xe-0/0/25 xe-0/0/26 xe-0/0/27	fte-0/1/6	xle-0/1/6
Q7	xe-0/0/28 xe-0/0/29 xe-0/0/30 xe-0/0/31	fte-0/1/7	xle-0/1/7

**Table 227: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (*continued*)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q8	xe-0/0/32	Not supported on this port	xle-0/1/8
	xe-0/0/33		
	xe-0/0/34		
	xe-0/0/35		
Q9	xe-0/0/36	Not supported on this port	xle-0/1/9
	xe-0/0/37		
	xe-0/0/38		
	xe-0/0/39		
Q10	xe-0/0/40	Not supported on this port	xle-0/1/10
	xe-0/0/41		
	xe-0/0/42		
	xe-0/0/43		
Q11	xe-0/0/44	Not supported on this port	xle-0/1/11
	xe-0/0/45		
	xe-0/0/46		
	xe-0/0/47		
Q12	xe-0/0/48	Not supported on this port	xle-0/1/12
	xe-0/0/49		
	xe-0/0/50		
	xe-0/0/51		
Q13	xe-0/0/52	Not supported on this port	xle-0/1/13
	xe-0/0/53		
	xe-0/0/54		
	xe-0/0/55		

**Table 227: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (continued)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)	40-Gigabit Ethernet Interfaces (On PIC 1)
Q14	xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59	Not supported on this port	xle-0/1/14
Q15	xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63	Not supported on this port	xle-0/1/15

**Table 228: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
0	xe-0/0/0	Not supported on this port
1	xe-0/0/1	Not supported on this port
2	xe-0/0/2	Not supported on this port
3	xe-0/0/3	Not supported on this port
4	xe-0/0/4	Not supported on this port
5	xe-0/0/5	Not supported on this port
6	xe-0/0/6	Not supported on this port
7	xe-0/0/7	Not supported on this port
8	xe-0/0/8	Not supported on this port
9	xe-0/0/9	Not supported on this port
10	xe-0/0/10	Not supported on this port
11	xe-0/0/11	Not supported on this port



**Table 228: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (*continued*)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
12	xe-0/0/12	Not supported on this port
13	xe-0/0/13	Not supported on this port
14	xe-0/0/14	Not supported on this port
15	xe-0/0/15	Not supported on this port
16	xe-0/0/16	Not supported on this port
17	xe-0/0/17	Not supported on this port
18	xe-0/0/18	Not supported on this port
19	xe-0/0/19	Not supported on this port
20	xe-0/0/20	Not supported on this port
21	xe-0/0/21	Not supported on this port
22	xe-0/0/22	Not supported on this port
23	xe-0/0/23	Not supported on this port
24	xe-0/0/24	Not supported on this port
25	xe-0/0/25	Not supported on this port
26	xe-0/0/26	Not supported on this port
27	xe-0/0/27	Not supported on this port
28	xe-0/0/28	Not supported on this port
29	xe-0/0/29	Not supported on this port
30	xe-0/0/30	Not supported on this port
31	xe-0/0/31	Not supported on this port
32	xe-0/0/32	Not supported on this port
33	xe-0/0/33	Not supported on this port

**Table 228: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (*continued*)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
34	xe-0/0/34	Not supported on this port
35	xe-0/0/35	Not supported on this port
36	xe-0/0/36	Not supported on this port
37	xe-0/0/37	Not supported on this port
38	xe-0/0/38	Not supported on this port
39	xe-0/0/39	Not supported on this port
40	xe-0/0/40	Not supported on this port
41	xe-0/0/41	Not supported on this port
42	xe-0/0/42	Not supported on this port
43	xe-0/0/43	Not supported on this port
44	xe-0/0/44	Not supported on this port
45	xe-0/0/45	Not supported on this port
46	xe-0/0/46	Not supported on this port
47	xe-0/0/47	Not supported on this port
48	xe-0/0/48:0 xe-0/0/48:1 xe-0/0/48:2 xe-0/0/48:3	et-0/1/0
49	xe-0/0/49:0 xe-0/0/49:1 xe-0/0/49:2 xe-0/0/49:3	et-0/1/1

**Table 228: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (*continued*)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 0)
50	xe-0/0/50:0 xe-0/0/50:1 xe-0/0/50:2 xe-0/0/50:3	et-0/1/2
51	xe-0/0/51:0 xe-0/0/51:1 xe-0/0/51:2 xe-0/0/51:3	et-0/1/3
52	xe-0/0/52:0 xe-0/0/52:1 xe-0/0/52:2 xe-0/0/52:3	et-0/1/4
53	xe-0/0/53:0 xe-0/0/53:1 xe-0/0/53:2 xe-0/0/53:3	et-0/1/5

**Table 229: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
0	xe-0/0/0	Not supported on this port	Not supported on this port
1	xe-0/0/1	Not supported on this port	Not supported on this port
2	xe-0/0/2	Not supported on this port	Not supported on this port
3	xe-0/0/3	Not supported on this port	Not supported on this port
4	xe-0/0/4	Not supported on this port	Not supported on this port

Table 229: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
5	xe-0/0/5	Not supported on this port	Not supported on this port
6	xe-0/0/6	Not supported on this port	Not supported on this port
7	xe-0/0/7	Not supported on this port	Not supported on this port
8	xe-0/0/8	Not supported on this port	Not supported on this port
9	xe-0/0/9	Not supported on this port	Not supported on this port
10	xe-0/0/10	Not supported on this port	Not supported on this port
11	xe-0/0/11	Not supported on this port	Not supported on this port
12	xe-0/0/12	Not supported on this port	Not supported on this port
13	xe-0/0/13	Not supported on this port	Not supported on this port
14	xe-0/0/14	Not supported on this port	Not supported on this port
15	xe-0/0/15	Not supported on this port	Not supported on this port
16	xe-0/0/16	Not supported on this port	Not supported on this port
17	xe-0/0/17	Not supported on this port	Not supported on this port
18	xe-0/0/18	Not supported on this port	Not supported on this port
19	xe-0/0/19	Not supported on this port	Not supported on this port
20	xe-0/0/20	Not supported on this port	Not supported on this port
21	xe-0/0/21	Not supported on this port	Not supported on this port
22	xe-0/0/22	Not supported on this port	Not supported on this port
23	xe-0/0/23	Not supported on this port	Not supported on this port
24	xe-0/0/24	Not supported on this port	Not supported on this port
25	xe-0/0/25	Not supported on this port	Not supported on this port
26	xe-0/0/26	Not supported on this port	Not supported on this port

**Table 229: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)**

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
27	xe-0/0/27	Not supported on this port	Not supported on this port
28	xe-0/0/28	Not supported on this port	Not supported on this port
29	xe-0/0/29	Not supported on this port	Not supported on this port
30	xe-0/0/30	Not supported on this port	Not supported on this port
31	xe-0/0/31	Not supported on this port	Not supported on this port
32	xe-0/0/32	Not supported on this port	Not supported on this port
33	xe-0/0/33	Not supported on this port	Not supported on this port
34	xe-0/0/34	Not supported on this port	Not supported on this port
35	xe-0/0/35	Not supported on this port	Not supported on this port
36	xe-0/0/36	Not supported on this port	Not supported on this port
37	xe-0/0/37	Not supported on this port	Not supported on this port
38	xe-0/0/38	Not supported on this port	Not supported on this port
39	xe-0/0/39	Not supported on this port	Not supported on this port
40	xe-0/0/40	Not supported on this port	Not supported on this port
41	xe-0/0/41	Not supported on this port	Not supported on this port
42	xe-0/0/42	Not supported on this port	Not supported on this port
43	xe-0/0/43	Not supported on this port	Not supported on this port
44	xe-0/0/44	Not supported on this port	Not supported on this port
45	xe-0/0/45	Not supported on this port	Not supported on this port
46	xe-0/0/46	Not supported on this port	Not supported on this port
47	xe-0/0/47	Not supported on this port	Not supported on this port

Table 229: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (*continued*)

Port Number	10-Gigabit Ethernet Interfaces (On PIC 0)	40-Gigabit Ethernet Interfaces (On PIC 1)	40-Gigabit Data Plane Uplink Interfaces (On PIC 1)
48	Not supported on this port	Not supported on this PIC	fte-0/1/0  <i>NOTE:</i> This interface is a fixed fte interface and cannot be changed to xle.
49	Not supported on this port	Not supported on this PIC	fte-0/1/1  <i>NOTE:</i> This interface is a fixed fte interface and cannot be changed to xle.
50	Not supported on this port	xle-0/1/2	fte-0/1/2  <i>NOTE:</i> By default, this interface is an fte interface but can be configured as an xle interface.
51	Not supported on this port	xle-0/1/3	fte-0/1/3  <i>NOTE:</i> By default, this interface is an fte interface but can be configured as an xle interface.
52	Not supported on this port	xle-0/1/4  <i>NOTE:</i> By default, this interface is an xle interface but can be configured as an fte interface.	fte-0/1/4
53	Not supported on this port	xle-0/1/5  <i>NOTE:</i> By default, this interface is an xle interface but can be configured as an fte interface.	fte-0/1/5

### Supported System Modes



*NOTE:* There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches depending on the system mode you configure. If you try to channelize ports that are restricted, the configuration is ignored.

The following system modes are available on the QFX5100-24Q switch:

- Default mode
- Mode-104-port
- Flexi-PIC mode
- Non-oversubscribed mode

See [Table 230 on page 2445](#) for more information regarding the supported system modes for your switch.

The following system modes are available on the QFX5100-96S switch:

- Default-mode
- Non-oversubscribed mode

See [Table 230 on page 2445](#) for more information regarding the supported system modes for your switch.

**Table 230: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software**

	Default-mode	Mode-104port	Flexi-pic-mode	Non-oversubscribed-mode
QFX5100-48S and QFX5100-48T	Not supported	Not supported	Not supported	Not supported
QFX5100-24Q	Supported  You do not need to configure the switch to be in this mode. On PIC 0, you can channelize all 24 40-Gbps QSFP+ ports. On PIC 1 and PIC 2, the 40-Gbps QSFP+ ports in the expansion modules are supported but cannot be channelized. In this mode, you can have one of two port combinations: 32 40-Gbps QSFP+ ports, or 96 10-Gigabit Ethernet ports plus 8 40-Gbps QSFP+ ports.	Supported  On PIC 0, all 24 40-Gbps QSFP+ ports are channelized by default, which provides 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in an expansion module on PIC 1 are supported. On PIC 1, ports 0 and 2 are channelized by default, and ports 1 and 3 are disabled. If 40-Gbps QSFP+ ports contained in an expansion module are detected on PIC 2, they are ignored.	Supported  On PIC 0, the first four ports (ports 0 through 3) cannot be channelized. 40-Gbps QSFP+ ports contained in expansion modules on PIC 1 and PIC 2 are supported but cannot be channelized.	Supported  All 24 40-Gbps QSFP+ ports on PIC 0 can be channelized to 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in the expansion modules on PIC 1 and PIC 2 are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.

**Table 230: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software (*continued*)**

	Default-mode	Mode-104port	Flexi-pic-mode	Non-oversubscribed-mode
QFX5100-96S	Supported  You do not need to configure the switch to be in this mode. On PIC 0, all 96 10-Gigabit Ethernet ports are supported. You can only channelize the 40-Gbps QSFP+ interfaces to 10-Gigabit Ethernet interfaces on ports 96 and 100. When you channelize the interfaces on ports 96 and 100, ports 97, 98, 99, 101, 102 and 103 are disabled.	Not supported	Not supported	Supported  On PIC 0, all 96 10-Gigabit Ethernet ports are supported. However, the eight 40-Gbps QSFP+ ports are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.

- Related Documentation**
- [Interfaces Overview on page 2389](#)
  - [Channelizing Interfaces on page 2608](#)
  - [Configuring the System Mode on page 2610](#)
  - [Understanding Interface Naming Conventions on page 2401](#)
  - *Rear Panel of a QFX3500 Device*
  - *Front Panel of a QFX3600 Device*



## Understanding Redundant Trunk Links

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.

To configure a redundant trunk link, create a redundant trunk group. The redundant trunk group is configured on the access switch and contains two links: a primary or active link, and a secondary link. If the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence.

Data traffic is forwarded only on the active link. Data traffic on the secondary link is dropped and shown as dropped packets when you issue the operational mode command **show interfaces *interface-name* extensive**.

While data traffic is blocked on the secondary link, Layer 2 control traffic is still permitted. For example, an LLDP session can be run between two switches on the secondary link.

Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You must disable RSTP on an interface if a redundant trunk group is configured on that interface. For example, in [Figure 48 on page 2448](#), in addition to disabling RSTP on the Switch 3 interfaces, you must also disable RSTP on the Switch 1 and Switch 2 interfaces connected to Switch 3. Spanning-tree protocols can, however, continue operating on other interfaces on those switches—for example on the link between Switch 1 and Switch 2.

[Figure 48 on page 2448](#) shows three switches in a basic topology for redundant trunk links. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2). Link 1 and Link 2 are in a redundant trunk group called group1. Link 1 is designated as the primary link. Traffic flows between Switch 3 in the access layer and Switch 1 in the distribution layer through Link 1. While Link 1 is active, Link 2 blocks traffic.

Figure 48: Redundant Trunk Group, Link 1 Active





Figure 49 on page 2448 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 49: Redundant Trunk Group, Link 2 Active



When Link 1 between Switch 1 and Switch 3 goes down, Link 2 takes over as the active link. Traffic between the access layer and the distribution layer is then automatically switched to Link 2 between Switch 1 and Switch 2.

**Related  
Documentation**

- [Example: Configuring Redundant Trunk Links for Faster Recovery](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578](#)

## Understanding Generic Routing Encapsulation

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets.

This topic describes:

- [Overview of GRE on page 2449](#)
- [GRE Tunneling on page 2449](#)
- [Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch on page 2451](#)
- [Configuration Limitations on page 2452](#)

### Overview of GRE

GRE encapsulates data packets and redirects them to a device that de-encapsulates them and routes them to their final destination. This allows the source and destination switches to operate as if they have a virtual point-to-point connection with each other (because the outer header applied by GRE is transparent to the encapsulated payload packet). For example, GRE tunnels allow routing protocols such as RIP and OSPF to forward data packets from one switch to another switch across the Internet. In addition, GRE tunnels can encapsulate multicast data streams for transmission over the Internet.

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see [“Configuration Limitations” on page 2452.](#))

As a *tunnel source router*, the switch encapsulates a payload packet for transport through the tunnel to a destination network. The payload packet is first encapsulated in a GRE packet, and then the GRE packet is encapsulated in a delivery protocol. The switch performing the role of a *tunnel remote router* extracts the tunneled packet and forwards the packet to its destination. Note that you can use one firewall term to terminate many GRE tunnels on a QFX5100 switch.

### GRE Tunneling

Data is routed by the system to the GRE endpoint over routes established in the route table. (These routes can be statically configured or dynamically learned by routing protocols such as RIP or OSPF.) When a data packet is received by the GRE endpoint, it is de-encapsulated and routed again to its destination address.

GRE tunnels are *stateless*—that is, the endpoint of the tunnel contains no information about the state or availability of the remote tunnel endpoint. Therefore, the switch

operating as a tunnel source router cannot change the state of the GRE tunnel interface to down if the remote endpoint is unreachable.

For details about GRE tunneling, see:

- [Encapsulation and De-Encapsulation on the Switch on page 2450](#)
- [Number of Source and Destination Tunnels Allowed on a Switch on page 2450](#)
- [Class of Service on GRE Tunnels on page 2450](#)
- [Applying Firewall Filters to GRE Traffic on page 2451](#)

### ***Encapsulation and De-Encapsulation on the Switch***

Encapsulation—A switch operating as a tunnel source router encapsulates and forwards GRE packets as follows:

1. When a switch receives a data packet (payload) to be tunneled, it sends the packet to the tunnel interface.
2. The tunnel interface encapsulates the data in a GRE packet and adds an outer IP header.
3. The IP packet is forwarded on the basis of the destination address in the outer IP header.

De-encapsulation—A switch operating as a tunnel remote router handles GRE packets as follows:

1. When the destination switch receives the IP packet from the tunnel interface, the outer IP header and GRE header are removed.
2. The packet is routed based on the inner IP header.

### ***Number of Source and Destination Tunnels Allowed on a Switch***

QFX5100 switches support as many as 512 GRE tunnels, including tunnels created with a firewall filter. That is, you can create a total of 512 GRE tunnels, regardless of which method you use.

EX switches support as many as 500 GRE tunnels between switches transmitting IPv4 or IPv6 payload packets over GRE. If a passenger protocol in addition to IPv4 and IPv6 is used, you can configure up to 333 GRE tunnels between the switches.

An EX switch can have a maximum of 20 tunnel source IP addresses configured, and each tunnel source IP can be configured with up to 20 destination IP addresses on a second switch. As a result, the two connected switches can have a maximum of 400 GRE tunnels. If the first switch is also connected to a third switch, the possible maximum number of tunnels is 500.

### ***Class of Service on GRE Tunnels***

When a network experiences congestion and delay, some packets might be dropped. Junos OS class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs and thereby set rules for packet loss. For details about CoS, see [Junos OS CoS for EX Series Switches Overview](#).

The following CoS components are available on a switch operating as a GRE tunnel source router or GRE tunnel remote router:

- At the GRE tunnel source—On a switch operating as a tunnel source router, you can apply CoS classifiers on an *ingress port* or on a *GRE port*, with the following results on CoS component support on tunneled packets:
  - Schedulers only—Based on the CoS classification on the ingress port, you can apply CoS schedulers on a GRE port of the switch to define output queues and control the transmission of packets through the tunnel after GRE encapsulation. However, you cannot apply CoS rewrite rules to these packets.
  - Schedulers and rewrite rules—Depending on the CoS classification on the GRE port, you can apply both schedulers and rewrite rules to the encapsulated packets transmitted through the tunnel.
- At the GRE tunnel endpoint—When the switch is a tunnel remote router, you can apply CoS classifiers on the GRE port and schedulers and rewrite rules on the egress port to control the transmission of a de-encapsulated GRE packet out from the egress port.

#### ***Applying Firewall Filters to GRE Traffic***

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch. (For details, see [Firewall Filters for EX Series Switches Overview](#).) Because of the encapsulation and de-encapsulation performed by GRE, you are constrained as to where you can apply a firewall filter to filter tunneled packets and which header will be affected. [Table 231 on page 2451](#) identifies these constraints.

**Table 231: Firewall Filter Application Points for Tunneled Packets**

Endpoint Type	Ingress Interface	Egress Interface
Source (encapsulating)	inner header	outer header
Remote (de-encapsulating)	Cannot filter packets on ingress interface	inner header

#### **Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch**

You can also use a firewall filter to de-encapsulate GRE traffic on a QFX5100 switch. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. See [“Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch” on page 5301](#) for information about how to configure a firewall filter for this purpose.

## Configuration Limitations

Table 232 on page 2452 lists features that are not supported with GRE.

**Table 232: Features Not Supported with GRE**

EX Switches	QFX Switches
MPLS over GRE tunnels	MPLS over GRE tunnels
GRE keepalives	GRE keepalives
GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets	GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets
BGP dynamic tunnels	BGP dynamic tunnels
Outer IP address must be IPv4	Outer IP address must be IPv4
Virtual routing instances	
Bidirectional Forwarding Detection (BFD) protocol over GRE distributed mode	
<p>OSPF limitation—Enabling OSPF on a GRE interface creates two equal-cost routes to the destination: one through the Ethernet network or uplink interface and the other through the tunnel interface. If data is routed through the tunnel interface, the tunnel might fail. To keep the interface operational, we recommend that you use a static route, disable OSPF on the tunnel interface, or configure the peer not to advertise the tunnel destination over the tunnel interface.</p>	

- Related Documentation**
- [Configuring Generic Routing Encapsulation Tunneling \(CLI Procedure\)](#)
  - [Configuring Generic Routing Encapsulation Tunneling on page 2600](#)
  - [Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch on page 5301](#)

## Understanding Ethernet OAM Link Fault Management

Juniper Networks Junos operating system (Junos OS) for Juniper Networks switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported on switches:

- Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote Fault Detection

Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection. The EX Series switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote Loopback Mode

Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

#### Related Documentation

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)

## Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups

You use resilient hashing to minimize flow remapping across members of a trunk/ECMP group in a load-balanced system. You can configure resilient hashing in link aggregation groups (LAGs) and in equal cost multipath (ECMP) groups.

- [Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing on page 2454](#)
- [Limitations and Caveats for Resilient Hashing on page 2455](#)
- [Resilient Hashing on LAGs on page 2455](#)
- [Resilient Hashing on ECMP on page 2456](#)

### Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing

---

Resilient hashing works in conjunction with the default static hashing algorithm. When members are added to or deleted from a trunk/ECMP group, the static hashing algorithm might remap destination paths. Resilient hashing distributes traffic across all members of a group by tracking the flow's member utilization. When a flow is affected by a member change, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table.

Resilient hashing thus provides the following benefits:

- Minimizes traffic-distribution imbalances among members of a trunk/ECMP group when members are added to or deleted from the group.
- Minimizes the impact on flows bound to unaffected members when a new member is added or an existing member is deleted from the group.

In normal hash-based load balancing, with the static hashing algorithm used alone, flows are assigned to members through the mathematical mod (%) operation. Any increase or decrease in the number of group members results in a complete remapping of flows to member IDs, as shown in the following example:

- Member ID = Hash (key) mod (number of members in group)
- Example:
  - Hash (key) = 10
  - $10 \bmod 5 = 0$  (member with ID 0 is selected for flow)
  - $10 \bmod 4 = 2$  (member with ID 2 is selected for the same flow when the number of members is decreased by 1)

Resilient hashing minimizes the destination path remapping when a member in the trunk/ECMP group is added or deleted.



When the flow is affected by a member change in the group, resilient hashing rebalances the flow by reprogramming the flow set table.

**Table 233: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from Trunk Groups**

Trunk Group Size	Normal (Static) Hashing Result	Resilient Hashing Result	Notes
4	Hash(10) % 4 = 2 Flow is assigned to member ID 2.	Flow is assigned to one of four group members based on flow set table entries.	Original trunk/ECMP group size is 4.
3	Hash(10) % 3 = 1 Flow is assigned to member ID 1.	Flow is assigned to same member as in the previous case.	Delete one member from original trunk/ECMP group. Trunk/ECMP group size is 3.
5	Hash(10) % 5 = 0 Flow is assigned to member ID 0.	There is minimal redistribution of flows from other members to this newly added member.	Add one member to original trunk group. Trunk/ECMP group size is 5.

### Limitations and Caveats for Resilient Hashing

Notice the following limitation and caveats for the resilient hashing feature:

- Resilient hashing applies only to unicast traffic.
- Resilient hashing supports a maximum of 1024 trunk groups, with each group having a maximum of 256 members.
- Resilient hashing does not guarantee that traffic distribution is even across all group members—it depends on the traffic pattern and on the organization of the resilient hashing flow set table in hardware. Resilient hashing *minimizes* remapping of flows to destination links when members are added to or deleted from the group.
- If resilient hashing is enabled on a trunk group or ECMP group and if **set forwarding-options enhanced-hash-key** with one of the options **hash-mode**, **inet**, **inet6**, or **layer2** is used, some flows might change destination links, because the new hash parameters might generate new hash indexes for the flows, and hence the new destination links.
- Resilient hashing is not supported on Virtual Chassis port (VCP) links.

### Resilient Hashing on LAGs

A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient hashing minimizes destination remapping behavior when a new member is added or deleted from the LAG.

A resilient hashing configuration on LAGs is per-aggregated-Ethernet-interface-based.

### Resilient Hashing on ECMP

---

An ECMP group for a route contains multiple next-hop equal cost addresses for the same destination in the routing table. (Routes of equal cost have the same preference and metric values.)

Junos OS uses the static hashing algorithm to choose one of the next-hop addresses in the ECMP group to install in the forwarding table. Resilient hashing enhances ECMPs by minimizing destination remapping behavior when a new member is added or deleted from the ECMP group.

A resilient hashing configuration on ECMP is global—it applies to all ECMP groups.

#### Related Documentation

- [Configuring Resilient Hashing for Trunk/ECMP Groups on page 2607](#)

## CHAPTER 34

# Configuration

- [Configuration Examples on page 2457](#)
- [Configuration Tasks on page 2585](#)
- [Configuration Tasks on page 2607](#)
- [Configuration Statements on page 2612](#)

### Configuration Examples

---

- [Example: Configuring Interfaces for Uplink Failure Detection on page 2457](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466](#)
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2493](#)
- [Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization on page 2530](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\) on page 2551](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)

### Example: Configuring Interfaces for Uplink Failure Detection

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate the failure information to the downlink interfaces. All of the network interface cards (NICs) on a server are configured as being either the primary link or the secondary link and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link to ensure that the traffic on the failed link is not dropped.

This example describes:

- [Requirements on page 2458](#)
- [Overview and Topology on page 2458](#)
- [Configuring Uplink Failure Detection on Both Switches on page 2459](#)
- [Verification on page 2460](#)

---

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.1 or later for the QFX Series
- Two QFX3500 switches
- Two aggregation switches
- One dual-homed server

---

## Overview and Topology

The topology in this example illustrates how to configure uplink failure detection on Switch A and Switch B. Switch A and Switch B are both configured with a link-to-monitor interface (the uplink interface to the aggregation switch) and a link-to-disable interface (the downlink interface to the server). For simplicity, only one group of link-to-monitor interfaces and link-to-disable interfaces is configured for each switch. The server is dual-homed to both Switch A and Switch B. In this scenario, if the link-to-monitor interface to Switch A is disabled, the server uses the link-to-monitor interface to Switch B instead.



**NOTE:** This example does not describe how to configure the dual-homed server or the aggregation switches. Please refer to the documentation for each of these devices for more information.

---

[Figure 45 on page 2392](#) illustrates a typical setup for uplink failure detection.

Figure 50: Uplink Failure Detection Configuration on Switches




Table 234 on page 2459 lists uplink failure settings for each QFX3500 switch.

Table 234: Settings for Uplink Failure Protection Example

Switch A	Switch B
<ul style="list-style-type: none"> <li>Group name: Group1</li> <li>Link-to-monitor interface: <b>xe-0/0/0</b></li> <li>Link-to-disable interface: <b>xe-0/0/1</b></li> </ul>	<ul style="list-style-type: none"> <li>Group name: Group2</li> <li>Link-to-monitor interface: <b>xe-0/0/0</b></li> <li>Link-to-disable interface: <b>xe-0/0/1</b></li> </ul>

### Configuring Uplink Failure Detection on Both Switches

To configure uplink failure detection on both switches, perform these tasks:

#### CLI Quick Configuration

To quickly configure uplink failure protection on Switch A and Switch B, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]
set uplink-failure-detection group group1
set uplink-failure-detection group group2
set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
set uplink-failure-detection group group1 link-to-disable xe-0/0/1
set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

#### Step-by-Step Procedure

To configure uplink failure protection on both switches:

- Specify a name for the uplink failure detection group on Switch A:
 

```
[edit protocols]
user@switch# set uplink-failure-detection group group1
```
- Add an uplink interface to the group on Switch A:
 

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
```

3. Add a downlink interface to the group on Switch A:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-disable xe-0/0/1
```
4. Specify a name for the uplink failure detection group on Switch B:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group2
```
5. Add an uplink interface to the group on Switch B:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
```
6. Add a downlink interface to the group on Switch B:  

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

**Results** Display the results of the configuration:

```
uplink-failure-detection {
  group {
    group1 {
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
    group2 {
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
  }
}
```

---

### Verification

To verify that uplink failure detection is working correctly, perform the following tasks on Switch A and Switch B:

- [Verifying That Uplink Failure Detection is Working Correctly on page 2460](#)

#### ***Verifying That Uplink Failure Detection is Working Correctly***

**Purpose** Verify that the switch disables the downlink interface when it detects an uplink failure.

- Action** 1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : xe-0/0/0*
Downlink             : xe-0/0/1*
Failure Action       : Inactive
```



**NOTE:** The asterisk (\*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : xe-0/0/0
Downlink             : xe-0/0/1
Failure Action       : Active
```

**Meaning** The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 2392](#)
  - [Configuring Interfaces for Uplink Failure Detection on page 2592](#)
  - [Verifying That Uplink Failure Detection Is Working Correctly](#)

## Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch

A QFX Series product allows you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. You can configure LAGs to connect a QFX Series product to other switches, like aggregation switches, servers, or routers. This example describes how to configure LAGs to connect a QFX3500, QFX3600, or QFX5100 switch to an aggregation switch.

- [Requirements on page 2462](#)
- [Overview and Topology on page 2462](#)
- [Configuration on page 2463](#)
- [Verification on page 2465](#)
- [Troubleshooting on page 2466](#)

---

### Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX3500 and QFX3600 switches, and Junos OS 13.2 or later for the QFX5100 switch.
- One QFX3500, QFX3600, or QFX5100 switch.

---

### Overview and Topology

In this example, the switch has one LAG comprising two 10-Gigabit Ethernet interfaces. This LAG is configured in port mode trunk so that the switch and the VLAN to which it has been assigned can send and receive traffic.

Configuring the Ethernet interfaces as LAGs has the following advantages:

- If one physical port is lost for any reason (a cable is unplugged or a switch port fails), the logical port transparently continues to function over the remaining physical port.
- Link Aggregation Control Protocol (LACP) can optionally be configured for link monitoring and automatic addition and deletion of individual links without user intervention.



**NOTE:** If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

---



The topology used in this example consists of one switch with a LAG configured between two of its 10-Gigabit Ethernet interfaces. The switch is connected to an aggregation switch.




Table 235 on page 2463 details the topology used in this configuration example.

**Table 235: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch**

Hostname	Base Hardware	Trunk Port
switch	QFX3500, QFX3600, or QFX5100 switch	ae0 is configured as a trunk port and combines the following two interfaces: xe-0/0/2 and xe-0/0/3 .

### Configuration

To configure a LAG between two 10-Gigabit Ethernet interfaces:

#### CLI Quick Configuration

To quickly configure a LAG between two 10-Gigabit Ethernet interfaces on a switch, copy the following commands and paste them into the switch terminal window:



**NOTE:** If you are configuring a LAG on the QFX5100 switch, use the interface-mode statement instead of the port-mode statement. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 43.

```
[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family ethernet-switching vlan members green
set interfaces xe-0/0/2 ether-options 802.ad ae0
set interfaces xe-0/0/3 ether-options 802.ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
```

#### Step-by-Step Procedure

To configure a LAG between a QFX Series switch and an aggregation switch:

- Specify the number of LAGs to be created on the switch:  
[edit chassis]  
user@switch# **set aggregated-devices ethernet device-count 1**
- Specify the number of links that need to be present for the ae0 LAG interface to be up:

- ```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 1
```
3. Specify the media speed of the ae0 link:
 

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed 10g
```
  4. Specify the members to be included within the aggregated Ethernet bundle:
 

```
[edit interfaces]
user@switch# set interfaces xe-0/0/2 ether-options 802.ad ae0
[edit interfaces]
user@switch# set interfaces xe-0/0/3 ether-options 802.ad ae0
```
  5. Assign a port mode of trunk to the ae0 link:



**NOTE:** If you are configuring a LAG on the QFX5100 switch, use the `interface-mode` statement instead of the `port-mode` statement. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

- ```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
or
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```
6. Assign the LAG to a VLAN:
 

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members green vlan-id 200
```
  7. (Optional): Designate one side of the LAG as active for LACP:
 

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active
```
  8. (Optional): Designate the interval and speed at which the interfaces send LACP packets:
 

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp periodic fast
```

## Results

Display the results of the configuration on a QFX3500 or QFX3600 switch:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
green {
  vlan-id 200;
}
}
interfaces {
```

```

ae0 {
  aggregated-ether-options {
    link-speed 10g;
    minimum-links 1;
  }
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members green;
      }
    }
  }
  xe-0/0/2 {
    ether-options {
      802.ad ae0;
    }
  }
  xe-0/0/3 {
    ether-options {
      802.ad ae0;
    }
  }
}

```

### Verification

To verify that switching is operational and one LAG has been created, perform these tasks:

- [Verifying That LAG ae0.0 Has Been Created on page 2465](#)
- [Verifying That LAG ae0 Has Been Created on page 2465](#)

#### *Verifying That LAG ae0.0 Has Been Created*

**Purpose** Verify that LAG ae0.0 has been created on the switch.

**Action** show interfaces ae0 terse

Interface	Admin	Link	Proto	Local	Remote
ae0	up	up			
ae0.0	up	up	eth-switch		

**Meaning** The output confirms that the ae0.0 link is up and shows the **family** and IP address assigned to this link.

#### *Verifying That LAG ae0 Has Been Created*

**Purpose** Verify that LAG ae0 has been created on the switch

**Action** show interfaces ae0 terse

Interface	Admin	Link	Proto	Local	Remote
ae0	up	down			
ae0.0	up	down	eth-switch		

**Meaning** The output shows that the **ae0.0** link is down.

### [Troubleshooting](#)

---

#### *Troubleshooting a LAG That Is Down*

**Problem** The **show interfaces terse** command shows that the LAG is **down**.

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.

**Related  
Documentation**

- [Configuring Link Aggregation on page 2593](#)
- [Verifying the Status of a LAG Interface on page 2750](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2751](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [show lacp statistics interfaces \(View\) on page 2875](#)

### **Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch**

QFX Series products allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. On a standalone switch, you can group up to 32 Ethernet interfaces to form a LAG. On a QFabric system, you can group up to 8 Ethernet interfaces to form a LAG. QFX Series products allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 2462](#):

- [Requirements on page 2467](#)
- [Overview and Topology on page 2467](#)
- [Configuring LACP for the LAG on the QFX Series on page 2467](#)

- [Verification on page 2468](#)
- [Troubleshooting on page 2469](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX3500 switch, Junos OS Release 12.1 or later for the QFX3600 switch, and Junos OS 13.2 or later for the QFX5100 switch.
- One QFX3500, QFX3600, or QFX5100 switch.

Before you configure LACP, be sure you have:

- Configured the ports on the switches as trunk ports.
- Configured the LAG.

## Overview and Topology

The topology in this example is exactly the same as the topology used in the [Configuring a LAG Between a QFX Switch and an Aggregation Switch](#) example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP mode can be either active or passive.



**NOTE:** If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

The interval can be fast (every second) or slow (every 30 seconds).

## Configuring LACP for the LAG on the QFX Series

To configure LACP for a QFX Series LAG, perform these tasks:

### CLI Quick Configuration

To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
```

**Step-by-Step Procedure**

To configure LACP for LAG ae0 :

1. Specify the aggregated Ethernet options for the LAG:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active periodic fast
```

**Results** Display the results of the configuration:

```
[edit interfaces]
user@switch# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

---

**Verification**

To verify that LACP packets are being exchanged, perform the following tasks:

- [Verifying the LACP Settings on page 2468](#)
- [Verifying That the LACP Packets Are Being Exchanged on page 2468](#)

**Verifying the LACP Settings**

**Purpose** Verify that LACP has been set up correctly.

**Action** Use the `show lacp interfaces interface-name` command to check that LACP has been enabled as active on one end.

```
user@switch> show lacp interfaces xe-0/0/2
```

Aggregated interface: ae0

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/2	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/0/2	Partner	No	Yes	No	No	No	Yes	Fast	Passive
LACP protocol:	Receive State		Transmit State		Mux State				
xe-0/0/2	Defaulted		Fast periodic		Detached				

**Meaning** The output indicates that LACP has been set up correctly and is active at one end.

**Verifying That the LACP Packets Are Being Exchanged**

**Purpose** Verify that LACP packets are being exchanged.

**Action** Use the **show interfaces aex statistics** command to display LACP information.

```
user@switch> show interfaces ae0 statistics
```

```
Physical interface: ae0, Enabled, Physical link is Down
Interface index: 153, SNMP ifIndex: 30
Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
Statistics      Packets      pps      Bytes      bps
Bundle:
  Input :           0           0           0           0
  Output:           0           0           0           0
Protocol inet
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

**Meaning** The output here shows that the link is down and that no PDUs are being exchanged.

### Troubleshooting

To troubleshoot a nonworking LACP link, perform these tasks:

- [Troubleshooting a Nonworking LACP Link on page 2469](#)

#### ***Troubleshooting a Nonworking LACP Link***

**Problem** The LACP link is not working.

**Solution** Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.
- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the **monitor traffic-interface lag-member detail** command.

- Related Documentation**
- [Configuring Link Aggregation on page 2593](#)
  - [Verifying the Status of a LAG Interface on page 2750](#)

- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2751](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [show lacp statistics interfaces \(View\) on page 2875](#)



## Example: Configuring Multichassis Link Aggregation



**NOTE:** Multichassis Link Aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI and QFX5100 standalone switches and EX4600 switches running Enhanced Layer 2 Software.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

The peers in an MC-LAG use an interchassis control link-protection link (ICL-PL) to replicate forwarding information across the peers. The Interchassis Control Protocol (ICCP) exchanges the control information between two MC-LAG switches. Additionally, ICCP propagates the operational state of MC-LAG members through the ICL-PL.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of an MC-LAG are two MC-LAG switches. Each of the switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.

- [Requirements on page 2471](#)
- [Overview on page 2472](#)
- [Configuration on page 2473](#)
- [Verification on page 2490](#)
- [Troubleshooting on page 2493](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2 or later for the QFX3500 and QFX3600 standalone switches and Junos OS Release 13.2X51-D10 or later for the QFX5100 standalone switches.
- Two QFX3500 or QFX3600 standalone switches, or two QFX5100 standalone switches.

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 2462](#).
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch” on page 2466](#).


## Overview

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, ICCP for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

## Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to a server. [Figure 51 on page 2472](#) shows the topology of this example.

**Figure 51: Configuring a Multichassis LAG Between Switch A and Switch B**



[Table 235 on page 2463](#) details the topology used in this configuration example.

**Table 236: Components of the Topology for Configuring a Multichassis LAG Between Two Switches**

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A	QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch	<b>ae0</b> is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of <b>ae0</b> : <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch A and <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch B.  <b>ae1</b> is configured as an MC-LAG, and the following two interfaces are part of <b>ae1</b> : <b>xe-0/0/44</b> on Switch A and <b>xe-0/0/46</b> on Switch B.
Switch B	QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch	

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.



**NOTE:** This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three different statements and one different option from the original CLI:

- The `port-mode` statement in the `[edit interfaces interface-name unit number family ethernet-switching]` hierarchy is not supported. Use the `interface-mode` statement instead.
- The `vlan` statement in the `[edit interfaces interface-name]` hierarchy is not supported. Use the `irb` statement instead.
- The `vlan.logical-interface-number` option in the `[edit vlans vlan-name l3-interface]` option is not supported. Use the `irb.logical-interface-number` option instead.
- The `service-id` statement in the `[edit switch-options]` hierarchy is required in the ELS CLI.

### Original CLI

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
```

```
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

## ELS

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
set switch-options service-id 10
```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.



**NOTE:** This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three different statements and one different option from the original CLI:

- The port-mode statement in the [edit interfaces *interface-name* unit *number* family ethernet-switching] hierarchy is not supported. Use the interface-mode statement instead.
- The vlan statement in the [edit interfaces *interface-name*] hierarchy is not supported. Use the irb statement instead.
- The vlan.logical-interface-number option in the [edit vlans *vlan-name* l3-interface] option is not supported. Use the irb.logical-interface-number option instead.
- The service-id statement in the [edit switch-options] hierarchy is required in the ELS CLI.

#### Original CLI

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

#### ELS

```
set chassis aggregated-devices ethernet device-count 2
```

```
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
set switch-options service-id 10
```

### ***Configuring MC-LAG on Two Switches***

#### **Step-by-Step Procedure**

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.  

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 2
```
2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.  

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
```
3. Configure a trunk interface between Switch A and Switch B.



**NOTE:** The `port-mode` statement is not supported on Enhanced Layer 2 Software (ELS). If you are running ELS, use the `interface-mode` statement.

Original CLI:

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
or
```

ELS:

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure a multichassis protection link between Switch A and Switch B.

Switch A:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Switch B:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

#### Step-by-Step Procedure

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a (BFD) session for ICCP on Switch A and Switch B.



**NOTE:** Configure at least 1000ms as the minimum receive interval.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
```

Switch B:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
```

3. Configure the peer IP address and minimum transmit interval for Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.



**NOTE:** Configure at least 1000ms as the transmit interval minimum interval.

Switch A:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
```

Switch B:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



**NOTE:** Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

Switch A:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



**NOTE:** By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
```

Switch B:



- ```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```
6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.
- ```
[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface vlan.500
[edit vlans]
user@switch# set v500 l3-interface irb.500
```



**NOTE:** The port-mode statement is not supported on Enhanced Layer 2 Software (ELS). If you are running ELS, use the interface-mode statement.

Original CLI:

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members
v500
```

or

ELS:

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk vlan members
v500
```

### Step-by-Step Procedure

To enable the MC-LAG interface:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.



**NOTE:** At least one end needs to be active. The other end can be either active or passive.

- ```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp active
```
2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.
- ```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```
3. Specify the same service ID on Switch A and Switch B.
- ELS:
- ```
[edit]
```

```
user@switch# set switch-options service-id 10
```

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

Switch B:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
```

5. Specify the operating mode of the MC-LAG on both Switch A and Switch B.



**NOTE:** Only active-active mode is supported at this time.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

6. Specify the status control for MC-LAG on Switch A and Switch B.



**NOTE:** You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

Switch B:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

7. Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.



**NOTE:** The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

8. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

10. Enable a VLAN on the MC-LAG on Switch A and Switch B.



**NOTE:** The `port-mode` statement is not supported on Enhanced Layer 2 Software (ELS). If you are running ELS, use the `interface-mode` statement.

Original CLI:

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
or
```

ELS:

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
[edit]
user@switch# set vlans v100 vlan-id 100
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

11. (Optional) Enable a private VLAN on the MC-LAG on Switch A and Switch B.

```
[edit]
user@switch# set vlans vlan100 pvlan isolation-vlan-id 200
extend-secondary-vlan-id
[edit]
user@switch# set vlans vlan100 interface ae0.0 pvlan-trunk
```

#### Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.
 

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```
2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:
 

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```
3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.



**NOTE:** The `ae1` interface is a downstream interface. This is why RSTP and `bpdu-block-on-edge` need to be configured.

```
[edit]
user@switch# set protocols rstp interface ae1.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.



**NOTE:** The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

[edit]

```
user@switch# set protocols rstp bpdu-block-on-edge
```

### Results

Display the results of the configuration on Switch A using the original CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 0;
    }
  }
}
```

```

        mode active-active;
        status-control active;
        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        port-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.2/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.2;
        peer 3.3.3.1 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.1 {
        interface ae0;
    }
}
}

```

```
vlan {  
  v100 {  
    vlan-id 100;  
  }  
  v500 {  
    vlan-id 500;  
    l3-interface vlan.500;  
  }  
}
```

Display the results of the configuration on Switch A using the ELS CLI.

```
chassis {  
  aggregated-devices {  
    ethernet {  
      device-count 2;  
    }  
  }  
}  
interfaces {  
  xe-0/0/12 {  
    ether-options {  
      802.3ad ae0;  
    }  
  }  
  xe-0/0/13 {  
    ether-options {  
      802.3ad ae0;  
    }  
  }  
  xe-0/0/44 {  
    ether-options {  
      802.3ad ae1;  
    }  
  }  
  ae0 {  
    unit 0 {  
      family ethernet-switching {  
        interface-mode trunk;  
        vlan {  
          members v500;  
        }  
      }  
    }  
  }  
  ae1 {  
    aggregated-ether-options {  
      lacp {  
        active;  
        system-id 00:01:02:03:04:05;  
        admin-key 3;  
      }  
      mc-ae {  
        mc-ae-id 3;  
        chassis-id 0;  
        mode active-active;  
      }  
    }  
  }  
}
```

```

        status-control active;
        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.2/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.2;
        peer 3.3.3.1 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.1 {
        interface ae0;
    }
}
}
switch-options {

```

```
    service-id 10;
  }
  vlans {
    v100 {
      vlan-id 100;
    }
    v500 {
      vlan-id 500;
      l3-interface irb.500;
    }
  }
}
```

Display the results of the configuration on Switch B using the original CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/46 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
      }
    }
  }
}
```



```

        chassis-id 1;
        mode active-active;
        status-control standby;
        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        port-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.1/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.1;
        peer 3.3.3.2 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.2 {
        interface ae0;
    }
}

```

```
}
vlands {
  v100 {
    vlan-id 100;
  }
  v500 {
    vlan-id 500;
    l3-interface vlan.500;
  }
}
```

Display the results of the configuration on Switch B using the ELS CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/46 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 1;
      }
    }
  }
}
```

```

        mode active-active;
        status-control standby;
        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.1/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.1;
        peer 3.3.3.2 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.2 {
        interface ae0;
    }
}
}

```

```
switch-options {  
  service-id 10;  
}  
vllans {  
  vl100 {  
    vl1an-id 100;  
  }  
  vl500 {  
    vl1an-id 500;  
    l3-interface irb.500;  
  }  
}
```

## Verification

---

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That ICCP Is Working on Switch A on page 2490](#)
- [Verifying That ICCP Is Working on Switch B on page 2490](#)
- [Verifying That LACP Is Active on Switch A on page 2491](#)
- [Verifying That LACP Is Active on Switch B on page 2491](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 2491](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 2492](#)
- [Verifying that MAC Learning Is Occurring on Switch A on page 2492](#)
- [Verifying that MAC Learning Is Occurring on Switch B on page 2492](#)

### *Verifying That ICCP Is Working on Switch A*

**Purpose** Verify that ICCP is running on Switch A.

**Action** [edit]  
user@switch# show iccp  
Redundancy Group Information for peer 3.3.3.1  
TCP Connection : Established  
Liveliness Detection : Up  
  
Client Application: MCSNOOPD  
  
Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

### *Verifying That ICCP Is Working on Switch B*

**Purpose** Verify that ICCP is running on Switch B.

**Action** show iccp  
[edit]  
user@switch# show iccp

```

Redundancy Group Information for peer 3.3.3.2
TCP Connection      : Established
Liveliness Detection : Up

```

```
Client Application: MCSNOOPD
```

```
Client Application: eswd
```

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

#### *Verifying That LACP Is Active on Switch A*

**Purpose** Verify that LACP is active on Switch A.

```

Action [edit]
user@switch# show lacp interfaces
Aggregated interface: ae1
  LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
  xe-0/0/46       Actor No   No   Yes   Yes  Yes  Yes   Fast    Active
  xe-0/0/46       Partner No   No   Yes   Yes  Yes  Yes   Fast    Active
  LACP protocol:   Receive State Transmit State Mux State
  xe-0/0/46       Current Fast periodic Collecting distributing

```

**Meaning** This output shows that Switch A is participating in LACP negotiation.

#### *Verifying That LACP Is Active on Switch B*

**Purpose** Verify that LACP is active on Switch B

```

Action [edit]
user@switch# show lacp interfaces
Aggregated interface: ae1
  LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
  xe-0/0/44       Actor No   No   Yes   Yes  Yes  Yes   Fast    Active
  xe-0/0/44       Partner No   No   Yes   Yes  Yes  Yes   Fast    Active
  LACP protocol:   Receive State Transmit State Mux State
  xe-0/0/44       Current Fast periodic Collecting distributing

```

**Meaning** This output shows that Switch B is participating in LACP negotiation.

#### *Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A*

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

**Action** [edit]  
user@switch# **show interfaces mc-ae**  
Member Link : ae1  
Current State Machine's State: mcae active state  
Local Status : active  
Local State : up  
Peer Status : active  
Peer State : up  
Logical Interface : ae1.0  
Topology Type : bridge  
Local State : up  
Peer State : up  
Peer Ip/MCP/State : 3.3.3.1 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch A is up and active.

***Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B***

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch B.

**Action** [edit]  
user@switch# **show interfaces mc-ae**  
Member Link : ae1  
Current State Machine's State: mcae active state  
Local Status : active  
Local State : up  
Peer Status : active  
Peer State : up  
Logical Interface : ae1.0  
Topology Type : bridge  
Local State : up  
Peer State : up  
Peer Ip/MCP/State : 3.3.3.2 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch B is up and active.

***Verifying that MAC Learning Is Occurring on Switch A***

**Purpose** Verify that MAC learning is working on Switch A.

**Action** [edit]  
user@switch# **show ethernet-switching table**  
Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries  

| VLAN | MAC address       | Type     | Age | Interfaces   |
|------|-------------------|----------|-----|--------------|
| V100 | *                 | Flood    | -   | All-members  |
| V100 | 00:10:94:00:00:05 | Learn(L) | 33  | ae0.0 (MCAE) |

**Meaning** The output shows four learned MAC addresses entries.

***Verifying that MAC Learning Is Occurring on Switch B***

**Purpose** Verify that MAC learning is working on Switch B.

**Action** [edit]  
 user@switch# **show ethernet-switching table**  
 Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries

| VLAN | MAC address       | Type     | Age | Interfaces    |
|------|-------------------|----------|-----|---------------|
| V100 | *                 | Flood    |     | - All-members |
| V100 | 00:10:94:00:00:05 | Learn(L) | 33  | ae0.0 (MCAE)  |

**Meaning** The output shows four learned MAC addresses entries.

## Troubleshooting

### *Troubleshooting a LAG That Is Down*

**Problem** The **show interfaces terse** command shows that the MC-LAG is **down**

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)

## Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP



**NOTE:** Multichassis Link Aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI and QFX5100 standalone switches running Enhanced Layer 2 Software (ELS).

There are two methods for enabling Layer 3 multicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 multicast MC-LAG is included in this example.

- [Requirements on page 2494](#)
- [Overview on page 2494](#)
- [Configuration on page 2496](#)
- [Verification on page 2529](#)

## Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX3500 and QFX3600 standalone switches and Junos OS Release 13.2X51-D10 or later for the QFX5100 standalone switches.
- Two QFX3500 or QFX3600 standalone switches, or two QFX5100 standalone switches.

Before you configure an MC-LAG for Layer 3 multicast using VRRP, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch”](#) on page 2462.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch”](#) on page 2466.

## Overview

---

In this example, you configure two MC-LAGs across two switches, consisting of two aggregated Ethernet interfaces (ae1 and ae2). To support the MC-LAG, create a third aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.



**NOTE:** Layer 3 connectivity is required for ICCP.

To complete the configuration, enable VRRP by completing the following steps for each MC-LAG:


- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group
- Configure Layer 3 connectivity between the VRRP groups

## Topology

The topology used in this example consists of two switches hosting two MC-LAGs—ae1 and ae2. The two switches are connected to a multicast source (Server 1) over the MC-LAG ae1, and a multicast receiver (Server 2) over the MC-LAG ae2. [Figure 52 on page 2495](#) shows the topology of this example.



Figure 52: Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP



g041361

Table 237 on page 2495 details the topology used in this configuration example.

Table 237: Components of the Topology for Configuring a Multichassis LAG for Layer 3 Multicast Using VRRP

| Hostname | Base Hardware  | Multichassis Link Aggregation Group   |
|----------|--|---|
| Switch A | QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch | <ul style="list-style-type: none"> <li>ae0 is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following two interfaces are part of ae0: xe-0/0/12 and xe-0/0/13 on Switch A and xe-0/0/12 and xe-0/0/13 on Switch B.</li> <li>ae1 is configured as an MC-LAG for the multicast source (Server 1), and the following two interfaces are part of ae1: xe-0/0/44 on Switch A and xe-0/0/46 on Switch B.</li> <li>ae2 is configured as an MC-LAG for the multicast receiver (Server 2), and the following two interfaces are part of ae2: xe-0/0/45 on Switch A and xe-0/0/47 on Switch B.</li> </ul> |
| Switch B | QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch |   |

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.



**NOTE:** This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three different statements and one different option from the original CLI:

- The `port-mode` statement in the `[edit interfaces interface-name unit number family ethernet-switching]` hierarchy is not supported. Use the `interface-mode` statement instead.
- The `vlan` statement in the `[edit interfaces interface-name]` hierarchy is not supported. Use the `irb` statement instead.
- The `vlan.logical-interface-number` option in the `[edit vlans vlan-name l3-interface]` option is not supported. Use the `irb.logical-interface-number` option instead.
- The `service-id` statement in the `[edit switch-options]` hierarchy is required in the ELS CLI.

### Original CLI:

```
set chassis aggregated-devices ethernet device-count 3
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces xe-0/0/45 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
```

```
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 virtual-address
  10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority 200
set interfaces vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 accept-data
set interfaces vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority 200
set interfaces vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface vlan.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval
  1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface vlan.100 priority 200
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 600
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
```

```
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

```
ELS: set chassis aggregated-devices ethernet device-count 3
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces xe-0/0/45 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set interfaces ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces irb unit 100 family inet address 10.1.1.11/24 vrrp-group 1 virtual-address
10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority 200
set interfaces irb unit 100 family inet address 10.1.1.11/24 vrrp-group 1 accept-data
set interfaces irb unit 200 family inet address 10.1.1.21/24 vrrp-group 2 virtual-address
10.1.1.2
set interfaces irb unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority 200
set interfaces irb unit 200 family inet address 10.1.1.21/24 vrrp-group 2 accept-data
set interfaces irb unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface irb.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval
1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
minimum-receive-interval 700
```

```

set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface vlan.100 priority 200
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 600
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
set switch-options service-id 10

```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

**Original CLI:**

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces xe-0/0/47 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk

```

```
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-aether mc-aether-id 4
set interfaces ae2 aggregated-ether-options mc-aether mode active-active
set interfaces ae2 aggregated-ether-options mc-aether status-control active
set interfaces ae2 aggregated-ether-options mc-aether init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching port-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 virtual-address
10.1.1.1
set interfaces vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority 150
set interfaces vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 accept-data
set interfaces vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 virtual-address
10.1.1.2
set interfaces vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority 150
set interfaces vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface vlan.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval
1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface vlan.100 priority 100
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
500
set protocols pim interface vlan.200 priority 500
set protocols pim interface vlan.200 dual-dr
```

```

set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0

```

```

ELS: set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces xe-0/0/47 ether-options 802.3ad ae2
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces ae2 aggregated-ether-options lacp active
set interfaces ae2 aggregated-ether-options lacp system-id 00:01:02:03:04:06
set interfaces ae2 aggregated-ether-options lacp admin-key 3
set interfaces ae2 aggregated-ether-options mc-ae mc-ae-id 4
set interfaces ae2 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae2 aggregated-ether-options mc-ae mode active-active
set interfaces ae2 aggregated-ether-options mc-ae status-control active
set interfaces ae2 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae2 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae2 unit 0 family ethernet-switching vlan members v200
set interfaces irb unit 100 family inet address 10.1.1.10/24 vrrp-group 1 virtual-address
  10.1.1.1
set interfaces irb unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 10.1.1.10/24 vrrp-group 1 accept-data
set interfaces irb unit 200 family inet address 10.1.1.20/24 vrrp-group 2 virtual-address
  10.1.1.2
set interfaces irb unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority 150
set interfaces irb unit 200 family inet address 10.1.1.20/24 vrrp-group 2 accept-data
set interfaces irb unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v200 vlan-id 200
set vlans v200 l3-interface irb.200
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.1

```

```
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval
  1000
set protocols igmp-snooping vlan all
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.100 bfd-liveness-detection transmit-interval
  threshold 500
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection
  minimum-receive-interval 700
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols ospf area 0.0.0.0 interface vlan.200 bfd-liveness-detection transmit-interval
  threshold 500
set protocols pim rp static address 1.0.0.3 group-ranges 239.0.0.0/8
set protocols pim interface vlan.100 priority 100
set protocols pim interface vlan.100 dual-dr
set protocols pim interface vlan.100 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.100 bfd-liveness-detection transmit-interval threshold
  500
set protocols pim interface vlan.200 priority 500
set protocols pim interface vlan.200 dual-dr
set protocols pim interface vlan.200 bfd-liveness-detection minimum-receive-interval
  700
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval
  minimum-interval 350
set protocols pim interface vlan.200 bfd-liveness-detection transmit-interval threshold
  500
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface ae2.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
set switch-options service-id 10
```

### *Configuring MC-LAG for Layer 3 Multicast Using VRRP on Two Switches*

#### **Step-by-Step Procedure**

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 3
```



2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

#### Switch A and Switch B

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
```

#### Switch A

```
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
user@switch# set xe-0/0/45 ether-options 802.3ad ae2
```

#### Switch B

```
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
user@switch# set xe-0/0/47 ether-options 802.3ad ae2
```

3. Configure ae0 as the trunk interface between Switch A and Switch B.

#### Switch A and Switch B

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```

#### Switch A and Switch B Using ELS

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```

4. Configure ae0 as the multichassis protection link between Switch A and Switch B.

#### Switch A

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

#### Switch B

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

#### Step-by-Step Procedure

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

#### Switch A

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

**Switch B**

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address, minimum receive interval, and minimum transmit interval for a Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

**Switch A**

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval
1000
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval
minimum-interval 1000
```

**Switch B**

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval
1000
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval
minimum-interval 1000
```

3. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



**NOTE:** Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

**Switch A**

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

**Switch B**

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```

4. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



**NOTE:** By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

**Switch A**

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip
10.207.64.233
```

#### Switch B

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip
10.207.64.232
```

5. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B.



**NOTE:** In ELS, use the *irb.logical-interface-number* instead.

#### Switch A and Switch B

```
[edit vlans]
user@switch# set v500 vlan-id 500
user@switch# set v500 l3-interface vlan.500
```

#### Switch A and Switch B Using ELS

```
[edit vlans]
user@switch# set v500 vlan-id 500
user@switch# set v500 l3-interface irb.500
```

#### Switch A and Switch B

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members v500
```

#### Switch A

```
[edit interfaces]
user@switch# set vlan unit 500 family inet address 3.3.3.2/24
```

#### Switch A Using ELS

```
[edit interfaces]
user@switch# set irbunit 500 family inet address 3.3.3.2/24
```

#### Switch B

```
[edit interfaces]
user@switch# set vlan unit 500 family inet address 3.3.3.1/24
```

#### Switch B Using ELS

```
[edit interfaces]
user@switch# set irb unit 500 family inet address 3.3.3.1/24
```

**Step-by-Step Procedure** To enable the ae1 and ae2 MC-LAG interfaces:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interfaces on Switch A and Switch B.



**NOTE:** At least one end needs to be active. The other end can be either active or passive.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options lacp active
```

```
user@switch# set ae2 aggregated-ether-options lacp active
```

2. Specify the same multichassis aggregated Ethernet (MC-AE) identification number for each MC-LAG peer on Switch A and Switch B.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```

```
user@switch# set ae2 aggregated-ether-options mc-ae mc-ae-id 4
```

3. Specify the same service ID on Switch A and Switch B.

**ELS:**

[edit]

```
set switch-options service-id 10
```

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

**Switch A**

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

```
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 0
```

**Switch B**

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 1
```

```
user@switch# set ae2 aggregated-ether-options mc-ae chassis-id 1
```

5. Specify the operating mode of the MC-LAGs on both Switch A and Switch B.



**NOTE:** Only active-active mode is supported at this time.

[edit interfaces]

```
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

```
user@switch# set ae2 aggregated-ether-options mc-ae mode active-active
```

6. Specify the status control for the MC-LAGs on Switch A and Switch B.



**NOTE:** You must configure status control on both Switch A and Switch B hosting the MC-LAGs. If one peer is in active mode, the other must be in standby mode.

#### Switch A

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
user@switch# set ae2 aggregated-ether-options mc-ae status-control active
```

#### Switch B

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
user@switch# set ae2 aggregated-ether-options mc-ae status-control standby
```

7. Specify the number of seconds by which the bring-up of the MC-LAG interfaces should be deferred after you reboot Switch A or Switch B.



**NOTE:** The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 420
user@switch# set ae2 aggregated-ether-options mc-ae init-delay-time 420
```

8. Specify the same LACP system ID for each MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
user@switch# set ae2 aggregated-ether-options lacp system-ID 00:01:02:03:04:06
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
user@switch# set ae2 aggregated-ether-options lacp admin-key 3
```

10. Enable a VLAN for each MC-LAG on Switch A and Switch B.

```
[edit vlans]
user@switch# set v100 vlan-id 100
user@switch# set v200 vlan-id 200
```

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
user@switch# set ae2 unit 0 family ethernet-switching vlan members v200
```

11. Configure ae1 and ae2 as trunk interfaces between Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
```

```
user@switch# set ae2 unit 0 family ethernet-switching port-mode trunk
```

ELS:

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ae2 unit 0 family ethernet-switching interface-mode trunk
```

#### Step-by-Step Procedure

To enable VRRP on the MC-LAGs on Switch A and Switch B:

1. Create a routed VLAN interface (RVI) for each MC-LAG, assign a virtual IP address that is shared between each switch in the VRRP groups, and assign an individual IP address for each switch in the VRRP groups.

##### Switch A

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1
virtual-address 10.1.1.1
user@switch# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2
virtual-address 10.1.1.2
```

##### Switch A Using ELS

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/24 vrrp-group 1
virtual-address 10.1.1.1
user@switch# set irb unit 200 family inet address 10.1.1.21/24 vrrp-group 2
virtual-address 10.1.1.2
```

##### Switch B

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1
virtual-address 10.1.1.1
user@switch# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2
virtual-address 10.1.1.2
```

##### Switch B Using ELS

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/24 vrrp-group 1
virtual-address 10.1.1.1
user@switch# set irb unit 200 family inet address 10.1.1.20/24 vrrp-group 2
virtual-address 10.1.1.2
```

2. Assign the priority for each switch in the VRRP groups:



**NOTE:** The switch configured with the highest priority is the master.

##### Switch A

```
[edit interfaces]
```

```

user@switch# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority
200
user@switch# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority
200

```

#### Switch A Using ELS

```

[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/24 vrrp-group 1 priority
200
user@switch# set irb unit 200 family inet address 10.1.1.21/24 vrrp-group 2 priority
200

```

#### Switch B

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority
150
user@switch# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority
150

```

#### Switch B Using ELS

```

[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.10/24 vrrp-group 1 priority
150
user@switch# set irb unit 200 family inet address 10.1.1.20/24 vrrp-group 2 priority
150

```

3. Enable the switch to accept all packets destined for the virtual IP address if it is the master in a VRRP group:

#### Switch A

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.11/24 vrrp-group 1
accept-data
user@switch# set vlan unit 200 family inet address 10.1.1.21/24 vrrp-group 2
accept-data

```

#### Switch A Using ELS

```

[edit interfaces]
user@switch# set irb unit 100 family inet address 10.1.1.11/24 vrrp-group 1 accept-data
user@switch# set irb unit 200 family inet address 10.1.1.21/24 vrrp-group 2
accept-data

```

#### Switch B

```

[edit interfaces]
user@switch# set vlan unit 100 family inet address 10.1.1.10/24 vrrp-group 1
accept-data
user@switch# set vlan unit 200 family inet address 10.1.1.20/24 vrrp-group 2
accept-data

```

#### Switch B Using ELS

```

[edit interfaces]

```

```
user@switch# set irb unit 100 family inet address 10.1.1.10/24 vrrp-group 1
accept-data
user@switch# set irb unit 200 family inet address 10.1.1.20/24 vrrp-group 2
accept-data
```

4. Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
user@switch# set v100 l3-interface vlan.100
user@switch# set v200 l3-interface vlan.200
```

ELS:

```
[edit interfaces]
user@switch# set v100 l3-interface irb.100
user@switch# set v200 l3-interface irb.200
```

#### Step-by-Step Procedure

To enable IGMP snooping:

1. Enable IGMP snooping for all VLANs on Switch A and Switch B.

```
[edit protocols]
user@switch# set igmp-snooping vlan all
```

#### Step-by-Step Procedure

To configure OSPF as the Layer 3 protocol:

1. Configure an OSPF area on Switch A and Switch B.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0
```

2. Assign the VLAN interfaces for the MC-LAGs as interfaces to the OSPF area on Switch A and Switch B.

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface vlan.100
user@switch# set interface vlan.200
```

3. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the OSPF interfaces on Switch A and Switch B.



**NOTE:** On a QFX5100 switch, the minimum transmit interval must be 1000 milliseconds or greater. Sub-second timers are not supported in Junos OS 13.2X51-D10 and later. If you configure the minimum transmit interval timer lower than 1000 milliseconds, the state of the MC-LAG can be affected.

---

```
[edit protocols ospf area 0.0.0.0]
user@switch# set interface vlan.100 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
minimum-interval 350
```



```

user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
threshold 500
user@switch# set interface vlan.200 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
minimum-interval 350
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
threshold 500

```

### Step-by-Step Procedure

To configure PIM as the multicast protocol:

1. Configure a static rendezvous point (RP) address on Switch A and Switch B.  

```

[edit protocols pim]
user@switch# set rp static address 1.0.0.3

```
2. Configure the address ranges of the multicast groups for which Switch A and Switch B can be a rendezvous point (RP).  

```

[edit protocols pim rp static address 1.0.0.3]
user@switch# set group-ranges 239.0.0.0/8

```
3. Enable PIM on the VLAN interfaces for the MC-LAGs on Switch A and Switch B.  

```

[edit protocols pim]
user@switch# set interface vlan.100 dual-dr
user@switch# set interface vlan.200 dual-dr

```
4. Configure each PIM interface's priority for being selected as the designated router (DR).

An interface with a higher priority value has a higher probability of being selected as the DR.

#### Switch A

```

[edit protocols pim]
user@switch# set interface vlan.100 priority 200
user@switch# set interface vlan.200 priority 600

```

#### Switch B

```

[edit protocols pim]
user@switch# set interface vlan.100 priority 100
user@switch# set interface vlan.200 priority 500

```

5. Configure the minimum receive interval, minimum transmit interval, and transmit interval threshold for a Bidirectional Forwarding Detection (BFD) session for the PIM interfaces on Switch A and Switch B.

```

[edit protocols pim]
user@switch# set interface vlan.100 bfd-liveness-detection
minimum-receive-interval 700
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
minimum-interval 350
user@switch# set interface vlan.100 bfd-liveness-detection transmit-interval
threshold 500
user@switch# set interface vlan.200 bfd-liveness-detection
minimum-receive-interval 700

```

```
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
minimum-interval 350
```

```
user@switch# set interface vlan.200 bfd-liveness-detection transmit-interval
threshold 500
```

**Step-by-Step  
Procedure**

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit protocols rstp]
user@switch# set interface all mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B.

```
[edit protocols rstp]
user@switch# set interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.



**NOTE:** The ae1 and ae2 interfaces are downstream interfaces. This is why RSTP and bpd-block-on-edge need to be configured.

```
[edit protocols rstp]
user@switch# set interface ae1.0 edge
user@switch# set interface ae2.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.



**NOTE:** The ae1 and ae2 interfaces are downstream interfaces. This is why RSTP and bpd-block-on-edge need to be configured.

```
[edit protocols rstp]
user@switch# set bpd-block-on-edge
```

**Results**

From configuration mode on Switch A, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

**Original CLI:**

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 3;
    }
  }
}
```

```

    }
  }
  interfaces {
    xe-0/0/12 {
      ether-options {
        802.3ad ae0;
      }
    }
    xe-0/0/13 {
      ether-options {
        802.3ad ae0;
      }
    }
    xe-0/0/44 {
      ether-options {
        802.3ad ae1;
      }
    }
    xe-0/0/45 {
      ether-options {
        802.3ad ae2;
      }
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
}

```

```
}
ae2 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:06;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 4;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
}
unit 0 {
  family ethernet-switching {
    port-mode trunk;
    vlan {
      members v200;
    }
  }
}
}
vlan {
  unit 100 {
    family inet {
      address 10.1.1.11/24 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.21/24 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 3.3.3.2/24;
    }
  }
}
}
protocols {
```

```

ospf {
  area 0.0.0.0 {
    interface vlan.100 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
    interface vlan.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 1.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
}
interface vlan.100 {
  priority 200;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {
  priority 600;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
}
iccp {

```

```
local-ip-addr 3.3.3.2;
peer 3.3.3.1 {
    session-establishment-hold-time 50;
    backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
    }
    liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
            minimum-interval 1000;
        }
    }
}
}
igmp-snooping {
    vlan all;
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface ae2.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.1 {
        interface ae0;
    }
}
}
vlangs {
    v100 {
        vlan-id 100;
        l3-interface vlan.100;
    }
    v200 {
        vlan-id 200;
        l3-interface vlan.200;
    }
    v500 {
        vlan-id 500;
        l3-interface vlan.500;
    }
}
}

ELS: chassis {
    aggregated-devices {
        ethernet {
```

```

        device-count 3;
    }
}
interfaces {
    xe-0/0/12 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/13 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/44 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-0/0/45 {
        ether-options {
            802.3ad ae2;
        }
    }
    ae0 {
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members v500;
                }
            }
        }
    }
    ae1 {
        aggregated-ether-options {
            lacp {
                active;
                system-id 00:01:02:03:04:05;
                admin-key 3;
            }
            mc-ae {
                mc-ae-id 3;
                chassis-id 0;
                mode active-active;
                status-control active;
                init-delay-time 240;
            }
        }
        unit 0 {
            family ethernet-switching {
                interface-mode trunk;
                vlan {
                    members v100;
                }
            }
        }
    }
}

```

```
    }
  }
}
ae2 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:06;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 4;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240;
    }
  }
}
unit 0 {
  family ethernet-switching {
    interface-mode trunk;
    vlan {
      members v200;
    }
  }
}
}
irb {
  unit 100 {
    family inet {
      address 10.1.1.1/24 {
        vrrp-group 1 {
          virtual-address 10.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 200 {
    family inet {
      address 10.1.1.2/24 {
        vrrp-group 2 {
          virtual-address 10.1.1.2;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 3.3.3.2/24;
    }
  }
}
```



```

}
protocols {
  ospf {
    area 0.0.0.0 {
      interface vlan.100 {
        bfd-liveness-detection {
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
    interface vlan.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 1.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
}
interface vlan.100 {
  priority 200;
  dual-dr;
  bfd-liveness-detection {
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {
  priority 600;
  dual-dr;
  bfd-liveness-detection {
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
}

```

```
}
iccp {
  local-ip-addr 3.3.3.2;
  peer 3.3.3.1 {
    session-establishment-hold-time 50;
    backup-liveness-detection {
      backup-peer-ip 10.207.64.233;
    }
    liveness-detection {
      minimum-receive-interval 1000;
      transmit-interval {
        minimum-interval 1000;
      }
    }
  }
}
}
igmp-snooping {
  vlan all;
}
}
rstp {
  interface ae0.0 {
    disable;
  }
  interface ae1.0 {
    edge;
  }
  interface ae2.0 {
    edge;
  }
  interface all {
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
}
multi-chassis {
  multi-chassis-protection 3.3.3.1 {
    interface ae0;
  }
}
}
switch-options {
  service-id 10;
}
}
vllans {
  v100 {
    vlan-id 100;
    l3-interface irb.100;
  }
  v200 {
    vlan-id 200;
    l3-interface irb.200;
  }
  v500 {
    vlan-id 500;
    l3-interfac irb.500;
  }
}
```

```
}
```

From configuration mode on Switch B, confirm your configuration by entering the **show chassis**, **show interfaces**, **show multi-chassis**, **show protocols**, and **show vlans** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Original CLI: chassis {
               aggregated-devices {
                 ethernet {
                   device-count 3;
                 }
               }
             }
            interfaces {
              xe-0/0/12 {
                ether-options {
                  802.3ad ae0;
                }
              }
              xe-0/0/13 {
                ether-options {
                  802.3ad ae0;
                }
              }
              xe-0/0/46 {
                ether-options {
                  802.3ad ae1;
                }
              }
              xe-0/0/47 {
                ether-options {
                  802.3ad ae2;
                }
              }
            }
            ae0 {
              unit 0 {
                family ethernet-switching {
                  port-mode trunk;
                  vlan {
                    members v500;
                  }
                }
              }
            }
            ae1 {
              aggregated-ether-options {
                lacp {
                  active;
                  system-id 00:01:02:03:04:05;
                  admin-key 3;
                }
                mc-ae {
                  mc-ae-id 3;
                  chassis-id 1;
                  mode active-active;
                }
              }
            }
          }
```

```
        status-control standby;
    }
}
unit 0 {
    family ethernet-switching {
        port-mode trunk;
        vlan {
            members v100;
        }
    }
}
ae2 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:06;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 4;
            chassis-id 1;
            mode active-active;
            status-control active;
            init-delay-time 240;
        }
    }
}
unit 0 {
    family ethernet-switching {
        port-mode trunk;
        vlan {
            members v200;
        }
    }
}
vlan {
    unit 100 {
        family inet {
            address 10.1.1.10/24 {
                vrrp-group 1 {
                    virtual-address 10.1.1.1;
                    priority 150;
                    accept-data;
                }
            }
        }
    }
}
unit 200 {
    family inet {
        address 10.1.1.20/24 {
            vrrp-group 2 {
                virtual-address 10.1.1.2;
                priority 150;
                accept-data;
            }
        }
    }
}
```

```

    }
  }
}
unit 500 {
  family inet {
    address 3.3.3.1/24;
  }
}
}
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface vlan.100 {
        bfd-liveness-detection {
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
    interface vlan.200 {
      bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
          minimum-interval 350;
          threshold 500;
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 1.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
}
interface vlan.100 {
  priority 100;
  dual-dr;
  bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {
  priority 500;

```

```
    dual-dr;
    bfd-liveness-detection { ## Warning: 'bfd-liveness-detection' is deprecated
        minimum-receive-interval 700;
        transmit-interval {
            minimum-interval 350;
            threshold 500;
        }
    }
}
}
iccp {
    local-ip-addr 3.3.3.1;
    peer 3.3.3.2 {
        session-establishment-hold-time 50;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.233;
        }
        liveness-detection {
            minimum-receive-interval 1000;
            transmit-interval {
                minimum-interval 1000;
            }
        }
    }
}
}
igmp-snooping {
    vlan all;
}
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface ae2.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.2 {
        interface ae0;
    }
}
}
vllans {
    v100 {
        vlan-id 100;
        l3-interface vllan.100;
    }
    v200 {
        vlan-id 200;
    }
}
```

```

        l3-interface vlan.200;
    }
    v500 {
        vlan-id 500;
        l3-interface vlan.500;
    }
}

ELS: chassis {
    aggregated-devices {
        ethernet {
            device-count 3;
        }
    }
}
interfaces {
    xe-0/0/12 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/13 {
        ether-options {
            802.3ad ae0;
        }
    }
    xe-0/0/46 {
        ether-options {
            802.3ad ae1;
        }
    }
    xe-0/0/47 {
        ether-options {
            802.3ad ae2;
        }
    }
}
ae0 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v500;
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            chassis-id 1;
        }
    }
}

```

```
        mode active-active;
        status-control standby;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members v100;
        }
    }
}
ae2 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:06;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 4;
            chassis-id 1;
            mode active-active;
            status-control active;
            init-delay-time 240;
        }
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members v200;
        }
    }
}
}
irb {
    unit 100 {
        family inet {
            address 10.1.1.10/24 {
                vrrp-group 1 {
                    virtual-address 10.1.1.1;
                    priority 150;
                    accept-data;
                }
            }
        }
    }
}
unit 200 {
    family inet {
        address 10.1.1.20/24 {
            vrrp-group 2 {
                virtual-address 10.1.1.2;
                priority 150;
                accept-data;
            }
        }
    }
}
```



```

    }
  }
}
unit 500 {
  family inet {
    address 3.3.3.1/24;
  }
}
}
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface vlan.100 {
        bfd-liveness-detection {
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
      interface vlan.200 {
        bfd-liveness-detection {
          minimum-receive-interval 700;
          transmit-interval {
            minimum-interval 350;
            threshold 500;
          }
        }
      }
    }
  }
}
pim {
  rp {
    static {
      address 1.0.0.3 {
        group-ranges {
          239.0.0.0/8;
        }
      }
    }
  }
}
interface vlan.100 {
  priority 100;
  dual-dr;
  bfd-liveness-detection {
    minimum-receive-interval 700;
    transmit-interval {
      minimum-interval 350;
      threshold 500;
    }
  }
}
interface vlan.200 {

```

```
    priority 500;
    dual-dr;
    bfd-liveness-detection {
        minimum-receive-interval 700;
        transmit-interval {
            minimum-interval 350;
            threshold 500;
        }
    }
}
}
}
iccp {
    local-ip-addr 3.3.3.1;
    peer 3.3.3.2 {
        session-establishment-hold-time 50;
        backup-liveness-detection {
            backup-peer-ip 10.207.64.233;
        }
        liveness-detection {
            minimum-receive-interval 1000;
            transmit-interval {
                minimum-interval 1000;
            }
        }
    }
}
}
igmp-snooping {
    vlan all;
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface ae2.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.2 {
        interface ae0;
    }
}
}
switch-options {
    service-id 10;
}
}
vllans {
    v100 {
        vlan-id 100;
    }
}
```

```

    l3-interface irb.100;
  }
  v200 {
    vlan-id 200;
    l3-interface irb.200;
  }
  v500 {
    vlan-id 500;
    l3-interface irb.500;
  }
}

```

### Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That Switch A is the Master Designated Router on page 2529](#)
- [Verifying That Switch B is the Backup Designated Router on page 2529](#)

#### *Verifying That Switch A is the Master Designated Router*

**Purpose** Verify that Switch A is the master designated router (DR).

**Action** From operational mode, enter the `show pim interfaces` command.

```
user@switch> show pim interfaces
```

Stat = Status, V = Version, NbrCnt = Neighbor Count,

S = Sparse, D = Dense, B = Bidirectional,

DR = Designated Router, P2P = Point-to-point link,

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name      | Stat | Mode | IP | V | State         | NbrCnt | JoinCnt(sg/*g) | DR | address   |
|-----------|------|------|----|---|---------------|--------|----------------|----|-----------|
| pim.32769 | Down | S    | 4  | 2 | P2P,NotCap    | 0      | 0/0            |    |           |
| vlan.100  | Up   | S    | 4  | 2 | DDR-DR,NotCap | 1      | 0/0            |    | 10.1.1.11 |
| vlan.200  | Up   | S    | 4  | 2 | DDR-DR,NotCap | 2      | 0/0            |    | 10.1.1.21 |

**Meaning** The DDR-DR state of the VLAN interfaces in the output shows that Switch A is the master designated router.

#### *Verifying That Switch B is the Backup Designated Router*

**Purpose** Verify that Switch B is the backup designated router (BDR).

**Action** From operational mode, enter the `show pim interfaces` command.

```
user@switch> show pim interfaces
```

Stat = Status, V = Version, NbrCnt = Neighbor Count,

S = Sparse, D = Dense, B = Bidirectional,

DR = Designated Router, P2P = Point-to-point link,

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

| Name      | Stat | Mode | IP | V | State          | NbrCnt | JoinCnt(sg/*g) | DR | address   |
|-----------|------|------|----|---|----------------|--------|----------------|----|-----------|
| pim.32769 | Down | S    | 4  | 2 | P2P,NotCap     | 0      | 0/0            |    |           |
| vlan.100  | Up   | S    | 4  | 2 | DDR-BDR,NotCap | 1      | 0/0            |    | 10.1.1.11 |
| vlan.200  | Up   | S    | 4  | 2 | DDR-BDR,NotCap | 2      | 0/0            |    | 10.1.1.21 |

**Meaning** The DDR-BDR state of the VLAN interfaces in the output shows that Switch B is the backup designated router.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)

## Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization



**NOTE:** Multichassis Link Aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI and QFX5100 standalone switches running Enhanced Layer 2 Software.



**NOTE:** Issuing a PING request on an MC-LAG with MAC synchronization enabled does not work.

There are 2 methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to synchronize the MAC addresses between the switches for the participating MC-LAG interfaces, or you can configure Virtual Router Redundancy Protocol (VRRP). The procedure to configure MAC address synchronization is included in this example. For more information on configuring VRRP for use in a Layer 3 unicast MC-LAG, see “[Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\)](#)” on page 2551.

- [Requirements on page 2530](#)
- [Overview on page 2531](#)
- [Configuration on page 2532](#)
- [Verification on page 2548](#)
- [Troubleshooting on page 2551](#)

### Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX3500 and QFX3600 standalone switches and Junos OS Release 13.2X51-D10 or later for the QFX5100 standalone switches.
- Two QFX3500 or QFX3600 standalone switches, or two QFX5100 standalone switches.

Before you configure an MC-LAG for Layer 3 unicast, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See “[Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch](#)” on page 2462.

- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch”](#) on page 2466.


## Overview

In this example, you configure an MC-LAG across two switches, consisting of two aggregated Ethernet interfaces, an interchassis control link-protection link (ICL-PL), multichassis protection link for the ICL-PL, ICCP for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers. Layer 3 connectivity is required for ICCP.

## Topology

The topology used in this example consists of two switches hosting an MC-LAG. The two switches are connected to a server. [Figure 51 on page 2472](#) shows the topology of this example.

**Figure 53: Configuring a Multichassis LAG Between Switch A and Switch B**



g041294

[Table 235 on page 2463](#) details the topology used in this configuration example.

Table 238: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

| Hostname | Base Hardware  | Multichassis Link Aggregation Group   |
|----------|--|---|
| Switch A | QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch | <b>ae0</b> is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of <b>ae0</b> : <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch A and <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch B.<br><br><b>ae1</b> is configured as an MC-LAG, and the following two interfaces are part of <b>ae1</b> : <b>xe-0/0/44</b> on Switch A and <b>xe-0/0/46</b> on Switch B. |
| Switch B | QFX3500 or QFX3600 standalone switch, or QFX5100 standalone switch |   |

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.



**NOTE:** This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three different statements and one different option from the original CLI:

- The `port-mode` statement in the `[edit interfaces interface-name unit number family ethernet-switching]` hierarchy is not supported. Use the `interface-mode` statement instead.
- The `vlan` statement in the `[edit interfaces interface-name]` hierarchy is not supported. Use the `irb` statement instead.
- The `vlan.logical-interface-number` option in the `[edit vlans vlan-name l3-interface]` option is not supported. Use the `irb.logical-interface-number` option instead.
- The `service-id` statement in the `[edit switch-options]` hierarchy is required in the ELS CLI.

#### Original CLI on Switch A:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
```

```

set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0

```

#### ELS on Switch A:

```

set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
set switch-options service-id 10

```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network

configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

#### Original CLI on Switch B:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

#### ELS on Switch B:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
```



```

set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
set switch-options service-id 10

```

### Configuring MC-LAG on Two Switches

#### Step-by-Step Procedure

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.

```

[edit chassis]
user@switch# set aggregated-devices ethernet device-count 2

```

2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.

```

[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1

```

3. Configure a trunk interface between Switch A and Switch B.

```

[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk

```

4. Configure a multichassis protection link between Switch A and Switch B.

Switch A:

```

[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0

```

Switch B:

```

[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0

```

**Step-by-Step  
Procedure**

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
```

3. Configure the peer IP address and minimum transmit interval for Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



**NOTE:** Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

---

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B:

```
[edit protocols]
```

- user@switch# **set iccp peer 3.3.3.2 session-establishment-hold-time 50**
5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



**NOTE:** By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```

6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B using the original CLI:

```
[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface vlan.500
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members v500
```

7. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B using ELS:

```
[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface irb.500
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk vlan members v500
```

**Step-by-Step Procedure**

To enable the MC-LAG interface:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.



**NOTE:** At least one end needs to be active. The other end can be either active or passive.

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options lacp active**

2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae mc-ae-id 3**

3. Specify the same service ID on Switch A and Switch B.

ELS:

[edit]

user@switch# **set switch-options service-id 10**

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A:

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae chassis-id 0**

Switch B:

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae chassis-id 1**

5. Specify the operating mode of the MC-LAG on both Switch A and Switch B.



**NOTE:** Only active-active mode is supported at this time.

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae mode active-active**

6. Specify the status control for MC-LAG on Switch A and Switch B.



**NOTE:** You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

Switch B:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae status-control standby
```

7. Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.



**NOTE:** The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

8. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-ID 00:01:02:03:04:05
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

10. Enable a VLAN on the MC-LAG on Switch A and Switch B using the original CLI:

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
[edit]
```

```
user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

11. Enable a VLAN on the MC-LAG on Switch A and Switch B using ELS:

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
[edit]
```

```
user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

### Step-by-Step Procedure

To enabled RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.

```
[edit]
```

- ```
user@switch# set protocols rstp interface ae1.0 edge
```
4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.

```
[edit]
```

```
user@switch# set protocols rstp bpdv-block-on-edge
```

### Results

Display the results of the configuration on Switch A using the original CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 0;
      mode active-active;
      status-control active;
    }
  }
}
```

```

        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        port-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.2/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.2;
        peer 3.3.3.1 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.1 {
        interface ae0;
    }
}
}
vlangs {
    v100 {

```

```
        vlan-id 100;
      }
    v500 {
      vlan-id 500;
      l3-interface vlan.500;
    }
  }
}
```

Display the results of the configuration on Switch A using ELS.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
    mc-ae {
      mc-ae-id 3;
      chassis-id 0;
      mode active-active;
      status-control active;
      init-delay-time 240
    }
  }
}
```



```

    }
  }
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
vlan {
  unit 500 {
    family inet {
      address 3.3.3.2/24;
    }
  }
}
}
protocols {
  iccp {
    local-ip-addr 3.3.3.2;
    peer 3.3.3.1 {
      session-establishment-hold-time 50;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 60;
        transmit-interval {
          minimum-interval 60;
        }
      }
    }
  }
}
rstp {
  interface ae0.0 {
    disable;
  }
  interface ae1.0 {
    edge;
  }
  interface all {
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
multi-chassis {
  multi-chassis-protection 3.3.3.1 {
    interface ae0;
  }
}
switch-options {
  service-id 10;
}

```

```
vlan {
  v100 {
    vlan-id 100;
  }
  v500 {
    vlan-id 500;
    l3-interface irb.500;
  }
}
```

Display the results of the configuration on Switch B using the original CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/46 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 1;
        mode active-active;
      }
    }
  }
}
```

```

        status-control standby;
        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        port-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.1/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.1;
        peer 3.3.3.2 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
    bpdu-block-on-edge;
}
}
multi-chassis {
    multi-chassis-protection 3.3.3.2 {
        interface ae0;
    }
}
}
vlangs {

```

```
v100 {  
  vlan-id 100;  
}  
v500 {  
  vlan-id 500;  
  l3-interface vlan.500;  
}  
}
```

Display the results of the configuration on Switch B using ELS.

```
chassis {  
  aggregated-devices {  
    ethernet {  
      device-count 2;  
    }  
  }  
}  
interfaces {  
  xe-0/0/12 {  
    ether-options {  
      802.3ad ae0;  
    }  
  }  
  xe-0/0/13 {  
    ether-options {  
      802.3ad ae0;  
    }  
  }  
  xe-0/0/46 {  
    ether-options {  
      802.3ad ae1;  
    }  
  }  
}  
ae0 {  
  unit 0 {  
    family ethernet-switching {  
      interface-mode trunk;  
      vlan {  
        members v500;  
      }  
    }  
  }  
}  
ae1 {  
  aggregated-ether-options {  
    lacp {  
      active;  
      system-id 00:01:02:03:04:05;  
      admin-key 3;  
    }  
    mc-ae {  
      mc-ae-id 3;  
      chassis-id 1;  
      mode active-active;  
      status-control standby;  
    }  
  }  
}
```

```

        init-delay-time 240
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 500 {
        family inet {
            address 3.3.3.1/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.1;
        peer 3.3.3.2 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
    rstp {
        interface ae0.0 {
            disable;
        }
        interface ae1.0 {
            edge;
        }
        interface all {
            mode point-to-point;
        }
        bpdu-block-on-edge;
    }
}
multi-chassis {
    multi-chassis-protection 3.3.3.2 {
        interface ae0;
    }
}
}
switch-options {
    service-id 10;
}

```

```
}
vllans {
  v100 {
    vlan-id 100;
  }
  v500 {
    vlan-id 500;
    l3-interface irb.500;
  }
}
```

---

### Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That ICCP Is Working on Switch A on page 2548](#)
- [Verifying That ICCP Is Working on Switch B on page 2548](#)
- [Verifying That LACP Is Active on Switch A on page 2549](#)
- [Verifying That LACP Is Active on Switch B on page 2549](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 2549](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 2550](#)
- [Verifying that MAC Learning Is Occurring on Switch A and Switch B on page 2550](#)

#### ***Verifying That ICCP Is Working on Switch A***

**Purpose** Verify that ICCP is running on Switch A.

**Action** [edit]  
user@switch# show iccp  
Redundancy Group Information for peer 3.3.3.1  
TCP Connection : Established  
Liveliness Detection : Up  
  
Client Application: MCSNOOPD  
  
Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

#### ***Verifying That ICCP Is Working on Switch B***

**Purpose** Verify that ICCP is running on Switch B.

**Action** show iccp  
  
[edit]  
user@switch# show iccp  
Redundancy Group Information for peer 3.3.3.2  
TCP Connection : Established  
Liveliness Detection : Up

Client Application: MCSNOOPD

Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

#### *Verifying That LACP Is Active on Switch A*

**Purpose** Verify that LACP is active on Switch A.

**Action** [edit]  
 user@switch# show lacp interfaces  
 Aggregated interface: ae1

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/46	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/46	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
LACP protocol:	Receive State	Transmit State	Mux State						
xe-0/0/46	Current	Fast periodic	Collecting distributing						

**Meaning** This output shows that Switch A is participating in LACP negotiation.

#### *Verifying That LACP Is Active on Switch B*

**Purpose** Verify that LACP is active on Switch B

**Action** [edit]  
 user@switch# show lacp interfaces  
 Aggregated interface: ae1

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/44	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/44	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
LACP protocol:	Receive State	Transmit State	Mux State						
xe-0/0/44	Current	Fast periodic	Collecting distributing						

**Meaning** This output shows that Switch B is participating in LACP negotiation.

#### *Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A*

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

**Action** [edit]  
user@switch# **show interfaces mc-ae**  
Member Link : ae1  
Current State Machine's State: mcae active state  
Local Status : active  
Local State : up  
Peer Status : active  
Peer State : up  
Logical Interface : ae1.0  
Topology Type : bridge  
Local State : up  
Peer State : up  
Peer Ip/MCP/State : 3.3.3.1 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch A is up and active.

***Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B***

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch B.

**Action** [edit]  
user@switch# **show interfaces mc-ae**  
Member Link : ae1  
Current State Machine's State: mcae active state  
Local Status : active  
Local State : up  
Peer Status : active  
Peer State : up  
Logical Interface : ae1.0  
Topology Type : bridge  
Local State : up  
Peer State : up  
Peer Ip/MCP/State : 3.3.3.2 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch B is up and active.

***Verifying that MAC Learning Is Occurring on Switch A and Switch B***

**Purpose** Verify that MAC learning is working on Switch A and B.

**Action** [edit]  
user@switch# **show ethernet-switching table**  
Ethernet-switching table: 10 entries, 4 learned, 0 persistent entries

VLAN	MAC address	Type	Age	Interfaces
v222	*	Flood	-	All-members
v222	00:00:5e:00:01:01	Static	-	Router
v222	00:10:94:00:00:05	Learn(L)	33	ae0.0 (MCAE)
v222	84:18:88:df:ac:ae	Learn(R)	0	ae2.0

**Meaning** The output shows four learned MAC addresses entries.



## Troubleshooting

---

### *Troubleshooting a LAG That Is Down*

**Problem** The `show interfaces terse` command shows that the MC-LAG is **down**

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2493](#)

### Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP)

There are two methods for enabling Layer 3 unicast functionality across a multichassis link aggregation group (MC-LAG). You can choose either to configure Virtual Router Redundancy Protocol (VRRP) or synchronize the MAC addresses for the Layer 3 interfaces of the switches participating in the MC-LAG. The procedure to configure VRRP for use in a Layer 3 unicast MC-LAG is included in this example. For more information on configuring MAC address synchronization, see *Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization*.

- [Requirements on page 2551](#)
- [Overview on page 2552](#)
- [Configuration on page 2553](#)
- [Verification on page 2572](#)
- [Troubleshooting on page 2577](#)

## Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 12.3 or later for the QFX Series
- Two QFX3500 or QFX3600 or QFX5100 switches

Before you configure an MC-LAG, be sure that you understand how to:

- Configure aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch”](#) on page 2462.
- Configure the Link Aggregation Control Protocol (LACP) on aggregated Ethernet interfaces on a switch. See [“Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch”](#) on page 2466.
- Configure Virtual Router Redundancy Protocol (VRRP) on a switch. See [“Configuring Basic VRRP Support”](#) on page 2285.

---

## Overview

In this example, you configure an MC-LAG across two switches by including interfaces from both switches in an aggregated Ethernet interface (ae1). To support the MC-LAG, create a second aggregated Ethernet interface (ae0) for the interchassis control link-protection link (ICL-PL). Configure a multichassis protection link for the ICL-PL, Interchassis Control Protocol (ICCP) for the peers hosting the MC-LAG, and Layer 3 connectivity between MC-LAG peers.



**NOTE:** Layer 3 connectivity is required for ICCP.

---

To complete the configuration, enable VRRP by completing the following steps:

- Create a routed VLAN interface (RVI)
- Create a VRRP group and assign a virtual IP address that is shared between each switch in the VRRP group
- Enable a member of a VRRP group to accept all packets destined for the virtual IP address if it is the master in the VRRP group
- Configure Layer 3 connectivity between the VRRP groups

## Topology

The topology used in this example consists of two switches hosting an MC-LAGs. The two switches are connected to a server. [Figure 54 on page 2553](#) shows the topology of this example.

Figure 54: Configuring a Multichassis LAG Between Switch A and Switch B




Table 235 on page 2463 details the topology used in this configuration example.

Table 239: Components of the Topology for Configuring a Multichassis LAG Between Two Switches

Hostname	Base Hardware	Multichassis Link Aggregation Group
Switch A	QFX3500, QFX3600, or QFX5100 switch	<b>ae0</b> is configured as an aggregated Ethernet interface, and is used as an ICL-PL. The following interfaces are part of <b>ae0</b> : <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch A and <b>xe-0/0/12</b> and <b>xe-0/0/13</b> on Switch B.  <b>ae1</b> is configured as an MC-LAG, and the following two interfaces are part of <b>ae1</b> : <b>xe-0/0/44</b> on Switch A and <b>xe-0/0/46</b> on Switch B.
Switch B		

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch A.



**NOTE:** This example shows how to configure MC-LAG using both the original CLI and Enhanced Layer 2 Software (ELS).

In ELS, there are three different statements and one different option from the original CLI:

- The port-mode statement in the [edit interfaces *interface-name* unit *number* family ethernet-switching] hierarchy is not supported. Use the interface-mode statement instead.
- The vlan statement in the [edit interfaces *interface-name*] hierarchy is not supported. Use the irb statement instead.
- The vlan.logical-interface-number option in the [edit vlans *vlan-name* l3-interface] option is not supported. Use the irb.logical-interface-number option instead.
- The service-id statement in the [edit switch-options] hierarchy is required in the ELS CLI.

#### Original CLI:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 100.1.1.1/24 vrrp-group 1 virtual-address 100.1.1.1
set interfaces vlan unit 100 family inet address 100.1.1.1/24 vrrp-group 1 priority 200
set interfaces vlan unit 100 family inet address 100.1.1.1/24 vrrp-group 1 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpd-block-on-edge
```

```
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

#### ELS:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/44 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 0
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control active
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 100 family inet address 100.1.1.1/24 vrrp-group 1 virtual-address 100.1.1.1
set interfaces irb unit 100 family inet address 100.1.1.1/24 vrrp-group 1 priority 200
set interfaces irb unit 100 family inet address 100.1.1.1/24 vrrp-group 1 accept-data
set interfaces irb unit 500 family inet address 3.3.3.2/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.2
set protocols iccp peer 3.3.3.1 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
set switch-options service-id 10
```

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and paste the commands into the CLI at the **[edit]** hierarchy level of Switch B.

#### Original CLI:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
```

```
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 virtual-address 100.1.1.1
set interfaces vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
set interfaces vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
set interfaces vlan unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface vlan.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface vlan.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

**ELS:**

```
set chassis aggregated-devices ethernet device-count 2
set interfaces xe-0/0/12 ether-options 802.3ad ae0
set interfaces xe-0/0/13 ether-options 802.3ad ae0
set interfaces xe-0/0/46 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae0 unit 0 family ethernet-switching vlan members v500 v100
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
set interfaces ae1 aggregated-ether-options lacp admin-key 3
set interfaces ae1 aggregated-ether-options mc-ae mc-ae-id 3
set interfaces ae1 aggregated-ether-options mc-ae chassis-id 1
set interfaces ae1 aggregated-ether-options mc-ae mode active-active
set interfaces ae1 aggregated-ether-options mc-ae status-control standby
set ae1 aggregated-ether-options mc-ae init-delay-time 240
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
set interfaces ae1 unit 0 family ethernet-switching vlan members v100
set interfaces irb unit 100 family inet address 100.1.1.10/24 vrrp-group 1 virtual-address 100.1.1.1
set interfaces irb unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
set interfaces irb unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
set interfaces irb unit 500 family inet address 3.3.3.1/24
set vlans v100 vlan-id 100
set vlans v100 l3-interface irb.100
set vlans v500 vlan-id 500
set vlans v500 l3-interface irb.500
set protocols iccp local-ip-addr 3.3.3.1
set protocols iccp peer 3.3.3.2 session-establishment-hold-time 50
set protocols iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.233
set protocols iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
set protocols iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
set protocols rstp interface ae0.0 disable
set protocols rstp interface ae1.0 edge
set protocols rstp interface all mode point-to-point
set protocols rstp bpdu-block-on-edge
set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
set switch-options service-id 10
```

### Configuring MC-LAG on Two Switches

#### Step-by-Step Procedure

To enable multichassis protection link between MC-LAG peers:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the number of LAGs on both Switch A and Switch B.  

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 2
```
2. Add member interfaces to the aggregated Ethernet interfaces on both Switch A and Switch B.  

```
[edit interfaces]
user@switch# set xe-0/0/12 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/13 ether-options 802.3ad ae0
[edit interfaces]
user@switch# set xe-0/0/44 ether-options 802.3ad ae1
[edit interfaces]
user@switch# set xe-0/0/46 ether-options 802.3ad ae1
```
3. Configure a trunk interface between Switch A and Switch B using the original CLI.  

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
```
4. Configure a trunk interface between Switch A and Switch B using ELS.  

```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```
5. Configure a multichassis protection link between Switch A and Switch B.  

Switch A:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.2 interface ae0
```

Switch B:

```
[edit]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

**Step-by-Step  
Procedure**

To enable ICCP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the local IP address to be in the ICCP connection on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.2
```

Switch B:

```
[edit protocols]
user@switch# set iccp local-ip-addr 3.3.3.1
```

2. Configure the peer IP address and minimum receive interval for a (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection minimum-receive-interval 1000
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
```

3. Configure the peer IP address and minimum transmit interval for Bidirectional Forwarding Detection (BFD) session for ICCP on Switch A and Switch B.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 liveness-detection transmit-interval minimum-interval 1000
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval 1000
```

4. (Optional) Configure the time during which an ICCP connection must succeed between MC-LAG peers on Switch A and Switch B.



**NOTE:** Configuring session establishment hold time helps in faster ICCP connection establishment. The recommended value is 50 seconds.

---

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 session-establishment-hold-time 50
```

Switch B:

```
[edit protocols]
```



- user@switch# **set iccp peer 3.3.3.2 session-establishment-hold-time 50**
5. (Optional) Configure the backup IP address to be used for backup liveness detection on both Switch A and Switch B.



**NOTE:** By default, backup liveness detection is not enabled. Configuring a backup IP address helps achieve sub-second traffic loss during a MC-LAG peer reboot.

Switch A:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.1 backup-liveness-detection backup-peer-ip 10.207.64.233
```

Switch B:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip 10.207.64.232
```

6. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B using the original CLI.

```
[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface vlan.500
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk vlan members v500 v100
```

7. Configure Layer 3 connectivity between the MC-LAG peers on both Switch A and Switch B using ELS.

```
[edit vlans]
user@switch# set v500 vlan-id 500
[edit vlans]
user@switch# set v500 l3-interface irb.500
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk vlan members v500 v100
```

**Step-by-Step Procedure**

To enable the MC-LAG interface:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable LACP on the MC-LAG interface on Switch A and Switch B.



**NOTE:** At least one end needs to be active. The other end can be either active or passive.

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options lacp active**

2. Specify the same multichassis aggregated Ethernet identification number on both MC-LAG peers on Switch A and Switch B.

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae mc-ae-id 3**

3. Specify the same service ID on Switch A and Switch B.

ELS:

[edit]

user@switch# **set switch-options service-id 10**

4. Specify a unique chassis ID for the MC-LAG on the MC-LAG peers on Switch A and Switch B.

Switch A:

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae chassis-id 0**

Switch B:

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae chassis-id 1**

5. Specify the operating mode of the MC-LAG on both Switch A and Switch B.



**NOTE:** Only active-active mode is supported at this time.

[edit interfaces]

user@switch# **set ae1 aggregated-ether-options mc-ae mode active-active**

6. Specify the status control for MC-LAG on Switch A and Switch B.



**NOTE:** You must configure status control on both Switch A and Switch B hosting the MC-LAG. If one peer is in active mode, the other must be in standby mode.

Switch A:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-aether-options status-control active
```

Switch B:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-aether-options status-control standby
```

7. Specify the number of seconds by which the bring-up of the MC-AE interface should be deferred after you reboot Switch A and Switch B.



**NOTE:** The recommended value for maximum VLAN configuration (for example, 4,000 VLANs) is 240 seconds. If IGMP snooping is enabled on all of the VLANs, the recommended value is 420 seconds.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-aether-options init-delay-time 240
```

8. Specify the same LACP system ID for the MC-LAG on Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

9. Specify the same LACP administration key on both Switch A and Switch B.

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

10. Enable a VLAN on the MC-LAG on Switch A and Switch B using the original CLI.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching port-mode trunk
[edit]
```

```
user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

11. Enable a VLAN on the MC-LAG on Switch A and Switch B using ELS.

```
[edit interfaces]
user@switch# set ae1 unit 0 family ethernet-switching interface-mode trunk
[edit]
```

```
user@switch# set vlans v100 vlan-id 100
```

```
[edit interfaces]
```

```
user@switch# set ae1 unit 0 family ethernet-switching vlan members v100
```

12. Enable VRRP on the MC-LAG on Switch A and Switch B:

- Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each switch in the VRRP group:

Switch A:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.1/24 vrrp-group 1
virtual-address 100.1.1.1
```

Switch B:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1
virtual-address 100.1.1.1
```

- Assign the priority for each switch in the VRRP group:



**NOTE:** The switch configured with the highest priority is the master.

Switch A:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 priority 200
```

Switch B:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
```

- Enable the switch to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

Switch A:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.11/24 vrrp-group 1 accept-data
```

Switch B:

```
[edit interfaces]
user@switch# set vlan unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
```

- Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
user@switch# set vlans v100 l3-interface vlan.100
```

13. Enable VRRP on the MC-LAG on Switch A and Switch B using ELS:

- Create a routed VLAN interface (RVI), assign a virtual IP address that is shared between each switch in the VRRP group, and assign an individual IP address for each switch in the VRRP group:

Switch A:

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 100.1.1.11/24 vrrp-group 1 virtual-address 100.1.1.1
```

Switch B:

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 100.1.1.10/24 vrrp-group 1 virtual-address 100.1.1.1
```

- Assign the priority for each switch in the VRRP group:



**NOTE:** The switch configured with the highest priority is the master.

Switch A:

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 100.1.1.11/24 vrrp-group 1 priority 200
```

Switch B:

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 100.1.1.10/24 vrrp-group 1 priority 150
```

- Enable the switch to accept all packets destined for the virtual IP address if it is the master in the VRRP group:

Switch A:

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 100.1.1.11/24 vrrp-group 1 accept-data
```

Switch B:

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 100.1.1.10/24 vrrp-group 1 accept-data
```

- Configure Layer 3 connectivity between Switch A and Switch B.

```
[edit interfaces]
user@switch# set irb v100 l3-interface irb.100
```

### Step-by-Step Procedure

To enable RSTP:

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Enable RSTP globally on all interfaces on Switch A and Switch B.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```

2. Disable RSTP on the ICL-PL interfaces on Switch A and Switch B:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

3. Configure the MC-LAG interfaces as edge ports on Switch A and Switch B.



**NOTE:** The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@switch# set protocols rstp interface ae1.0 edge
```

4. Enable BPDU blocking on all interfaces except for the ICL-PL interfaces on Switch A and Switch B.



**NOTE:** The ae1 interface is a downstream interface. This is why RSTP and bpdu-block-on-edge need to be configured.

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

### Results

Display the results of the configuration on Switch A using the original CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
  ae0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v500;
        }
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 0;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members v100;
        }
      }
    }
  }
}
```

```

}
vlan {
  unit 100 {
    family inet {
      address 100.1.1.1/24 {
        vrrp-group 1 {
          virtual-address 100.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 3.3.3.2/24;
    }
  }
}
protocols {
  iccp {
    local-ip-addr 3.3.3.2;
    peer 3.3.3.1 {
      session-establishment-hold-time 50;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
  }
  rstp {
    interface ae0.0 {
      disable;
    }
    interface ae1.0 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
    bpdu-block-on-edge;
  }
}
multi-chassis {
  multi-chassis-protection 3.3.3.1 {
    interface ae0;
  }
}
vlans {
  v100 {

```

```
        vlan-id 100;
        l3-interface vlan.100;
    }
    v500 {
        vlan-id 500;
        l3-interface vlan.500;
    }
}
```

Display the results of the configuration on Switch A using ELS.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      system-id 00:01:02:03:04:05;
      admin-key 3;
    }
  }
  mc-ae {
    mc-ae-id 3;
    chassis-id 0;
    mode active-active;
    status-control active;
  }
}
```



```

        init-delay-time 240;
    }
}
unit 0 {
    family ethernet-switching {
        interface-mode trunk;
        vlan {
            members v100;
        }
    }
}
vlan {
    unit 100 {
        family inet {
            address 100.1.1.1/24 {
                vrrp-group 1 {
                    virtual-address 100.1.1.1;
                    priority 200;
                    accept-data;
                }
            }
        }
    }
    unit 500 {
        family inet {
            address 3.3.3.2/24;
        }
    }
}
}
protocols {
    iccp {
        local-ip-addr 3.3.3.2;
        peer 3.3.3.1 {
            session-establishment-hold-time 50;
            backup-liveness-detection {
                backup-peer-ip 10.207.64.233;
            }
            liveness-detection {
                minimum-receive-interval 1000;
                transmit-interval {
                    minimum-interval 1000;
                }
            }
        }
    }
}
rstp {
    interface ae0.0 {
        disable;
    }
    interface ae1.0 {
        edge;
    }
    interface all {
        mode point-to-point;
    }
}

```

```
    }
    bpdv-block-on-edge;
  }
}
multi-chassis {
  multi-chassis-protection 3.3.3.1 {
    interface ae0;
  }
}
switch-options {
  service-id 10;
}
vlangs {
  v100 {
    vlan-id 100;
    l3-interface irb.100;
  }
  v500 {
    vlan-id 500;
    l3-interface irb.500;
  }
}
```

Display the results of the configuration on Switch B using the original CLI.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
      802.3ad ae1;
    }
  }
}
ae0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v500;
      }
    }
  }
}
```

```

    }
  }
  ae1 {
    aggregated-ether-options {
      lacp {
        active;
        system-id 00:01:02:03:04:05;
        admin-key 3;
      }
      mc-ae {
        mc-ae-id 3;
        chassis-id 1;
        mode active-active;
        status-control active;
        init-delay-time 240;
      }
    }
  }
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members v100;
      }
    }
  }
}
vlan {
  unit 100 {
    family inet {
      address 100.1.1.10/24 {
        vrrp-group 1 {
          virtual-address 100.1.1.1;
          priority 200;
          accept-data;
        }
      }
    }
  }
  unit 500 {
    family inet {
      address 3.3.3.1/24;
    }
  }
}
}
protocols {
  iccp {
    local-ip-addr 3.3.3.1;
    peer 3.3.3.2 {
      session-establishment-hold-time 50;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {

```

```
        minimum-interval 1000;
      }
    }
  }
  rstp {
    interface ae0.0 {
      disable;
    }
    interface ae1.0 {
      edge;
    }
    interface all {
      mode point-to-point;
    }
    bpdu-block-on-edge;
  }
}
multi-chassis {
  multi-chassis-protection 3.3.3.2 {
    interface ae0;
  }
}
vllans {
  v100 {
    vlan-id 100;
    l3-interface vlan.100;
  }
  v500 {
    vlan-id 500;
    l3-interface vlan.500;
  }
}
```

Display the results of the configuration on Switch B using ELS.

```
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  xe-0/0/12 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/13 {
    ether-options {
      802.3ad ae0;
    }
  }
  xe-0/0/44 {
    ether-options {
```

```

        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v500;
            }
        }
    }
}
ae1 {
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:01:02:03:04:05;
            admin-key 3;
        }
        mc-ae {
            mc-ae-id 3;
            chassis-id 1;
            mode active-active;
            status-control active;
            init-delay-time 240;
        }
    }
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members v100;
            }
        }
    }
}
vlan {
    unit 100 {
        family inet {
            address 100.1.1.10/24 {
                vrrp-group 1 {
                    virtual-address 100.1.1.1;
                    priority 200;
                    accept-data;
                }
            }
        }
    }
    unit 500 {
        family inet {
            address 3.3.3.1/24;
        }
    }
}
}

```

```
protocols {
  iccp {
    local-ip-addr 3.3.3.1;
    peer 3.3.3.2 {
      session-establishment-hold-time 50;
      backup-liveness-detection {
        backup-peer-ip 10.207.64.233;
      }
      liveness-detection {
        minimum-receive-interval 1000;
        transmit-interval {
          minimum-interval 1000;
        }
      }
    }
  }
}
rstp {
  interface ae0.0 {
    disable;
  }
  interface ae1.0 {
    edge;
  }
  interface all {
    mode point-to-point;
  }
  bpdu-block-on-edge;
}
}
multi-chassis {
  multi-chassis-protection 3.3.3.2 {
    interface ae0;
  }
}
switch-options {
  service-id 10;
}
}
vllans {
  v100 {
    vlan-id 100;
    l3-interface irb.100;
  }
  v500 {
    vlan-id 500;
    l3-interface irb.500;
  }
}
}
```

---

## Verification

To verify that the MC-LAG group has been created and is working properly, perform these tasks:

- [Verifying That ICCP Is Working on Switch A on page 2573](#)
- [Verifying That ICCP Is Working on Switch B on page 2573](#)

- [Verifying That LACP Is Active on Switch A on page 2574](#)
- [Verifying That LACP Is Active on Switch B on page 2574](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A on page 2574](#)
- [Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B on page 2574](#)
- [Verifying that MAC Learning Is Occurring on Switch A on page 2575](#)
- [Verifying that MAC Learning Is Occurring on Switch B on page 2575](#)
- [Verifying that Switch A is the Master in the VRRP Group on page 2576](#)
- [Verifying that Switch B is the Backup Member in the VRRP Group on page 2576](#)
- [Verifying that the Virtual IP Address is Attached to an Individual Address on Switch A on page 2577](#)
- [Verifying that the Virtual IP Address is Attached to an Individual Address on Switch B on page 2577](#)

### ***Verifying That ICCP Is Working on Switch A***

**Purpose** Verify that ICCP is running on Switch A.

**Action** [edit]  
 user@switch# **show iccp**  
 Redundancy Group Information for peer 3.3.3.1  
     TCP Connection : Established  
     Liveliness Detection : Up  
  
 Client Application: MCSNOOPD  
  
 Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

### ***Verifying That ICCP Is Working on Switch B***

**Purpose** Verify that ICCP is running on Switch B.

**Action** **show iccp**  
  
 [edit]  
 user@switch# **show iccp**  
 Redundancy Group Information for peer 3.3.3.2  
     TCP Connection : Established  
     Liveliness Detection : Up  
  
 Client Application: MCSNOOPD  
  
 Client Application: eswd

**Meaning** This output shows that the TCP connection between the peers hosting the MC-LAG is up, liveness detection is up, and MCSNOOPD and ESWD client applications are running.

**Verifying That LACP Is Active on Switch A**

**Purpose** Verify that LACP is active on Switch A.

**Action** [edit]  
user@switch# show lacp interfaces  
Aggregated interface: ae1

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/46	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/46	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active

LACP protocol:                      Receive State      Transmit State                      Mux State  
xe-0/0/46                                      Current      Fast periodic      Collecting      distributing

**Meaning** This output shows that Switch A is participating in LACP negotiation.

**Verifying That LACP Is Active on Switch B**

**Purpose** Verify that LACP is active on Switch B

**Action** [edit]  
user@switch# show lacp interfaces  
Aggregated interface: ae1

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/0/44	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-0/0/44	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active

LACP protocol:                      Receive State      Transmit State                      Mux State  
xe-0/0/44                                      Current      Fast periodic      Collecting      distributing

**Meaning** This output shows that Switch B is participating in LACP negotiation.

**Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch A**

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch A.

**Action** [edit]  
user@switch# show interfaces mc-ae

```
Member Link                : ae1
Current State Machine's State: mcae active state
Local Status                : active
Local State                  : up
Peer Status                  : active
Peer State                   : up
  Logical Interface          : ae1.0
  Topology Type               : bridge
  Local State                 : up
  Peer State                  : up
  Peer Ip/MCP/State           : 3.3.3.1 ae0.0 up
```

**Meaning** This output shows that the MC-AE interface on Switch A is up and active.

**Verifying That the MC-AE and ICL-PL Interfaces Are Up on Switch B**

**Purpose** Verify that the MC-AE and ICL-PL interfaces are up on Switch B.



**Action** [edit]  
 user@switch# **show interfaces mc-ae**  
 Member Link : ae1  
 Current State Machine's State: mcae active state  
 Local Status : active  
 Local State : up  
 Peer Status : active  
 Peer State : up  
 Logical Interface : ae1.0  
 Topology Type : bridge  
 Local State : up  
 Peer State : up  
 Peer Ip/MCP/State : 3.3.3.2 ae0.0 up

**Meaning** This output shows that the MC-AE interface on Switch B is up and active.

#### *Verifying that MAC Learning Is Occurring on Switch A*

**Purpose** Verify that MAC learning is working on Switch A.

**Action** [edit]  
 user@switch# **show ethernet-switching table**  
 Ethernet-switching table: 6 entries, 1 learned, 0 persistent entriesC

VLAN	MAC address	Type	Age	Interfaces
v100	*	Flood		- All-members
v100	00:00:5e:00:01:01	Static		- Router
v100	78:fe:3d:5a:07:42	Static		- Router
v100	78:fe:3d:5b:ad:c2	Learn(R)	0	ae0.0
v500	*	Flood		- All-members
v500	78:fe:3d:5a:07:42	Static		- Router

**Meaning** The output shows two static MAC address in VLAN v100 and one static MAC address in VLAN v500. These addresses belong to the Layer 3 RVI addresses on both Switch A and Switch B that you configured in the MC-LAG. The ICL-PL interface configured on the VRRP master member learned the VLAN v100 Learn (R) MAC address of the VRRP backup member.

#### *Verifying that MAC Learning Is Occurring on Switch B*

**Purpose** Verify that MAC learning is working on Switch B.

**Action** [edit]user@switch# **show ethernet-switching table**

Ethernet-switching table: 7 entries, 1 learned, 0 persistent entries

VLAN	MAC address	Type	Age	Interfaces
v100	*	Flood		- All-members
v100	00:00:5e:00:01:01	Static		- Router
v100	78:fe:3d:5a:07:42	Learn(R)	0	ae0.0
v100	78:fe:3d:5b:ad:c2	Static		- Router
v200	78:fe:3d:5b:ad:c2	Static		- Router
v500	*	Flood		- All-members
v500	78:fe:3d:5b:ad:c2	Static		- Router

**Meaning** The output shows two static MAC address in VLAN v100 and one static MAC address in VLAN v500. These addresses belong to the Layer 3 RVI addresses on both Switch A and Switch B that you configured in the MC-LAG. The ICL-PL interface configured on the VRRP backup member learned the VLAN v100 Learn (R) MAC address of the VRRP master member.

*Verifying that Switch A is the Master in the VRRP Group*

**Purpose** Verify that Switch A is the master member in the VRRP group.

**Action** [edit]user@switch# **show vrrp**

Interface	State	Group	VR state	VR Mode	Timer	Type	Address
vlan.100	up	1	master	Active	A 0.605	lcl	100.1.1.11
						vip	100.1.1.1

**Meaning** The output shows that Switch A is the master member in the VRRP group.

*Verifying that Switch B is the Backup Member in the VRRP Group*

**Purpose** Verify that Switch B is the backup member in the VRRP group.

**Action** [edit]user@switch# **show vrrp**

Interface	State	Group	VR state	VR Mode	Timer	Type	Address
vlan.100	up	1	backup	Active	A 0.605	lcl	100.1.1.10
						vip	100.1.1.1

**Meaning** The output shows that Switch B is the backup member in the VRRP group.

*Verifying that the Virtual IP Address is Attached to an Individual Address on Switch A*

**Action** [edit]  
 user@switch# run show interfaces terse vlan

Interface	Admin	Link	Proto	Local	Remote
vlan	up	up			
vlan.100	up	up	inet	100.1.1.1/24	100.1.1.11/24
vlan.500	up	up	inet	3.3.3.2/24	

**Meaning** The output shows that the virtual IP address (100.1.1.1/24) is bound to the individual IP address (100.1.1.11/24) on Switch A.

*Verifying that the Virtual IP Address is Attached to an Individual Address on Switch B*

**Action** [edit]  
 user@switch# run show interfaces terse vlan

Interface	Admin	Link	Proto	Local	Remote
vlan	up	up			
vlan.100	up	up	inet	100.1.1.1/24	100.1.1.10/24
vlan.500	up	up	inet	3.3.3.1/24	

**Meaning** The output shows that the virtual IP address (100.1.1.1/24) is bound to the individual IP address (100.1.1.10/24) on Switch B.

**Troubleshooting***Troubleshooting a LAG That Is Down*

**Problem** The show interfaces terse command shows that the MC-LAG is down

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that the MC-LAG is part of family Ethernet switching (Layer 2 LAG).
- Verify that the MC-LAG member is connected to the correct MC-LAG member at the other end.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)

## Example: Configuring Redundant Trunk Links for Faster Recovery



**NOTE:** This example uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Example: Configuring Redundant Trunk Links for Faster Recovery*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 43.

You can manage network convergence by configuring both a primary link and a secondary link on a switch; this is called a redundant trunk group (RTG). If the primary link in a redundant trunk group fails, it passes its known MAC address locations to the secondary link, which automatically takes over after one minute.

This example describes how to create a redundant trunk group with a primary and a secondary link:

- [Requirements on page 2578](#)
- [Overview and Topology on page 2578](#)
- [Disabling RSTP on Switches 1 and 2 on page 2581](#)
- [Configuring Redundant Trunk Links on Switch 3 on page 2581](#)
- [Verification on page 2582](#)

### Requirements

---

This example uses the following hardware and software components:

- Two EX Series or QFX Series distribution switches
- One EX Series or QFX Series access switch
- The appropriate software release for your platform:
  - For EX Series switches: Junos OS Release 13.2X50-D10 or later
  - For the QFX Series: Junos OS Release 13.2X50-D15 or later

Before you configure the redundant trunk links network on the access and distribution switches, be sure you have:

- Configured interfaces ge-0/0/9 and ge-0/0/10 on the access switch, Switch 3, as trunk interfaces. .
- Configured one trunk interface on each distribution switch, Switch 1 and Switch 2.
- Connected the three switches as shown in the topology for this example (see [Figure 55 on page 2580](#)).

### Overview and Topology

---

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for trunk interface network recovery. When a trunk

interface fails, data traffic is routed to another trunk interface after one minute, thereby keeping network convergence time to a minimum.

This example shows the configuration of a redundant trunk group that includes one primary link (and its interface) and one unspecified link (and its interface) that serves as the secondary link.

A second type of redundant trunk group, not illustrated in the example, consists of two unspecified links (and their interfaces); in this case, neither of the links is primary. The software selects an active link by comparing the port numbers of the two links and activating the link with the higher port number. For example, if the two link interfaces use interfaces ge-0/1/0 and ge-0/1/1, the software activates ge-0/1/1. (In the interface names, the final number is the port number.)

The two links in a redundant trunk group generally operate the same way, whether they are configured as primary/unspecified or unspecified/unspecified. Data traffic initially passes through the active link but is blocked on the inactive link. While data traffic is blocked on the secondary link, note that Layer 2 control traffic is still permitted if the link is active. For example, an LLDP session can be run between two switches on the secondary link. If the active link either goes down or is disabled administratively, it broadcasts a list of its known MAC addresses for data traffic; the other link immediately picks up and adds the MAC addresses to its address table, becomes active, and begins forwarding traffic.

The one difference in operation between the two types of redundant trunk groups occurs when a primary link is active, goes down, is replaced by the secondary link, and then reactivates. When a primary link is re-enabled like this while the secondary link is active, the primary link waits 2 minutes (you can change the time interval by using the preempt cutover timer to accommodate your network) and then takes over as the active link. In other words, the primary link has priority and is always activated if it is available. This differs from the behavior of two unspecified links, both of which act as equals. Because the unspecified links are equal, the active link remains active until it either goes down or is disabled administratively; this is the only time that the other unspecified link learns the MAC addresses and immediately becomes active.

The example given here illustrates a primary/unspecified configuration for a redundant trunk group because that configuration gives you more control and is more commonly used.



**NOTE:** Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You will need to disable RSTP on the two distribution switches in the example, Switch 1 and Switch 2. Spanning-tree protocols can, however, continue operating in other parts of the network—for example, between the distribution switches and also in links between distribution switches and the enterprise core.

Figure 55 on page 2580 displays an example topology containing three switches. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer.

Switch 3 is connected to the distribution layer through trunk interfaces ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2).

Table 240 on page 2580 lists the components used in this redundant trunk group.

Because RSTP and RTGs cannot operate simultaneously on a switch, you disable RSTP on Switch 1 and Switch 2 in the first configuration task, and you disable RSTP on Switch 3 in the second task.

The second configuration task creates a redundant trunk group called example 1 on Switch 3. The trunk interfaces ge-0/0/9.0 and ge-0/0/10.0 are the two links configured in the second configuration task. You configure the trunk interface ge-0/0/9.0 as the primary link. You configure the trunk interface ge-0/0/10.0 as an unspecified link, which becomes the secondary link by default.

Figure 55: Topology for Configuring the Redundant Trunk Links




Table 240: Components of the Redundant Trunk Link Topology

Property	Settings
Switch hardware	<ul style="list-style-type: none"> <li>Switch 1—1 EX Series or QFX Series distribution switch</li> <li>Switch 2—1 EX Series or QFX Series distribution switch</li> <li>Switch 3—1 EX Series or QFX Series access switch</li> </ul>
Trunk interfaces	On Switch 3 (access switch): ge-0/0/9.0 and ge-0/0/10.0
Redundant trunk group	example1

### Disabling RSTP on Switches 1 and 2

To disable RSTP on Switch 1 and Switch 2, perform this task on each switch:

**CLI Quick Configuration** To quickly disable RSTP on Switch 1 and Switch 2, copy the following command and paste it into each switch terminal window:

```
[edit]
set protocols rstp disable
```

**Step-by-Step Procedure** To disable RSTP on Switch 1 and Switch 2:

1. Disable RSTP on Switch 1 and Switch 2:
- ```
[edit]
user@switch# set protocols rstp disable
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
protocols {
  rstp {
    disable;
  }
}
```

### Configuring Redundant Trunk Links on Switch 3

To configure redundant trunk links on Switch 3, perform this task:

**CLI Quick Configuration** To quickly configure the redundant trunk group example1 on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp disable
set switch-options redundant-trunk-group group example1 interface ge-0/0/9.0 primary
set switch-options redundant-trunk-group group example1 interface ge-0/0/10.0
set redundant-trunk-group group example1 preempt-cutover-timer 60
```

**Step-by-Step Procedure** Configure the redundant trunk group example1 on Switch 3.

1. Turn off RSTP:
 

```
[edit]
user@switch# set protocols rstp disable
```
2. Name the redundant trunk group example1 while configuring trunk interface ge-0/0/9.0 as the primary link and ge-0/0/10 as an unspecified link to serve as the secondary link:
 

```
[edit switch-options]
user@switch# set redundant-trunk-group group example1 interface ge-0/0/9.0 primary
user@switch# set redundant-trunk-group group example1 interface ge-0/0/10.0
```
3. (Optional) Change the time interval (from the default 120 seconds) that a re-enabled primary link waits to take over for an active secondary link:
 

```
[edit switch-options]
user@switch# set redundant-trunk-group group example1 preempt-cutover-timer 60
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
switch-options
  redundant-trunk-group {
    group example1 {
      preempt-cutover-timer 60;
      interface ge-0/0/9.0 {
        primary;
      }
      interface ge-0/0/10.0;
    }
  }
protocols {
  rstp {
    disable;
  }
}
```

---

### Verification

To confirm that the configuration is set up correctly, perform this task:

- [Verifying That a Redundant Trunk Group Was Created on page 2582](#)

#### *Verifying That a Redundant Trunk Group Was Created*

**Purpose** Verify that the redundant trunk group example1 has been created on Switch 1 and that trunk interfaces are members of the redundant trunk group.

**Action** List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group
```

| Group name | Interface   | State  | Time of last flap | Flap count |
|------------|-------------|--------|-------------------|------------|
| example1   | ge-0/0/9.0  | Up/Pri | Never             | 0          |
|            | ge-0/0/10.0 | Up     | Never             | 0          |

**Meaning** The **show redundant-trunk-group** command lists all redundant trunk groups configured on the switch as well as the interface names and their current states (up or down for an unspecified link, and up or down and primary for a primary link). For this configuration example, the output shows that the redundant trunk group example1 is configured on the switch. The **Up** beside the interfaces indicates that both link cables are physically connected. The **Pri** beside trunk interface ge-0/0/9.0 indicates that it is configured as the primary link.

**Related Documentation**

- [Understanding Redundant Trunk Links on page 2447](#)



## Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches

Junos OS for switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet interface:

- [Requirements on page 2583](#)
- [Overview and Topology on page 2583](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 1 on page 2583](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 2 on page 2584](#)
- [Verification on page 2585](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later for EX Series switches
- Two EX3200 or EX4200 switches connected directly

### Overview and Topology

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two EX4200 switches connected directly. Before you begin configuring Ethernet OAM LFM on two switches, connect the two switches directly through a trunk interface.

### Configuring Ethernet OAM Link Fault Management on Switch 1

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/0
set interface ge-0/0/0 link-discovery active
set interface ge-0/0/0 pdu-interval 800
set interface ge-0/0/0 remote-loopback
```

**Step-by-Step Procedure** To configure Ethernet OAM LFM on switch 1:

1. Enable IEEE 802.3ah OAM support on an interface:
 

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface (OAM LFM) ge-0/0/0
```

2. Specify that the interface initiates the discovery process by configuring the link discovery mode to **active**:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch1# set interface ge-0/0/0 link-discovery active
```
3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800 on switch 1:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch1# set interface pdu-interval 800
```
4. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Ensure that the remote DTE supports remote loopback mode. To set the remote DTE in loopback mode  

```
[edit protocols oam ethernet link-fault-management]  
user@switch1# set interface ge-0/0/0.0 remote-loopback
```

### Results

Check the results of the configuration:

```
[edit]  
user@switch1# show  
  
protocols {  
  oam {  
    ethernet {  
      link-fault-management {  
        interface ge-0/0/0 {  
          pdu-interval 800;  
          link-discovery active;  
          remote-loopback;  
        }  
      }  
    }  
  }  
}
```

---

### Configuring Ethernet OAM Link Fault Management on Switch 2

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM on switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management ]  
set interface ge-0/0/1  
set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Step-by-Step Procedure** To configure Ethernet OAM LFM on switch 2:

1. Enable OAM on the peer interface on switch 2:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch2# set interface ge-0/0/1
```
2. Enable remote loopback support for the local interface:  

```
[edit protocols oam ethernet link-fault-management]  
user@switch2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Results** Check the results of the configuration:

```
[edit]
```

```

user@switch2# show

protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/1 {
          negotiation-options {
            allow-remote-loopback;
          }
        }
      }
    }
  }
}

```

### Verification

#### *Verifying That OAM LFM Has Been Configured Properly*

**Purpose** Verify that OAM LFM has been configured properly.

**Action** Use the `show oam ethernet link-fault-management` command:

```
user@switch1#show oam ethernet link-fault-management
```

### Sample Output

```

Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported

```

**Meaning** When the output displays the MAC address and the discover state is **Send Any**, it means that OAM LFM has been configured properly.

**Related Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)
- [Understanding Ethernet OAM Link Fault Management on page 2452](#)

### Configuration Tasks

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)
- [Configuring Aggregated Ethernet LACP on page 2589](#)
- [Configuring Ethernet Loopback Capability on page 2589](#)
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2590](#)
- [Configuring Interfaces for Uplink Failure Detection on page 2592](#)

- [Configuring a Layer 3 Logical Interface on page 2593](#)
- [Configuring Link Aggregation on page 2593](#)
- [Configuring Local Link Bias \(CLI Procedure\) on page 2596](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)
- [Configuring Generic Routing Encapsulation Tunneling on page 2600](#)
- [Configuring the LPM Table With Junos OS 13.2x51-D10 on page 2602](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)
- [Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up on page 2606](#)
- [Configuring Resilient Hashing for Trunk/ECMP Groups on page 2607](#)

## Configuring Gigabit and 10-Gigabit Ethernet Interfaces

Devices include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
- Configures Storm Control on all 10-Gigabit Ethernet network interfaces
- Provides basic Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP) configuration

This topic describes:

- [Configuring Port Mode on page 2586](#)
- [Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces on page 2587](#)
- [Configuring the Speed of Gigabit Ethernet Copper SFP Interfaces on page 2588](#)
- [Configuring the IP Options on page 2588](#)

### Configuring Port Mode

---

If you are connecting a switch to other switches and to routers on the LAN, you need to assign the interface to a logical port and you need to configure the logical port as a trunk port.

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode on the original CLI:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
port-mode trunk
```

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode on the Enhanced Layer 2 software (ELS):

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
interface-mode trunk
```

### Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces

Devices include a factory default configuration that enables Gigabit Ethernet and 10-Gigabit Ethernet and interfaces with applicable link settings.

The following default configurations are available on Gigabit Ethernet interfaces:

- The speed for Gigabit Ethernet interfaces is set to 1000 Mbps by default. The speed for 1-Gigabit Ethernet Copper SFP interfaces is 1 Gbps by default.
- Gigabit Ethernet interfaces operate in full-duplex mode.
- Autonegotiation is not supported.

To enable autonegotiation, issue the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



**NOTE:** Do not use the **set interface *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- The speed for 10-Gigabit Ethernet interfaces is set to 10 Gbps by default. The speed cannot be configured.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is not supported.

To enable autonegotiation, issue the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



**NOTE:** Do not use the **set interface *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

The **ether-options** statement enables you to modify the following options:

- **802.3ad**—Specify an aggregated Ethernet bundle for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
- **autonegotiation**—Enable or disable autonegotiation of flow control, link mode, and speed for Gigabit Ethernet interfaces.
- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic** for Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.

To set **ether-options** for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces:

```
[edit]
user@switch# set interfaces interface-name ether-options
```

---

### Configuring the Speed of Gigabit Ethernet Copper SFP Interfaces

You can configure the speed of Gigabit Ethernet copper SFP interfaces on the EX4600 and QFX5100 devices. The default speed is 1Gbps.



**NOTE:** Autonegotiation is not supported on EX4600 and QFX5100 devices.

---

1. Configure the speed of the interface:

```
[edit]
user@switch# set interfaces interface-name speed speed
```

For example, to configure a speed of 100Mbps on the **ge-0/1/0** interface:

```
[edit]
user@switch# set interfaces ge-0/1/0 speed 100m
```

2. To delete the speed of the interface:

```
[edit]
user@switch# delete interfaces interface-name speed speed
```

For example, to delete a speed of 100Mbps on the **ge-0/1/0** interface:

```
[edit]
user@switch# delete interfaces ge-0/1/0 speed 100m
```

---

### Configuring the IP Options

To specify an IP address for the logical unit:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

#### Related Documentation

- [Monitoring Interface Status and Traffic on page 335](#)
- [show interfaces xe on page 2852](#)
- [show interfaces ge-](#)
- [speed on page 2732](#)
- [Understanding Interface Naming Conventions on page 2401](#)

## Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic interval
```

For example, to specify the interval as fast, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic fast
```

## Configuring Ethernet Loopback Capability

To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]
user@switch# delete interfaces interface-name ether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**

**Related  
Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)

## Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Juniper Networks EX Series and QFX Series switches use a hashing algorithm to determine how to forward traffic over a Link Aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

Configuring the fields used by the hashing algorithm is useful in scenarios where most of the traffic entering the bundle is similar and the traffic needs to be managed in the LAG bundle. For instance, if the only difference in the IP packets for all incoming traffic is the source and destination IP address, you can tune the hashing algorithm to make hashing decisions more efficiently by configuring the algorithm to make hashing decisions using only those fields.

- [Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing on page 2590](#)
- [Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing on page 2591](#)
- [Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing on page 2591](#)

### Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing

To configure the hashing algorithm to use fields in the Layer 2 header for hashing:

1. Configure the hash mode to Layer 2 header:

```
[edit forwarding-options enhanced-hash-key]  
user@switch# set hash-mode layer2-header
```

The default hash mode is Layer 2 payload. Therefore, this step must be performed if you have not previously configured the hash mode.

2. Configure the fields in the Layer 2 header that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]  
user@switch# set layer2 {no-destination-mac-address | no-ether-type |  
no-source-mac-address | vlan-id}
```

By default, the hashing algorithm uses the values in the destination MAC address, Ethertype, and source MAC address fields in the header to hash traffic on the LAG.



You can configure the hashing algorithm to not use the values in these fields by configuring **no-destination-mac-address**, **no-ether-type**, or **no-source-mac-address**.

You can also configure the hashing algorithm to include the VLAN ID field in the header by configuring the **vlan-id** option.

If you want the hashing algorithm to not use the Ethertype field for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set layer2 no-ether-type
```

### Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing

To configure the hashing algorithm to use fields in the IP payload for hashing:

1. Configure the hash mode to Layer 2 payload:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-payload
```

The IP payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IP payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet {no-ipv4-destination-address | no-ipv4-source-address |
no-l4-destination-port | no-l4-source-port | no-protocol | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and protocol fields and instead hash traffic based only on the IPv4 source and destination addresses:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet no-l4-destination-port no-l4-source-port no-protocol
```

### Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing

To configure the hashing algorithm to use fields in the IPv6 payload for hashing:

1. Configure the hash mode to Layer 2 payload:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set hash-mode layer2-payload
```

The IPv6 payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IPv6 payload that the hashing algorithm uses for hashing:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet6 {no-ipv6-destination-address | no-ipv6-source-address |
no-l4-destination-port | no-l4-source-port | no-next-header | vlan-id}
```

For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and the Next Header fields and instead hash traffic based only on the IPv6 source and IPv6 destination address fields only:

```
[edit forwarding-options enhanced-hash-key]
user@switch# set inet6 no-l4-destination-port no-l4-source-port no-next-header
```

#### Related Documentation

- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396](#)

- *Understanding Aggregated Ethernet Interfaces and LACP*

## Configuring Interfaces for Uplink Failure Detection

You can configure uplink failure detection to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure information to downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Follow these configuration guidelines:

- Configure an interface in only one group.
- Configure a maximum of eight groups for each switch.
- Configure a maximum of eight uplinks to monitor and a maximum of 48 downlinks to disable in each group.
- Configure physical links and logical links in separate groups.

To configure uplink failure detection on a switch:

1. Specify a name for an uplink failure detection group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Repeat Step 2 for each uplink interface you add to the group.

4. Add a downlink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

5. Repeat Step 4 for each downlink interface you add to the group.



**NOTE:** After you have configured an uplink failure detection group, use the `show uplink-failure-detection group (Uplink Failure Detection) group-name` command to verify that all interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.

### Related Documentation

- [Overview of Uplink Failure Detection on page 2392](#)
- [Example: Configuring Interfaces for Uplink Failure Detection on page 2457](#)
- [Verifying That Uplink Failure Detection Is Working Correctly](#)

## Configuring a Layer 3 Logical Interface

Devices use Layer 3 logical interfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. Layer 3 logical interfaces route traffic between subnets.

To configure Layer 3 logical interfaces, enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See *Configuring VLANs*.

To configure Layer 3 logical interfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]
user@switch# set unit logical-unit-number vlan-id vlan-id-number
```

### Related Documentation

- [Understanding Layer 3 Logical Interfaces on page 2408](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 2750](#)

## Configuring Link Aggregation

Use the link aggregation feature to aggregate one or more links to form a virtual link or aggregation group. The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases link availability.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregation group.



**NOTE:** On QFX5100 and EX4600 standalone switches and on QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.

1. [Creating an Aggregated Ethernet Interface on page 2594](#)
2. [Configuring the VLAN Name and VLAN ID Number on page 2594](#)
3. [Configuring Aggregated Ethernet LACP on page 2594](#)

## Creating an Aggregated Ethernet Interface

---

To create an aggregated Ethernet interface:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count device-count
```

For example, to specify 5:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count
```

2. Specify the minimum number of links for the aggregated Ethernet interface (aex), that is, the defined bundle, to be labeled “up”:



**NOTE:** By default only one link must be up for the bundle to be labeled “up”.

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links minimum-links
```

For example, to specify 5:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links 5
```

3. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed link-speed
```

For example, to specify 10g:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed 10g
```

4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interface-name ether-options 802.3ad aex
user@switch# set interface-name ether-options 802.3ad aex
```

## Configuring the VLAN Name and VLAN ID Number

---

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
```

For example, 100.

## Configuring Aggregated Ethernet LACP

---

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
```

```
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

```
[edit interfaces]
```

```
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
```

```
user@switch# set aex aggregated-ether-options lacp periodic interval
```

For example, to specify the interval as fast, execute the following command:

```
[edit interfaces]
```

```
user@switch# set aex aggregated-ether-options lacp periodic fast
```

#### Related Documentation

- [Understanding Interface Naming Conventions on page 2401](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
- [Verifying the Status of a LAG Interface on page 2750](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2751](#)
- [show lacp statistics interfaces \(View\) on page 2875](#)

## Configuring Local Link Bias (CLI Procedure)

Local link bias is used to conserve bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF on a different member link in the LAG bundle.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.

To enable local link bias on a LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias
where aex is the name of the aggregated Ethernet link bundle.
```

For instance, to enable local link bias on aggregated Ethernet interface ae0:

```
[edit]
user@switch# set interface ae0 aggregated-ether-options local-bias
```

### Related Documentation

- [Understanding Local Link Bias on page 2408](#)

## Configuring Multichassis Link Aggregation



**NOTE:** Multichassis Link Aggregation (MC-LAG) is supported on QFX3500 and QFX3600 standalone switches running the original CLI, and on QFX5100 and EX4600 standalone switches running Enhanced Layer 2 Software.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running the Spanning Tree Protocol (STP).

The MC-LAG switches use the Interchassis Control Protocol (ICCP) to exchange the control information between two MC-LAG switches.

On one end of an MC-LAG is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to detect the MC-LAG. On the other side of MC-LAG are two MC-LAG switches. Each of the switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregated Ethernet interface or multichassis aggregated Ethernet interface group.

Perform the following steps on each switch that is hosting an MC-LAG:

1. Specify the same multichassis aggregated Ethernet identification number for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae mc-ae-id number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae mc-ae-id 3
```

2. Specify a unique chassis ID for the MC-LAG that the aggregated Ethernet interface belongs to on each switch.

```
[edit interfaces]
user@switch# set aeX aggregated-ether-options mc-ae chassis-id number
```

For example:

```
[edit interfaces]
user@switch# set ae1 aggregated-ether-options mc-ae chassis-id 0
```

3. Specify the mode of the MC-LAG the aggregated Ethernet interface belongs to.



**NOTE:** Only active-active mode is supported at this time.

```
[edit interfaces]
```

```
user@switch# set aeX aggregated-ether-options mc-ae mode mode
```

For example:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae mode active-active
```

4. Specify whether the aggregated Ethernet interface participating in the MC-LAG is primary or secondary. Primary is **active**, and secondary is **standby**.



**NOTE:** You must configure status control on both switches hosting the MC-LAG. If one switch is in active mode, the other must be in standby mode.

```
[edit interfaces]
```

```
user@switch# set aeX aggregated-ether-options mc-ae status-control (active | standby)
```

For example:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options mc-ae status-control active
```

5. Specify the same LACP system ID on each switch.

```
[edit interfaces]
```

```
user@switch# set aeX aggregated-ether-options lacp system-id mac-address
```

For example:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp system-id 00:01:02:03:04:05
```

6. Specify the same LACP administration key on each switch.

```
[edit interfaces]
```

```
user@switch# set aeX aggregated-ether-options lacp admin-key number
```

For example:

```
[edit interfaces]
```

```
user@switch# set ae1 aggregated-ether-options lacp admin-key 3
```

7. Configure ICCP by doing the following on each switch hosting the MC-LAG:

- a. Configure the local IP address to be used by all switches hosting the MC-LAG.

```
[edit protocols]
```

```
user@switch# set iccp local-ip-addr local-ip-address
```

For example:

```
[edit protocols]
```

```
user@switch# set iccp local-ip-addr 3.3.3.1
```

- b. (Optional) Configure the IP address of the switch and the time during which an ICCP connection must succeed between the switches hosting the MC-LAG.

Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

```
[edit protocols]
```

```
user@switch# set iccp peer peer-ip-address session-establishment-hold-time seconds
```

For example:

```
[edit protocols]
```

```
user@switch# set iccp peer 3.3.3.2 session-establishment-hold-time 50
```



- c. (Optional) Configure the IP address to be used for backup liveness detection:



**NOTE:** By default, backup liveness detection is not enabled. Configure backup liveness detection if you require minimal traffic loss during a reboot. Backup liveness detection helps achieve sub-second traffic loss during an MC-LAG reboot.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address backup-liveness-detection backup-peer-ip
ip-address
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 backup-liveness-detection backup-peer-ip
10.207.64.232
```

- d. Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.



**NOTE:** Configuring the minimum receive interval is required to enable BFD.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection minimum-receive-interval
seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection minimum-receive-interval 1000
```

- e. Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.

```
[edit protocols]
user@switch# set iccp peer peer-ip-address liveness-detection transmit-interval
minimum-interval seconds
```

For example:

```
[edit protocols]
user@switch# set iccp peer 3.3.3.2 liveness-detection transmit-interval minimum-interval
1000
```

8. Configure a multichassis protection link between the switches.

```
[edit]
user@switch# set multi-chassis multi-chassis-protection peer-ip-address interface
interface-name
```

For example:

```
[edit protocols]
user@switch# set multi-chassis multi-chassis-protection 3.3.3.1 interface ae0
```

9. If you are using ELS, configure the **service-id** on both switches.

The **service-id** must be the same number on both switches.

```
[edit]
user@switch# set switch-options service-id number
```

For example:

```
[edit]
user@switch# set switch-options service-id 10
```

10. Enable RSTP globally on all interfaces.

```
[edit]
user@switch# set protocols rstp interface all mode point-to-point
```

11. Disable RSTP on the ICL-PL interfaces on both switches.

```
[edit]
user@switch# set protocols rstp interface interface-name disable
```

For example:

```
[edit]
user@switch# set protocols rstp interface ae0.0 disable
```

12. Configure the MC-LAG interfaces as edge ports on both switches.

```
set protocols rstp interface interface-name
```

For example:

```
[edit]
user@switch# set protocols rstp interface ae1.0
```

13. Enable BPDU block on all interfaces except for the ICL-PL interfaces on both switches.

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

For example:

```
[edit]
user@switch# set protocols rstp bpdu-block-on-edge
```

#### Related Documentation

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol \(VRRP\) on page 2551](#)
- [Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2493](#)

## Configuring Generic Routing Encapsulation Tunneling

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets. GRE tunneling is accomplished through tunnel endpoints that encapsulate or de-encapsulate traffic.

You can also use a firewall filter to de-encapsulate GRE traffic on a QFX5100 switch. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation.

For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. For more information on this feature, see [“Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch” on page 5301](#).

This topic describes:

1. [Configuring a GRE Tunnel on page 2601](#)

## Configuring a GRE Tunnel

To configure a GRE tunnel interface:

1. Create a GRE interface with a unit number and address:

[edit interfaces]

user@switch# set gr-0/0/0 unit *number* family inet *address*



**NOTE:** The base name of the interface must be gr-0/0/0.

This is a pseudo interface, and the address you specify can be any IP address. The routing table must specify **gr-0/0/0.x** as the outgoing interface for any packets that will be tunneled.

If you configure a GRE interface on a QFX5100 switch that is a member of a Virtual Chassis and later change the Virtual Chassis member number of the switch, the name of the GRE interface does not change in any way (because it is a pseudo interface). For example, if you change the member number from **0** to **5**, the GRE interface name does *not* change from **gr-0/0/0.x** to **gr-5/0/0.x**.

2. Specify the tunnel source address for the logical interface:

[edit interfaces]

user@switch# set gr-0/0/0 unit *number* tunnel source *source-address*

3. Specify the destination address:

[edit interfaces]

user@switch# set gr-0/0/0 unit *number* tunnel destination *destination-address*

The destination address must be reachable through static or dynamic routing. If you use static routing, you must get the destination MAC address (for example, by using **ping**) before user traffic can be forwarded through the tunnel.

### Related Documentation

- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 2752](#)
- [Understanding Generic Routing Encapsulation on page 2449](#)
- [Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch on page 5301](#)

## Configuring the LPM Table With Junos OS 13.2x51-D10

In addition to choosing a profile, you can further optimize memory allocation for LPM table entries by configuring how many IPv6 addresses with prefixes in the range /65 through /127 you want to store. If you want to use more than 16 IPv6 addresses with prefixes in this range, you must enter and commit the following statement:

[edit]

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix value
```

in which *value* can be a value in the range 1 through 128. Each increment adds support for 16 IPv6 addresses with prefixes between /65 and /127, for a maximum of 2048 such addresses (16 x 128 = 2048). The system supports 16 of these addresses by default, so to increase the number of supported addresses, you must enter a value of 2 or greater. For example, if you enter **2**, the system will support 32 IPv6 addresses with prefixes in the range /65 through /127.



**NOTE:** When you configure the `num-65-127-prefix` value, all the data interfaces on the switch restart. The management interfaces are unaffected.

The LPM table is shared, and each increment that you add for IPv6 addresses with prefixes in the range /65 through /127 reduces the number of forwarding table entries that are available for IPv4 addresses and IPv6 addresses with prefixes less than /65.

[Table 110 on page 1700](#) provides examples of valid combinations that the LPM table can store, also using the **l2-profile-one** profile. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

**Table 241: Example LPM Table Combinations Using l2-profile-one With Junos OS 13.2X51-D10**

| IPv4 entries | IPv6 Entries (prefix <= 64) | IPv6 Entries (prefix >= 65) | num-65-127-prefix |
|--------------|-----------------------------|-----------------------------|-------------------|
| 16K          | 0K                          | 16                          | 1 (default)       |
| 0K           | 8K                          | 16                          | 1 (default)       |
| 8K           | 4K                          | 16                          | 1 (default)       |
| 4K           | 4K                          | 1K                          | <b>64</b>         |
| 2K           | 5K                          | 1K                          | <b>64</b>         |
| 0K           | 6K                          | 1K                          | <b>64</b>         |
| 4K           | 2K                          | 2K                          | <b>128</b>        |
| 2K           | 3K                          | 2K                          | <b>128</b>        |
| 0K           | 4K                          | 2K                          | <b>128</b>        |

[Table 242 on page 2604](#) provides examples of valid combinations that the LPM table can store when you use the **lpm-profile** profile. As before, each row represents a case in which the table is full and cannot accommodate any more entries.

Table 242: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10

| IPv4 entries | IPv6 Entries (prefix <= 64) | IPv6 Entries (prefix >= 65) | num-65-127-prefix |
|--------------|-----------------------------|-----------------------------|-------------------|
| 128K         | 0K                          | 16                          | 1 (default)       |
| 0K           | 8K                          | 16                          | 1 (default)       |
| 8K           | 4K                          | 16                          | 1 (default)       |
| 4K           | 4K                          | 1K                          | 64                |
| 2K           | 5K                          | 1K                          | 64                |
| 0K           | 6K                          | 1K                          | 64                |
| 4K           | 2K                          | 2K                          | 128               |
| 2K           | 3K                          | 2K                          | 128               |
| 0K           | 4K                          | 2K                          | 128               |

**Related Documentation**

- [Understanding the Unified Forwarding Table](#)
- [Configuring the Unified Forwarding Table on page 1697](#)

## Configuring Ethernet OAM Link Fault Management (CLI Procedure)

Ethernet OAM link fault management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across point-to-point Ethernet links either directly or through repeaters.

To configure Ethernet OAM LFM using the CLI:

1. Enable IEEE 802.3ah OAM support on an interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name
```



**NOTE:** You can configure Ethernet OAM LFM on aggregated interfaces.



**NOTE:** The remaining steps are optional. You can choose which of these features to configure for Ethernet OAM LFM on your switch.

2. Specify whether the interface or the peer initiates the discovery process by configuring the link discovery mode to **active** or **passive** (**active** = interface initiates; **passive** = peer initiates):

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name link-discovery active
```

3. Configure a periodic OAM PDU-sending interval (in milliseconds) for fault detection:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-interval interval
```

4. Specify the number of OAM PDUs that an interface can miss before the link between peers is considered down:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-threshold threshold-value
```

5. Configure event threshold values on an interface for the local errors that trigger the sending of link event TLVs:

- Set the threshold value (in seconds) for sending frame-error events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-error count
```

- Set the threshold value (in seconds) for sending frame-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period count
```

- Set the threshold value (in seconds) for sending frame-period-summary events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period-summary count
```

- Set the threshold value (in seconds) for sending symbol-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds symbol-period count
```



**NOTE:** You can disable the sending of link event TLVs.

To disable the sending of link event TLVs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name negotiation-options no-allow-link-events
```

6. Create an action profile to define event fault flags and thresholds to be taken when the link fault event occurs. Then apply the action profile to one or more interfaces. (You can also apply multiple action profiles to a single interface.)

- a. Name the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name
```

- b. Specify actions to be taken by the system when the link fault event occurs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name action syslog
user@switch# set action-profile profile-name action link-down
```

- c. Specify events for the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name event link-adjacency-loss
```



**NOTE:** For each action profile, you must specify at least one link event and one action. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all actions are executed. You can set a low threshold for a specific action such as logging the error and set a high threshold for another action such as system logging.

7. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Set the remote DTE in loopback mode (the remote DTE must support remote-loopback mode) and then enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name remote-loopback
user@switch# set interface interface-name negotiation-options allow-remote-loopback
```

#### Related Documentation

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)
- [Understanding Ethernet OAM Link Fault Management on page 2452](#)

## Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up

In an MC-LAG topology, a link without Link Access Control Protocol (LACP) configuration remains down and cannot be accessed by the provider edge (PE) devices in the topology.

To ensure that the peer with limited LACP capability is up and accessible on the MC-LAG network, configure one of the aggregated Ethernet links or interfaces on a PE device to be up by using the **force-up** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

You can configure the *force-up* feature on either the PE device in active mode or in standby mode, or on both. However, in order to prevent duplicate traffic and packet drops, you can configure the force-up feature only on one aggregated Ethernet link of the PE device in an MC-LAG topology. If multiple aggregated Ethernet links are up on the PE device that has the force-up feature configured, then the device selects the link based on the LACP port ID and port priority. The port with the lowest priority is given preference. In case of two ports with the same priority, the one with the lowest port ID is given preference.



**NOTE:** If LACP comes up partially in the MC-LAG topology—that is, it comes up on one of the PE devices and does not come up on the other PE devices—the force-up feature is disabled.

#### Related Documentation

- [Understanding Multichassis Link Aggregation](#)
- [Configuring Multichassis Link Aggregation](#)
- [force-up](#)
- [force-up on page 2660](#)



## Configuring Resilient Hashing for Trunk/ECMP Groups

You use resilient hashing to minimize flow remapping across members of a trunk/ECMP group in a load-balanced system. You can configure resilient hashing in link aggregation groups (LAGs) and in equal cost multipath (ECMP) sets.

This topic includes:

1. [Configuring Resilient Hashing on LAGs on page 2607](#)
2. [Configuring Resilient Hashing on ECMP Groups on page 2607](#)

---

### Configuring Resilient Hashing on LAGs

To enable resilient hashing for a LAG:

- Configure resilient hashing on the aggregated Ethernet interface:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options resilient-hash
```

---

### Configuring Resilient Hashing on ECMP Groups

To enable resilient hashing for ECMP groups:

- Configure resilient hashing for ECMP:

```
[edit forwarding-options]
user@switch# set enhanced-hash-key ecmp-resilient-hash
```

#### Related Documentation

- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 2454](#)

---

## Configuration Tasks

- [Channelizing Interfaces on page 2608](#)
- [Configuring the System Mode on page 2610](#)

## Channelizing Interfaces

The QFX3500, QFX3600, QFX5100, and EX4600 switches provide 40-Gbps QSFP+ ports that can be channelized. Channelization allows you to configure 40-Gbps QSFP+ ports to operate as four 10-Gigabit Ethernet (xe) interfaces. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. By default, the four 40-Gbps QSFP+ ports operate as 40-Gigabit Ethernet (et) ports. When an et port is channelized to four xe ports, a colon is used to signify the four separate channels. For example, on a QFX3500 standalone switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are *xe-0/1/2:0*, *xe-0/1/2:1*, *xe-0/1/2:2*, and *xe-0/1/2:3*.

By default, the 40-Gbps QSFP+ ports on EX4600 and QFX5100 switches are channelized automatically (auto-channelized) if any of the four channels on a 40-Gbps QSFP+ port receive data, unless you have configured channelization either at the chassis level or at the port level. Auto-channelization is not supported on interfaces contained in expansion modules or on Virtual Chassis ports.

You can disable auto-channelization by including the **disable-auto-speed-detection** statement at the **[edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high) channel-speed]** hierarchy.

There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches, depending on the system mode you enable. If you try to channelize ports that are restricted, the configuration is ignored. See [“Configuring the System Mode” on page 2610](#) for more information.



**CAUTION:** The Packet Forwarding Engine on the switch is restarted when you configure or delete a port. As a result, you might experience packet loss on the device. When you channelize a 40-Gbps QSFP+ port on the master of a Virtual Chassis, traffic might be disrupted on the master as well as on the line card members, and a mastership switchover occurs.

The following steps describe how to configure a block of ports or an individual port to operate as 10-Gigabit Ethernet ports.

1. To configure a block of 40-Gigabit Ethernet (et) ports to operate as 10-Gigabit Ethernet ports, specify a port range and channel speed:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port-range port-range-low port-range-high channel-speed speed
```

For example, to configure ports 0 through 3 on PIC 1 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 1]
user@switch# set port-range 0 3 channel-speed 10g
```

2. To configure an individual 40-Gigabit Ethernet (et) port to operate as 10-Gigabit Ethernet (xe) ports, specify a port number and channel speed:

```
[edit chassis fpc 0 pic 0]
user@switch# set port port-number channel-speed speed
```

For example, to configure port 2 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port 2 channel-speed 10g
```

3. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

4. To return a range of ports to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-range port-range-low port-range-high channel-speed speed
```

For example, to return ports 0 through 3 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-range 0 3 channel-speed 10g
```

5. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

6. To return a port to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port port-number channel-speed speed
```

For example, to return port 2 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port 2 channel-speed 10g
```

7. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

The following steps describe how to disable auto-channelization at the port level.

1. To disable auto-channelization at the port level, include the **disable** statement:

```
[edit]
user@switch# set chassis fpc slot-number pic pic-number (port port-number |
  port-range port-range-low port-range-high) channel-speed
  disable-auto-speed-detection
```

For example, to disable auto-channelization for one port:

```
[edit]
```

```
user@switch# set chassis fpc 0 pic 0 port 2 channel-speed  
disable-auto-speed-detection
```

For example, to disable auto-channelization for a range of ports:

```
[edit]  
user@switch# set chassis fpc 0 pic 0 port-range 2 4 channel-speed  
disable-auto-speed-detection
```

2. Review your configuration and issue the **commit** command.

```
[edit]  
user@switch# commit  
commit complete
```

#### Related Documentation

- [Configuring the System Mode on page 2610](#)
- [channel-speed on page 2634](#)
- [fpc on page 2661](#)
- [pic on page 2723](#)

## Configuring the System Mode

You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. By default, all QSFP+ interfaces are auto-channelized. Auto-channelization is not supported on interfaces contained in expansion modules or on Virtual Chassis ports. To disable auto-channelization, see [“Channelizing Interfaces” on page 2608](#) for more information.



**NOTE:** When you request the system mode change, you must reboot for the system mode to take effect.



**CAUTION:** The Packet Forwarding Engine on the switch is restarted when you issue system mode changes. As a result, you might experience packet loss on the switch.

The following system modes are available on the QFX5100-24Q switch:

- Default-mode
- Mode-104-port
- Flexi-PIC mode
- Non-oversubscribed mode

The following system modes are available on the QFX5100-96S switch:

- Default-mode
- Non-oversubscribed mode

See [Table 243 on page 2611](#) for more information regarding the supported system modes for your switch.

**Table 243: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software**

|             | Default-mode   | Mode-104port   | Flexi-pic-mode  | Non-oversubscribed-mode   |
|-------------|--|--|---|---|
| QFX5100-48S | Not supported  | Not supported  | Not supported   | Not supported   |
| QFX5100-24Q | Supported<br><br>You do not need to configure the switch to be in this mode. On PIC 0, you can channelize all 24 40-Gbps QSFP+ ports. On PIC 1 and PIC 2, the 40-Gbps QSFP+ ports in the expansion modules are supported but cannot be channelized. In this mode, you can have one of two port combinations: 32 40-Gbps QSFP+ ports, or 96 10-Gigabit Ethernet ports plus 8 40-Gbps QSFP+ ports. | Supported<br><br>On PIC 0, all 24 40-Gbps QSFP+ ports are channelized by default, which provides 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in an expansion module on PIC 1 are supported. On PIC 1, ports 0 and 2 are channelized by default, and ports 1 and 3 are disabled. If 40-Gbps QSFP+ ports contained in an expansion module are detected on PIC 2, they are ignored. | Supported<br><br>On PIC 0, the first four ports (ports 0 through 3) cannot be channelized. 40-Gbps QSFP+ ports contained in expansion modules on PIC 1 and PIC 2 are supported but cannot be channelized. | Supported<br><br>All 24 40-Gbps QSFP+ ports on PIC 0 can be channelized to 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in the expansion modules on PIC 1 and PIC 2 are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode. |
| QFX5100-96S | Supported<br><br>You do not need to configure the switch to be in this mode. On PIC 0, all 96 10-Gigabit Ethernet ports are supported. You can only channelize the 40-Gbps QSFP+ interfaces to 10-Gigabit Ethernet interfaces on ports 96 and 100. When you channelize the interfaces on ports 96 and 100, ports 97, 98, 99, 101, 102 and 103 are disabled.                                      | Not supported  | Not supported   | Supported<br><br>On PIC 0, all 96 10-Gigabit Ethernet ports are supported. However, the eight 40-Gbps QSFP+ ports are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode.  |

The following steps describe how to change the system mode.

1. To change the system mode, issue the following operational command:

```
{master:0}
```

```
root> request chassis system-mode mode
```

For example:

```
{master:0}
```

```
root> request chassis system-mode non-oversubscribed-mode
```

2. To return to the default mode (default-mode), issue the following operational command:

```
{master:0}
```

```
root> request chassis system-mode default-mode
```

3. To see which system mode is configured, issue the following operational command:

```
{master:0}
```

```
root> show chassis system-mode
```

#### Related Documentation

- [Understanding Interface Naming Conventions on page 2401](#)
- [Understanding Port Ranges and System Modes on page 2421](#)
- [Channelizing Interfaces on page 2608](#)

## Configuration Statements

---

- [\[edit interfaces et\] Configuration Statement Hierarchy on the QFX Series on page 2616](#)
- [802.3ad on page 2622](#)
- [action \(OAM LFM\) on page 2623](#)
- [action-profile on page 2624](#)
- [address on page 2625](#)
- [aggregated-devices on page 2627](#)
- [aggregated-ether-options on page 2628](#)
- [alarm \(chassis\) on page 2630](#)
- [allow-remote-loopback on page 2631](#)
- [authentication-key \(ICCP\) on page 2631](#)
- [auto-negotiation on page 2632](#)
- [backup-liveness-detection on page 2633](#)
- [backup-peer-ip on page 2633](#)
- [channel-speed on page 2634](#)
- [chassis on page 2635](#)
- [chassis-id on page 2636](#)
- [configured-flow-control on page 2637](#)
- [container-devices on page 2638](#)
- [craft-lockout on page 2639](#)
- [description \(Interfaces\) on page 2640](#)

- [destination \(Tunnels\) on page 2641](#)
- [detection-time \(Liveness Detection\) on page 2641](#)
- [device-count on page 2642](#)
- [disk-failure-action on page 2642](#)
- [ecmp-resilient-hash on page 2643](#)
- [enhanced-hash-key on page 2644](#)
- [ethernet on page 2645](#)
- [ethernet \(OAM LFM\) on page 2646](#)
- [ethernet \(Alarm\) on page 2648](#)
- [ethernet-switching on page 2649](#)
- [ether-options on page 2650](#)
- [eui-64 on page 2651](#)
- [event \(OAM LFM\) on page 2651](#)
- [event-thresholds on page 2652](#)
- [family on page 2653](#)
- [fibre-channel \(Alarm\) on page 2656](#)
- [filter on page 2657](#)
- [flow-control on page 2659](#)
- [force-up on page 2660](#)
- [fpc on page 2661](#)
- [frame-error on page 2662](#)
- [frame-period on page 2662](#)
- [frame-period-summary on page 2663](#)
- [gratuitous-arp-reply on page 2663](#)
- [group on page 2664](#)
- [group \(Redundant Trunk Groups\) on page 2665](#)
- [hash-mode on page 2666](#)
- [hold-time \(Physical Interface\) on page 2668](#)
- [iccp on page 2670](#)
- [irb \(Interfaces\) on page 2672](#)
- [inet \(interfaces\) on page 2675](#)
- [inet \(enhanced-hash-key\) on page 2676](#)
- [inet6 \(interfaces\) on page 2677](#)
- [inet6 \(enhanced-hash-key\) on page 2678](#)
- [interface \(Multichassis Protection\) on page 2679](#)
- [interface \(OAM LFM\) on page 2680](#)
- [interface \(Redundant Trunk Groups\) on page 2681](#)

- [interface-mode](#) on page 2682
- [interface-range](#) on page 2684
- [interfaces](#) on page 2686
- [lacp \(802.3ad\)](#) on page 2693
- [lacp \(Aggregated Ethernet\)](#) on page 2694
- [layer2 \(enhanced-hash-key\)](#) on page 2695
- [link-adjacency-loss](#) on page 2696
- [link-discovery](#) on page 2697
- [link-down](#) on page 2697
- [link-event-rate](#) on page 2698
- [link-fault-management](#) on page 2699
- [link-to-disable](#) on page 2700
- [link-to-monitor](#) on page 2700
- [link-down](#) on page 2701
- [link-mode](#) on page 2702
- [link-speed](#) on page 2703
- [liveness-detection](#) on page 2704
- [local-bias](#) on page 2705
- [local-ip-addr \(ICCP\)](#) on page 2705
- [loopback \(Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet\)](#) on page 2706
- [management-ethernet \(Alarm\)](#) on page 2706
- [member](#) on page 2707
- [member-range](#) on page 2707
- [mc-ae](#) on page 2708
- [mc-ae-id](#) on page 2709
- [minimum-interval \(Liveness Detection\)](#) on page 2709
- [minimum-links](#) on page 2710
- [minimum-receive-interval \(Liveness Detection\)](#) on page 2710
- [mode \(QFX Series\)](#) on page 2711
- [multi-chassis](#) on page 2711
- [multi-chassis-protection](#) on page 2712
- [multiplier \(Liveness Detection\)](#) on page 2712
- [mtu](#) on page 2713
- [negotiation-options](#) on page 2714
- [no-adaptation \(Liveness Detection\)](#) on page 2714
- [no-allow-link-events](#) on page 2715
- [no-gratuitous-arp-request](#) on page 2715



- [oam on page 2716](#)
- [on-disk-failure on page 2718](#)
- [on-loss-of-keepalives on page 2719](#)
- [pdu-interval on page 2720](#)
- [pdu-threshold on page 2720](#)
- [peer \(ICCP\) on page 2721](#)
- [peer \(Multichassis\) on page 2722](#)
- [periodic on page 2722](#)
- [pic on page 2723](#)
- [preempt-cutover-timer on page 2724](#)
- [redundancy \(Graceful Switchover\) on page 2725](#)
- [redundant-trunk-group on page 2726](#)
- [remote-loopback on page 2727](#)
- [resilient-hash on page 2727](#)
- [rx-buffers on page 2728](#)
- [routing-engine on page 2729](#)
- [service-id on page 2730](#)
- [session-establishment-hold-time on page 2730](#)
- [source on page 2731](#)
- [speed on page 2732](#)
- [status-control on page 2732](#)
- [symbol-period on page 2733](#)
- [syslog \(OAM LFM\) on page 2733](#)
- [targeted-broadcast on page 2734](#)
- [threshold \(Detection Time\) on page 2734](#)
- [traceoptions \(ICCP\) on page 2735](#)
- [transmit-interval \(Liveness Detection\) on page 2737](#)
- [traceoptions \(Individual Interfaces\) on page 2738](#)
- [traceoptions \(OAM LFM\) on page 2739](#)
- [traps on page 2740](#)
- [tunnel on page 2741](#)
- [tunnel-port on page 2741](#)
- [tx-buffers on page 2742](#)
- [unit on page 2744](#)
- [uplink-failure-detection on page 2745](#)
- [version \(Liveness Detection\) on page 2745](#)

- [vlan-id on page 2746](#)
- [vlan-tagging on page 2746](#)

## [edit interfaces et] Configuration Statement Hierarchy on the QFX Series

This topic lists supported and unsupported configuration statements in the **[edit interfaces et]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific QFX Series platforms, see *QFX Series Virtual Chassis Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit interfaces et\] Hierarchy Level on page 2616](#)
- [Unsupported Statements in the \[edit interfaces et\] Hierarchy Level on page 2620](#)

### Supported Statements in the [edit interfaces et] Hierarchy Level

The following hierarchy shows the **[edit interfaces et]** configuration statements supported on EX Series switches.

```
interfaces {
  et-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    encapsulation type;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
          port-priority number;
        }
      }
    }
    ethernet-switch-profile {
      tag-protocol-id [tpids];
    }
    (flow-control | no-flow-control);
    (loopback | no-loopback);
    no-auto-mdix;
  }
  flexible-vlan-tagging;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  native-vlan-id
```

```

no-gratuitous-arp-request;
traceoptions {
    flag flag;
}
(traps | no-traps);
unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    encapsulation type;
    family ccc;
    filter {
        group group-number;
        input filter-name;
        input-list [filter-names];
        output filter-name;
        output-list [filter-names];
    }
    policer {
        input policer-name;
        output policer-name;
    }
}
family ethernet-switching {
    filter {
        input filter-name;
        output filter-name;
    }
    interface-mode (access | trunk);
    recovery-timeout seconds;
    storm-control profile-name;
    vlan {
        members (vlan-name | [-vlan-names] | all);
    }
}
family inet {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
    }
}

```

```
(preempt | no-preempt) {
    hold-time seconds;
}
priority number;
track {
    interface interface-name {
        priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [addresses];
vrrp-inherit-from {
    active-group group-number;
    active-interface interface-name;
}
}
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
    }
}
```

```

    priority number;
    track {
        interface interface-name {
            priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {

    input filter-name;

    output filter-name;

}
mtu bytes;
nd6-stale-time time;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
input-vlan-map action;
output-vlan-map action;
proxy-arp (restricted | unrestricted);
swap-by-poppush;
(traps | no-traps);
vlan-id vlan-id-number;
vlan-id-list [vlan-id vlan-id-vlan-id];
}
vlan-tagging;
}
}

```

### Unsupported Statements in the [edit interfaces et] Hierarchy Level

All statements in the [edit interfaces et] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

**Table 244: Unsupported [edit interfaces et] Configuration Statements for the QFX Series**

| Statement                 | Hierarchy   |
|---------------------------|---|
| passive-monitor-mode      | [edit interfaces et]                                |
| stacked-vlan-tagging      | [edit interfaces et]                                |
| asynchronous-notification | [edit interfaces et ether-options]                  |
| ignore-l3-incompletes     | [edit interfaces et ether-options]                  |
| mpls                      | [edit interfaces et ether-options]                  |
| source-address-filter     | [edit interfaces et ether-options]                  |
| source-filtering          | [edit interfaces et ether-options]                  |
| no-source-filtering       | [edit interfaces et ether-options]                  |
| accept-source-mac         | [edit interfaces et unit]                           |
| layer2-policer            | [edit interfaces et unit]                           |
| native-inner-vlan-id      | [edit interfaces et unit]                           |
| vlan-id-range             | [edit interfaces et unit]                           |
| vlan-tags                 | [edit interfaces et unit]                           |
| mpls                      | [edit interfaces et unit family]                    |
| tcc                       | [edit interfaces et unit family]                    |
| vpls                      | [edit interfaces et unit family]                    |
| bridge-domain-type        | [edit interfaces et unit family ethernet-switching] |
| inner-vlan-id-list        | [edit interfaces et unit family ethernet-switching] |
| vlan-rewrite              | [edit interfaces et unit family ethernet-switching] |
| policer                   | [edit interfaces et unit family inet]               |

Table 244: Unsupported [edit interfaces et] Configuration Statements for the QFX Series  
(continued)

| Statement           | Hierarchy   |
|---------------------|---|
| sampling            | [edit interfaces et unit family inet]                                     |
| service             | [edit interfaces et unit family inet]                                     |
| targeted-broadcast  | [edit interfaces et unit family inet]                                     |
| unnumbered-address  | [edit interfaces et unit family inet]                                     |
| bandwidth-threshold | [edit interfaces et unit family inet address vrrp-group track interface]  |
| service             | [edit interfaces et unit family inet6]                                    |
| bandwidth-threshold | [edit interfaces et unit family inet6 address vrrp-group track interface] |
| group               | [edit interfaces et unit family inet6 filter]                             |

**Related Documentation** • [QFX Series Virtual Chassis Software Features Overview](#)

## 802.3ad

---

**Syntax**    802.3ad aex;  
              lACP {  
                  force-up;  
                  (primary | backup);  
              }  
              port-priority;  
              }

**Hierarchy Level**    [edit [interfaces interface-name ether-options](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Specify the aggregated Ethernet logical interface number.



**NOTE:** The port-priority statement is not supported on QFabric systems.

---

**Options**    aex—Aggregated Ethernet logical interface number.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Link Aggregation on page 2593](#)
- [Configuring Aggregated Ethernet LACP on page 2589](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2393](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 1234](#)
- *Junos OS Network Interfaces Library for Routing Devices*



---

## action (OAM LFM)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>action {<br/>    syslog;<br/>    link-down;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | <p>Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure)</a> on page 2604</li></ul>                          |

## action-profile

---

**Syntax**    `action-profile profile-name;`  
              `action {`  
                  `syslog;`  
                  `link-down;`  
                  `}`  
              `event {`  
                  `link-adjacency-loss;`  
                  `link-event-rate {`  
                      `frame-error count;`  
                      `frame-period count;`  
                      `frame-period-summary count;`  
                      `symbol-period count;`  
                      `}`  
                  `}`  
              `}`

**Hierarchy Level**    `[edit protocols oam ethernet link-fault-management]`

**Release Information**    Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description**    Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.

The remaining statements are explained separately.

**Options**    *profile-name*—Name of the action profile.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**    • [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)

## address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcidlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority-number number;
            track {
                priority-cost seconds;
                priority-hold-time interface-name {
                    interface priority;
                    bandwidth-threshold bits-per-second {
                        priority;
                    }
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family *family*],  
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*  
 family *family*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the interface address.

**Options** *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

The remaining statements are explained separately.



**NOTE:** The `edit logical-systems` hierarchy is not available on QFabric systems.

---

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the Protocol Family*
- *Junos OS Administration Library for Routing Devices*
- *family*
- *negotiate-address*
- *unnumbered-address (Ethernet)*

## aggregated-devices

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>aggregated-devices {   ethernet {     device-count <i>number</i>;   } }</pre>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis</a> ],<br>[edit <a href="#">chassis</a> node-group <i>name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure properties for aggregated devices on the switch.<br><br>The remaining statements are explained separately.   |
| <b>Default</b>                  | Aggregated devices are disabled.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2393</a></li> <li>• <a href="#">Configuring Link Aggregation on page 2593</a></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## aggregated-ether-options

---

**Syntax**    aggregated-ether-options {  
              **configured-flow-control** {  
                  rx-buffers (on | off);  
                  tx-buffers (on | off);  
              }  
              ethernet-switch-profile {  
                  tag-protocol-id;  
                  (fcoe-lag | no-fcoe-lag);  
                  (flow-control | no-flow-control);  
                  **lACP mode** {  
                      admin-key *key*;  
                      periodic *interval*;  
                      system-id *mac-address*;  
                      force-up;  
                  }  
              }  
              (link-protection | no-link-protection);  
              link-speed *speed*;  
              local-bias;  
              (loopback | no-loopback);  
              mc-ae {  
                  chassis-id *chassis-id*;  
                  mc-ae-id *mc-ae-id*;  
                  mode (active-active);  
                  status-control (active | standby);  
              }  
              minimum-links *number*;  
              rebalance-periodic;  
              resilient-hash;  
              source-address-filter *filter*;  
              (source-filtering | no-source-filtering);  
          }

**Hierarchy Level**    [edit **interfaces** *aex*]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statements **fcoe-lag** and **no-fcoe-lag** introduced in Junos OS Release 13.2X52-D10 for the QFX Series.  
Statements **force-up**, **lACP**, and **resilient-hash** introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

**Description**    Configure properties specific to a specific aggregated Ethernet interface.

The statements are explained separately.

**Default**    Options are not enabled.

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related  
Documentation**

- [Understanding Aggregated Ethernet Interfaces and LACP on page 2393](#)
- [Configuring Aggregated Ethernet LACP on page 2589](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## alarm (chassis)

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>alarm {<br/>    interface-type {<br/>        alarm-name (ignore   red   yellow);<br/>    }<br/>}</pre>  |
| Hierarchy Level          | [edit chassis],<br>[edit chassis interconnect-device <i>name</i> ],<br>[edit chassis node-group <i>name</i> ]  |
| Release Information      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for the ACX Series.  |
| Description              | <p>Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the <b>RED ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the contact on the craft interface or LCD screen. Yellow alarm conditions light the <b>YELLOW ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the craft interface or LCD screen.</p> <p>To configure more than one alarm, include multiple <i>alarm-name</i> lines.</p> |
| Options                  | <p><i>alarm-name</i>—Alarm condition. For a list of conditions, see <i>System-Wide Alarms and Alarms for Each Interface Type</i>.</p> <p><i>ignore</i>—The specified alarm condition does not set off any alarm.</p> <p><i>interface-type</i>—Type of interface on which you are configuring the alarm: <b>atm</b>, <b>ethernet</b>, <b>sonet</b>, or <b>t3</b>.</p> <p><b>red</b>—The specified alarm condition sets off a red alarm.</p> <p><b>yellow</b>—The specified alarm condition sets off a yellow alarm.</p>   |
| Required Privilege Level | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Alarms on page 7191</a></li><li>• <a href="#">Chassis Conditions That Trigger Alarms</a></li><li>• <a href="#">Chassis Alarm Messages on a QFX3500 Device on page 7192</a></li><li>• <a href="#">Interface Alarm Messages on page 7195</a></li></ul>   |



## allow-remote-loopback


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | allow-remote-loopback;  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the EX Series switch.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul> |

## authentication-key (ICCP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | authentication-key <i>key</i> ;  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp</a> <a href="#">peer</a> < <i>peer-IP-address</i> >],<br>[edit protocols <a href="#">iccp</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.  |
| <b>Description</b>              | Specify the authentication key (password). The QFX3500 and MX Series device uses this password to verify the authenticity of packets sent from the peers hosting a multichassis link aggregation group (MC-LAG). Peer-level authentication takes precedence over global-level authentication.<br><br>Interchassis Control Protocol (ICCP) uses MD5 authentication. |
| <b>Options</b>                  | <b>key</b> —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |

## auto-negotiation

---

|  |   |
|--|---|
| <b>Syntax</b>  | (auto-negotiation   no-auto-negotiation);   |
| <b>Hierarchy Level</b>   | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>   | <p>Explicitly enable or disable autonegotiation.</p> <ul style="list-style-type: none"><li>• <b>auto-negotiation</b>—Enable autonegotiation.</li><li>• <b>no-auto-negotiation</b>—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options.</li></ul> |
| <b>Default</b>   | Autonegotiation is automatically enabled for Gigabit Ethernet interfaces. Autonegotiation is not an option for 10-Gigabit Ethernet interfaces. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.   |
| <hr/> <div> <b>NOTE:</b> Autonegotiation is not supported on QFX5100 devices.</div> <hr/> |   |
| <b>Required Privilege Level</b>  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">speed on page 2732</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>   |

## backup-liveness-detection

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>backup-liveness-detection {<br/>    <a href="#">backup-peer-ip ip4-address</a></code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2R1 for EX Series switches.   |
| <b>Description</b>              | Backup liveness detection determines the peer status (whether it is up or down) by exchanging keep alive messages (UDP-based packets) over the management link between the two Interchassis Control Protocol (ICCP) peers. When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |

## backup-peer-ip

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>backup-peer-ip <i>ip4-address</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer backup-liveness-detection</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 13.2R1 for EX Series switches. |
| <b>Description</b>              | Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration.                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                          |

## channel-speed

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | channel-speed (10g; disable-auto-speed-detection) ;  |
| <b>Hierarchy Level</b>          | [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> (port <i>port-number</i>   port-range <i>port-range-low</i> <i>port-range-high</i> )]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.  |
| <b>Description</b>              | (QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Enable the specified port on the Physical Interface Card (PIC) to perform in the specified channel speed. Additionally, you can disable auto-speed detection. |
| <b>Default</b>                  | 40g (40-Gigabit Ethernet).   |
| <b>Options</b>                  | <b>10g</b> —Set the channel speed to 10g (10-Gigabit Ethernet).<br><br><b>disable-auto-speed-detection</b> —Disable auto-speed detection.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Channelizing Interfaces on page 2608</a></li></ul>   |

## chassis

```
Syntax  chassis {
        routing-engine {
        redundancy {
            failover {
                on-disk-failure {
                disk-failure-action (halt | reboot);
                }
                on-loss-of-keepalives;
            }
            graceful-switchover;
        }
        aggregated-devices {
            ethernet {
                device-count number;
            }
            alarm {
                interface-type {
                    alarm-name (red | yellow | ignore);
                }
            }
        }
        forwarding-options profile-name {
            num-65-127-prefix value
        }
        fpc slot {
            auto-speed-detection disable
            pic pic-number{
                port port-number{
                    tunnel-port port-number tunnel-services;
                    channel-speed speed;
                }
                port-range port-range-low port-range-high {
                    channel-speed speed;
                }
            }
        }
        maximum-ecmp next-hops;
    }
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation** • [Configuring Link Aggregation on page 2593](#)


---

## chassis-id

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>chassis-id <i>chassis-id</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aggregated-ether-options mc-aes</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  |
| <b>Description</b>              | Specify the chassis ID of the multichassis aggregated Ethernet interface device. LACP uses the chassis ID to calculate the port number of the multichassis link aggregation group (MC-LAG) physical member links. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |

## configured-flow-control

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | configured-flow-control {<br><b>rx-buffers</b> (on   off);<br><b>tx-buffers</b> (on   off);<br>}   |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]   |
| <b>Description</b>              | <p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the <b>rx-buffers</b> and the <b>tx-buffers</b> values when you configure asymmetric flow control.</p> <p>Use the <b>flow-control</b> and <b>no-flow-control</b> statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <hr/> <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> <hr/> |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.   |
| <b>Options</b>                  | The statements are explained separately.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">congestion-notification-profile on page 6220</a></li> <li>• <a href="#">flow-control on page 2659</a></li> </ul>  |

## container-devices

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>container-devices {<br/>    device-count <i>number</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | <pre>[edit chassis]<br/>[edit chassis interconnect-device <i>name</i>]<br/>[edit chassis node-group <i>name</i>]</pre>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for QFX Series switches.   |
| <b>Description</b>              | Specify the container devices configuration. The <b>number</b> option specifies the number of sequentially numbered container interfaces, from <b>ci0</b> to <b>ci127</b> maximum. |
| <b>Options</b>                  | <b>number</b> —Number of container devices.<br><b>Range:</b> 1 through 128   |
| <b>Required Privilege Level</b> | <b>chassis</b> —To view this statement in the configuration.<br><b>chassis-control</b> —To add this statement to the configuration.  |



## craft-lockout

```

Syntax  craft-lockout {
        alarm {
            interface-type {
                link-down (red | yellow | ignore);
            }
        }
        container-devices {
            device-count number;
        }
        fpc slot {
            pic pic-number {
                fibre-channel {
                    port-range {
                        port-range-low port-range-high;
                    }
                }
            }
        }
        routing-engine {
            on-disk-failure {
                disk-failure-action (halt | reboot);
            }
        }
    }

```

**Hierarchy Level** [edit chassis interconnect-device]

**Release Information** Statement introduced in Junos Release 11.3 for the QFX Series.

**Description** Disable the physical operation of the craft interface front panel.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring the Junos OS to Disable the Physical Operation of the Craft Interface*

## description (Interfaces)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>description text;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit <a href="#">interfaces</a> interface-name],</code><br><code>[edit <a href="#">interfaces</a> interface-name unit logical-unit-number],</code><br><code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]</code>   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.  |
| <b>Description</b>              | <p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the <b>show interfaces</b> commands, and is also exposed in the <b>ifAlias</b> Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>   |
| <b>Options</b>                  | <b>text</b> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.   |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Interface Description</i></li><li>• <i>Adding a Logical Unit Description to the Configuration</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <i>Using DHCP Relay Agent Option 82 Information</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li><li>• <i>Example: Connecting Access Switches to a Distribution Switch</i></li></ul> |

## destination (Tunnels)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>destination address;</code>  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> ],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.  |
| <b>Description</b>              | For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.  |
| <b>Options</b>                  | <b>address</b> —Address of the remote side of the connection.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Interface Address</i></li> <li>• <i>point-to-point</i></li> </ul>  |

## detection-time (Liveness Detection)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>detection-time {<br/>    <i>milliseconds</i>;<br/>}</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <b>iccp</b> peer <b>liveness-detection</b> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.                            |
| <b>Description</b>              | The Bidirectional Forwarding Detection (BFD) timers are adaptive and can be adjusted to be faster or slower.<br><br>The remaining statement is explained separately. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |

## device-count

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>device-count <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis aggregated-devices ethernet</a> ],<br>[edit <a href="#">chassis node-group <i>name</i> aggregated-devices ethernet</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure the number of aggregated Ethernet logical devices available to the switch.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation on page 2593</a></li><li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462</a></li></ul> |

## disk-failure-action

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>disk-failure-action (halt   reboot);</code>   |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">redundancy on-disk-failure</a> ]<br>[edit chassis routing-engine on-disk-failure]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure the Routing Engine to halt or reboot when the Routing Engine hard disk fails.   |
| <b>Options</b>                  | <b>halt</b> —Specify the Routing Engine to halt.<br><b>reboot</b> —Specify the Routing Engine to reboot.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">graceful-switchover on page 2309</a></li><li>• <a href="#">Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors</a></li><li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li><li>• <a href="#">High Availability Features for EX Series Switches Overview</a></li></ul> |

---

## ecmp-resilient-hash

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ecmp-resilient-hash;   |
| <b>Hierarchy Level</b>          | [edit forwarding-options enhanced-hash-key]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.   |
| <b>Description</b>              | Enable resilient hashing for ECMP groups, to minimize remapping of destination paths.  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Resilient Hashing for Trunk/ECMP Groups on page 2607</a></li></ul> |

## enhanced-hash-key

---

**Syntax**    enhanced-hash-key {  
              ecmp-resilient-hash;  
              fabric-load-balance {  
                  flowlet {  
                      inactivity-interval *interval*;  
                  }  
              per-packet;  
          }  
              hash-mode {  
                  layer2-header;  
                  layer2-payload;  
              }  
              inet {  
                  no-ipv4-destination-address;  
                  no-ipv4-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-protocol;  
                  vlan-id;  
              }  
              inet6 {  
                  no-ipv6-destination-address;  
                  no-ipv6-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-next-header;  
                  vlan-id;  
              }  
              layer2 {  
                  no-destination-mac-address;  
                  no-ether-type;  
                  no-source-mac-address;  
                  vlan-id;  
              }  
          }  
      }

**Hierarchy Level**    [edit forwarding-options]

**Release Information**    Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.  
                              Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.  
                              The **fabric-load-balance** statement introduced in Junos OS Release 14.1X53-D10.

**Description**    Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The remaining statements are explained separately.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2590</a></li> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396</a></li> </ul> |

## ethernet

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ethernet {<br>device-count <i>number</i> ;<br>}   |
| <b>Hierarchy Level</b>          | [edit chassis aggregated-devices],<br>[edit chassis node-group aggregated-devices]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure properties for aggregated Ethernet devices on the switch.<br><br>The remaining statement is explained separately.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2593</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## ethernet (OAM LFM)

---

```
Syntax  ethernet {
        connectivity-fault-management {
            action-profile profile-name {
                action {
                    interface-down;
                }
                default-actions {
                    interface-down;
                }
                event {
                    adjacency-loss;
                }
            }
        }
        esp-traceoptions {
            file filename <files number> <no-stamp> <replace> <size size> <world-readable |
                no-world-readable>;
            flag (all | error | esp | interface | krt | lib | normal | task | timer);
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            mip-half-function (none | default | explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s | 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    priority
                    remote-mep mep-id {
                        action-profile profile-name;
                        sla-iterator-profile profile-name {
                            data-tlv-size size;
                            iteration-count count-value;
                            priority priority-value;
                        }
                    }
                }
            }
            short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
        }
    }
    performance-monitoring {
```



```

sla-iterator-profiles {
  profile-name {
    calculation-weight {
      delay delay-value;
      delay-variation delay-variation-value;
    }
    cycle-time cycle-time-value;
    iteration-period iteration-period-value;
    measurement-type two-way-delay;
    passive;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}

```

```
    }  
  }  
}
```

|                                 |   |
|---------------------------------|---|
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.<br><b>connectivity-fault-management</b> introduced in Junos OS Release 10.2 for EX Series switches.  |
| <b>Description</b>              | Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces on switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.<br><br>The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583</a></li><li>• <a href="#">Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li><li>• <a href="#">Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)</a></li></ul> |

---

## ethernet (Alarm)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ethernet {<br><a href="#">link-down</a> (red   yellow   ignore);<br>}  |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">alarm</a> ],<br>[edit chassis interconnect-device <i>name</i> <a href="#">alarm</a> ],<br>[edit chassis node-group <i>name</i> <a href="#">alarm</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure alarms for an Ethernet interface.  |
| <b>Options</b>                  | The remaining statement is explained separately.—  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Alarms on page 7191</a></li><li>• <a href="#">Interface Alarm Messages on page 7195</a></li></ul>              |

## ethernet-switching

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> ethernet-switching {   filter {     group <i>filter-group-number</i>;     input <i>filter-name</i>;     input-list [ <i>filter-names</i> ];     output <i>filter-name</i>;     output-list [ <i>filter-names</i> ];   }   <b>interface-mode</b> (access   trunk);   <b>recovery-timeout</b> <i>seconds</i>;   storm-control <i>profile-name</i>;   vlan {     members (<i>vlan-name</i>   [<i>-vlan-names</i>]   all);   } } </pre> |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>ge-chassis/slot/port unit logical-unit-number</i> ] family   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Configure Ethernet switching protocol family information for the logical interface.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li> <li>• <a href="#">JUNOS Software Network Interfaces Configuration Guide</a></li> </ul>  |

## ether-options

---

**Syntax** ether-options {  
    802.3ad aex {  
        lACP {  
            force-up;  
            (primary | backup);  
        }  
    }  
    (auto-negotiation | no-auto-negotiation);  
    configured-flow-control {  
        rx-buffers (on | off);  
        tx-buffers (on | off);  
    }  
    (flow-control | no-flow-control);  
    link-mode mode;  
    (loopback | no-loopback);  
    speed (auto-negotiation | no-auto-negotiation);  
}

**Hierarchy Level** [edit [interfaces](#) *interface-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure **ether-options** properties for a Gigabit Ethernet or 10-Gigabit Ethernet interface.  
  
The statements are explained separately.

**Default** Enabled.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## eui-64

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | eui-64;  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series. |
| <b>Description</b>              | For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Interface Address</a></li> </ul>  |

## event (OAM LFM)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>event {   link-adjacency-loss;   link-event-rate {     frame-error count;     frame-period count;     frame-period-summary count;     symbol-period count;   } }</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">action-profile</a> <i>profile-name</i> ]                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Configure link events in an action profile for Ethernet OAM link fault management (LFM).<br><br>The remaining statements are explained separately.                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul>                             |

## event-thresholds

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>event-thresholds {<br/>    frame-error count;<br/>    frame-period count;<br/>    frame-period-summary count;<br/>    symbol-period count;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface</a> <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | <p>Configure threshold limit values for link events in periodic OAM PDUs.</p> <p>The remaining statements are explained separately.</p>                     |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li></ul>               |

## family

```
Syntax  family {
    ethernet-switching {
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        interface-mode (access | trunk);
        recovery-timeout seconds;
        storm-control profile-name;
        vlan {
            members (vlan-name [ -vlan-names ] | all);
        }
    }
    inet {
        accounting {
            destination-class-usage;
            source-class-usage {
                input;
                output;
            }
        }
    }
    address ipv4-address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        preferred;
        primary;
        vrrp-group group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            advertisements-threshold number;
            authentication-key key;
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority number;
            track {
                interface interface-name {
                    priority-cost number;
                }
                priority-hold-time seconds;
                route ip-address/mask routing-instance instance-name priority-cost cost;
            }
            virtual-address [ addresses ];
            vrrp-inherit-from {
                active-group group-number;
                active-interface interface-name;
            }
        }
    }
}
```

```
    }
  }
  filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
  }
  mtu bytes;
  no-neighbor-learn;
  no-redirects;
  primary;
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
    }
  }
}
inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address address {
  eui-64;
  ndp ip-address (mac | multicast-mac) mac-address <publish>;
  preferred;
  primary;
  vrrp-inet6-group group-id {
    accept-data | no-accept-data;
    advertisements-threshold number;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    inet6-advertise-interval milliseconds;
    preempt | no-preempt {
      hold-time seconds;
    }
  }
  priority number;
  track {
    interface interface-name {
      priority-cost number;
    }
    priority-hold-time seconds;
    route ip-address/mask routing-instance instance-name priority-cost cost;
  }
  virtual-inet6-address [addresses];
  virtual-link-local-address ipv6-address;
  vrrp-inherit-from {
    active-group group-name;
    active-interface interface-name;
```



```

    }
  }
}
(dad-disable | no-dad-disable);
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
mtu bytes;
nd6-stale-time time;
no-neighbor-learn;
no-redirects;
policer {
  input policer-name;
  output policer-name;
}
rpf-check {
  fail-filter filter-name;
  mode {
    loose;
  }
}
mpls {
  filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
  }
  mtu bytes;
}
}
}

```

|                            |   |
|----------------------------|---|
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces interface-name unit logical-unit-number</a> ],<br>[edit <a href="#">interfaces interface-range interface-name unit logical-unit-number family</a> ]  |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>         | Configure protocol family information for the logical interface on the QFX Series product.  |
| <b>Default</b>             | Access interfaces on the QFX Series are set to <b>family ethernet-switching</b> by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting first.<br><br>You must configure a logical interface to be able to use the physical device. |

**Options** Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- 10-Gigabit Ethernet (**xe**)

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)
- [Configuring Link Aggregation on page 2593](#)
- [Configuring IRB Interfaces on page 1675](#)
- *Junos OS Network Interfaces Library for Routing Devices*

---

## fibre-channel (Alarm)

---

**Syntax** fibre-channel {  
    [link-down](#) (red | yellow | ignore);  
}

**Hierarchy Level** [edit chassis [alarm](#)],  
[edit chassis interconnect-device *name* [alarm](#)],  
[edit chassis node-group *name* [alarm](#)]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure alarms for a Fibre Channel interface.

**Options** The remaining statement is explained separately.—

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Alarms on page 7191](#)
- [Interface Alarm Messages on page 7195](#)

## filter

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>filter {   group <i>filter-group-number</i>;   input <i>filter-name</i>;   input-list [ <i>filter-names</i> ];   output <i>filter-name</i>;   output-list [ <i>filter-names</i> ]; }</pre>  |
| <b>Hierarchy Level</b>     | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</pre>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>                                      |
| <b>Description</b>         | <p>Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the <b>family ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>mpls</b>, or <b>vpls</b> only.</p> |




**NOTE:** On QFX3500 and QFX3600 switches running Enhanced Layer 2 Software, VPLS is not supported.

|                                 |  |
|---------------------------------|--|
| <b>Options</b>                  | <p><b>group <i>filter-group-number</i></b>—Define an interface to be part of a filter group.<br/> <b>Range:</b> 1 through 255</p> <p><b>input <i>filter-name</i></b>—Name of one filter to evaluate when packets are received on the interface.</p> <p><b>output <i>filter-name</i></b>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Applying a Filter to an Interface</i></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> <li>• <i>Junos OS Administration Library for Routing Devices</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> </ul> |

- *Configuring Firewall Filters (CLI Procedure)*
- *Configuring Firewall Filters and Policers for VPLS*
- *family*
- *family*

## flow-control

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | (flow-control   no-flow-control);  |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> <ul style="list-style-type: none"> <li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li> <li>• <b>no-flow-control</b>—Disable flow control.</li> </ul> |
| <b>Default</b>                  | Flow control is disabled.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">configured-flow-control on page 2637</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>  |

## force-up

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | force-up;   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> 802.3ad lacp;<br>[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">aggregated-ether-options</a> lacp;   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure the state of the interface as up when the peer has limited LACP capability. You can also Configure the peer interface (in MC-LAG) to remain up even with limited LACP capability.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2393</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2589</a></li><li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li><li>• </li></ul> |

## fpc

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>fpc slot {   auto-speed-detection disable;   pic <i>pic-number</i> {     tunnel-port <i>port-number</i> tunnel-services;     port <i>port-number</i> {       channel-speed (<i>speed</i> disable-auto-speed-detection) ;     }     port-range <i>port-range-low</i> <i>port-range-high</i> {       channel-speed (<i>speed</i> disable-auto-speed-detection);     }   } }</pre> |
| <b>Hierarchy Level</b>          | [edit chassis]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.</p> <p>For generic routing encapsulation (GRE) tunneling, use the <b>tunnel-port</b> statement to specify the port that you want to convert to a GRE tunnel port.</p>  |
| <b>Options</b>                  | <p><b>slot</b>—Number of the FPC slot. For QFX3500 and QFX3600 devices, the slot number is always 0.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show chassis fpc on page 639</a></li> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul>   |

## frame-error

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>frame-error count;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | <p>Configure the threshold value for sending frame error events or taking the action specified in the action profile.</p> <p>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for frame error events.</p> <p><b>Range:</b> 1 through 100 seconds</p> <p><b>Default:</b> 1 second</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li></ul>  |

## frame-period

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>frame-period count;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | <p>Configure the number of frame errors within the last N frames that has exceeded a threshold.</p> <p>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.</p> |
| <b>Options</b>                  | <p><i>count</i>—Threshold count in seconds for frame error events.</p> <p><b>Range:</b> 1 through 100 seconds</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li></ul>  |



## frame-period-summary

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>frame-period-summary count;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management event link-event-rate</a> ],<br>[edit protocols <a href="#">oam ethernet link-fault-management interface interface-name event-thresholds</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | <p>Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period.</p> |
| <b>Options</b>                  | <p><code>count</code>—Threshold count in seconds for frame period summary error events.</p> <p><b>Range:</b> 1 through 100 seconds</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul>   |

## gratuitous-arp-reply

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>(gratuitous-arp-reply   no-gratuitous-arp-reply);</code>  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ],<br>[edit interfaces <a href="#">interface-range interface-range-name</a> ]                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Enable processing of ARP updates received via gratuitous ARP reply messages.  |
| <b>Default</b>                  | Updating of the ARP cache is disabled on all Ethernet interfaces.   |
| <b>Options</b>                  | <p><code>gratuitous-arp-reply</code>—Update the ARP cache.</p> <p><code>no-gratuitous-arp-reply</code>—Do not update the ARP cache.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>      |

## group

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>group group-name {<br/>    link-to-monitor interface-name;<br/>    link-to-disable interface-name;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit protocols uplink-failure-detection]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure a group of uplink and downlink interfaces for uplink failure detection.  |
| <b>Options</b>                  | <p><i>group-name</i>—Name of the uplink failure detection group.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2392</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2592</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2457</a></li></ul> |

## group (Redundant Trunk Groups)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>group <i>name</i> {     interface <i>interface-name</i> &lt;primary&gt;;     interface <i>interface-name</i>;     preempt-cutover-timer <i>seconds</i>; }</pre>  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with ELS:<br/>[edit switch-options <b>redundant-trunk-group</b>]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options <b>redundant-trunk-group</b>]</li> </ul>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| <b>Description</b>              | Create a redundant trunk group.   |
| <b>Options</b>                  | <p><b>name</b>—The name of the redundant trunk group. The group name must start with a letter and can consist of letters, numbers, dashes, and underscores.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring Redundant Trunk Links for Faster Recovery</i></li> <li><a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578</a></li> <li><a href="#">Understanding Redundant Trunk Links on page 2447</a></li> </ul>  |

## hash-mode

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>hash-mode {<br/>    layer2-header;<br/>    layer2-payload;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.   |
| <b>Description</b>              | <p>Select the mode for the hashing algorithm.</p> <p>The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.</p> <p>For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.</p> <p>The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to <b>layer2-payload</b> if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to <b>layer2-header</b> if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.</p> <p>If the hash mode is set to <b>layer2-payload</b>, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the <b>set forwarding-options enhanced-hash-key inet</b> statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the <b>set forwarding-options enhanced-hash-key inet6</b> statement.</p> <p>If the hash mode is set to <b>layer2-header</b>, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the <b>set forwarding-options enhanced-hash-key layer2</b> statement.</p> |
| <b>Default</b>                  | layer2-payload   |
| <b>Options</b>                  | <p><b>layer-2-payload</b>—Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.</p> <p><b>layer-2-header</b>—Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.</p>  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |

**Related  
Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2590](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396](#)
- [enhanced-hash-key on page 2644](#)
- [inet on page 2676](#)
- [inet6 on page 2678](#)
- [layer2 on page 2695](#)

## hold-time (Physical Interface)

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | hold-time up <i>milliseconds</i> down <i>milliseconds</i> ;   |
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces</a> <i>interface-name</i> ],<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> ]   |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 10.4R5 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.   |
| <b>Description</b>         | Specify the <b>hold-time</b> value to use to damp shorter interface transitions milliseconds.<br>When an interface goes from up to down, it is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly, an interface is not advertised as being up until it has remained up for the hold-time period. |



### NOTE:

- We recommend that you configure the hold-time value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for hold-time depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.
  - The hold-time option is not available for controller interfaces.
- 

|                                 |   |
|---------------------------------|---|
| <b>Default</b>                  | Interface transitions are not damped.   |
| <b>Options</b>                  | <b>down <i>milliseconds</i></b> —Hold time to use when an interface transitions from up to down.<br>Junos OS advertises the transition within 100 milliseconds of the time value you specify.<br><b>Range:</b> 0 through 4,294,967,295<br><b>Default:</b> 0 (interface transitions are not damped)<br><br><b>up <i>milliseconds</i></b> —Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.<br><b>Range:</b> 0 through 4,294,967,295<br><b>Default:</b> 0 (interface transitions are not damped) |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>advertise-interval</i></li><li>• <i>interfaces (for EX Series switches)</i></li><li>• <i>Physical Interface Damping Overview</i></li></ul>   |

- *Damping Shorter Physical Interface Transitions*
- *Damping Longer Physical Interface Transitions*

## iccp

```
Syntax  iccp {
    authentication-key string;
    local-ip-addr local-ip-addr;
    peer ip-address{
        authentication-key string;
        backup-liveness-detection {
            backup-peer-ip ip-address;
        }
        liveness-detection {
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (1 | automatic);
        }
        local-ip-addr ipv4-address;
        session-establishment-hold-time seconds;
    }
    session-establishment-hold-time seconds;
    traceoptions {
        file <filename> <files number> <match regular-expression> <microsecond-stamp>
        <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

**Hierarchy Level** [edit protocols]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Configure Interchassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.



**NOTE:** Backup liveness detection is not supported on MX Series routers.

The remaining statement are explained separately.



|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

## irb (Interfaces)

---

```
Syntax  irb {
    accounting-profile name;
    description text;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        encapsulation type;
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage {
                    input;
                    output;
                }
            }
        }
        address ipv4-address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            preferred;
            primary;
            vrrp-group group-number {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                advertisements-threshold number;
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
                (preempt | no-preempt) {
                    hold-time seconds;
                }
            }
            priority number;
            track {
                interface interface-name {
                    bandwidth-threshold bandwidth;
                    priority-cost number;
                }
                priority-hold-time seconds;
                route ip-address/mask routing-instance instance-name priority-cost cost;
            }
        }
        virtual-address [ addresses ];
    }
}
```

```

        vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
        }
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;

```

```

    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

Hierarchy Level [edit interfaces *interface-name*

|                                 |  |
|---------------------------------|--|
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3R2 for EX Series switches.<br><b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series.   |
| <b>Description</b>              | Configure the properties of a specific integrated bridging and routing (IRB) interface.<br><br>The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">[edit interfaces] Hierarchy Level</a></li> <li>• <a href="#">[edit interfaces] Configuration Statement Hierarchy on EX Series Switches</a></li> </ul> |

## inet (interfaces)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>inet {     address <i>address</i> {         primary;         filter input <i>filter-name</i>;         filter output <i>filter-name</i>;         targeted-broadcast;     } }</pre> |
| <b>Hierarchy Level</b>          | <a href="#">[edit interfaces interface-name unit logical-unit-number family]</a> ,<br><a href="#">[edit interfaces interface-range interface-name unit logical-unit-number family]</a> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure the primary IP address for the logical interface.  |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.  |
| <b>Options</b>                  | The remaining statements are explained separately.—  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li> </ul>  |

## inet (enhanced-hash-key)

---

**Syntax**    `inet {  
              no-ipv4-destination-address;  
              no-ipv4-source-address;  
              no-l4-destination-port;  
              no-l4-source-port;  
              no-protocol;  
              vlan-id;  
          }`

**Hierarchy Level**    [edit forwarding-options [enhanced-hash-key](#)]

**Release Information**    Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

**Description**    Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.

When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the **inet** statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to **layer2-payload**. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

**Default**    The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:

- IP destination address
- IP source address
- Layer 4 destination port
- Layer 4 source port
- Protocol

**Options**    **no-ipv4-destination-address**—Exclude the IPv4 destination address field from the hashing algorithm.

**no-ipv4-source-address**—Exclude the IPv4 source address field from the hashing algorithm.

**no-l4-destination-port**—Exclude the Layer 4 destination port field from the hashing algorithm.

**no-l4-source-port**—Exclude the Layer 4 source port field from the hashing algorithm.

**no-protocol**—Exclude the protocol field from the hashing algorithm.

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

**Required Privilege** interface—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2590](#)
  - [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396](#)
  - [enhanced-hash-key on page 2644](#)
  - [hash-mode on page 2666](#)
  - [inet6 on page 2678](#)

## inet6 (interfaces)

**Syntax**

```
inet6 {
    address address {
        eui-64
        preferred
        primary;
        filter input filter-name;
        filter output filter-name;
    }
}
```

**Hierarchy Level** [edit [interfaces interface-name unit logical-unit-number](#) family],  
[edit [interfaces interface-range interface-name unit logical-unit-number](#) family]

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Configure the primary IP address for the logical interface.

**Default** You must configure a logical interface to be able to use the physical device.

**Options** The remaining statements are explained separately.—

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)

## inet6 (enhanced-hash-key)

---

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>inet6 {<br/>    no-ipv6-destination-address;<br/>    no-ipv6-source-address;<br/>    no-l4-destination-port;<br/>    no-l4-source-port;<br/>    no-next-header;<br/>    vlan-id;<br/>}</pre>  |
| <b>Hierarchy Level</b>     | [edit forwarding-options <a href="#">enhanced-hash-key</a> ]   |
| <b>Release Information</b> | Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.   |
| <b>Description</b>         | <p>Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.</p> <p>When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-payload</b> statement.</p> |
| <b>Default</b>             | <p>The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:</p> <ul style="list-style-type: none"><li>• IP destination address</li><li>• IP source address</li><li>• Layer 4 destination port</li><li>• Layer 4 source port</li><li>• Next header</li></ul>  |
| <b>Options</b>             | <p><b>no-ipv6-destination-address</b>—Exclude the IPv6 destination address field from the hashing algorithm.</p> <p><b>no-ipv6-source-address</b>—Exclude the IPv6 source address field from the hashing algorithm.</p> <p><b>no-l4-destination-port</b>—Exclude the Layer 4 destination port field from the hashing algorithm.</p>  |



**no-l4-source-port**—Exclude the Layer 4 source port field from the hashing algorithm.

**no-next-header**—Exclude the Next Header field from the hashing algorithm.

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

**Required Privilege** interface—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2590](#)
  - [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396](#)
  - [enhanced-hash-key on page 2644](#)
  - [hash-mode on page 2666](#)
  - [inet on page 2676](#)

## interface (Multichassis Protection)

**Syntax** interface *interface-name*;

**Hierarchy Level** [edit [multi-chassis multi-chassis-protection peer](#)]

**Release Information** Statement introduced in Junos OS Release 9.6 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two switches hosting a multichassis link aggregation group (MC-LAG) use this link to pass Interchassis Control Protocol (ICCP) and data traffic.

**Required Privilege** interface—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.

## interface (OAM LFM)

---

**Syntax**    `interface interface-name {  
          link-discovery (active | passive);  
          pdu-interval interval;  
          pdu-threshold threshold-value;  
          remote-loopback;  
          event-thresholds {  
              frame-error count;  
              frame-period count;  
              frame-period-summary count;  
              symbol-period count;  
          }  
          negotiation-options {  
              allow-remote-loopback;  
              no-allow-link-events;  
          }  
          }  
          }`

**Hierarchy Level**    [edit protocols [oam](#) [ethernet](#) [link-fault-management](#)]

**Release Information**    Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description**    Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

**Options**    *interface-name*—Name of the interface to be enabled for IEEE 802.3ah OAM link fault management (LFM) support.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.



**Related Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)

## interface (Redundant Trunk Groups)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> &lt;primary&gt;; interface <i>interface-name</i>;</pre>  |
| <b>Hierarchy Level</b>          | <p>For platforms with ELS:</p> <pre>[edit switch-options <b>redundant-trunk-group</b> <i>group name</i>]</pre> <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <b>redundant-trunk-group</b> <i>group name</i>]</pre>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over as the primary link without waiting for normal STP convergence.</p>   |
| <b>Options</b>                  | <p><b>interface <i>interface-name</i></b>—A logical interface or an aggregated interface containing multiple ports.</p> <p><b>primary</b>—(Optional) Specify one of the interfaces in the redundant group as the primary link. The interface without this option is the secondary link in the redundant group. If a link is not specified as <b>primary</b>, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are <b>ge-0/1/0</b> and <b>ge-0/1/1</b>, the software assigns <b>ge-0/1/1</b> as the active link.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery</a></li> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578</a></li> <li>• <a href="#">Understanding Redundant Trunk Links on page 2447</a></li> </ul>   |

## interface-mode

|                          |  |
|--------------------------|--|
| Syntax                   | interface-mode (access   trunk);   |
| Hierarchy Level          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching],<br>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]   |
| Release Information      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.   |
| Description              | <p> <b>NOTE:</b> This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see <a href="#">port-mode</a>. For ELS details, see “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 43.</p> <p>(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the <b>trunk</b> option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the <b>vlan-id</b> or <b>vlan-id-list</b> statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the <b>access</b> option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the <b>vlan-id</b> statement.</p> <p> <b>NOTE:</b> On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure <b>interface-mode</b> and <b>irb</b> for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> |
| Options                  | <p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p>   |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| Related Documentation    | <ul style="list-style-type: none"> <li>Configuring a Logical Interface for Access Mode</li> <li>Configuring a Logical Interface for Trunk Mode</li> </ul>  |

- *Example: Connecting Access Switches to a Distribution Switch*

## interface-range

**Syntax** `interface-range interface-range-name {`  
     `disable;`  
     `description text;`  
     `ether-options {`  
         `802.3ad aex {`  
             `lacp {`  
                 `force-up;`  
             `}`  
         `}`  
     `(auto-negotiation | no-auto-negotiation);`  
     `(flow-control | no-flow-control);`  
     `link-mode mode;`  
     `speed (auto-negotiation | speed);`  
     `}`  
     `hold-time milliseconds down milliseconds;`  
     `member interface-name;`  
     `member-range starting-interface-name to ending-interface-name;`  
     `mtu bytes;`  
     `unit logical-unit-number {`  
         `description text;`  
         `disable;`  
         `family family-name {...}`  
         `(traps | no traps);`  
         `vlan-id vlan-id-number;`  
     `}`  
     `}`

**Hierarchy Level** [edit [interfaces](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX series.

**Description** Group interfaces that share a common configuration profile.



**NOTE:** The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces.

**Options** `interface-range-name`—Name of the interface range.



**NOTE:** You can use regular expressions and wildcards to specify the interfaces in the member range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

**Required Privilege Level** `interface`—To view this statement in the configuration.  
     `interface-control`—To add this statement to the configuration.

- Related Documentation**
- [Understanding Interface Ranges on page 2406](#)
  - [Interfaces Overview on page 2389](#)
  - [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)
  - *Junos OS Network Interfaces Library for Routing Devices*

## interfaces

---

```
Syntax  interfaces {
        aex {
            disable;
            aggregated-ether-options {
                configured-flow-control {
                    rx-buffers (on | off);
                    tx-buffers (on | off);
                }
            }
            (fcoe-lag | no-fcoe-lag);
            (flow-control | no-flow-control);
            lacp mode {
                admin-key key;
                force-up;
                periodic interval;
                system-id mac-address;
            }
            link-speed speed;
            local-bias;
            loopback;
            no-loopback;
            minimum-links number;
        }
        mc-ae {
            chassis-id chassis-id;
            mc-ae-id mc-ae-id;
            mode (active-active);
            status-control (active | standby);
        }
        description text;
        gratuitous-arp-reply | no-gratuitous-arp-reply
        hold-time down milliseconds up milliseconds;
        mtu bytes;
        no-gratuitous-arp-request;
        traceoptions;
        (traps | no traps);
        unit logical-unit-number {
            disable;
            description text;
            family {
                ethernet-switching {
                    filter input filter-name;
                    filter output filter-name;
                    native-vlan-id vlan-id;
                    port-mode mode;
                    reflective-relay;
                    vlan {
                        members [ (all | names | vlan-ids) ];
                    }
                }
            }
            inet {
                address address {
                    primary;
```



```

    }
    filter input filter-name;
    filter output filter-name;
    primary;
    targeted-broadcast;
  }
  (traps | no traps);
  vlan-id vlan-id-number;
}
vlan-tagging;
}
interface-range interface-range-name {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
  disable;
  description text;
  family family-name {...}
  (traps | no traps);
  vlan-id vlan-id-number;
}
}
lo0 {
  disable;
  description text;
  hold-time milliseconds down milliseconds;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    disable;
    description text;
    family {
      inet {
        address address {
          primary;
        }
      }
      filter input filter-name;
    }
  }
}

```

```
        filter output filter-name;  
        primary;  
        targeted-broadcast;  
    }  
    (traps | no traps);  
}  
}  
mex {  
    disable;  
    description text;  
    hold-time milliseconds down milliseconds;  
    (gratuitous-arp-reply | no-gratuitous-arp-reply);  
    no-gratuitous-arp-request;  
    traceoptions;  
    traps;  
    unit logical-unit-number {  
        disable;  
        description text;  
        family {  
            ethernet-switching {  
                filter input filter-name;  
                filter output filter-name;  
                native-vlan-id vlan-id;  
                port-mode mode;  
                reflective-relay;  
            }  
            vlan {  
                members [ (all | names | vlan-ids) ];  
            }  
        }  
        inet {  
            address address {  
                primary;  
                filter input filter-name;  
                filter output filter-name;  
                primary;  
                targeted-broadcast;  
            }  
        }  
    }  
    traps;  
    vlan-id vlan-id-number;  
}  
vlan-tagging;  
vlan {  
    disable;  
    description text;  
    (gratuitous-arp-reply | no-gratuitous-arp-reply);  
    hold-time milliseconds down milliseconds;  
    mtu bytes;  
    no-gratuitous-arp-request;  
    traceoptions;  
    (traps | no traps);  
    unit logical-unit-number {  
        description text;  
        disable;  
        family {  
            inet {
```

```

        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
    }
    (traps | no traps);
}
}
fc-0/0/port {
    fibrechannel-options {
        bb-sc-n;
        (loopback | no-loopback);
        speed (auto-negotiation | 2g | 4g | 8g);
    }
    unit logical-unit-number {
        disable;
        description text;
        family {
            fibre-channel {
                port-mode np-port;
            }
        }
        (traps | no traps);
    }
}
ge-0/0/port {
    disable;
    description text;
    ether-options {
        802.3ad aex {
            lacp {
                force-up;
                primary;
            }
        }
    }
    (auto-negotiation | no-auto-negotiation);
    configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
    }
    (flow-control | no-flow-control);
    link-mode mode;
    loopback;
    no-loopback;
    speed (auto-negotiation | speed);
}
gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
    description text;
    disable;

```

```
family {
  ethernet-switching {
    filter input filter-name;
    filter output filter-name;
    native-vlan-id vlan-id;
    port-mode mode;
    reflective-relay;
    vlan {
      members [ (all | names | vlan-ids) ];
    }
  }
  inet {
    address address {
      primary;
    }
    filter input filter-name;
    filter output filter-name;
    primary;
    targeted-broadcast;
  }
  (traps | no traps);
  vlan-id vlan-id-number;
}
vlan-tagging;
}
vrrp-group group-id {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bits-per-second priority-cost priority;
      priority-cost priority;
    }
    priority-hold-time seconds;
    route prefix/prefix-length routing-instance instance-name priority-cost priority;
  }
}
virtual-address [ addresses ];
}
xe-0/0/port {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
        (primary | backup);
      }
    }
  }
}
```

```

configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
}
(flow-control | no-flow-control);
loopback;
no-loopback;
}
(gratuitous-arp-reply | no-gratuitous-arp-reply)
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
    disable;
    description text;
    family {
        ethernet-switching {
            filter input filter-name;
            filter output filter-name;
            native-vlan-id vlan-id;
            port-mode mode;
            reflective-relay;
            vlan {
                members [ (all | names | vlan-ids) ];
            }
        }
        fibre-channel {
            port-mode (f-port | np-port);
        }
        inet {
            address address {
                primary;
            }
            filter input filter-name;
            filter output filter-name;
            primary;
            targeted-broadcast;
        }
        (traps | no traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}
}

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

|                                 |  |
|---------------------------------|--|
| <b>Description</b>              | <p>Configure the interfaces on the QFX Series.</p> <p>Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.</p>   |
| <b>Options</b>                  | <p><b>aex</b>—Configure an aggregated Ethernet interface.</p> <p><b>xe-0/0/</b><i>port</i>/<b>—</b>Configure a 10-Gigabit Ethernet interface.</p> <p><b>ge-0/0/</b><i>port</i>/<b>—</b>Configure a Gigabit Ethernet interface.</p> <p><b>fc-0/0/</b><i>port</i>/<b>—</b>Configure a Fibre Channel interface.</p> <p><b>meX</b>/<b>—</b>Configure a management interface.</p> <p><b>mc-ae</b>—Configure a multichassis aggregated Ethernet (MC-AE) interface.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Interfaces Overview on page 2389</a></li><li>• <a href="#">Understanding Interface Ranges on page 2406</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <a href="#">Configuring Link Aggregation on page 2593</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2593</a></li></ul>  |

---

## lacp (802.3ad)

---

**Syntax**    `lacp {  
              force-up;  
              (primary | backup);  
              port-priority;  
          }`

**Hierarchy Level**    [edit [interfaces interface-name ether-options 802.3ad](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.



**NOTE:** The port-priority statement is not supported on QFabric systems.

---

**Required Privilege Level**    interface—To view this statement in the configuration.  
                                  interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring Link Aggregation on page 2593](#)
- [Configuring Aggregated Ethernet LACP on page 2589](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2393](#)

## larp (Aggregated Ethernet)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>larp (active   passive) {<br/>    admin-key <i>key</i>;<br/>    periodic <i>interval</i><br/>    system-ID <i>mac-address</i>;<br/>    force-up;<br/>}</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name</a> <a href="#">aggregated-ether-options</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation on page 2593</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2589</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2393</a></li></ul> |



## layer2 (enhanced-hash-key)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> layer2 {     no-destination-mac-address;     no-ether-type;     no-source-mac-address;     vlan-id; } </pre>  |
| <b>Hierarchy Level</b>     | [edit forwarding-options <b>enhanced-hash-key</b> ]   |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p>   |
| <b>Description</b>         | <p>Select the fields in the Layer 2 header that are used by the hashing algorithm to make hashing decisions.</p> <p>When traffic enters a link aggregation group (LAG) bundle, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links. The hashing algorithm always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When traffic is exiting a device that has enabled ECMP, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to forward traffic to the next hop device.</p> <p>The hashing algorithm only inspects the fields in the Layer 2 header when the hash mode is set to Layer 2 header. You can set the hash mode to Layer 2 header using the <b>set forwarding-options enhanced-hash-key hash-mode layer2-header</b> statement.</p> |
| <b>Default</b>             | <p>The hash mode of the hashing algorithm is set to Layer 2 payload, by default. When the hash mode is set to Layer 2 payload, the hashing algorithm does not use fields in the Layer 2 header to make hashing decisions.</p> <p>The following fields are used by the hashing algorithm when the hash mode of the hashing algorithm is set to Layer 2 header, by default:</p> <ul style="list-style-type: none"> <li>• Destination MAC address</li> <li>• Ethertype</li> <li>• Source MAC address</li> </ul>  |
| <b>Options</b>             | <p><b>no-destination-mac-address</b>—Exclude the destination MAC address field from the hashing algorithm.</p> <p><b>no-ether-type</b>—Exclude the Ethertype field from the hashing algorithm.</p> <p><b>no-source-mac-address</b>—Exclude the source MAC address field from the hashing algorithm.</p>   |

**vlan-id**—Include the VLAN ID field in the hashing algorithm.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 2590](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396](#)
- [enhanced-hash-key on page 2644](#)
- [hash-mode on page 2666](#)

---

## link-adjacency-loss

---

**Syntax** link-adjacency-loss;

**Hierarchy Level** [edit protocols [oam ethernet link-fault-management action-profile event](#)]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description** Configure **loss of adjacency** event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the [action](#) statement.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)

## link-discovery

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | link-discovery (active   passive);   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management interface interface-name</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs.         |
| <b>Options</b>                  | <p><i>active</i>—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.</p> <p><i>passive</i>—In passive mode, the peer initiates the discovery process.</p> <p>Once the discovery process is initiated, both sides participate in discovery.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul>  |

## link-down

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | link-down;  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile action</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Mark the interface as down for transit traffic.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul> |

## link-event-rate

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>link-event-rate {<br/>    frame-error <i>count</i>;<br/>    frame-period <i>count</i>;<br/>    frame-period-summary <i>count</i>;<br/>    symbol-period <i>count</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile event</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | <p>Configure the number of link fault management (LFM) events per second.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li></ul>  |

## link-fault-management

```
Syntax  link-fault-management {
        action-profile profile-name;
        action {
            syslog;
            link-down;
        }
        event {
            link-adjacency-loss;
            link-event-rate {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
        }
        interface interface-name {
            link-discovery (active | passive);
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
            event-thresholds {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
            negotiation-options {
                allow-remote-loopback;
                no-allow-link-events;
            }
        }
    }
```

**Hierarchy Level** [edit protocols [oam](#) [ethernet](#)]

**Release Information** Statement introduced in Junos OS Release 9.4 for EX Series switches.

**Description** Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)

## link-to-disable

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>link-to-disable <i>interface-name</i>;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of eight downlink interfaces in a group.  |
| <b>Options</b>                  | <i>interface-name</i> —Name of the downlink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2392</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2592</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2457</a></li></ul> |

## link-to-monitor

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>link-to-monitor <i>interface-name</i>;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of eight uplink interfaces in a group.  |
| <b>Options</b>                  | <i>interface-name</i> —Name of the uplink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface.  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Uplink Failure Detection on page 2392</a></li><li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2592</a></li><li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2457</a></li></ul> |

---

## link-down

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | link-down (red   yellow   ignore);   |
| <b>Hierarchy Level</b>          | [edit chassis <a href="#">alarm ethernet</a> ],<br>[edit chassis <a href="#">alarm fibre-channel</a> ],<br>[edit chassis interconnect-device <i>name</i> <a href="#">alarm ethernet</a> ],<br>[edit chassis node-group <i>name</i> <a href="#">alarm fibre-channel</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Specify either red, yellow, or ignore to display when the link is down.  |
| <b>Options</b>                  | <p><b>red</b>—Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.</p> <p><b>yellow</b>—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.</p> <p><b>ignore</b>—Suppresses or ignores the alarm.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    |  |

## link-mode

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>link-mode mode;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Set the device's link-connection characteristic.   |
| <b>Default</b>                  | The <b>full-duplex</b> mode is enabled.  |
| <b>Options</b>                  | <p><b>mode</b>—Link characteristic:</p> <ul style="list-style-type: none"><li>• <b>full-duplex</b>—Connection is full duplex.</li><li>• <b>half-duplex</b>—Connection is half duplex.</li><li>• <b>automatic</b>—Link mode is negotiated.</li></ul> <p>If <b>no-auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select only <b>full-duplex</b> or <b>half-duplex</b>. If <b>auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select any mode.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>   |



## link-speed

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | link-speed <i>speed</i> ;  |
| <b>Hierarchy Level</b>          | [edit interfaces aex <a href="#">aggregated-ether-options</a> ]  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | For aggregated Ethernet interfaces only, set the required link speed.  |
| <b>Options</b>                  | <p><b>speed</b>—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p>On QFX5100 standalone switches, you can configure <b>mixed</b> as the link speed. The <b>mixed</b> option allows you to configure mixed rate aggregated Ethernet bundles on a QFX5100 standalone switch with link speeds of 40G and 10G only. Load balancing will not work if you configure link speeds that are not supported.</p> <p>Aggregated Ethernet links on the QFX Series can have one of the following speed values:</p> <ul style="list-style-type: none"> <li>• <b>100g</b>—Links are 100 Gbps.</li> <li>• <b>100m</b>—Links are 100 Mbps.</li> <li>• <b>10g</b>—Links are 10 Gbps.</li> <li>• <b>1g</b>—Links are 1 Gbps.</li> <li>• <b>40g</b>—Links are 40 Gbps.</li> <li>• <b>50g</b>—Links are 50 Gbps.</li> <li>• <b>80g</b>—Links are 80 Gbps.</li> <li>• <b>8g</b>—Links are 8 Gbps.</li> <li>• <b>0c192</b>—Links are OC-192.</li> <li>• <b>mixed</b>—Links are 10 Gbps and 40Gbps.</li> </ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2593</a></li> </ul>  |

## liveness-detection

---

**Syntax**

```
liveness-detection {  
  detection-time {  
    threshold milliseconds;  
  }  
  minimum-interval milliseconds;  
  minimum-receive-interval milliseconds;  
  multiplier number;  
  no-adaptation;  
  transmit-interval {  
    minimum-interval milliseconds;  
    threshold milliseconds;  
  }  
  version (1 | automatic);  
}
```

**Hierarchy Level** [edit protocols *iccp* *peer*]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Enable Bidirectional Forwarding Detection (BFD). BFD enables rapid detection of communication failures between peers.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## local-bias

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>local-bias;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit interfaces aex aggregated-ether-options]</code>  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices.  |
| <b>Description</b>              | <p>Enable local link bias for all links in the aggregated Ethernet interface.</p> <p>Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.</p> <p>You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Local Link Bias (CLI Procedure) on page 2596</a></li> <li>• <a href="#">Understanding Local Link Bias on page 2408</a></li> </ul>   |

## local-ip-addr (ICCP)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>local-ip-addr <i>local-ip-address</i>;</code>   |
| <b>Hierarchy Level</b>          | <p><code>[edit protocols <a href="#">iccp</a>],</code><br/> <code>[edit protocols <a href="#">iccp</a> <a href="#">peer</a> <i>peer-IP-address</i>]</code></p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.0 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 12.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> |
| <b>Description</b>              | Specify the local IP address of the interchassis link (ICL) interface that Interchassis Control Protocol (ICCP) uses to communicate to the peers that host a multichassis link aggregation group (MC-LAG).                          |
| <b>Options</b>                  | <b><i>local-ip-address</i></b> —Default local IP address to be used by all peers.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |

## loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | (loopback   no-loopback);  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> aggregated-ether-options],<br>[edit interfaces <i>interface-name</i> ether-options] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | For aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.            |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet Loopback Capability on page 2589</a></li></ul>    |

## management-ethernet (Alarm)

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | management-ethernet {<br>link-down (red   yellow   ignore);<br>}  |
| <b>Hierarchy Level</b>     | [edit chassis alarm],<br>[edit chassis interconnect-device <i>name</i> alarm],<br>[edit chassis node-group <i>name</i> alarm] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.2 for the QFX Series.   |
| <b>Description</b>         | Configure alarms for a management Ethernet interface.   |



**NOTE:** If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.

---

|                                 |   |
|---------------------------------|---|
| <b>Options</b>                  | The remaining statement is explained separately.—   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Alarms on page 7191</a></li><li>• <a href="#">Interface Alarm Messages on page 7195</a></li></ul> |

## member

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>member interface-name;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-range interface-range-name</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Specify the name of the member interface belonging to an interface range on the QFX Series switch.  |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li> <li>• <a href="#">Interfaces Overview on page 2389</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## member-range

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>member-range starting-interface-name ending-interface-name;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-range interface-range-name</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.   |
| <b>Options</b>                  | <i>starting interface-name ending interface-name</i> —Name of the first member and the name of the last member in the interface sequence.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Interface Ranges on page 2406</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li> <li>• <a href="#">Interfaces Overview on page 2389</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul> |

## mc-ae

**Syntax** `mc-ae {  
     chassis-id chassis-id;  
     mc-ae-id mc-ae-id;  
     mode (active-active);  
     status-control (active | standby);  
 }`

**Hierarchy Level** [edit [interfaces aggregated-ether-options](#)]

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

**Description** Specify the multichassis aggregated Ethernet interface configuration.

**Options** **chassis-id**—Specify the chassis ID for LACP to calculate the port number of the MC-LAG physical member links.

**mc-ae-id**—Specify the identification number of MC-LAG device. The two MC-LAG QFX3500 devices that manage a given MC-LAG must have the same **mc-lag-id**.

**mode (active | active)**—Specify that the MC-LAG is in active-active mode. In this mode, if a member interface of the MC-LAG goes down, traffic can still be forwarded to the QFX3500 devices hosting the MC-LAG using the interchassis link-protection link (ICL-PL). The links from the client-device connected to both of the QFX3500 devices will remain active. Only active-active mode is supported at this time.

**status-control (active | standby)**—Specify if a peer is in active or standby mode. In active mode, the peer is considered the primary device, and in standby mode it is considered the secondary device. If the ICL-PL goes down, the peer in standby mode will bring its member links to standby state. If the Interchassis Control Protocol (ICCP) connection goes down, the peer in standby mode will change the LACP system ID to the default value on its member links.



**NOTE:** You cannot have both peers hosting an MC-LAG be in active or standby mode. One peer must be in active mode, and the other peer must be in standby mode.

**Required Privilege Level** interface—To view this statement in the configuration.  
 interface-control—To add this statement to the configuration.

## mc-ae-id

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>mc-ae-id <i>mc-ae-id</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aggregated-ether-options mc-ae</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  |
| <b>Description</b>              | Specify the multichassis aggregated Ethernet (MC-AE) identification number of the MC-AE that a given aggregated Ethernet interface belongs to. The two peers that host a given multichassis link aggregation group (MC-LAG) must have the same multichassis aggregated Ethernet ID. |
| <b>Options</b>                  | <b>Range:</b> 1 through 65535.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |

## minimum-interval (Liveness Detection)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>minimum-interval <i>milliseconds</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.   |
| <b>Description</b>              | Configure simultaneously the minimum interval at which the peer transmits liveness detection requests and the minimum interval at which the peer expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately by using the <b>transmit-interval</b> <b>minimal-interval</b> and <b>minimum-receive-interval</b> statements, respectively. |
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).<br><b>Range:</b> 1 through 255,000  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |

## minimum-links

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>minimum-links <i>number</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit interfaces <code>aex</code> <a href="#">aggregated-ether-options</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | For an aggregated Ethernet interface, set the minimum number of links that must be up for the bundle to be labeled up.                              |
| <b>Options</b>                  | <i>number</i> —Number of links.<br><b>Range:</b> 1 through 8<br><b>Default:</b> 1   |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation on page 2593</a></li></ul>   |

## minimum-receive-interval (Liveness Detection)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>minimum-receive-interval <i>milliseconds</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <code>iccp</code> <code>peer</code> <a href="#">liveness-detection</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| <b>Description</b>              | Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session.  |
| <b>Options</b>                  | <i>milliseconds</i> —Specify the minimum interval value.<br><b>Range:</b> 1 through 255,000  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.  |



## mode (QFX Series)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>mode active-active ;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aggregated-ether-options mc-ae</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  |
| <b>Description</b>              | Configure the multichassis link aggregation group (MC-LAG) to be in active-active mode. In active-active mode, all of the members of the MC-LAG will be active on both routing or switching devices. Only active-active mode is supported at this time. |
| <b>Options</b>                  | <b>active-active</b> —Specify that all of the members of the MC-LAG will be active on both routing or switching devices.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    |   |

## multi-chassis

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>multi-chassis {   multi-chassis-protection peer-ip-address {     interface interface-name;   } }</pre>   |
| <b>Hierarchy Level</b>          | [edit]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                                   |
| <b>Description</b>              | Configure an interchassis link-protection link (ICL-PL) between the two peers that host a multichassis link aggregation group (MC-LAG). You can configure either an aggregated Ethernet interface or a 10-Gigabit Ethernet interface to be an ICL-PL. |
| <b>Options</b>                  | <b>interface interface-name</b> —Specify the logical interface name of the peer.<br><br>The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |

## multi-chassis-protection

---


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>multi-chassis-protection <i>peer-ip-address</i> {<br/>    <b>interface</b> <i>interface-name</i>;<br/>}</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">multi-chassis</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.   |
| <b>Description</b>              | Configure multichassis link protection between the two peers that host a multichassis link aggregation group (MC-LAG). If the Interchassis Control Protocol (ICCP) connection is up and the interchassis link (ICL) comes up, the peer configured as standby brings up the multichassis aggregated Ethernet (MC-AE) interfaces shared with the peer. Multichassis protection must be configured on one interface for each peer.<br><br>The remaining statements are explained separately. |
| <b>Options</b>                  | <b>interface <i>interface-name</i></b> —Specify the logical interface name of the peer.<br><br>The remaining statements are explained separately.   |
| <b>Required Privilege Level</b> | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.   |

## multiplier (Liveness Detection)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>multiplier <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp</a> peer <a href="#">liveness-detection</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.   |
| <b>Description</b>              | Configure the number of liveness detection requests not received by the peer before Bidirectional Forwarding Detection (BFD) declares the peer is down. |
| <b>Options</b>                  | <b>number</b> —Maximum allowable number of liveness detection requests missed by the peer.<br><b>Range:</b> 1 through 255<br><b>Default:</b> 3          |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.                     |
| <b>Related Documentation</b>    |   |

## mtu

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>mtu bytes;</code>   |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> ],<br>[edit <b>interfaces</b> <i>interface-range</i> <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On a QFX3500, QFX3600, and QFX5100 switch, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.</p> <p>Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:</p> <ul style="list-style-type: none"> <li>• <b>For LAG interfaces</b>—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (<b>aex</b>) automatically configures the jumbo MTU size on the member links.</li> <li>• <b>For RVIs</b>—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named <b>vlan</b>. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the <b>vlan</b> interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named <b>vlan</b> (the RVI). On a QFX5100 switch, jumbo frames on the RVI are configured on the basis of the interface MTU.</li> </ul> |
|                                 | <div>  <p><b>CAUTION:</b> Setting or deleting the jumbo MTU size on the RVI (the <b>vlan</b> interface) while the switch is transmitting packets might result in dropped packets.</p> </div>   |
| <b>Options</b>                  | <p><b>bytes</b> —MTU size.</p> <p><b>Range:</b> 64 through 9216 bytes</p> <p><b>Default:</b> 1514 bytes</p>   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>  |

- Related Documentation**
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586](#)
  - *Junos OS Network Interfaces Library for Routing Devices*

---

## negotiation-options

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>negotiation-options {<br/>    allow-remote-loopback;<br/>    no-allow-link-events;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam ethernet link-fault-management interface</b> <i>interface-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | <p>Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | • <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a>   |

---

## no-adaptation (Liveness Detection)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>no-adaptation;</pre>  |
| <b>Hierarchy Level</b>          | [edit protocols <b>iccp peer liveness-detection</b> ]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.0 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 12.2 for the QFX Series.</p> |
| <b>Description</b>              | Configure Bidirectional Forwarding Detection (BFD) sessions to not adapt to changing network conditions.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                       |
| <b>Related Documentation</b>    |  |

## no-allow-link-events

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-allow-link-events;   |
| <b>Hierarchy Level</b>          | [edit protocols <b>oam</b> <b>ethernet link-fault-management interface</b> <i>interface-name</i> <b>negotiation-options</b> ]                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Disable the sending of link event TLVs.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul> |

## no-gratuitous-arp-request

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-gratuitous-arp-request;  |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> ],<br>[edit <b>interfaces interface-range</b> <i>interface-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs). |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IRB Interfaces on page 1675</a></li> </ul>   |

## oam

---

```
Syntax  oam {
    ethernet {
        connectivity-fault-management {
            action-profile profile-name {
                action {
                    interface-down;
                }
                default-actions {
                    interface-down;
                }
                event {
                    adjacency-loss;
                }
            }
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            mip-half-function (none | default | explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s | 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    remote-mep mep-id {
                        action-profile profile-name;
                    }
                }
            }
        }
    }
    performance-monitoring {
        sla-iterator-profiles {
            profile-name {
                calculation-weight {
                    delay delay-value;
                    delay-variation delay-variation-value;
                }
                cycle-time cycle-time-value;
                iteration-period iteration-period-value;
                measurement-type two-way-delay;
                passive;
            }
        }
    }
}
```

```

    }
  }
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
}
}

```

|                          |  |
|--------------------------|--|
| Hierarchy Level          | [edit protocols]   |
| Release Information      | Statement introduced in Junos OS Release 9.4 for EX Series switches.<br><b>connectivity-fault-management</b> introduced in Junos OS Release 10.2 for EX Series switches.   |
| Description              | Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.<br><br>The remaining statements are explained separately. |
| Required Privilege Level | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |

- Related Documentation**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)
  - *Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches*
  - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)
  - *Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)*

---

## on-disk-failure

---

- Syntax** `on-disk-failure {  
    disk-failure-action (halt | reboot);  
}`
- Hierarchy Level** `[edit chassis redundancy]  
[edit chassis routing-engine]`
- Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.
- Description** Instruct the router to halt or reboot if it detects hard disk errors on the Routing Engine.
- Options** The remaining statement is explained separately.
- Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.
- Related Documentation**
- [graceful-switchover on page 2309](#)
  - *Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors*
  - *Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)*
  - *High Availability Features for EX Series Switches Overview*



## on-loss-of-keepalives

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | on-loss-of-keepalives;   |
| <b>Hierarchy Level</b>          | [edit chassis <b>redundancy</b> failover]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Instruct the backup router to take mastership if it detects a loss of keepalive signal from the master Routing Engine.   |
| <b>Default</b>                  | <p>The <b>on-loss-of-keepalives</b> statement must be included at the [edit chassis <b>redundancy failover</b>] hierarchy level for failover to occur.</p> <p>When the <b>on-loss-of-keepalives</b> statement is included but graceful Routing Engine switchover <i>is not</i> configured, failover occurs after 300 seconds (5 minutes).</p> <p>When the <b>on-loss-of-keepalives</b> statement is included and graceful Routing Engine switchover <i>is</i> configured, the keepalive signal is automatically enabled and the failover time is set to 2 seconds.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">graceful-switchover on page 2309</a></li> <li>• <i>keepalive-time</i></li> <li>• <i>Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</i></li> <li>• <i>High Availability Features for EX Series Switches Overview</i></li> </ul>   |

## pdu-interval

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>pdu-interval <i>interval</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | Specify the periodic OAM PDU sending interval for fault detection. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.   |
| <b>Options</b>                  | <i>interval</i> —Periodic OAM PDU sending interval.<br><b>Range:</b> 400 through 1000 milliseconds<br><b>Default:</b> 1000 milliseconds  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583</a></li><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li></ul> |

## pdu-threshold

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>pdu-threshold <i>threshold-value</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> ]                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.   |
| <b>Description</b>              | Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces. |
| <b>Options</b>                  | <i>threshold-value</i> —Number of PDUs missed before declaring the peer lost.<br><b>Range:</b> 3 through 10 PDUs<br><b>Default:</b> 3 PDUs   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li></ul>                                      |

## peer (ICCP)

**Syntax**

```
peer ip-address {
    authentication-key string;
    backup-liveness-detection {
        backup-peer-ip ip-address;
    }
    liveness-detection {
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (1 | automatic);
    }
    local-ip-address ipv4-address;
    session-establishment-hold-time seconds;
}
```

**Hierarchy Level** [edit protocols [iccp](#)]

**Release Information** Statement introduced in Junos OS Release 10.0 for MX Series routers.  
Statement introduced in Junos OS Release 12.2 for the QFX Series.  
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

**Description** Configure the peers that host a multichassis link aggregation group (MC-LAG). You must configure Interchassis Control Protocol (ICCP) for both peers that host the MC-LAG.



**NOTE:** Backup liveness detection is not supported on MX Series routers.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## peer (Multichassis)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>peer ip-address {<br/>    interface interface-name;<br/>}</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">multi-chassis</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.   |
| <b>Description</b>              | Configure the IP address of the peer that is part of the interchassis link-protection link (ICL-PL). If the Interchassis Control Connection Protocol (ICCP) is up and the interchassis link (ICL) comes up, the peer configured as standby will bring up the MC-AE interfaces shared with the active peer specified by the <b>peer</b> statement. You must specify the physical interface of the peer. |
| <b>Options</b>                  | <code>interface interface-name</code> —Specify the logical interface name of the peer.   |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.  |

## periodic

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>periodic (fast   slow);</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aex aggregated-ether-options lacp</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure the interval for periodic transmission of LACP packets.  |
| <b>Default</b>                  | <code>fast</code>  |
| <b>Options</b>                  | <code>interval</code> —Interval at which to periodically transmit LACP packets: <ul style="list-style-type: none"><li>• <code>fast</code>—Receive packets every second. This is the default.</li><li>• <code>slow</code>—Receive packets every 30 seconds.</li></ul> |
| <b>Required Privilege Level</b> | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2393</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>  |

## pic

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> pic <i>pic-number</i> {     tunnel-port <i>port-number</i> tunnel-services;     port <i>port-number</i> {         channel-speed (<i>speed</i> disable-auto-speed-detection) ;     }     port-range <i>port-range-low</i> <i>port-range-high</i> {         channel-speed (<i>speed</i> disable-auto-speed-detection) ;     } } </pre>   |
| <b>Hierarchy Level</b>          | [edit chassis fpc slot]  |
| <b>Release Information</b>      | Option <b>channel-speed</b> introduced in Junos OS Release 13.2 for the QFX Series.  |
| <b>Description</b>              | (QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Configure a specific port or a range of ports to operate as 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.   |
| <b>Options</b>                  | <p><b>pic <i>pic-number</i></b>—(QFX3500 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify <b>1</b> to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.</p> <p>(QFX3600 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify <b>0</b> to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.</p> <p><b>port <i>physical-port-number</i></b>—Port number on which you want to configure the port type.</p> <p><b>port-range-low</b>—Lowest-numbered port in the range of ports.</p> <p><b>port-range-high</b>—Highest-numbered port in the range of ports.</p> <p><b>channel-speed (<i>speed</i>  disable-auto-speed-detection)</b> —Configure <i>10g</i> for 10-Gigabit Ethernet type ports, and configure <i>disable-auto-speed-detection</i> to disable auto-channelization.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Channelizing Interfaces on page 2608</a></li> </ul>   |

## preempt-cutover-timer

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>preempt-cutover-timer seconds;</code>  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For platforms with ELS:<br/>[edit switch-options <b>redundant-trunk-group</b> group name]</li><li>For platforms without ELS:<br/>[edit ethernet-switching-options <b>redundant-trunk-group</b> group name]</li></ul>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 43 for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| <b>Description</b>              | Change the length of time that a re-enabled primary link waits to take over from an active secondary link in a redundant trunk group.  |
| <b>Default</b>                  | If you do not change the time with the <b>preempt-cutover-timer</b> statement, a re-enabled primary link takes over from the active secondary link after 120 seconds.  |
| <b>Options</b>                  | <p><b>seconds</b>—Number of seconds that the primary link waits to take over from the active secondary link.</p> <p><b>Range:</b> 1 through 600 seconds</p>  |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>Example: Configuring Redundant Trunk Links for Faster Recovery</i></li><li><a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578</a></li><li><a href="#">Understanding Redundant Trunk Links on page 2447</a></li></ul>   |

## redundancy (Graceful Switchover)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> redundancy {   failover {     on-disk-failure;     on-loss-of-keepalives;   }   graceful-switchover; } </pre>  |
| <b>Hierarchy Level</b>          | [edit chassis]   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Enable redundant Routing Engines on a Virtual Chassis with two or more member switches or on a Virtual Chassis Fabric, on a standalone EX6200 or EX8200 switch with more than one Routing Engine.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Default</b>                  | Redundancy is enabled for the Routing Engines.   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">graceful-switchover on page 2309</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 2270</a></li> <li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2268</a></li> <li>• <a href="#">Installing Software on an EX Series Switch with Redundant Routing Engines (CLI Procedure)</a></li> <li>• <a href="#">High Availability Features for EX Series Switches Overview</a></li> </ul> |

## redundant-trunk-group

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>redundant-trunk-group {<br/>  group name {<br/>    interface interface-name &lt;primary&gt;;<br/>    interface interface-name;<br/>    preempt-cutover-timer seconds;<br/>  }<br/>}</pre>   |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>• For platforms with ELS:<br/>[edit switch-options]</li><li>• For platforms without ELS:<br/>[edit ethernet-switching-options]</li></ul>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10 (ELS). (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| <b>Description</b>              | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal spanning-tree protocol convergence.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Redundant Trunk Links for Faster Recovery</i></li><li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578</a></li><li>• <a href="#">Understanding Redundant Trunk Links on page 2447</a></li></ul>   |



## remote-loopback

---


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | remote-loopback;  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam</a> <a href="#">ethernet</a> <a href="#">link-fault-management</a> <a href="#">interface</a> <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.           |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul> |

## resilient-hash

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | resilient-hash;  |
| <b>Hierarchy Level</b>          | [edit interfaces <a href="#">aex</a> <a href="#">aggregated-ether-options</a> ]]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.   |
| <b>Description</b>              | Enable resilient hashing for a LAG, to minimize remapping of destination paths.  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Resilient Hashing for Trunk/ECMP Groups on page 2607</a></li> </ul> |

## rx-buffers

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | rx-buffers (on   off);  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces interface-name ether-options configured-flow-control</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | <p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
|                                 | <p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p>   |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.  |
| <b>Options</b>                  | <b>on   off</b> —Enable or disable an interface to generate and send Ethernet PAUSE messages.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2659</a></li> <li>• <a href="#">tx-buffers on page 2742</a></li> </ul>  |

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## routing-engine

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> routing-engine {   on-disk-failure {     disk-failure-action (halt   reboot);   } } </pre>   |
| <b>Hierarchy Level</b>          | <p>[edit chassis]<br/> [edit chassis interconnect-device <i>name</i>],<br/> [edit chassis node-group <i>name</i>]</p>  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure a Routing Engine to halt or reboot automatically when a hard disk error occurs. A hard disk error may cause a Routing Engine to enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. Rebooting or halting prevents this. |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.<br/> interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring the Junos OS to Enable a Routing Engine to Reboot on Hard Disk Errors</i></li> <li>• <i>Junos OS High Availability Library for Routing Devices</i></li> </ul>  |

## service-id

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>service-id number;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit switch-options]</code><br><code>[edit vlans <i>vlan-name</i>]</code>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>              | Specify a service identifier for each multichassis aggregated Ethernet interface that belongs to a link aggregation group (LAG).                                   |
| <b>Options</b>                  | <b>number</b> —A number that identifies a particular service.<br><b>Range:</b> 1 through 65535   |
| <b>Required Privilege Level</b> | <b>system</b> —To view this statement in the configuration.<br><b>system control</b> —To add this statement to the configuration.                                  |

## session-establishment-hold-time

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>session-establishment-hold-time <i>seconds</i>;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit protocols <a href="#">iccp peer</a>],</code><br><code>[edit protocols <a href="#">iccp</a>]</code>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches. |
| <b>Description</b>              | Specify the time during which an Interchassis Control Protocol (ICCP) connection must be established between peers.  |
| <b>Options</b>                  | <b>seconds</b> —Time (in seconds) within which a successful ICCP connection must be established.   |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.  |


## source

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>source <i>source-address</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>              | Specify the source address of the tunnel.   |
| <b>Default</b>                  | If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.   |
| <b>Options</b>                  | <b><i>source-address</i></b> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul>  |

## speed

---

|  |  |
|--|--|
| <b>Syntax</b>  | (speed 100m   1g);   |
| <b>Hierarchy Level</b>   | [edit <a href="#">interfaces</a> <i>interface-name</i> ]   |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>   | Configure the speed of the interface. On QFX5100 devices using 1-Gigabit Ethernet Copper SFP interfaces, you can configure the speed to be 100 Mbps. To return to the default speed of 1 Gbps, delete the <b>100m</b> . statement at the <b>[edit interfaces <i>interface-name</i> speed]</b> CLI hierarchy. |
| <div> <b>NOTE:</b> Autonegotiation is not supported on QFX5100 devices.</div> |  |
| <b>Default</b>   | The speed for 1-Gigabit Ethernet Copper SFP interfaces is set to 1 Gbps by default, but you can configure the speed to be 100 Mbps. The speed for 10-Gigabit Ethernet interfaces is set to 10 Gbps by default and cannot be configured to operate in a different speed.                                      |
| <b>Options</b>   | <ul style="list-style-type: none"><li>• <b>100m</b>—100 Mbps</li><li>• <b>1g</b>—1 Gbps</li></ul>  |
| <b>Required Privilege Level</b>  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">auto-negotiation on page 2632</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>                                       |

## status-control

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | status-control (active   standby);  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces aggregated-ether-options</a> <i>mc-ae</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.                              |
| <b>Description</b>              | Specify whether a peer hosting a multichassis link aggregation group (MC-LAG) is primary or secondary. Primary is considered active, and secondary is considered standby. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |

## symbol-period

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>symbol-period count;</code>   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ; <a href="#">event link-event-rate</a> ] ,<br>[edit protocols <a href="#">oam ethernet link-fault-management interface</a> <i>interface-name</i> <a href="#">event-thresholds</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Configure the threshold for sending symbol period events or taking the action specified in the action profile.<br><br>Symbol code errors occur on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period. You cannot configure the default value to a different value. |
| <b>Options</b>                  | <i>count</i> —Threshold count in seconds for symbol period events.<br><b>Range:</b> 1 through 100 seconds   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul>   |

## syslog (OAM LFM)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>syslog;</code>  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">oam ethernet link-fault-management action-profile</a> <i>profile-name</i> ; <a href="#">action</a> ]                |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.              |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul> |


## targeted-broadcast

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | targeted-broadcast;   |
| <b>Hierarchy Level</b>          | [edit <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> family inet],<br>[edit <b>interfaces</b> <i>interface-range</i> <i>interface-range-name</i> <b>unit</b> <i>logical-unit-number</i> family inet]                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Specify whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.                       |
| <b>Default</b>                  | When this statement is not included, broadcast packets are sent to the Routing Engine only.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IP Directed Broadcast on an EX Series Switch</i></li><li>• <i>Configuring IP Directed Broadcast (CLI Procedure)</i></li><li>• <i>Understanding IP Directed Broadcast for EX Series Switches</i></li></ul> |

## threshold (Detection Time)

---

|  |  |
|--|--|
| <b>Syntax</b>  | threshold <i>milliseconds</i> ;  |
| <b>Hierarchy Level</b>   | [edit protocols <b>iccp</b> <i>peer</i> <b>liveness-detection</b> <i>detection-time</i> ]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.  |
| <b>Description</b>   | Specify the threshold for the adaptation of the detection time for a Bidirectional Forwarding Detection (BFD) session. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent. |
| <div> <b>NOTE:</b> The threshold time must be greater than or equal to the <b>minimum-interval</b> or the <b>minimum-receive-interval</b> values.</div> |  |
| <b>Options</b>   | <b>milliseconds</b> — Value for the detection time adaptation threshold.<br><b>Range:</b> 1 through 255,000  |
| <b>Required Privilege Level</b>  | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |



## traceoptions (ICCP)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; }</pre>  |
| <b>Hierarchy Level</b>     | [edit <a href="#">protocols iccp</a> ]   |
| <b>Release Information</b> | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.  |
| <b>Description</b>         | Set Interchassis Control Protocol (ICCP) tracing options.  |
| <b>Default</b>             | Tracing operations are disabled.   |
| <b>Options</b>             | <p><b>disable</b>—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. By default, the log file is stored in <b>/var/log/</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 1 trace file only</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The following are the tracing options:</p> <ul style="list-style-type: none"> <li><b>all</b>—All tracing operations</li> <li><b>config-internal</b>—Trace configuration internals.</li> <li><b>general</b>—Trace general events.</li> <li><b>normal</b>—All normal events.</li> </ul> <p><b>Default:</b> If you do not specify this option, only unusual or abnormal operations are traced.</p> <ul style="list-style-type: none"> <li><b>parse</b>—Trace configuration parsing.</li> <li><b>policy</b>—Trace policy operations and actions.</li> </ul> |

- **route**—Trace routing table changes.
- **state**—Trace state transitions.
- **task**—Trace protocol task processing.
- **timer**—Trace protocol task timer processing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.        |
|                                 | routing-control—To add this statement to the configuration. |

---


## transmit-interval (Liveness Detection)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>transmit-interval {<br/>    minimum-interval <i>milliseconds</i>;<br/>    threshold <i>milliseconds</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit protocols <i>iccp</i> peer <i>liveness-detection</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.0 for MX Series routers.<br>Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.  |
| <b>Description</b>              | <p>Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |

## traceoptions (Individual Interfaces)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>traceoptions {<br/>    flag <i>flag</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The <b>traceoptions</b> statement for interfaces does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system <b>syslog</b> file in the directory <b>/var/log</b>.</p>   |
|                                 | <div> <b>NOTE:</b> The <b>traceoptions</b> statement is not supported on the QFX3000 QFabric system.</div>   |
| <b>Default</b>                  | If you do not include this statement, no interface-specific tracing operations are performed.   |
| <b>Options</b>                  | <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"><li>• <b>all</b>—All interface tracing operations</li><li>• <b>event</b>—Interface events</li><li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li><li>• <b>media</b>—Interface media changes</li><li>• <b>q921</b>—ISDN Q.921 frames</li><li>• <b>q931</b>—ISDN Q.931 frames</li></ul> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Tracing Operations of an Individual Router or Switch Interface</i></li></ul>   |

## traceoptions (OAM LFM)

**Syntax** traceoptions {  
     file *filename* <files *number*> <match *regex*> <size *size*> <world-readable |  
         no-world-readable>;  
     flag *flag* ;  
     no-remote-trace;  
 }

**Release Information** Statement introduced in JUNOS Release 10.2 for EX Series switches.

**Description** Configure tracing options the link fault management.

**Options** file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory */var/log*.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000

**Default:** 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **action-profile**—Trace action profile invocation events.
- **all**—Trace all events.
- **configuration**—Trace configuration events.
- **protocol**—Trace protocol processing events.
- **routing socket**—Trace routing socket events.

match—(Optional) Refine the output to log only those lines that match the given regular expression.

no-world-readable—(Optional) Restrict file access to the user who created the file.

no-remote-trace—(Optional) Disable the remote trace.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**world-readable**—(Optional) Enable unrestricted file access.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 2604](#)

---

## traps

---

**Syntax** (traps | no-traps);

**Hierarchy Level** [edit interfaces *interface-name*],  
[edit interfaces *interface-name* *unit* *logical-unit-number*],  
[edit interfaces *interface-range* *interface-range-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Enable or disable the sending of SNMP notifications when the state of the connection changes.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Enabling or Disabling SNMP Notifications on Physical Interfaces*
- *Enabling or Disabling SNMP Notifications on Logical Interfaces*

## tunnel

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>tunnel {   destination destination-address;   source source-address;   ttl number; }</pre>  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.   |
| <b>Description</b>              | <p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul>   |


## tunnel-port

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | tunnel-port <i>port-number</i> tunnel-services;  |
| <b>Hierarchy Level</b>          | [edit chassis fpc <i>slot</i> pic <i>pic-number</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series. |
| <b>Description</b>              | Configure the port number for generic routing encapsulation (GRE) tunneling.   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li> </ul>             |

## tx-buffers

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | tx-buffers (on   off);   |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
|                                 | <div> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</div> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p>  |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.   |
| <b>Options</b>                  | <b>on   off</b> —Enable or disable an interface to respond to an Ethernet PAUSE message.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">flow-control on page 2659</a></li><li>• <a href="#">rx-buffers on page 2728</a></li></ul>  |



- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## unit

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>unit <i>logical-unit-number</i> {<br/>    family {<br/>        ethernet-switching {<br/>            filter input <i>filter-name</i>;<br/>            filter output <i>filter-name</i>;<br/>            native-vlan-id <i>vlan-id</i>;<br/>            port-mode <i>mode</i>;<br/>            vlan {<br/>                members [ (all   <i>names</i>   <i>vlan-ids</i>) ];<br/>            }<br/>        }<br/>        fibre-channel {<br/>            port-mode (f-port   np-port);<br/>        }<br/>        inet {<br/>            address <i>address</i> {<br/>                primary;<br/>            }<br/>            filter input <i>filter-name</i>;<br/>            filter output <i>filter-name</i>;<br/>            primary;<br/>            targeted-broadcast;<br/>        }<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces <i>interface-name</i></a> ],<br>[edit <a href="#">interfaces <i>interface-range</i> <i>interface-range-name</i></a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.   |
| <b>Default</b>                  | You must configure a logical interface to be able to use the physical device.   |
| <b>Options</b>                  | <b><i>logical-unit-number</i></b> —Number of the logical unit.<br><b>Range:</b> 0 through 16,384<br><br>The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <a href="#">Configuring Link Aggregation on page 2593</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>  |

## uplink-failure-detection

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>uplink-failure-detection {   group group-name {     link-to-monitor interface-name;     link-to-disable interface-name;   } }</pre>   |
| <b>Hierarchy Level</b>          | [edit protocols]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failure information to the downlink interfaces.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Uplink Failure Detection on page 2392</a></li> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2592</a></li> <li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2457</a></li> </ul> |

## version (Liveness Detection)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | version (1   automatic);   |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">iccp peer liveness-detection</a> ]   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.0 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 12.2 for the QFX Series.</p> |
| <b>Description</b>              | Configure the Bidirectional Forwarding Detection (BFD) protocol version to detect.   |
| <b>Options</b>                  | <p>1—Use BFD protocol version 1.</p> <p><b>automatic</b>—Autodetect the BFD protocol version.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                       |

## vlan-id

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>vlan-id <i>vlan-id-number</i>;</code>   |
| <b>Hierarchy Level</b>     | [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]            |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>         | For 10-Gigabit Ethernet and aggregated Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface. |



**NOTE:** The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

|                                 |   |
|---------------------------------|---|
| <b>Options</b>                  | <i>vlan-id-number</i> —Valid VLAN identifier.<br><b>Range:</b> 1 through 4094   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-tagging on page 1813</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2593</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul> |

## vlan-tagging

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>vlan-tagging;</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ]<br>[edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-range-name</i> ]            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.  |
| <b>Default</b>                  | VLAN tagging is disabled by default.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">vlan-id on page 2746</a></li><li>• <a href="#">Configuring a Layer 3 Logical Interface on page 2593</a></li></ul> |

# Administration

- [Routine Monitoring on page 2747](#)
- [Monitoring Commands on page 2753](#)

## Routine Monitoring

---

- [Monitoring System Process Information on page 2747](#)
- [Monitoring System Properties on page 2748](#)
- [Monitoring Interface Status and Traffic on page 2749](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 2750](#)
- [Verifying the Status of a LAG Interface on page 2750](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2751](#)
- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 2752](#)

## Monitoring System Process Information

**Purpose** View the processes running on the device.

**Action** To view the software processes running on the device:  
[edit system]

user@switch> [show system processes](#)

**Meaning** [Table 22 on page 333](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

**Table 245: Summary of System Process Information Output Fields**

| Field | Values                     |
|-------|----------------------------|
| PID   | Identifier of the process. |
| Name  | Owner of the process.      |

Table 245: Summary of System Process Information Output Fields (*continued*)

| Field              | Values   |
|--------------------|--|
| State              | Current state of the process.                            |
| CPU Load           | Percentage of the CPU that is being used by the process. |
| Memory Utilization | Amount of memory that is being used by the process.      |
| Start Time         | Time of day when the process started.                    |

- Related Documentation**
- [Monitoring System Properties on page 334](#)
  - [show system uptime on page 1137](#)

## Monitoring System Properties

**Purpose** View system properties such as the name, IP address, and resource usage.

**Action** To monitor system properties in the CLI, enter the following commands:

- [show system uptime](#)
- [show system users](#)
- [show system storage](#)

**Meaning** [Table 23 on page 334](#) summarizes key output fields in the system properties display.

Table 246: Summary of Key System Properties Output Fields

| Field                      | Values  | Additional Information                                 |
|----------------------------|---|--|
| <b>General Information</b> |   |  |
| Serial Number              | Serial number of device.  |  |
| Junos OS Version           | Version of Junos OS active on the switch, including whether the software is for domestic or export use. | Export software is for use outside the USA and Canada. |
| Hostname                   | Name of the device.   |  |
| IP Address                 | IP address of the device.   |  |
| Loopback Address           | Loopback address.   |  |
| Domain Name Server         | Address of the domain name server.  |  |
| Time Zone                  | Time zone on the device.  |  |

Table 246: Summary of Key System Properties Output Fields (*continued*)

| Field                          | Values   | Additional Information   |
|--------------------------------|--|--|
| <b>Time</b>                    |  |  |
| Current Time                   | Current system time, in Coordinated Universal Time (UTC).  |  |
| System Booted Time             | Date and time when the device was last booted and how long it has been running.  |  |
| Protocol Started Time          | Date and time when the protocols were last started and how long they have been running.  |  |
| Last Configured Time           | Date and time when a configuration was last committed. This field also shows the name of the user who issued the last <b>commit</b> command. |  |
| Load Average                   | CPU load average for 1, 5, and 15 minutes.   |  |
| <b>Storage Media</b>           |  |  |
| Internal Flash Memory          | Usage details of internal flash memory.  |  |
| External Flash Memory          | Usage details of external USB flash memory.  |  |
| <b>Logged in Users Details</b> |  |  |
| User                           | Username of any user logged in to the switch.  |  |
| Terminal                       | Terminal through which the user is logged in.  |  |
| From                           | System from which the user has logged in. A hyphen indicates that the user is logged in through the console.                                 |  |
| Login Time                     | Time when the user logged in.  | This is the <b>user@switch</b> field in <b>show system users</b> command output. |
| Idle Time                      | How long the user has been idle.   |  |

- Related Documentation**
- [Monitoring System Process Information on page 333](#)
  - [show system processes on page 1051](#)

## Monitoring Interface Status and Traffic

**Purpose** View interface status to monitor interface bandwidth utilization and traffic statistics.

**Action** • To view interface status for all the interfaces, enter [show interfaces xe](#).

- To view status and statistics for a specific interface, enter **show interfaces xe interface-name**.
- To view status and traffic statistics for all interfaces, enter either **show interfaces xe detail** or **show interfaces xe extensive**.

**Meaning** For details about output from the CLI commands, see [show interfaces xe](#).

## Verifying That Layer 3 Logical Interfaces Are Working

**Purpose** After configuring Layer 3 logical interfaces, verify that they are set up properly and transmitting data.

- Action** 1. To determine if you have successfully created the logical interfaces and the links are up:

```
[edit interfaces]
user@switch> show interfaces interface-name terse
```

| Interface      | Admin | Link | Proto | Local      | Remote |
|----------------|-------|------|-------|------------|--------|
| ge-0/0/0       | up    | up   |       |            |        |
| ge-0/0/0.0     | up    | up   | inet  | 1.1.1.1/24 |        |
| ge-0/0/0.1     | up    | up   | inet  | 2.1.1.1/24 |        |
| ge-0/0/0.2     | up    | up   | inet  | 3.1.1.1/24 |        |
| ge-0/0/0.3     | up    | up   | inet  | 4.1.1.1/24 |        |
| ge-0/0/0.4     | up    | up   | inet  | 5.1.1.1/24 |        |
| ge-0/0/0.32767 | up    | up   |       |            |        |

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine if packets were transmitted correctly on the logical interface VLANs:

```
user@switch> ping ip-address
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

**Meaning** The output confirms that the logical interfaces have been created and the links are up.

**Related Documentation**

- [Configuring a Layer 3 Logical Interface on page 2593](#)

## Verifying the Status of a LAG Interface

**Purpose** Verify that a link aggregation group (LAG) (**ae0**) has been created on the switch.

- Action** To verify that the **ae0** LAG has been created:

```
[edit interfaces]
show interfaces ae0 terse
```

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|-------|-------|--------|
|-----------|-------|------|-------|-------|--------|



```

ae0                up      up
ae0.0              up      up      inet      10.10.10.2/24

```

**Meaning** The output confirms that the **ae0** link is up and shows the family and IP address assigned to this link.

- Related Documentation**
- [Configuring Link Aggregation on page 2593](#)
  - [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2751](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466](#)
  - [show lacp statistics interfaces \(View\) on page 2875](#)

## Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. [Verifying the LACP Setup on page 2751](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 2751](#)

### Verifying the LACP Setup

**Purpose** Verify that the LACP has been set up correctly.

**Action** To verify that LACP has been enabled as active on one end:

```

user@switch>show lacp interfaces xe-0/0/0
Aggregated interface: ae0
LACP state:
xe-0/1/0      Actor No Yes No No No Yes Fast Active
xe-0/1/0      PartnerNo Yes No No No Yes Fast Passive
LACP protocol: Receive State Transmit State Mux State
xe-0/1/0      Defaulted Fast periodic Detached

```

**Meaning** This example shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, one side must be set as active in order for the bundled link to be up.

### Verifying That LACP Packets Are Being Exchanged

**Purpose** Verify that LACP packets are being exchanged between interfaces.

**Action** Use the `show lacp statistics interfaces interface-name` command to display LACP BPDU exchange information.

```
show lacp statistics interfaces ae0
```

```
Aggregated interface: ae0
```

| LACP Statistics: | LACP Rx | LACP Tx | Unknown Rx | Illegal Rx |
|------------------|---------|---------|------------|------------|
| xe-0/0/2         | 1352    | 2035    | 0          | 0          |
| xe-0/0/3         | 1352    | 2056    | 0          | 0          |

**Meaning** The output here shows that the link is up and that PDUs are being exchanged.

- Related Documentation**
- [Configuring Link Aggregation on page 2593](#)
  - [Verifying the Status of a LAG Interface on page 2750](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466](#)
  - [show lacp statistics interfaces \(View\) on page 2875](#)

## Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

**Purpose** Verify that the generic routing encapsulation (GRE) interface is sending tunneled traffic.

**Action** Display status information about the specified GRE interface by using the command `show interfaces`.

```
user@switch> show interfaces gr-0/0/0.0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
Interface index: 132, SNMP ifIndex: 26
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
Flags: Point-To-Point SNMP-Traps 16384
IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
Input packets : 0
Output packets: 0
Protocol inet, MTU: 1476
Flags: None
Addresses, Flags: Is-Primary
Local: 1.10.1.1
```

**Meaning** The output indicates that the GRE interface gr-0/0/0 is up. The output displays the name of the physical interface and the traffic statistics for this interface---the number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.

**Related Documentation** • *Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)*

## Monitoring Commands

---

- [monitor interface](#)
- [show forwarding-options enhanced-hash-key](#)
- [show iccp](#)
- [show interfaces diagnostics optics](#)
- [show interfaces ge](#)
- [show interfaces \(GRE\)](#)
- [show interfaces irb](#)
- [show interfaces mc-ae](#)
- [show interfaces queue](#)
- [show interfaces xe](#)
- [show lacp interfaces](#)
- [show lacp statistics interfaces \(View\)](#)
- [show oam ethernet link-fault-management](#)
- [show redundant-trunk-group](#)
- [show uplink-failure-detection](#)

## monitor interface

**Syntax**    `monitor interface`  
               `<interface-name> | traffic <detail>>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                               Command introduced in Junos OS Release 9.0 for EX Series switches.  
                               Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.



**NOTE:** This command is not supported on the QFX3000 QFabric system.

**Options**    **none**—Display real-time statistics for all interfaces.

**detail**—(Optional) With traffic option only, display detailed output.

**interface-name**—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

**traffic**—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

**Additional Information**    The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the **c** key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type in the [CLI Explorer](#). To control the output of the **monitor interface** command while it is running, use the keys listed in [Table 247 on page 2754](#). The keys are not case-sensitive.

**Table 247: Output Control Keys for the monitor interface Command**

| Key | Action   |
|-----|--|
| c   | Clears (returns to zero) the delta counters since <b>monitor interface</b> was started. This does not clear the accumulative counter. To clear the accumulative counter, use the <b>clear interfaces interval</b> command. |
| f   | Freezes the display, halting the display of updated statistics and delta counters.   |
| i   | Displays information about a different interface. The command prompts you for the name of a specific interface.  |

**Table 247: Output Control Keys for the monitor interface Command** *(continued)*

| Key      | Action   |
|----------|--|
| n        | Displays information about the next interface. The <b>monitor interface</b> command displays the physical or logical interfaces in the same order as the <b>show interfaces terse</b> command. |
| q or Esc | Quits the command and returns to the command prompt.   |
| t        | Thaws the display, resuming the update of the statistics and delta counters.   |

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 248 on page 2755](#). The keys are not case-sensitive.

**Table 248: Output Control Keys for the monitor interface traffic Command**

| Key      | Action   |
|----------|--|
| b        | Displays the statistics in units of bits and bits per second (bps).  |
| c        | Clears (return to 0) the delta counters in the <b>Current Delta</b> column. The statistics counters are not cleared. |
| d        | Displays the <b>Current Delta</b> column (instead of the rate column) in Bps or packets per second (pps).            |
| p        | Displays the statistics in units of packets and packets per second (pps).  |
| q or Esc | Quits the command and returns to the command prompt.   |
| r        | Displays the rate column (instead of the <b>Current Delta</b> column) in Bps and pps.                                |

**Required Privilege Level** trace

**List of Sample Output** [monitor interface \(Physical\) on page 2757](#)  
[monitor interface \(OTN Interface\) on page 2758](#)  
[monitor interface \(MX2020 Routers with MPC6E and OTN MICInterface\) on page 2759](#)  
[monitor interface \(Logical\) on page 2760](#)  
[monitor interface \(QFX3500 Switch\) on page 2760](#)  
[monitor interface traffic on page 2761](#)  
[monitor interface traffic \(QFX3500 Switch\) on page 2761](#)  
[monitor interface traffic detail \(QFX3500 Switch\) on page 2762](#)

**Output Fields** [Table 249 on page 2756](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 249: monitor interface Output Fields

| Field Name               | Field Description  | Level of Output |
|--------------------------|--|-----------------|
| <b>routerl</b>           | Hostname of the router.  | All levels      |
| <b>Seconds</b>           | How long the monitor interface command has been running or how long since you last cleared the counters.   | All levels      |
| <b>Time</b>              | Current time (UTC).  | All levels      |
| <b>Delay x/y/z</b>       | Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> <li>• <b>x</b>—Time taken for the last polling (in milliseconds).</li> <li>• <b>y</b>—Minimum time taken across all pollings (in milliseconds).</li> <li>• <b>z</b>—Maximum time taken across all pollings (in milliseconds).</li> </ul>  | All levels      |
| <b>Interface</b>         | Short description of the interface, including its name, status, and encapsulation.   | All levels      |
| <b>Link</b>              | State of the link: <b>Up</b> , <b>Down</b> , or <b>Test</b> .  | All levels      |
| <b>Current delta</b>     | Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.   | All levels      |
| <b>Local Statistics</b>  | (Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| <b>Remote Statistics</b> | (Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>  | All levels      |

Table 249: monitor interface Output Fields (*continued*)

| Field Name         | Field Description  | Level of Output |
|--------------------|--|-----------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | All levels      |
| Description        | With the <b>traffic</b> option, displays the interface description configured at the <b>[edit interfaces <i>interface-name</i>]</b> hierarchy level.   | detail          |

## Sample Output

### monitor interface (Physical)

```

user@host> monitor interface so-0/0/0
router1                               Seconds: 19                      Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
    Input packets:                6045 (0 pps)
    Input bytes:                  6290065 (0 bps)
    Output packets:               10376 (0 pps)
    Output bytes:                 10365540 (0 bps)
Encapsulation statistics:
    Input keepalives:             1901
    Output keepalives:            1901
    NCP state: Opened
    LCP state: Opened
Error statistics:
    Input errors:                 0
    Input drops:                 0
    Input framing errors:         0
    Policed discards:            0
    L3 incompletes:              0
    L2 channel errors:           0
    L2 mismatch timeouts:        0
    Carrier transitions:          1
    Output errors:               0
    Output drops:                0
    Aged packets:                0
Active alarms : None
Active defects: None
SONET error counts/seconds:
    LOS count                    1
    LOF count                    1
    SEF count                    1
    ES-S                         0
    SES-S                        0
SONET statistics:
    BIP-B1                      458871

```

```

BIP-B2                      460072          [0]
REI-L                      465610          [0]
BIP-B3                      458978          [0]
REI-P                      458773          [0]

```

## Received SONET overhead:

```

F1      : 0x00 J0      : 0x00 K1      : 0x00
K2      : 0x00 S1      : 0x00 C2      : 0x00
C2(cmp) : 0x00 F2      : 0x00 Z3      : 0x00
Z4      : 0x00 S1(cmp) : 0x00

```

## Transmitted SONET overhead:

```

F1      : 0x00 J0      : 0x01 K1      : 0x00
K2      : 0x00 S1      : 0x00 C2      : 0xcf
F2      : 0x00 Z3      : 0x00 Z4      : 0x00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

## monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0
```

```
Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
```

## Traffic statistics:

```

Input bytes:                0 (0 bps)
Output bytes:               0 (0 bps)
Input packets:              0 (0 pps)
Output packets:             0 (0 pps)

```

## Error statistics:

```

Input errors:               0
Input drops:                0
Input framing errors:       0
Policed discards:          0
L3 incompletes:             0
L2 channel errors:          0
L2 mismatch timeouts:       0
Carrier transitions:         5
Output errors:              0
Output drops:               0
Aged packets:               0

```

Active alarms : None

Active defects: None

## Input MAC/Filter statistics:

```

Unicast packets            0
Broadcast packets          0
Multicast packets          0
Oversized frames           0
Packet reject count        0
DA rejects                 0
SA rejects                 0

```

## Output MAC/Filter Statistics:

```

Unicast packets            0
Broadcast packets          0
Multicast packets          0
Packet pad count           0
Packet error count         0

```

## OTN Link 0

```

OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
OTN OC - Seconds
LOS                        2

```



```

      LOF                                     9
OTN OTU - FEC Statistics
  Corr err ratio                             N/A
  Corr bytes                                 0
  Uncorr words                               0
OTN OTU - Counters
  BIP                                         0
  BBE                                         0
  ES                                          0
  SES                                         0
  UAS                                         422
OTN ODU - Counters
  BIP                                         0
  BBE                                         0
  ES                                          0
  SES                                         0
  UAS                                         422
OTN ODU - Received Overhead    APSPPC 0-3:      0

```

### monitor interface (MX2020 Routers with MPC6E and OTN MICInterface)

```

user@host> monitor interface xe-3/0/0
host name                               Seconds: 67
Time: 23:46:46
Delay: 0/0/13

Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                             0 (0 bps)
  Output bytes:                             0 (0 bps)
  Input packets:                           0 (0 pps)
  Output packets:                           0 (0 pps)
Error statistics:
  Input errors:                             0
  Input drops:                              0
  Input framing errors:                     0
  Policed discards:                         0
  L3 incompletes:                           0
  L2 channel errors:                        0
  L2 mismatch timeouts:                     0
  Carrier transitions:                       3
  Output errors:                             0
  Output drops:                             0
  Aged packets:                             0
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
  LOS                                       0
  LOF                                       0
OTN OTU - FEC Statistics
  Corr err ratio                             N/A
  Corr bytes                                 0
  Uncorr words                               0
OTN OTU - Counters
  BIP                                         0
  BBE                                         0
  ES                                          0
  SES                                         0
  UAS                                         0
OTN ODU - Counters
  BIP                                         0

```

```

BBE                                0                                [0]
ES                                0                                [0]
SES                                0                                [0]
UAS                                0                                [0]
OTN ODU - Received Overhead       [0]
APSPCC 0-3:                       00 00 00 00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (Logical)

```

user@host> monitor interface so-1/0/0.0
host name                Seconds: 16                Time: 15:33:39
                                                                Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:
Input bytes:              0                                [0]
Output bytes:             0                                [0]
Input packets:            0                                [0]
Output packets:           0                                [0]
Remote statistics:
Input bytes:              0 (0 bps)                       [0]
Output bytes:             0 (0 bps)                       [0]
Input packets:            0 (0 pps)                       [0]
Output packets:           0 (0 pps)                       [0]
Traffic statistics:
Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

### monitor interface (QFX3500 Switch)

```

user@switch> monitor interface ge-0/0/0
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:
Input bytes:              0 (0 bps)                       [0]
Output bytes:             0 (0 bps)                       [0]
Input packets:            0 (0 pps)                       [0]
Output packets:           0 (0 pps)                       [0]
Error statistics:
Input errors:             0                                [0]
Input drops:              0                                [0]
Input framing errors:     0                                [0]
Policed discards:        0                                [0]
L3 incompletes:           0                                [0]
L2 channel errors:       0                                [0]
L2 mismatch timeouts:    0                                [0]
Carrier transitions:      0                                [0]
Output errors:            0                                [0]
Output drops:             0                                [0]
Aged packets:             0                                [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
Unicast packets           0                                [0]
Broadcast packets         0 Multicast packet             [0]

```

Interface warnings:  
 o Outstanding LINK alarm

### monitor interface traffic

```
user@host> monitor interface traffic
host name                               Seconds: 15                               Time: 12:31:09
```

| Interface | Link | Input packets | (pps) | Output packets | (pps) |
|-----------|------|---------------|-------|----------------|-------|
| so-1/0/0  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/0  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/1  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/2  | Down | 0             | (0)   | 0              | (0)   |
| so-1/1/3  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/0  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/1  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/2  | Down | 0             | (0)   | 0              | (0)   |
| t3-1/2/3  | Down | 0             | (0)   | 0              | (0)   |
| so-2/0/0  | Up   | 211035        | (1)   | 36778          | (0)   |
| so-2/0/1  | Up   | 192753        | (1)   | 36782          | (0)   |
| so-2/0/2  | Up   | 211020        | (1)   | 36779          | (0)   |
| so-2/0/3  | Up   | 211029        | (1)   | 36776          | (0)   |
| so-2/1/0  | Up   | 189378        | (1)   | 36349          | (0)   |
| so-2/1/1  | Down | 0             | (0)   | 18747          | (0)   |
| so-2/1/2  | Down | 0             | (0)   | 16078          | (0)   |
| so-2/1/3  | Up   | 0             | (0)   | 80338          | (0)   |
| at-2/3/0  | Up   | 0             | (0)   | 0              | (0)   |
| at-2/3/1  | Down | 0             | (0)   | 0              | (0)   |

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

### monitor interface traffic (QFX3500 Switch)

```
user@switch> monitor interface traffic
switch                               Seconds: 7                               Time: 16:04:37
```

| Interface | Link | Input packets | (pps) | Output packets | (pps) |
|-----------|------|---------------|-------|----------------|-------|
| ge-0/0/0  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/1  | Up   | 392187        | (0)   | 392170         | (0)   |
| ge-0/0/2  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/3  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/4  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/5  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/6  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/7  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/8  | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/9  | Up   | 392184        | (0)   | 392171         | (0)   |
| ge-0/0/10 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/11 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/12 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/13 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/14 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/15 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/16 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/17 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/18 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/19 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/20 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/21 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/22 | Up   | 392172        | (0)   | 392187         | (0)   |
| ge-0/0/23 | Up   | 392185        | (0)   | 392173         | (0)   |

|       |      |   |     |         |     |
|-------|------|---|-----|---------|-----|
| vcp-0 | Down | 0 |     | 0       |     |
| vcp-1 | Down | 0 |     | 0       |     |
| ae0   | Down | 0 | (0) | 0       | (0) |
| bme0  | Up   | 0 |     | 1568706 |     |

### monitor interface traffic detail (QFX3500 Switch)

user@switch> monitor interface traffic detail  
switch

Seconds: 74

Time: 16:03:02

| Interface<br>Description | Link | Input packets | (pps) | Output packets | (pps) |
|--------------------------|------|---------------|-------|----------------|-------|
| ge-0/0/0                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/1                 | Up   | 392183        | (0)   | 392166         | (0)   |
| ge-0/0/2                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/3                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/4                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/5                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/6                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/7                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/8                 | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/9                 | Up   | 392181        | (0)   | 392168         | (0)   |
| ge-0/0/10                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/11                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/12                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/13                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/14                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/15                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/16                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/17                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/18                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/19                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/20                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/21                | Down | 0             | (0)   | 0              | (0)   |
| ge-0/0/22                | Up   | 392169        | (0)   | 392184         | (1)   |
| ge-0/0/23                | Up   | 392182        | (0)   | 392170         | (0)   |
| vcp-0                    | Down | 0             |       | 0              |       |
| vcp-1                    | Down | 0             |       | 0              |       |
| ae0                      | Down | 0             | (0)   | 0              | (0)   |
| bme0                     | Up   | 0             |       | 1568693        |       |

## show forwarding-options enhanced-hash-key

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show forwarding-options enhanced-hash-key</b>   |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p><b>Fabric Load Balancing Options</b> output fields introduced in Junos OS Release 14.1X53-D10.</p>  |
| <b>Description</b>              | <p>Display information about which packet fields are used by the hashing algorithm to make hashing decisions.</p> <p>You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the <b>forwarding-options enhanced-hash-key</b> statement.</p>   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2590</a></li> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396</a></li> <li>• <a href="#">enhanced-hash-key on page 2644</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 2764</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 2765</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 2765</a></p>   |
| <b>Output Fields</b>            | <p><a href="#">Table 250 on page 2763</a> lists the output fields for the <b>show forwarding-options enhanced-hash-key</b> command. Output fields are listed in the approximate order in which they first appear.</p>  |

**Table 250: show forwarding-options enhanced-hash-key Output Fields**

| Field Name                   | Field Description  |
|------------------------------|--|
| <b>Hash-Mode</b>             | Current hash mode: Layer 2 header or Layer 2 payload.  |
| <b>Protocol</b>              | Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.              |
| <b>Destination L4 Port</b>   | Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.   |
| <b>Source L4 Port</b>        | Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.        |
| <b>Destination IPv4 Addr</b> | Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No. |

Table 250: show forwarding-options enhanced-hash-key Output Fields (*continued*)

| Field Name                     | Field Description  |
|--------------------------------|--|
| <b>Source IPv4 Addr</b>        | Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.  |
| <b>Vlan id</b>                 | Indicates whether the Vlan id field is or is not used by the hashing algorithm: Yes or No.   |
| <b>Next Hdr</b>                | Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.  |
| <b>Destination IPv6 Addr</b>   | Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.   |
| <b>Source IPv6 Addr</b>        | Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.  |
| <b>Ether Type</b>              | Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.  |
| <b>Destination MAC Address</b> | Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.   |
| <b>Source MAC Address</b>      | Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.  |
| <b>Load Balancing Method</b>   | Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet.<br><br>The load balancing method is flowlet by default, and can be configured using the <a href="#">fabric-load-balance</a> statement. |
| <b>Fabric Link Scale</b>       | Indicates the fabric link scale, in mbps.  |
| <b>Inactivity Interval</b>     | Indicates the fabric load balance inactivity interval, in microseconds (us).<br><br>The inactivity interval is 16 microseconds by default, and can be configured using the <a href="#">inactivity-interval</a> statement.            |
| <b>Hash Region Size/Trunk</b>  | Indicates the hash region size, in buckets per fabric trunk.   |

## Sample Output

### show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
-----
```

```
Hash-Mode                               :layer2-payload
```

```
inet Hash settings-
```

```
-----
```

```
inet packet fields
```

```
Protocol                               : Yes
Destination L4 Port                    : Yes
Source L4 Port                         : Yes
Destination IPv4 Addr                  : Yes
Source IPv4 Addr                       : Yes
Vlan id                               : No
```

```
inet6 Hash settings-
```

```
-----
```

```
inet6 packet fields
```

```
Next Hdr                              : Yes
Destination L4 Port                    : Yes
Source L4 Port                         : Yes
Destination IPv6 Addr                  : Yes
Source IPv6 Addr                       : Yes
Vlan id                               : No
```

#### show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
```

```
-----
```

```
Hash-Mode                               : layer2-header
```

```
layer2 Hash settings-
```

```
-----
```

```
layer2 packet fields
```

```
Ether Type                            : Yes
Destination MAC Address                : Yes
Source MAC Address                     : Yes
VLAN ID                               : No
```

#### show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>
```

```
Fabric Load Balancing Options
```

```
-----
```

```
Load Balancing Method : Flowlet
Fabric Link Scale      : 40960 (mbps)
Inactivity Interval    : 16 (us)
Hash Region Size/Trunk : 1024 (buckets)
```

## show iccp

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show iccp</b> <brief   detail><br><b>logical-system</b> [ <i>system-name</i>   all]   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for the MX Series.<br>Support for logical systems added in Junos OS Release 14.1.  |
| <b>Description</b>              | Display Interchassis Control Protocol (ICCP) information about the multichassis link aggregation group (MC-LAG) peers, including the state of the TCP connection, Bidirectional Forwarding Detection protocol, backup liveness peer status, and MCSNOOPD, LACPD, and ESWD applications.  |
| <b>Options</b>                  | <p><b>logical-system</b> [<i>system-name</i>   all]—(Optional) Display information for a specified logical system or all systems.</p> <p><b>none</b>—Display ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.</p> <p><b>brief</b>—Display brief ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.</p> <p><b>detail</b>—Display detailed ICCP information about the MC-LAG peers, including the state of the TCP connection and Bidirectional Forwarding Detection protocol, and MCSNOOPD, LACPD, and ESWD applications.</p> |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">iccp on page 2670</a></li> <li>• <a href="#">Understanding Multichassis Link Aggregation on page 2411</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show iccp on page 2767</a>   |
| <b>Output Fields</b>            | <a href="#">Table 251 on page 2766</a> lists the output fields for the <b>show iccp</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 251: show iccp

| Field Name                            | Field Description   |
|---------------------------------------|---|
| Redundancy Group Information for peer | Aggregated Ethernet interface name.   |
| TCP Connection                        | Specifies if the TCP connection between the peers hosting the MC-LAG is up or down.                     |
| Liveness Detection                    | Specifies if liveness detection, also known as Bidirectional Forwarding Detection (BFD), is up or down. |



Table 251: show iccp (*continued*)

| Field Name                     | Field Description   |
|--------------------------------|---|
| Client Application             | Specifies information regarding the state of the MCSNOOPD and ESWD client applications.                             |
| Redundancy Group IDs<br>Joined | Denotes the redundancy group unique identifier that is associated for the particular client application or process. |

## Sample Output

### show iccp

```
user@switch> show iccp
Redundancy Group Information for peer 3.3.3.2
  TCP Connection      : Established
  Liveliness Detection : Up

Client Application: MCSNOOPD

Client Application: eswd
```

## show interfaces diagnostics optics

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>show interfaces diagnostics optics <i>interface-name</i></code>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  |
| <b>Description</b>              | <p>Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series or QFX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information. For a list of transceivers supported on EX Series switches and their specifications, including DOM support, see <i>Pluggable Transceivers Supported on EX Series Switches</i>.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p> |
| <b>Options</b>                  | <i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed: <i>ge-fpc/pic/port</i> , <i>xe-fpc/pic/port</i> , or <i>et-fpc/pic/port</i> .   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Monitoring Interface Status and Traffic</i></li> <li>• <a href="#">Monitoring Interface Status and Traffic on page 335</a></li> <li>• <i>Installing a Transceiver in an EX Series Switch</i></li> <li>• <i>Installing a Transceiver in a QFX Series Device</i></li> <li>• <i>Removing a Transceiver from an EX Series Switch</i></li> <li>• <i>Removing a Transceiver from a QFX Series Device</i></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul>  |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 2775</a></p> <p><a href="#">show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 2776</a></p> <p><a href="#">show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 2777</a></p> <p><a href="#">show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver) on page 2778</a></p> <p><a href="#">show interfaces diagnostics optics et-4/1/0 (CFP Transceiver) on page 2779</a></p>   |
| <b>Output Fields</b>            | <a href="#">Table 49 on page 942</a> lists the output fields for the <b>show interfaces diagnostics optics</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 252: show interfaces diagnostics optics Output Fields

| Field Name         | Field Description                            |
|--------------------|--|
| Physical interface | Displays the name of the physical interface. |

Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name   | Field Description  |
|--|--|
| <b>Laser bias current</b>  | Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents. |
| <b>Laser output power</b><br>(Not available for QSFP+ transceivers)                                  | Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).   |
| <b>Laser temperature</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)               | Displays the laser temperature, in Celsius and Fahrenheit.   |
| <b>Module temperature</b>  | Displays the temperature, in Celsius and Fahrenheit.   |
| <b>Module voltage</b><br>(Not available for XFP transceivers)  | Displays the voltage, in Volts.  |
| <b>Laser rx power</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)                  | Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).   |
| <b>Receiver signal average optical power</b><br>(Not available for XFP, QSFP+, and CFP transceivers) | Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).  |
| <b>Laser bias current high alarm</b>   | Displays whether the laser bias power setting high alarm is <b>On</b> or <b>Off</b> .  |
| <b>Laser bias current low alarm</b>  | Displays whether the laser bias power setting low alarm is <b>On</b> or <b>Off</b> .   |
| <b>Laser bias current high warning</b>   | Displays whether the laser bias power setting high warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser bias current low warning</b>  | Displays whether the laser bias power setting low warning is <b>On</b> or <b>Off</b> .   |
| <b>Laser output power high alarm</b><br>(Not available for QSFP+ transceivers)                       | Displays whether the laser output power high alarm is <b>On</b> or <b>Off</b> .  |
| <b>Laser output power low alarm</b><br>(Not available for QSFP+ transceivers)                        | Displays whether the laser output power low alarm is <b>On</b> or <b>Off</b> .   |
| <b>Laser output power high warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the laser output power high warning is <b>On</b> or <b>Off</b> .  |

Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name  | Field Description   |
|---|---|
| <b>Laser output power low warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the laser output power low warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser temperature high alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays whether the laser temperature high alarm is <b>On</b> or <b>Off</b> .    |
| <b>Laser temperature low alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays whether the laser temperature low alarm is <b>On</b> or <b>Off</b> .     |
| <b>Laser temperature high warning</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the laser temperature high warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser temperature low warning</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays whether the laser temperature low warning is <b>On</b> or <b>Off</b> .   |
| <b>Module temperature high alarm</b><br>(Not available for QSFP+ transceivers)                      | Displays whether the module temperature high alarm is <b>On</b> or <b>Off</b> .   |
| <b>Module temperature low alarm</b><br>(Not available for QSFP+ transceivers)                       | Displays whether the module temperature low alarm is <b>On</b> or <b>Off</b> .    |
| <b>Module temperature high warning</b><br>(Not available for QSFP+ transceivers)                    | Displays whether the module temperature high warning is <b>On</b> or <b>Off</b> . |
| <b>Module temperature low warning</b><br>(Not available for QSFP+ transceivers)                     | Displays whether the module temperature low warning is <b>On</b> or <b>Off</b> .  |
| <b>Module voltage high alarm</b><br>(Not available for XFP and QSFP+ transceivers)                  | Displays whether the module voltage high alarm is <b>On</b> or <b>Off</b> .       |
| <b>Module voltage low alarm</b><br>(Not available for XFP and QSFP+ transceivers)                   | Displays whether the module voltage low alarm is <b>On</b> or <b>Off</b> .        |
| <b>Module voltage high warning</b><br>(Not available for XFP and QSFP+ transceivers)                | Displays whether the module voltage high warning is <b>On</b> or <b>Off</b> .     |
| <b>Module voltage low warning</b><br>(Not available for XFP and QSFP+ transceivers)                 | Displays whether the module voltage low warning is <b>On</b> or <b>Off</b> .      |

Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| <b>Laser rx power high alarm</b><br>(Not available for QSFP+ and CFP transceivers)                          | Displays whether the receive laser power high alarm is <b>On</b> or <b>Off</b> .   |
| <b>Laser rx power low alarm</b><br>(Not available for QSFP+ and CFP transceivers)                           | Displays whether the receive laser power low alarm is <b>On</b> or <b>Off</b> .  |
| <b>Laser rx power high warning</b><br>(Not available for QSFP+ and CFP transceivers)                        | Displays whether the receive laser power high warning is <b>On</b> or <b>Off</b> .   |
| <b>Laser rx power low warning</b><br>(Not available for QSFP+ and CFP transceivers)                         | Displays whether the receive laser power low warning is <b>On</b> or <b>Off</b> .  |
| <b>Laser bias current high alarm threshold</b><br>(Not available for QSFP+ transceivers)                    | Displays the vendor-specified threshold for the laser bias current high alarm.   |
| <b>Module not ready alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)                      | Displays whether the module not ready alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module has an operational fault. |
| <b>Module low power alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                 | Displays whether the module low power alarm is <b>On</b> or <b>Off</b> .   |
| <b>Module initialization incomplete alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the module initialization incomplete alarm is <b>On</b> or <b>Off</b> .   |
| <b>Module fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                     | Displays whether the module fault alarm is <b>On</b> or <b>Off</b> .   |
| <b>PLD Flash initialization fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays whether the PLD Flash initialization fault alarm is <b>On</b> or <b>Off</b> .   |
| <b>Power supply fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)               | Displays whether the power supply fault alarm is <b>On</b> or <b>Off</b> .   |
| <b>Checksum fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                   | Displays whether the checksum fault alarm is <b>On</b> or <b>Off</b> .   |
| <b>Tx laser disabled alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                | Displays whether the Tx laser disabled alarm is <b>On</b> or <b>Off</b> .  |

Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name   | Field Description  |
|--|--|
| <b>Module power down alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers) | Displays whether the module power down alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , module is in a limited power mode, low for normal operation.                            |
| <b>Tx data not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers) | Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is <b>On</b> or <b>Off</b> .  |
| <b>Tx not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)      | Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is <b>On</b> or <b>Off</b> .   |
| <b>Tx laser fault alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)    | Laser fault condition. Displays whether the Tx laser fault alarm is <b>On</b> or <b>Off</b> .  |
| <b>Tx CDR loss of lock alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)    | Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .                   |
| <b>Rx not ready alarm</b><br>(Not available for SFP, SFP+, QSFP+, and CFP transceivers)      | Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is <b>On</b> or <b>Off</b> .  |
| <b>Rx loss of signal alarm</b><br>(Not available for SFP and SFP+ transceivers)              | Receive loss of signal alarm. When the output is <b>On</b> , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is <b>On</b> or <b>Off</b> . |
| <b>Rx CDR loss of lock alarm</b><br>(Not available for SFP, SFP+, and QSFP+ transceivers)    | Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .   |
| <b>Laser bias current low alarm threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser bias current low alarm.  |
| <b>Laser bias current high warning threshold</b><br>(Not available for QSFP+ transceivers)   | Displays the vendor-specified threshold for the laser bias current high warning.   |
| <b>Laser bias current low warning threshold</b><br>(Not available for QSFP+ transceivers)    | Displays the vendor-specified threshold for the laser bias current low warning.  |
| <b>Laser output power high alarm threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the laser output power high alarm.   |
| <b>Laser output power low alarm threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser output power low alarm.  |

Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name   | Field Description  |
|--|--|
| <b>Laser output power high warning threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the laser output power high warning. |
| <b>Laser output power low warning threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the laser output power low warning.  |
| <b>Module temperature high alarm threshold</b><br>(Not available for QSFP+ transceivers)       | Displays the vendor-specified threshold for the module temperature high alarm.   |
| <b>Module temperature low alarm threshold</b><br>(Not available for QSFP+ transceivers)        | Displays the vendor-specified threshold for the module temperature low alarm.    |
| <b>Module temperature high warning threshold</b><br>(Not available for QSFP+ transceivers)     | Displays the vendor-specified threshold for the module temperature high warning. |
| <b>Module temperature low warning threshold</b><br>(Not available for QSFP+ transceivers)      | Displays the vendor-specified threshold for the module temperature low warning.  |
| <b>Module voltage high alarm threshold</b><br>(Not available for XFP and QSFP+ transceivers)   | Displays the vendor-specified threshold for the module voltage high alarm.       |
| <b>Module voltage low alarm threshold</b><br>(Not available for XFP and QSFP+ transceivers)    | Displays the vendor-specified threshold for the module voltage low alarm.        |
| <b>Module voltage high warning threshold</b><br>(Not available for XFP and QSFP+ transceivers) | Displays the vendor-specified threshold for the module voltage high warning.     |
| <b>Module voltage low warning threshold</b><br>(Not available for XFP and QSFP+ transceivers)  | Displays the vendor-specified threshold for the module voltage low warning.      |
| <b>Laser rx power high alarm threshold</b><br>(Not available for QSFP+ transceivers)           | Displays the vendor-specified threshold for the laser rx power high alarm.       |
| <b>Laser rx power low alarm threshold</b><br>(Not available for QSFP+ transceivers)            | Displays the vendor-specified threshold for the laser rx power low alarm.        |
| <b>Laser rx power high warning threshold</b><br>(Not available for QSFP+ transceivers)         | Displays the vendor-specified threshold for the laser rx power high warning.     |

Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| <b>Laser rx power low warning threshold</b><br>(Not available for QSFP+ transceivers)                         | Displays the vendor-specified threshold for the laser rx power low warning.                                |
| <b>Laser temperature high alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays the vendor-specified threshold for the laser temperature high alarm, in Celsius and Fahrenheit.   |
| <b>Laser temperature low alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays the vendor-specified threshold for the laser temperature low alarm, in Celsius and Fahrenheit.    |
| <b>Laser temperature high warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays the vendor-specified threshold for the laser temperature high warning, in Celsius and Fahrenheit. |
| <b>Laser temperature low warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays the vendor-specified threshold for the laser temperature low warning, in Celsius and Fahrenheit.  |
| <b>SOA bias current high alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)    | Displays the vendor-specified threshold for SOA bias current high alarm.                                   |
| <b>SOA bias current low alarm threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)     | Displays the vendor-specified threshold for SOA bias current low alarm.                                    |
| <b>SOA bias current high warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)  | Displays the vendor-specified threshold for SOA bias current high warning.                                 |
| <b>SOA bias current low warning threshold</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)   | Displays the vendor-specified threshold for SOA bias current low warning.                                  |
| <b>Laser receiver power high alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers)                 | Displays whether the laser receiver power high alarm is <b>On</b> or <b>Off</b> .                          |
| <b>Laser receiver power low alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers)                  | Displays whether the laser receiver power low alarm is <b>On</b> or <b>Off</b> .                           |
| <b>Laser receiver power high warning</b><br>(Not available for SFP, SFP+, and XFP transceivers)               | Displays whether the laser receiver power high warning is <b>On</b> or <b>Off</b> .                        |
| <b>Laser receiver power low warning</b><br>(Not available for SFP, SFP+, and XFP transceivers)                | Displays whether the laser receiver power low warning is <b>On</b> or <b>Off</b> .                         |



Table 252: show interfaces diagnostics optics Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| <b>Laser receiver power</b><br>(Not available for SFP, SFP+, and XFP transceivers)                  | Displays the laser receiver power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm). |
| <b>Tx loss of signal functionality alarm</b><br>(Not available for SFP, SFP+, and XFP transceivers) | Displays whether the Tx loss of signal functionality alarm is <b>On</b> or <b>Off</b> .      |
| <b>APD supply fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)         | Displays whether the APD supply fault alarm is <b>On</b> or <b>Off</b> .                     |
| <b>TEC fault alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)                | Displays whether the TEC fault alarm is <b>On</b> or <b>Off</b> .                            |
| <b>Wavelength unlocked alarm</b><br>(Not available for SFP, SFP+, XFP, and QSFP+ transceivers)      | Displays whether the Wavelength unlocked alarm is <b>On</b> or <b>Off</b> .                  |

## Sample Output

### show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```

user@switch> show interfaces diagnostics optics ge-0/1/0
Physical interface: ge-0/1/0
  Laser bias current           : 5.444 mA
  Laser output power          : 0.3130 mW / -5.04 dBm
  Module temperature          : 36 degrees C / 97 degrees F
  Module voltage              : 3.2120 V
  Receiver signal average optical power : 0.3840 mW / -4.16 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm     : Off
  Module voltage low alarm      : Off
  Module voltage high warning   : Off
  Module voltage low warning    : Off
  Laser rx power high alarm     : Off
  Laser rx power low alarm      : Off
  Laser rx power high warning   : Off
  Laser rx power low warning    : Off
  Laser bias current high alarm threshold : 15.000 mA
  Laser bias current low alarm threshold  : 1.000 mA
  Laser bias current high warning threshold : 12.000 mA

```

```

Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6300 mW / -2.01 dBm
Laser output power low alarm threshold : 0.0660 mW / -11.80 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0780 mW / -11.08 dBm
Module temperature high alarm threshold : 109 degrees C / 228 degrees F
Module temperature low alarm threshold : -29 degrees C / -20 degrees F
Module temperature high warning threshold : 103 degrees C / 217 degrees F
Module temperature low warning threshold : -13 degrees C / 9 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7939 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0157 mW / -18.04 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
  Laser bias current : 4.968 mA
  Laser output power : 0.4940 mW / -3.06 dBm
  Module temperature : 27 degrees C / 81 degrees F
  Module voltage : 3.2310 V
  Receiver signal average optical power : 0.0000
  Laser bias current high alarm : Off
  Laser bias current low alarm : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm : Off
  Module voltage low alarm : Off
  Module voltage high warning : Off
  Module voltage low warning : Off
  Laser rx power high alarm : Off
  Laser rx power low alarm : On
  Laser rx power high warning : Off
  Laser rx power low warning : On
  Laser bias current high alarm threshold : 10.500 mA
  Laser bias current low alarm threshold : 2.000 mA
  Laser bias current high warning threshold : 9.000 mA
  Laser bias current low warning threshold : 2.500 mA
  Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
  Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
  Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
  Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F

```

```

Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver)

```
user@switch> show interfaces diagnostics optics xe-0/1/0
```

```
Physical interface: xe-0/1/0
```

```

Laser bias current                : 8.029 mA
Laser output power                 : 0.6430 mW / -1.92 dBm
Module temperature                 : 4 degrees C / 39 degrees F
Laser rx power                    : 0.0012 mW / -29.21 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Module temperature high alarm      : Off
Module temperature low alarm       : Off
Module temperature high warning    : Off
Module temperature low warning     : Off
Laser rx power high alarm          : Off
Laser rx power low alarm           : On
Laser rx power high warning        : Off
Laser rx power low warning         : On
Module not ready alarm             : On
Module power down alarm            : Off
Tx data not ready alarm            : Off
Tx not ready alarm                 : Off
Tx laser fault alarm               : Off
Tx CDR loss of lock alarm          : Off
Rx not ready alarm                 : On
Rx loss of signal alarm            : On
Rx CDR loss of lock alarm          : On
Laser bias current high alarm threshold : 13.000 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 12.000 mA
Laser bias current low warning threshold : 3.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold  : 0.1650 mW / -7.83 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 90 degrees C / 194 degrees F
Module temperature low alarm threshold  : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold    : 0.8912 mW / -0.50 dBm
Laser rx power low alarm threshold     : 0.0912 mW / -10.40 dBm
Laser rx power high warning threshold  : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold   : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
  Module temperature                : 33 degrees C / 92 degrees F
  Module voltage                    : 3.3060 V
Lane 0
  Laser bias current                : 7.182 mA
  Laser receiver power              : 0.743 mW / -1.29 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning   : Off
  Laser receiver power high alarm   : Off
  Laser receiver power low alarm    : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning  : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off
Lane 1
  Laser bias current                : 7.326 mA
  Laser receiver power              : 0.752 mW / -1.24 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning   : Off
  Laser receiver power high alarm   : Off
  Laser receiver power low alarm    : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning  : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off
Lane 2
  Laser bias current                : 7.447 mA
  Laser receiver power              : 0.790 mW / -1.03 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning   : Off
  Laser receiver power high alarm   : Off
  Laser receiver power low alarm    : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning  : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off
Lane 3
  Laser bias current                : 7.734 mA
  Laser receiver power              : 0.768 mW / -1.15 dBm
  Laser bias current high alarm     : Off
  Laser bias current low alarm      : Off
  Laser bias current high warning   : Off
  Laser bias current low warning   : Off
  Laser receiver power high alarm   : Off
  Laser receiver power low alarm    : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning  : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm           : Off

```

## Sample Output

### show interfaces diagnostics optics et-4/1/0 (CFP Transceiver)

```

user@switch> show interfaces diagnostics optics et-4/1/0
Physical interface: et-4/1/0
  Module temperature                : 38 degrees C / 101 degrees F
  Module voltage                    : 3.2500 V
  Module temperature high alarm     : Off
  Module temperature low alarm      : Off
  Module temperature high warning   : Off
  Module temperature low warning    : Off
  Module voltage high alarm         : Off
  Module voltage low alarm          : Off
  Module voltage high warning       : Off
  Module voltage low warning        : Off
  Module not ready alarm            : Off
  Module low power alarm            : Off
  Module initialization incomplete alarm : Off
  Module fault alarm                : Off
  PLD Flash initialization fault alarm : Off
  Power supply fault alarm          : Off
  Checksum fault alarm              : Off
  Tx laser disabled alarm           : Off
  Tx loss of signal functionality alarm : Off
  Tx CDR loss of lock alarm         : Off
  Rx loss of signal alarm           : Off
  Rx CDR loss of lock alarm         : Off
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold : 3.5000 V
  Module voltage low alarm threshold : 3.0990 V
  Module voltage high warning threshold : 3.4000 V
  Module voltage low warning threshold : 3.2000 V
  Laser bias current high alarm threshold : 250.000 mA
  Laser bias current low alarm threshold : 37.500 mA
  Laser bias current high warning threshold : 225.000 mA
  Laser bias current low warning threshold : 50.000 mA
  Laser output power high alarm threshold : 3.9800 mW / 6.00 dBm
  Laser output power low alarm threshold : 0.4670 mW / -3.31 dBm
  Laser output power high warning threshold : 3.5480 mW / 5.50 dBm
  Laser output power low warning threshold : 0.5240 mW / -2.81 dBm
  Laser rx power high alarm threshold : 3.5481 mW / 5.50 dBm
  Laser rx power low alarm threshold : 0.0616 mW / -12.10 dBm
  Laser rx power high warning threshold : 3.1622 mW / 5.00 dBm
  Laser rx power low warning threshold : 0.0691 mW / -11.61 dBm
  Laser temperature high alarm threshold : 67 degrees C / 153 degrees F
  Laser temperature low alarm threshold : 35 degrees C / 95 degrees F
  Laser temperature high warning threshold : 62 degrees C / 144 degrees F
  Laser temperature low warning threshold : 40 degrees C / 104 degrees F
  SOA bias current high alarm threshold : 0.000 mA
  SOA bias current low alarm threshold : 0.000 mA
  SOA bias current high warning threshold : 0.000 mA
  SOA bias current low warning threshold : 0.000 mA
Lane 0
  Laser bias current                : 131.684 mA
  Laser output power                 : 1.002 mW / 0.01 dBm
  Laser temperature                  : 54 degrees C / 128 degrees F
  Laser receiver power               : 0.497 mW / -3.03 dBm

```

```

Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Laser temperature high alarm       : Off
Laser temperature low alarm        : Off
Laser temperature high warning     : Off
Laser temperature low warning      : Off
Laser receiver power high alarm    : Off
Laser receiver power low alarm     : Off
Laser receiver power high warning  : Off
Laser receiver power low warning   : Off
Tx loss of signal functionality alarm : Off
Rx CDR loss of lock alarm          : Off
Rx loss of signal alarm            : Off
Rx CDR loss of lock alarm          : Off
APD supply fault alarm             : Off
TEC fault alarm                   : Off
Wavelength unlocked alarm          : Off

Lane 1
Laser bias current                  : 122.345 mA
Laser output power                  : 1.002 mW / 0.01 dBm
Laser temperature                   : 51 degrees C / 124 degrees F
Laser receiver power                : 0.611 mW / -2.14 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off
Laser output power low alarm        : Off
Laser output power high warning     : Off
Laser output power low warning      : Off
Laser temperature high alarm        : Off
Laser temperature low alarm         : Off
Laser temperature high warning      : Off
Laser temperature low warning       : Off
Laser receiver power high alarm     : Off
Laser receiver power low alarm      : Off
Laser receiver power high warning   : Off
Laser receiver power low warning    : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm           : Off
Rx loss of signal alarm             : Off
Rx CDR loss of lock alarm           : Off
APD supply fault alarm              : Off
TEC fault alarm                    : Off
Wavelength unlocked alarm           : Off

Lane 2
Laser bias current                  : 112.819 mA
Laser output power                  : 1.000 mW / 0.00 dBm
Laser temperature                   : 50 degrees C / 122 degrees F
Laser receiver power                : 0.540 mW / -2.67 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off

```

```

Laser output power low alarm           : Off
Laser output power high warning        : Off
Laser output power low warning         : Off
Laser temperature high alarm           : Off
Laser temperature low alarm            : Off
Laser temperature high warning         : Off
Laser temperature low warning          : Off
Laser receiver power high alarm        : Off
Laser receiver power low alarm         : Off
Laser receiver power high warning      : Off
Laser receiver power low warning       : Off
Tx loss of signal functionality alarm  : Off
Tx CDR loss of lock alarm              : Off
Rx loss of signal alarm                : Off
Rx CDR loss of lock alarm              : Off
APD supply fault alarm                 : Off
TEC fault alarm                       : Off
Wavelength unlocked alarm              : Off

Lane 3
Laser bias current                     : 100.735 mA
Laser output power                     : 1.002 mW / 0.01 dBm
Laser temperature                      : 50 degrees C / 122 degrees F
Laser receiver power                   : 0.637 mW / -1.96 dBm
Laser bias current high alarm          : Off
Laser bias current low alarm           : Off
Laser bias current high warning        : Off
Laser bias current low warning         : Off
Laser output power high alarm          : Off
Laser output power low alarm           : Off
Laser output power high warning        : Off
Laser output power low warning         : Off
Laser temperature high alarm           : Off
Laser temperature low alarm            : Off
Laser temperature high warning         : Off
Laser temperature low warning          : Off
Laser receiver power high alarm        : Off
Laser receiver power low alarm         : Off
Laser receiver power high warning      : Off
Laser receiver power low warning       : Off
Tx loss of signal functionality alarm  : Off
Tx CDR loss of lock alarm              : Off
Rx loss of signal alarm                : Off
Rx CDR loss of lock alarm              : Off
APD supply fault alarm                 : Off
TEC fault alarm                       : Off
Wavelength unlocked alarm              : Off

```

## show interfaces ge

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show interfaces <i>device-name:type-fpc/pic/port</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;descriptions&gt;</code><br><code>&lt;media&gt;</code><br><code>&lt;routing-instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;snmp-index <i>snmp-index</i>&gt;</code><br><code>&lt;statistics&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Display status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.   |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b><i>device-name:type-fpc/pic/port</i></b>—The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing-instance or display all routing-instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Monitoring Interface Status and Traffic on page 335</a></li><li>• <a href="#">Troubleshooting Network Interfaces on page 1234</a></li><li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1234</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li></ul>   |
| <b>List of Sample Output</b>    | <a href="#">show interfaces on page 2790</a><br><a href="#">show interfaces brief on page 2790</a><br><a href="#">show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled) on page 2790</a><br><a href="#">show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled) on page 2791</a>   |



[show interfaces extensive \(Symmetric Flow Control and Autonegotiation Enabled\) on page 2792](#)

[show interfaces extensive \(Asymmetric Flow Control and Autonegotiation Enabled\) on page 2794](#)

[show interfaces terse on page 2796](#)

[show interfaces terse \(QFabric Systems\) on page 2796](#)

**Output Fields** [Table 253 on page 2783](#) lists the output fields for the **show interfaces ge** command. Output fields are listed in the approximate order in which they appear.

**Table 253: show interfaces ge Output Fields**

| Field Name                     | Field Description  | Level of Output               |
|--------------------------------|--|-------------------------------|
| <b>Physical Interface</b>      |  |                               |
| <b>Physical interface</b>      | Name of the physical interface.  | All levels                    |
| <b>Enabled</b>                 | State of the interface: <b>Enabled</b> or <b>Disabled</b> .  | All levels                    |
| <b>Interface index</b>         | Index number of the physical interface, which reflects its initialization sequence.  | <b>detail extensive none</b>  |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.  | <b>detail extensive none</b>  |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.  | <b>detail extensive</b>       |
| <b>Description</b>             | Optional user-specified description.   | <b>brief detail extensive</b> |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.  | All levels                    |
| <b>MTU</b>                     | Maximum transmission unit size on the physical interface. The default is 1514.   | All levels                    |
| <b>Speed</b>                   | Speed at which the interface is running.   | All levels                    |
| <b>Loopback</b>                | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .   | All levels                    |
| <b>Source filtering</b>        | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .   | All levels                    |
| <b>Flow control</b>            | Flow control status: <b>Enabled</b> or <b>Disabled</b> .<br><br><i>NOTE:</i> This field is only displayed if asymmetric flow control is not configured.  | <b>detail extensive</b>       |
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ):<br><br><ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/><b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/><b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul><br><i>NOTE:</i> This field is only displayed if asymmetric flow control is configured. | <b>detail extensive</b>       |

Table 253: show interfaces ge Output Fields (*continued*)

| Field Name                     | Field Description  | Level of Output              |
|--------------------------------|--|------------------------------|
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .  | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>  | All levels                   |
| <b>Device flags</b>            | Information about the physical device.   | All levels                   |
| <b>Interface flags</b>         | Information about the interface.   | All levels                   |
| <b>Link flags</b>              | Information about the link.  | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.   | <b>detail extensive none</b> |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.  | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.  | <b>detail extensive none</b> |
| <b>Hardware address</b>        | MAC address of the hardware.   | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .   | <b>detail extensive none</b> |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.  | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled on the switch.</p> | <b>detail extensive</b>      |

Table 253: show interfaces ge Output Fields (*continued*)

| Field Name          | Field Description   | Level of Output  |
|---------------------|---|------------------|
| <b>Input errors</b> | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b> |

Table 253: show interfaces ge Output Fields (*continued*)

| Field Name                              | Field Description  | Level of Output              |
|---|--|------------------------------|
| <b>Output errors</b>                    | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>             |
| <b>Egress queues</b>                    | Total number of egress queues supported on the specified interface.  | <b>detail extensive</b>      |
| <b>Queue counters (Egress )</b>         | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>  | <b>detail extensive</b>      |
| <b>Queue Number</b>                     | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.   | <b>detail extensive</b>      |
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>   | <b>detail extensive none</b> |

Table 253: show interfaces ge Output Fields (*continued*)

| Field Name               | Field Description  | Level of Output  |
|--------------------------|--|------------------|
| <b>MAC statistics</b>    | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b> |
| <b>Filter Statistics</b> | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.  | <b>extensive</b> |

Table 253: show interfaces ge Output Fields (*continued*)

| Field Name                             | Field Description   | Level of Output       |
|--|---|-----------------------|
| Autonegotiation information            | <p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> <li>• <b>Link partner speed</b>—Speed of the link partner.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive             |
| Packet Forwarding Engine configuration | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>  | extensive             |
| Logical Interface                      |   |                       |
| Logical interface                      | Name of the logical interface.  | All levels            |
| Index                                  | Index number of the logical interface, which reflects its initialization sequence.  | detail extensive none |
| SNMP ifIndex                           | SNMP interface index number for the logical interface.  | detail extensive none |
| Generation                             | Unique number for use by Juniper Networks technical support only.   | detail extensive      |
| Flags                                  | Information about the logical interface.  | All levels            |

Table 253: show interfaces ge Output Fields (*continued*)

| Field Name                     | Field Description   | Level of Output              |
|--------------------------------|---|------------------------------|
| <b>Encapsulation</b>           | Encapsulation on the logical interface.   | All levels                   |
| <b>Protocol</b>                | Protocol family.  | <b>detail extensive none</b> |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.  | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b> | If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.   | <b>extensive</b>             |
| <b>Local statistics</b>        | Number and rate of bytes and packets destined to and from the switch.   | <b>extensive</b>             |
| <b>Transit statistics</b>      | Number and rate of bytes and packets transiting the switch.   | <b>extensive</b>             |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b>      |
| <b>Route Table</b>             | Route table in which the logical interface address is located. For example, 0 refers to the routing table <b>inet.0</b> .   | <b>detail extensive none</b> |
| <b>Input Filters</b>           | Names of any input filters applied to this interface.   | <b>detail extensive</b>      |
| <b>Output Filters</b>          | Names of any output filters applied to this interface.  | <b>detail extensive</b>      |
| <b>Flags</b>                   | Information about protocol family flags.<br><br>If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled. | <b>detail extensive</b>      |
| <b><i>protocol-family</i></b>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.  | <b>brief</b>                 |
| <b>Flags</b>                   | Information about the address flags.  | <b>detail extensive none</b> |
| <b>Destination</b>             | IP address of the remote side of the connection.  | <b>detail extensive none</b> |
| <b>Local</b>                   | IP address of the logical interface.  | <b>detail extensive none</b> |
| <b>Broadcast</b>               | Broadcast address of the logical interlace.   | <b>detail extensive none</b> |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```
user@switch> show interfaces ge-0/0/9
Physical interface: ge-0/0/9, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 21
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
  Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 22)
  Flags: SNMP-Traps
  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch
  Flags: None
```

### show interfaces brief

```
user@switch> show interfaces ge-0/0/9 brief
Physical interface: ge-0/0/9, Enabled, Physical link is Down
  Description: voice priority and tcp and icmp traffic rate-limiting filter at i
  ngress port
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface ge-0/0/9.0
  Flags: Device-Down SNMP-Traps Encapsulation: ENET2
  eth-switch
```

### show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled)

```
user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
  Interface index: 193, SNMP ifIndex: 206, Generation: 196
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
```



```

Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped   : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes   : 0 0 bps
  Output bytes  : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets: 0
  Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets

  0 best-effort      0              0              0
  1 assured-forw     0              0              0
  5 expedited-fo     0              0              0
  7 network-cont     0              0              0

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes   : 0
  Output bytes  : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes   : 0 0 bps
  Output bytes  : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,

```

#### show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
Interface index: 193, SNMP ifIndex: 206, Generation: 196
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on ,
Auto-negotiation: Enabled,

```

```

Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped   : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0                      0 bps
Output bytes  : 0                      0 bps
Input packets : 0                      0 pps
Output packets: 0                      0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters:

```

|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 1 assured-forw | 0              | 0                   | 0               |
| 5 expedited-fo | 0              | 0                   | 0               |
| 7 network-cont | 0              | 0                   | 0               |

```

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Local statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Transit statistics:
Input bytes   : 0                      0 bps
Output bytes  : 0                      0 bps
Input packets : 0                      0 pps
Output packets: 0                      0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,,

```

#### show interfaces extensive (Symmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
Interface index: 49164, SNMP ifIndex: 574, Generation: 142

```

```

Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
Last flapped   : 2011-02-25 00:45:03 UTC (22:42:48 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets : 0          0 pps
Output packets: 0          0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:

```

|                | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort  | 0              | 0                   | 0               |
| 2 no-loss      | 0              | 0                   | 0               |
| 3 fcoe         | 0              | 0                   | 0               |
| 7 network-cont | 0              | 0                   | 0               |

```

Queue number:      Mapped forwarding classes
0                  best-effort
2                  no-loss
3                  fcoe
7                  network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:

```

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |

```

VLAN tagged frames          0
Code violations              0
MAC Priority Flow Control Statistics:
  Priority : 0                0          0
  Priority : 1                0          0
  Priority : 2                0          0
  Priority : 3                0          0
  Priority : 4                0          0
  Priority : 5                0          0
  Priority : 6                0          0
  Priority : 7                0          0
Filter statistics:
  Input packet count          0
  Input packet rejects        0
  Input DA rejects            0
  Input SA rejects            0
  Output packet count         0
  Output packet pad count     0
  Output packet error count   0
  CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth          Buffer Priority
Limit
    0 best-effort             75          750000000    75          0          low
none
    7 network-control         5           500000000    5           0          low
none
    8 mcast-be                15          1500000000   15          0          low
none
    11 mcast-nc               5           500000000    5           0          low
none

```

#### show interfaces extensive (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
  Interface index: 49164, SNMP ifIndex: 574, Generation: 142
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on
  Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
  Last flapped   : 2011-02-25 00:45:03 UTC (22:42:48 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0          0 bps
  Output bytes: 0          0 bps
  Input packets: 0        0 pps

```

```

Output packets:                0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                0                0                0
  2 no-loss                    0                0                0
  3 fcoe                       0                0                0
  7 network-cont               0                0                0

Queue number:      Mapped forwarding classes
  0                best-effort
  2                no-loss
  3                fcoe
  7                network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
  Total octets              Receive      Transmit
  Total packets             0          0
  Unicast packets           0          0
  Broadcast packets         0          0
  Multicast packets         0          0
  CRC/Align errors          0          0
  FIFO errors               0          0
  MAC control frames        0          0
  MAC pause frames          0          0
  Oversized frames          0
  Jabber frames             0
  Fragment frames           0
  VLAN tagged frames        0
  Code violations           0
MAC Priority Flow Control Statistics:
  Priority : 0              0          0
  Priority : 1              0          0
  Priority : 2              0          0
  Priority : 3              0          0
  Priority : 4              0          0
  Priority : 5              0          0
  Priority : 6              0          0
  Priority : 7              0          0
Filter statistics:
  Input packet count        0
  Input packet rejects      0
  Input DA rejects          0
  Input SA rejects          0

```

```

Output packet count                                0
Output packet pad count                            0
Output packet error count                          0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link Partner:
  Link mode: Full-duplex, Flow control: None, Remote fault: OK,
  Link partner Speed: 1000 Mbps
Local resolution:
  Flow control: enable PAUSE transmit and Disable PAUSE receive, Remote
fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      0 best-effort        75      750000000    75      0      low
none
      7 network-control    5      50000000    5      0      low
none
      8 mcast-be           15     150000000   15      0      low
none
      11 mcast-nc          5      50000000    5      0      low
none

```

#### show interfaces terse

```

user@switch> show interfaces ge-0/0/12 terse
Interface      Admin Link Proto  Local      Remote
ge-0/0/12      up    up

```

#### show interfaces terse (QFabric Systems)

```

user@switch> show interfaces node1:ge-0/0/0 terse
Physical interface: node1:ge-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 2884086
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
Interface flags: Internal: 0x4000
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped    : Never
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)

```

## show interfaces (GRE)


|  |  |
|--|--|
| <b>Syntax</b>  | <pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>   |
| <b>Release Information</b>   | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>   |
| <b>Description</b>   | Display status information about the specified generic routing encapsulation (GRE) interface.  |
| <b>Options</b>   | <p><b><i>interface-type</i></b>—On M Series and T Series routers and EX Series switches, the interface type is <i>gr-fpc/pic/port</i>.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified output level of interface information.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <div>  <p><b>NOTE:</b> You can configure generic routing encapsulation (GRE) interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the <i>Junos OS MPLS Applications Library for Routing Devices</i> and the <i>Junos OS, Release 14.1</i>.</p> </div> |  |
| <b>Required Privilege Level</b>  | view   |
| <b>List of Sample Output</b>   | <p><a href="#">show interfaces (GRE) on page 2801</a></p> <p><a href="#">show interfaces brief (GRE) on page 2801</a></p> <p><a href="#">show interfaces detail (GRE) on page 2801</a></p> <p><a href="#">show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 2802</a></p> <p><a href="#">show interfaces extensive (GRE) on page 2803</a></p>   |
| <b>Output Fields</b>   | <p><a href="#">Table 254 on page 2798</a> lists the output fields for the <b>show interfaces (GRE)</b> command. Output fields are listed in the approximate order in which they appear.</p>  |

Table 254: GRE show interfaces Output Fields

| Field Name                     | Field Description  | Level of Output              |
|--------------------------------|--|------------------------------|
| <b>Physical Interface</b>      |  |                              |
| <b>Physical interface</b>      | Name of the physical interface.  | All levels                   |
| <b>Enabled</b>                 | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .   | All levels                   |
| <b>Interface index</b>         | Physical interface's index number, which reflects its initialization sequence.   | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.  | <b>detail extensive none</b> |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.  | <b>detail extensive</b>      |
| <b>Type</b>                    | Type of interface.   | All levels                   |
| <b>Link-level type</b>         | Encapsulation used on the physical interface.  | All levels                   |
| <b>MTU</b>                     | MTU size on the physical interface.  | All levels                   |
| <b>Speed</b>                   | Speed at which the interface is running.   | All levels                   |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.  | <b>detail extensive</b>      |
| <b>Device Flags</b>            | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .   | All levels                   |
| <b>Interface Flags</b>         | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .  | All levels                   |
| <b>Input rate</b>              | Input rate in bits per second (bps) and packets per second (pps).  | None specified               |
| <b>Output rate</b>             | Output rate in bps and pps.  | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.  | <b>detail extensive</b>      |
| <b>Traffic statistics</b>      | <p>The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>      |
| <b>Logical Interface</b>       |  |                              |
| <b>Logical interface</b>       | Name of the logical interface.   | All levels                   |
| <b>Index</b>                   | Logical interface index number, which reflects its initialization sequence.  | <b>detail extensive none</b> |



Table 254: GRE show interfaces Output Fields (*continued*)

| Field Name                     | Field Description  | Level of Output       |
|--------------------------------|--|-----------------------|
| SNMP ifIndex                   | Logical interface SNMP interface index number.   | detail extensive none |
| Generation                     | Unique number for use by Juniper Networks technical support.   | detail extensive      |
| Flags                          | <p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under <i>Common Output Fields Description</i>. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>Reassemble-Pkts</b>—If the <b>Flags</b> field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation.</li> </ul>                               | All levels            |
| IP-Header                      | <p>IP header of the logical interface. If the <b>tunnel key</b> statement is configured, this information is included in the <b>IP Header</b> entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>df</b>—If the <b>IP-Header</b> field includes this string immediately following the 16 bits of identification information (that is, if <b>:df:</b> displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation.</li> </ul> | All levels            |
| Encapsulation                  | Encapsulation on the logical interface.  | All levels            |
| Copy-tos-to-outer-ip-header    | <p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel.</li> <li>• <b>Off</b>—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header.</li> </ul> <p><b>NOTE:</b> EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always <b>Off</b> in switch output.</p>  | detail extensive      |
| Gre keepalives configured      | <p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p><b>NOTE:</b> EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always <b>Off</b> in switch output.</p>   | detail extensive      |
| Gre keepalives adjacency state | Status of the other end of the GRE tunnel: <b>Up</b> or <b>Down</b> . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up.  | detail extensive      |
| Input packets                  | Number of packets received on the logical interface.   | None specified        |
| Output packets                 | Number of packets transmitted on the logical interface.  | None specified        |

Table 254: GRE show interfaces Output Fields (*continued*)

| Field Name                    | Field Description   | Level of Output              |
|-------------------------------|---|------------------------------|
| <b>Traffic statistics</b>     | <p>Rate of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input rate</b>—Rate of bits and packets received on the interface.</li> <li>• <b>Output rate</b>—Rate of bits and packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>      |
| <b>Local statistics</b>       | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.  | <b>detail extensive</b>      |
| <b>Transit statistics</b>     | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.  | <b>detail extensive none</b> |
| <b>Protocol</b>               | Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , or <b>mpls</b> .   | <b>detail extensive none</b> |
| <b><i>protocol-family</i></b> | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.  | <b>brief</b>                 |
| <b>MTU</b>                    | MTU size on the logical interface.  | <b>detail extensive none</b> |
| <b>Generation</b>             | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b>      |
| <b>Route table</b>            | Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .  | <b>detail extensive</b>      |
| <b>Flags</b>                  | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .  | <b>detail extensive none</b> |
| <b>Addresses, Flags</b>       | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .   | <b>detail extensive none</b> |
| <b>Destination</b>            | IP address of the remote side of the connection.  | <b>detail extensive none</b> |
| <b>Local</b>                  | IP address of the logical interface.  | <b>detail extensive none</b> |
| <b>Broadcast</b>              | Broadcast address of the logical interface.   | <b>detail extensive none</b> |
| <b>Generation</b>             | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b>      |

## Sample Output

### show interfaces (GRE)

```

user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
  Flags: None
  Addresses, Flags: Is-Primary
    Local: 1.10.1.1

```

### show interfaces brief (GRE)

```

user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
  IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  inet 10.100.0.1/30
  mpls

```

### show interfaces detail (GRE)

```

user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0

```

```

Output packets:                0
Local statistics:
Input bytes :                  0
Output bytes :                  0
Input packets:                 0
Output packets:                0
Transit statistics:
Input bytes :                  0          0 bps
Output bytes :                  0          0 bps
Input packets:                 0          0 pps
Output packets:                0          0 pps
Protocol inet, MTU: 1476, Generation: 12, Route table: 0
Flags: None
Addresses, Flags: Is-Primary
Destination: Unspecified, Local: 1.10.1.1, Broadcast: Unspecified,
Generation: 15

```

### show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

```

user@switch> show interfaces gr-2/0/15 detail
Physical interface: gr-2/0/15, Enabled, Physical link is Up
Interface index: 195, SNMP ifIndex: 846, Generation: 198
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:0f:d2, Hardware address: 00:1f:12:38:0f:d2
Device flags    : Present Running
Interface flags: Point-To-Point SNMP-Traps
Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
Traffic statistics:
Input bytes :          5600636          0 bps
Output bytes :          5600636          0 bps
Input packets:          20007          0 pps
Output packets:          20007          0 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 180.20.30.2:180.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
Input bytes :          5600886
Output bytes :          2881784
Input packets:          20010
Output packets:          10018
Local statistics:
Input bytes :           398
Output bytes :           264
Input packets:           5
Output packets:           3
Transit statistics:
Input bytes :          5600488          0 bps
Output bytes :          2881520          0 bps
Input packets:          20005          0 pps
Output packets:          10015          0 pps

```

```

Protocol inet, Generation: 159, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 90.90.90/24, Local: 90.90.90.10, Broadcast: 90.90.90.255,
  Generation: 144

```

```

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)

```

```

Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 160.20.40.2:160.20.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down

```

```
Traffic statistics:
```

```

Input bytes :          260
Output bytes :        2880148
Input packets:           4
Output packets:       10002

```

```
Local statistics:
```

```

Input bytes :          112
Output bytes :           0
Input packets:           2
Output packets:           0

```

```
Transit statistics:
```

```

Input bytes :          148          0 bps
Output bytes :        2880148        0 bps
Input packets:           2          0 pps
Output packets:       10002          0 pps

```

```
Protocol inet, Generation: 171, Route table: 0
```

```
Flags: None
```

```
Addresses, Flags: Is-Preferred Is-Primary
```

```

  Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
  Generation: 160

```

### show interfaces extensive (GRE)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 2801](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 2802](#).

## show interfaces irb

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>show interfaces irb &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance <i>instance-name</i>&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>  |
| <b>Release Information</b>      | <p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series</p>  |
| <b>Description</b>              | Display integrated routing and bridging interfaces information.   |
| <b>Options</b>                  | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance <i>instance-name</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the interface with the specified SNMP index.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Additional Information</b>   | Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.  |
| <b>Required Privilege Level</b> | view  |
| <b>List of Sample Output</b>    | <p><a href="#">show interfaces irb extensive on page 2808</a></p> <p><a href="#">show interfaces irb snmp-index on page 2809</a></p>  |
| <b>Output Fields</b>            | <a href="#">Table 255 on page 2804</a> lists the output fields for the <b>show interfaces irb</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 255: show interfaces irb Output Fields**

| Field Name                | Field Description   | Level of Output |
|---------------------------|---|-----------------|
| <b>Physical Interface</b> |   |                 |
| <b>Physical interface</b> | Name of the physical interface.   | All levels      |
| <b>Enabled</b>            | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | All levels      |

Table 255: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description  | Level of Output                    |
|--------------------------------|--|------------------------------------|
| <b>Proto</b>                   | Protocol configured on the interface.  | <b>terse</b>                       |
| <b>Interface index</b>         | Physical interface index number, which reflects its initialization sequence.   | <b>detail extensive none</b>       |
| <b>SNMP ifIndex</b>            | SNMP index number for the physical interface.  | <b>detail extensive none</b>       |
| <b>Type</b>                    | Physical interface type.   | <b>detail extensive none</b>       |
| <b>Link-level type</b>         | Encapsulation being used on the physical interface.  | <b>detail extensive brief none</b> |
| <b>MTU</b>                     | MTU size on the physical interface.  | <b>detail extensive brief none</b> |
| <b>Clocking</b>                | Reference clock source: <b>Internal</b> or <b>External</b> . Always unspecified on IRB interfaces.   | <b>detail extensive brief</b>      |
| <b>Speed</b>                   | Speed at which the interface is running. Always unspecified on IRB interfaces.   | <b>detail extensive brief</b>      |
| <b>Device flags</b>            | Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .   | <b>detail extensive brief none</b> |
| <b>Interface flags</b>         | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .  | <b>detail extensive brief none</b> |
| <b>Link type</b>               | Physical interface link type: <b>full duplex</b> or <b>half duplex</b> .   | <b>detail extensive none</b>       |
| <b>Link flags</b>              | Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .   | <b>detail extensive none</b>       |
| <b>Physical Info</b>           | Physical interface information.  | All levels                         |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.  | <b>detail extensive</b>            |
| <b>Current address</b>         | Configured MAC address.  | <b>detail extensive none</b>       |
| <b>Hardware address</b>        | MAC address of the hardware.   | <b>detail extensive none</b>       |
| <b>Alternate link address</b>  | Backup address of the link.  | <b>detail extensive</b>            |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> . | <b>detail extensive none</b>       |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.  | <b>detail extensive</b>            |

Table 255: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description   | Level of Output         |
|--------------------------------|---|-------------------------|
| <b>Traffic statistics</b>      | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>   | <b>detail extensive</b> |
| <b>IPv6 transit statistics</b> | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>   | <b>detail extensive</b> |
| <b>Input errors</b>            | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>           | <b>detail extensive</b> |
| <b>Output errors</b>           | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the DPC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>detail extensive</b> |

---

#### Logical Interface

---



Table 255: show interfaces irb Output Fields (*continued*)

| Field Name                     | Field Description   | Level of Output                 |
|--------------------------------|---|---------------------------------|
| <b>Logical interface</b>       | Name of the logical interface.  | All levels                      |
| <b>Index</b>                   | Index number of the logical interface (which reflects its initialization sequence).   | <b>detail extensive</b><br>none |
| <b>SNMP ifIndex</b>            | SNMP interface index number of the logical interface.   | <b>detail extensive</b><br>none |
| <b>Generation</b>              | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b>         |
| <b>Flags</b>                   | Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .   | <b>detail extensive</b>         |
| <b>Encapsulation</b>           | Encapsulation on the logical interface.   | <b>detail extensive</b>         |
| <b>Bandwidth</b>               | Speed at which the interface is running.  | <b>detail extensive</b>         |
| <b>Routing Instance</b>        | Routing instance IRB is configured under.   | <b>detail extensive</b>         |
| <b>Bridging Domain</b>         | Bridging domain IRB is participating in.  | <b>detail extensive</b>         |
| <b>Traffic statistics</b>      | Number and rate of bytes and packets received and transmitted on the logical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>   | <b>detail extensive</b>         |
| <b>IPv6 transit statistics</b> | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> | <b>detail extensive</b>         |
| <b>Local statistics</b>        | Statistics for traffic received from and transmitted to the Routing Engine.   | <b>detail extensive</b>         |
| <b>Transit statistics</b>      | Statistics for traffic transiting the router.   | <b>detail extensive</b>         |
| <b>Protocol</b>                | Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under <i>Common Output Fields Description</i> .  | <b>detail extensive</b>         |
| <b>MTU</b>                     | Maximum transmission unit size on the logical interface.  | <b>detail extensive</b>         |
| <b>Maximum labels</b>          | Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.   | <b>detail extensive</b><br>none |

Table 255: show interfaces irb Output Fields (*continued*)

| Field Name              | Field Description   | Level of Output         |
|-------------------------|---|-------------------------|
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b> |
| <b>Route table</b>      | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.   | <b>detail extensive</b> |
| <b>Addresses, Flags</b> | Information about address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .                 | <b>detail extensive</b> |
| <b>Policer</b>          | The policer that is to be evaluated when packets are received or transmitted on the interface.  | <b>detail extensive</b> |
| <b>Flags</b>            | Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> . | <b>detail extensive</b> |

## Sample Output

### show interfaces irb extensive

```

user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  IPv6 transit statistics:
    Input bytes   : 0
    Output bytes  : 0
    Input packets : 0
    Output packets: 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0

```

```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 154, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255,
    Generation: 155
Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
  Flags: Is-Primary
  Policer: Input: __default_arp_policer

```

#### show interfaces irb snmp-index

```

user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags : Present Running
  Interface flags: SNMP-Traps
  Link type : Full-Duplex
  Link flags : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped : Never
    Input packets : 0
    Output packets: 0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1500
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
  Protocol multiservice, MTU: 1500
    Flags: Is-Primary

```

## show interfaces mc-ae

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show interfaces mc-ae id <i>identifier</i> unit <i>number</i></b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.   |
| <b>Description</b>              | On peers with multi-chassis aggregated Ethernet ( <b>mc-aeX</b> ) interfaces, use this command to display information about the <b>mc-aeX</b> interfaces.  |
| <b>Options</b>                  | <b>identifier</b> —(Optional) Name of the multichassis aggregated Ethernet interface.<br><b>number</b> —(Optional) Specify the logical interface by unit number.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Multichassis Link Aggregation on page 2411</a> (QFX Series Switches)</li> <li>• <a href="#">Understanding Multichassis Link Aggregation</a> (EX Series Switches)</li> <li>• <a href="#">Configuring Multichassis Link Aggregation on page 2597</a> (QFX Series Switches)</li> <li>• <a href="#">Configuring Multichassis Link Aggregation</a> (EX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation on page 2471</a> (QFX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation with Layer 3 MAC Address Synchronization on page 2530</a> (QFX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast using MAC Address Synchronization</a> (QFX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using Virtual Router Redundancy Protocol (VRRP) on page 2551</a> (QFX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Unicast Using VRRP on EX9200 Switches</a> (EX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on page 2493</a> (QFX Series Switches)</li> <li>• <a href="#">Example: Configuring Multichassis Link Aggregation for Layer 3 Multicast Using VRRP on EX9200 Switches</a> (EX Series Switches)</li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show interfaces mc-ae on page 2811</a>   |
| <b>Output Fields</b>            | <a href="#">Table 256 on page 2810</a> lists the output fields for the <b>show interfaces mc-ae</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 256: show interfaces mc-ae Output Fields**

| Output Field Name                    | Field Description   |
|--------------------------------------|---|
| <b>Current State Machine's State</b> | Specifies the state of the MC-LAG initialization state machine. |

Table 256: show interfaces mc-ae Output Fields (*continued*)

| Output Field Name        | Field Description  |
|--------------------------|--|
| <b>Member Link</b>       | Specifies the identifiers of the configured multichassis link aggregated interface members.  |
| <b>Local Status</b>      | Specifies the status of the local link: <b>active</b> or <b>standby</b> .  |
| <b>Peer Status</b>       | Specifies the status of the peer link: <b>active</b> or <b>standby</b> .   |
| <b>Peer State</b>        | Specifies the status of the local and peer links in an <b>active/active</b> MC-LAG configuration   |
| <b>Logical Interface</b> | Specifies the identifier and unit of the AE interface.   |
| <b>Topology Type</b>     | Specifies the bridge configured on the AE.   |
| <b>Local State</b>       | Specifies if the local device is up or down.   |
| <b>Peer State</b>        | Specifies if the peer device is up or down.  |
| <b>Peer Ip/MCP/State</b> | Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet (MC-AE) interfaces that are part of the peer. |

## Sample Output

### show interfaces mc-ae

```

user@host> show interfaces mc-ae ae1 512
Member Link           : ae0
Current State Machine's State: mcae active state
Local Status          : active
Local State           : up
Peer Status           : active
Peer State            : up
    Logical Interface  : ae0.0
    Topology Type      : bridge
    Local State        : up
    Peer State         : up
    Peer Ip/MCP/State  : 3.3.3.2 ae1.0 up

```

## show interfaces queue

---

**Syntax**    show interfaces queue  
              <aggregate | remaining-traffic>  
              <both-ingress-egress>  
              <egress>  
              <forwarding-class *forwarding-class*>  
              <ingress>  
              <interface-name *interface-name*>  
              <l2-statistics>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              **both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              **l2-statistics** option introduced in Junos OS Release 12.1.

**Description**    Display class-of-service (CoS) queue information for physical interfaces.

**Options**    **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

**remaining-traffic**—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 257 on page 2813](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 257 on page 2813](#).

Table 257: Layer 2 Overhead, Transmitted Packets/Bytes

| Protocol       | Fragmentation       |                                   | LFI |
|----------------|---------------------|-----------------------------------|-----|
|                | First fragmentation | Second to <i>n</i> fragmentations |     |
|                | Bytes               | Bytes                             |     |
| MLPPP (Long)   | 13                  | 12                                | 8   |
| MLPPP (short)  | 11                  | 10                                | 8   |
| MLFR (FRF15)   | 12                  | 10                                | 8   |
| MFR (FRF16)    | 10                  | 8                                 | -   |
| MCMLPPP(Long)  | 13                  | 12                                | -   |
| MCMLPPP(Short) | 11                  | 10                                | -   |

## Layer 2 Statistics—Fragmentation Overhead Calculation

## MLPPP/MC-MLPPP Overhead details:

=====

## Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

## Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

## MLFR (FRF15) Overhead details:

=====

## Fragment 1:

```

Framereley header         : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

## Fragments 2 ...n :

```

Framereley header         : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

## MFR (FRF16) Overhead details:

=====

```
Fragment 1:
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS    : 4 bytes

Fragments 2 ...n :
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  HDLC flag and FCS    : 4 bytes
```

## Overhead with LFI

```
MLPPP(Long & short sequence):
=====
  Outer PPP header      : 4 bytes
  HDLC flag and FCS     : 4 bytes
```

```
MLFR (FRF15):
=====
  Framereelay header    : 2 bytes
  Control,NLPID         : 2 bytes
  HDLC flag and FCS     : 4 bytes
```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

## Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.



On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

**Required Privilege Level**

view

**List of Sample Output**

[show interfaces queue \(Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC\) on page 2820](#)  
[show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 2821](#)  
[show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 2823](#)  
[show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 2823](#)  
[show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 2824](#)  
[show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 2828](#)  
[show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 2831](#)  
[show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 2833](#)  
[show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 2834](#)  
[show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 2835](#)  
[show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 2838](#)  
[show interfaces queue \(QFX Series\) on page 2848](#)  
[show interfaces queue l2-statistics \(lsq interface\) on page 2849](#)  
[show interfaces queue lsq \(lsq-ifd\) on page 2849](#)

**Output Fields** Table 258 on page 2816 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

**Table 258: show interfaces queue Output Fields**

| Field Name  | Field Description   |
|---|---|
| Physical interface  | Name of the physical interface.   |
| Enabled   | State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .  |
| Interface index   | Physical interface's index number, which reflects its initialization sequence.  |
| SNMP ifindex  | SNMP index number for the interface.  |
| Forwarding classes supported  | Total number of forwarding classes supported on the specified interface.  |
| Forwarding classes in use   | Total number of forwarding classes in use on the specified interface.   |
| Ingress queues supported  | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.   |
| Ingress queues in use   | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.  |
| Output queues supported   | Total number of output queues supported on the specified interface.   |
| Output queues in use  | Total number of output queues in use on the specified interface.  |
| Egress queues supported   | Total number of egress queues supported on the specified interface.   |
| Egress queues in use  | Total number of egress queues in use on the specified interface.  |
| Queue counters (Ingress)  | CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> |
| Burst size  | (Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.   |
| The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command: |   |
| Queue   | Queue number.   |
| Forwarding classes  | Forwarding class name.  |

Table 258: show interfaces queue Output Fields (*continued*)

| Field Name                  | Field Description   |
|-----------------------------|---|
| <b>Queued Packets</b>       | <p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p> |
| <b>Queued Bytes</b>         | <p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 259 on page 2819</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>   |
| <b>Transmitted Packets</b>  | <p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 2812</a></p>  |
| <b>Transmitted Bytes</b>    | <p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 259 on page 2819</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 2812</a></p>                                       |
| <b>Tail-dropped packets</b> | Number of packets dropped because of tail drop.   |
| <b>RL-dropped packets</b>   | <p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>   |
| <b>RL-dropped bytes</b>     | <p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>   |

Table 258: show interfaces queue Output Fields (*continued*)

| Field Name          | Field Description   |
|---------------------|---|
| RED-dropped packets | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| RED-dropped bytes   | <p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 259 on page 2819</a>.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> <li>(J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority bytes dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority bytes dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>                 |

Byte counts vary by interface hardware. [Table 259 on page 2819](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware.

[Table 259 on page 2819](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 259: Byte Count by Interface Hardware

| Interface Hardware               | Output Level                | Byte Count Includes  | Comments   |
|----------------------------------|-----------------------------|--|--|
| Gigabit Ethernet IQ and IQE PICs | Interface                   | <p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>  | <p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p> |
|                                  | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>   | —  |
| Non-IQ PIC                       | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted.</li> </ul> | <p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>  |

Table 259: Byte Count by Interface Hardware (*continued*)

| Interface Hardware                                   | Output Level                | Byte Count Includes   | Comments   |
|--|-----------------------------|---|--|
| IQ and IQE PICs with a SONET/SDH interface           | Interface                   | <p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>   | The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.   |
|  | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>  | For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.                      |
| Non-IQ PIC with a SONET/SDH interface                | Interface                   | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul> | For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP). |
| Interfaces configured with Frame Relay Encapsulation | Interface                   | The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.   |  |
| 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs        | Interface                   | <p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>  | The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.   |
| 4-port 1G IQ2 and IQ2-E PICs                         | Packet forwarding component | Queued: 478 bytes of Layer 3 packet.  | —  |
| 8-port 1G IQ2 and IQ2-E PICs                         |                             | Transmitted: 478 bytes of Layer 3 packet.   |  |

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
  Transmitted:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
    Tail-dropped packets :              0              0 pps
    RL-dropped packets  :          205050862        602295 pps
    RL-dropped bytes    :          13595326612      327648832 bps
    RED-dropped packets :              0              0 pps
      Low              :              0              0 pps
      Medium-low       :              0              0 pps
      Medium-high      :              0              0 pps
      High             :              0              0 pps
    RED-dropped bytes   :              0              0 bps
      Low              :              0              0 bps
      Medium-low       :              0              0 bps
      Medium-high      :              0              0 bps
      High             :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps
```

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :              5              0 pps
    Bytes            :              242              0 bps
  Transmitted:
    Packets          :              5              0 pps
    Bytes            :              242              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets :              0              0 pps
    RED-dropped bytes   :              0              0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets          :          42603765          595484 pps
```

```

Bytes          :          5453281920          609776496 bps
Transmitted:
Packets        :          42603765          595484 pps
Bytes          :          5453281920          609776496 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets        :          45          0 pps
Bytes          :          3930          0 bps
Transmitted:
Packets        :          45          0 pps
Bytes          :          3930          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 4, Forwarding classes: af11
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 5, Forwarding classes: ef11
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets        :          31296413          437436 pps
Bytes          :          4005940864          447935200 bps
Transmitted:
Packets        :          31296413          437436 pps
Bytes          :          4005940864          447935200 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps

```



```

Transmitted:
Packets      :                0                0 pps
Bytes        :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps

```

#### show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :                5240762                3404 pps
    Bytes        :            3020710354            15934544 bps
  Transmitted:
    Packets      :                5240762                3404 pps
    Bytes        :            3020710354            15934544 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                0                0 pps
    RED-dropped bytes :                0                0 bps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                0                0 pps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :                2480391                1650 pps
    Bytes        :            1304685666            6945704 bps
  Transmitted:
    Packets      :                2478740                1650 pps
    Bytes        :            1303817240            6945704 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                1651                0 pps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                1651                0 pps
    RED-dropped bytes :                868426                0 bps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                868426                0 pps

```

#### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:

```

```

Packets      :      13      0 pps
Bytes        :      622      0 bps
Transmitted:
Packets      :      13      0 pps
Bytes        :      622      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: af1
Queued:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Transmitted:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Transmitted:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Transmitted:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps

```

#### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets      :      148450735      947295 pps
Bytes        :      8016344944      409228848 bps
Transmitted:
Packets      :      76397439      487512 pps
Bytes        :      4125461868      210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets :      72053285      459783 pps
Low          :      72053285      459783 pps
Medium-low   :      0      0 pps
Medium-high  :      0      0 pps
High         :      0      0 pps
RED-dropped bytes  :      3890877444      198626472 bps

```

```

Low : 3890877444 198626472 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 410278257 473940 pps
Bytes : 22156199518 204742296 bps
Transmitted:
Packets : 4850003 4033 pps
Bytes : 261900162 1742256 bps
Tail-dropped packets : Not Available
RED-dropped packets : 405425693 469907 pps
Low : 405425693 469907 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 21892988124 203000040 bps
Low : 21892988124 203000040 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```

```

Queued:
  Packets      :          76605230          485376 pps
  Bytes       :          5209211400        264044560 bps
Transmitted:
  Packets      :          76444631          484336 pps
  Bytes       :          5198235612        263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets :          160475          1040 pps
  Low         :          160475          1040 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :          10912300        565760 bps
  Low           :          10912300        565760 bps
  Medium-low    :              0              0 bps
  Medium-high   :              0              0 bps
  High          :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Transmitted:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Tail-dropped packets : Not Available
RED-dropped packets :              0              0 pps
  Low         :              0              0 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :              0              0 bps
  Low         :              0              0 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          4836136          3912 pps
  Bytes       :          333402032        2139056 bps
Transmitted:
  Packets      :          3600866          1459 pps
  Bytes       :          244858888        793696 bps
Tail-dropped packets : Not Available
RED-dropped packets :          1225034          2450 pps
  Low         :          1225034          2450 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :          83302312        1333072 bps
  Low         :          83302312        1333072 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Transmitted:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Tail-dropped packets : Not Available

```

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
| Low                 | : | 0 | 0 pps |
| Medium-low          | : | 0 | 0 pps |
| Medium-high         | : | 0 | 0 pps |
| High                | : | 0 | 0 pps |
| RED-dropped bytes   | : | 0 | 0 bps |
| Low                 | : | 0 | 0 bps |
| Medium-low          | : | 0 | 0 bps |
| Medium-high         | : | 0 | 0 bps |
| High                | : | 0 | 0 bps |

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

##### Queued:

|         |   |            |               |
|---------|---|------------|---------------|
| Packets | : | 77059796   | 486384 pps    |
| Bytes   | : | 3544750624 | 178989576 bps |

##### Transmitted:

|                      |   |            |               |
|----------------------|---|------------|---------------|
| Packets              | : | 77059797   | 486381 pps    |
| Bytes                | : | 3544750670 | 178988248 bps |
| Tail-dropped packets | : | 0          | 0 pps         |
| RED-dropped packets  | : | 0          | 0 pps         |
| Low                  | : | 0          | 0 pps         |
| Medium-low           | : | 0          | 0 pps         |
| Medium-high          | : | 0          | 0 pps         |
| High                 | : | 0          | 0 pps         |
| RED-dropped bytes    | : | 0          | 0 bps         |
| Low                  | : | 0          | 0 bps         |
| Medium-low           | : | 0          | 0 bps         |
| Medium-high          | : | 0          | 0 bps         |
| High                 | : | 0          | 0 bps         |

Queue: 1, Forwarding classes: expedited-forwarding

##### Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

##### Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 2, Forwarding classes: assured-forwarding

##### Queued:

|         |   |           |             |
|---------|---|-----------|-------------|
| Packets | : | 4846580   | 3934 pps    |
| Bytes   | : | 222942680 | 1447768 bps |

##### Transmitted:

|                      |   |           |             |
|----------------------|---|-----------|-------------|
| Packets              | : | 4846580   | 3934 pps    |
| Bytes                | : | 222942680 | 1447768 bps |
| Tail-dropped packets | : | 0         | 0 pps       |
| RED-dropped packets  | : | 0         | 0 pps       |
| Low                  | : | 0         | 0 pps       |
| Medium-low           | : | 0         | 0 pps       |
| Medium-high          | : | 0         | 0 pps       |

```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
  Transmitted:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps

```

```

    RED-dropped bytes      :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                :                7055              1 pps
    Bytes                  :            451552              512 bps
  Transmitted:
    Packets                :                7055              1 pps
    Bytes                  :            451552              512 bps
    Tail-dropped packets : Not Available
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :                1031              0 pps
    Bytes                  :            143292              0 bps
  Transmitted:
    Packets                :                1031              0 pps
    Bytes                  :            143292              0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes      :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes      :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes      :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                :                77009             11 pps
    Bytes                  :            6894286             7888 bps
  Transmitted:
    Packets                :                77009             11 pps
    Bytes                  :            6894286             7888 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes      :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps

```

## Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

## Queued:

|         |   |        |       |
|---------|---|--------|-------|
| Packets | : | 1031   | 0 pps |
| Bytes   | : | 147328 | 0 bps |

## Transmitted:

|                      |   |        |       |
|----------------------|---|--------|-------|
| Packets              | : | 1031   | 0 pps |
| Bytes                | : | 147328 | 0 bps |
| Tail-dropped packets | : | 0      | 0 pps |
| RED-dropped packets  | : | 0      | 0 pps |
| Low, non-TCP         | : | 0      | 0 pps |
| Low, TCP             | : | 0      | 0 pps |
| High, non-TCP        | : | 0      | 0 pps |
| High, TCP            | : | 0      | 0 pps |
| RED-dropped bytes    | : | 0      | 0 bps |
| Low, non-TCP         | : | 0      | 0 bps |
| Low, TCP             | : | 0      | 0 bps |
| High, non-TCP        | : | 0      | 0 bps |
| High, TCP            | : | 0      | 0 bps |

Queue: 1, Forwarding classes: expedited-forwarding

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low, non-TCP         | : | 0 | 0 pps |
| Low, TCP             | : | 0 | 0 pps |
| High, non-TCP        | : | 0 | 0 pps |
| High, TCP            | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low, non-TCP         | : | 0 | 0 bps |
| Low, TCP             | : | 0 | 0 bps |
| High, non-TCP        | : | 0 | 0 bps |
| High, TCP            | : | 0 | 0 bps |

Queue: 2, Forwarding classes: assured-forwarding

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low, non-TCP         | : | 0 | 0 pps |
| Low, TCP             | : | 0 | 0 pps |
| High, non-TCP        | : | 0 | 0 pps |
| High, TCP            | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low, non-TCP         | : | 0 | 0 bps |
| Low, TCP             | : | 0 | 0 bps |
| High, non-TCP        | : | 0 | 0 bps |
| High, TCP            | : | 0 | 0 bps |

Queue: 3, Forwarding classes: network-control

## Queued:

|         |   |          |          |
|---------|---|----------|----------|
| Packets | : | 94386    | 12 pps   |
| Bytes   | : | 13756799 | 9568 bps |

## Transmitted:



|                      |   |          |          |
|----------------------|---|----------|----------|
| Packets              | : | 94386    | 12 pps   |
| Bytes                | : | 13756799 | 9568 bps |
| Tail-dropped packets | : | 0        | 0 pps    |
| RED-dropped packets  | : | 0        | 0 pps    |
| Low, non-TCP         | : | 0        | 0 pps    |
| Low, TCP             | : | 0        | 0 pps    |
| High, non-TCP        | : | 0        | 0 pps    |
| High, TCP            | : | 0        | 0 pps    |
| RED-dropped bytes    | : | 0        | 0 bps    |
| Low, non-TCP         | : | 0        | 0 bps    |
| Low, TCP             | : | 0        | 0 bps    |
| High, non-TCP        | : | 0        | 0 bps    |
| High, TCP            | : | 0        | 0 bps    |

### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                254                0 pps
    Bytes        :            16274                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps

```

```

    RED-dropped bytes      :                0          0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                3          0 pps
    Bytes                  :               126          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :             80564692          0 pps
    Bytes                  :          3383717100          0 bps
  Transmitted:
    Packets                :             80564692          0 pps
    Bytes                  :          3383717100          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :             80564685          0 pps
    Bytes                  :          3383716770          0 bps
  Transmitted:
    Packets                :             80564685          0 pps

```

```

Bytes : 3383716770 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Transmitted:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 288 0 pps
Bytes : 18450 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
  Transmitted:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

#### show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
  Transmitted:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :      0      0 pps
    Low          :      0      0 pps

```

```

Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps

```

```

      High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 109355853 471736 pps
    Bytes : 7436199152 256627968 bps
  Transmitted:
    Packets : 109355852 471736 pps
    Bytes : 7436198640 256627968 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps

```

```
Transmitted:
Packets      :                0                0 pps
Bytes        :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low        :                0                0 pps
  Medium-low :                0                0 pps
  Medium-high:                0                0 pps
  High       :                0                0 pps
RED-dropped bytes :                0                0 bps
  Low        :                0                0 bps
  Medium-low :                0                0 bps
  Medium-high:                0                0 bps
  High       :                0                0 bps
```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```
user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

  Interface index: 192, SNMP ifIndex: 1948

  Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
  Lam

  Forwarding classes: 16 supported, 9 in use

  Egress queues: 8 supported, 8 in use

  Queue: 0, Forwarding classes: DEFAULT

  Queued:

    Packets      :                214886                13449 pps
    Bytes        :                9884756                5164536 bps

  Transmitted:

    Packets      :                214886                13449 pps
    Bytes        :                9884756                5164536 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
      Low        :                0                0 pps
      Medium-low :                0                0 pps
      Medium-high:                0                0 pps
      High       :                0                0 pps
    RED-dropped bytes :                0                0 bps
      Low        :                0                0 bps
      Medium-low :                0                0 bps
```



|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 bps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

Queue: 1, Forwarding classes: REALTIME

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 pps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 pps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 pps |
|------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 bps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 bps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 bps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

Queue: 2, Forwarding classes: PRIVATE

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| Medium-low        | : | 0 | 0 pps |
| Medium-high       | : | 0 | 0 pps |
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 3, Forwarding classes: CONTROL

Queued:

|         |   |      |       |
|---------|---|------|-------|
| Packets | : | 60   | 0 pps |
| Bytes   | : | 4560 | 0 bps |

Transmitted:

|                      |   |      |       |
|----------------------|---|------|-------|
| Packets              | : | 60   | 0 pps |
| Bytes                | : | 4560 | 0 bps |
| Tail-dropped packets | : | 0    | 0 pps |
| RED-dropped packets  | : | 0    | 0 pps |
| Low                  | : | 0    | 0 pps |
| Medium-low           | : | 0    | 0 pps |
| Medium-high          | : | 0    | 0 pps |
| High                 | : | 0    | 0 pps |
| RED-dropped bytes    | : | 0    | 0 bps |
| Low                  | : | 0    | 0 bps |
| Medium-low           | : | 0    | 0 bps |
| Medium-high          | : | 0    | 0 bps |
| High                 | : | 0    | 0 bps |

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |

|                   |   |   |       |
|-------------------|---|---|-------|
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

#### Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

##### Queued:

|         |   |          |             |
|---------|---|----------|-------------|
| Packets | : | 371365   | 23620 pps   |
| Bytes   | : | 15597330 | 7936368 bps |

##### Transmitted:

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Packets              | : | 371365   | 23620 pps   |
| Bytes                | : | 15597330 | 7936368 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |
| High                 | : | 0        | 0 pps       |
| RED-dropped bytes    | : | 0        | 0 bps       |
| Low                  | : | 0        | 0 bps       |
| Medium-low           | : | 0        | 0 bps       |
| Medium-high          | : | 0        | 0 bps       |
| High                 | : | 0        | 0 bps       |

Queue: 1, Forwarding classes: REALTIME

##### Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|                                       |   |   |       |
|---------------------------------------|---|---|-------|
| Bytes                                 | : | 0 | 0 bps |
| Transmitted:                          |   |   |       |
| Packets                               | : | 0 | 0 pps |
| Bytes                                 | : | 0 | 0 bps |
| Tail-dropped packets                  | : | 0 | 0 pps |
| RED-dropped packets                   | : | 0 | 0 pps |
| Low                                   | : | 0 | 0 pps |
| Medium-low                            | : | 0 | 0 pps |
| Medium-high                           | : | 0 | 0 pps |
| High                                  | : | 0 | 0 pps |
| RED-dropped bytes                     | : | 0 | 0 bps |
| Low                                   | : | 0 | 0 bps |
| Medium-low                            | : | 0 | 0 bps |
| Medium-high                           | : | 0 | 0 bps |
| High                                  | : | 0 | 0 bps |
| Queue: 2, Forwarding classes: PRIVATE |   |   |       |
| Queued:                               |   |   |       |
| Packets                               | : | 0 | 0 pps |
| Bytes                                 | : | 0 | 0 bps |
| Transmitted:                          |   |   |       |
| Packets                               | : | 0 | 0 pps |
| Bytes                                 | : | 0 | 0 bps |
| Tail-dropped packets                  | : | 0 | 0 pps |
| RED-dropped packets                   | : | 0 | 0 pps |
| Low                                   | : | 0 | 0 pps |
| Medium-low                            | : | 0 | 0 pps |
| Medium-high                           | : | 0 | 0 pps |
| High                                  | : | 0 | 0 pps |
| RED-dropped bytes                     | : | 0 | 0 bps |
| Low                                   | : | 0 | 0 bps |
| Medium-low                            | : | 0 | 0 bps |

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 bps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

Queue: 3, Forwarding classes: CONTROL

Queued:

|         |   |       |       |
|---------|---|-------|-------|
| Packets | : | 32843 | 0 pps |
|---------|---|-------|-------|

|       |   |         |        |
|-------|---|---------|--------|
| Bytes | : | 2641754 | 56 bps |
|-------|---|---------|--------|

Transmitted:

|         |   |       |       |
|---------|---|-------|-------|
| Packets | : | 32843 | 0 pps |
|---------|---|-------|-------|

|       |   |         |        |
|-------|---|---------|--------|
| Bytes | : | 2641754 | 56 bps |
|-------|---|---------|--------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 pps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 pps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 pps |
|------|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| RED-dropped bytes | : | 0 | 0 bps |
|-------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 bps |
|-----|---|---|-------|

|            |   |   |       |
|------------|---|---|-------|
| Medium-low | : | 0 | 0 bps |
|------------|---|---|-------|

|             |   |   |       |
|-------------|---|---|-------|
| Medium-high | : | 0 | 0 bps |
|-------------|---|---|-------|

|      |   |   |       |
|------|---|---|-------|
| High | : | 0 | 0 bps |
|------|---|---|-------|

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

Transmitted:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

|       |   |   |       |
|-------|---|---|-------|
| Bytes | : | 0 | 0 bps |
|-------|---|---|-------|

|                      |   |   |       |
|----------------------|---|---|-------|
| Tail-dropped packets | : | 0 | 0 pps |
|----------------------|---|---|-------|

|                     |   |   |       |
|---------------------|---|---|-------|
| RED-dropped packets | : | 0 | 0 pps |
|---------------------|---|---|-------|

|     |   |   |       |
|-----|---|---|-------|
| Low | : | 0 | 0 pps |
|-----|---|---|-------|

|                   |   |   |       |
|-------------------|---|---|-------|
| Medium-low        | : | 0 | 0 pps |
| Medium-high       | : | 0 | 0 pps |
| High              | : | 0 | 0 pps |
| RED-dropped bytes | : | 0 | 0 bps |
| Low               | : | 0 | 0 bps |
| Medium-low        | : | 0 | 0 bps |
| Medium-high       | : | 0 | 0 bps |
| High              | : | 0 | 0 bps |

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:



|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |

High : 0 0 bps

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available

```

|                        |   |       |
|------------------------|---|-------|
| Total-dropped packets: | 0 | 0 pps |
| Total-dropped bytes :  | 0 | 0 bps |

### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           1           0 pps
    Bytes        :        1001           0 bps
  Transmitted:
    Packets      :           5           0 pps
    Bytes        :        1062           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :           1           0 pps
    Bytes        :        1500           0 bps
  Transmitted:
    Packets      :           6           0 pps
    Bytes        :       1573           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :           1           0 pps
    Bytes        :         512           0 bps
  Transmitted:
    Packets      :           3           0 pps
    Bytes        :         549           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
=====

```

### show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0

```

## Queue: 0, Forwarding classes: be

## Queued:

|         |   |          |             |
|---------|---|----------|-------------|
| Packets | : | 55576    | 1206 pps    |
| Bytes   | : | 29622008 | 5145472 bps |

## Transmitted:

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Packets              | : | 55576    | 1206 pps    |
| Bytes                | : | 29622008 | 5145472 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RL-dropped packets   | : | 0        | 0 pps       |
| RL-dropped bytes     | : | 0        | 0 bps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |
| High                 | : | 0        | 0 pps       |
| RED-dropped bytes    | : | 0        | 0 bps       |
| Low                  | : | 0        | 0 bps       |
| Medium-low           | : | 0        | 0 bps       |
| Medium-high          | : | 0        | 0 bps       |
| High                 | : | 0        | 0 bps       |

## Queue: 1, Forwarding classes: ef

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RL-dropped packets   | : | 0 | 0 pps |
| RL-dropped bytes     | : | 0 | 0 bps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 2, Forwarding classes: af

## Queued:

|         |   |   |       |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes   | : | 0 | 0 bps |

## Transmitted:

|                      |   |   |       |
|----------------------|---|---|-------|
| Packets              | : | 0 | 0 pps |
| Bytes                | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |
| RL-dropped packets   | : | 0 | 0 pps |
| RL-dropped bytes     | : | 0 | 0 bps |
| RED-dropped packets  | : | 0 | 0 pps |
| Low                  | : | 0 | 0 pps |
| Medium-low           | : | 0 | 0 pps |
| Medium-high          | : | 0 | 0 pps |
| High                 | : | 0 | 0 pps |
| RED-dropped bytes    | : | 0 | 0 bps |
| Low                  | : | 0 | 0 bps |
| Medium-low           | : | 0 | 0 bps |
| Medium-high          | : | 0 | 0 bps |
| High                 | : | 0 | 0 bps |

## Queue: 3, Forwarding classes: nc

|                      |   |          |             |
|----------------------|---|----------|-------------|
| Queued:              |   |          |             |
| Packets              | : | 22231    | 482 pps     |
| Bytes                | : | 11849123 | 2057600 bps |
| Transmitted:         |   |          |             |
| Packets              | : | 22231    | 482 pps     |
| Bytes                | : | 11849123 | 2057600 bps |
| Tail-dropped packets | : | 0        | 0 pps       |
| RL-dropped packets   | : | 0        | 0 pps       |
| RL-dropped bytes     | : | 0        | 0 bps       |
| RED-dropped packets  | : | 0        | 0 pps       |
| Low                  | : | 0        | 0 pps       |
| Medium-low           | : | 0        | 0 pps       |
| Medium-high          | : | 0        | 0 pps       |
| High                 | : | 0        | 0 pps       |
| RED-dropped bytes    | : | 0        | 0 bps       |
| Low                  | : | 0        | 0 bps       |
| Medium-low           | : | 0        | 0 bps       |
| Medium-high          | : | 0        | 0 bps       |
| High                 | : | 0        | 0 bps       |

## show interfaces xe

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>show interfaces <i>device-name:type-fpc/pic/port</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;descriptions&gt;</code><br><code>&lt;media&gt;</code><br><code>&lt;routing-instance (all   <i>instance-name</i>)&gt;</code><br><code>&lt;snmp-index <i>snmp-index</i>&gt;</code><br><code>&lt;statistics&gt;</code>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.   |
| <b>Options</b>                  | <p><b><i>device-name:type-fpc/pic/port</i></b>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p> |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Monitoring Interface Status and Traffic on page 335</a></li><li>• <a href="#">Troubleshooting Network Interfaces on page 1234</a></li><li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 1234</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li></ul>  |
| <b>List of Sample Output</b>    | <a href="#">show interfaces on page 2860</a><br><a href="#">show interfaces (Asymmetric Flow Control) on page 2861</a><br><a href="#">show interfaces brief on page 2861</a><br><a href="#">show interfaces detail on page 2861</a><br><a href="#">show interfaces detail (Asymmetric Flow Control) on page 2863</a><br><a href="#">show interfaces extensive on page 2864</a><br><a href="#">show interfaces extensive (Asymmetric Flow Control) on page 2866</a>   |

[show interfaces terse on page 2868](#)

[show interfaces \(QFabric System\) on page 2868](#)

**Output Fields** Table 260 on page 2853 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

**Table 260: show interfaces xe Output Fields**

| Field Name  | Field Description   | Level of Output              |
|---|---|------------------------------|
| <b>Physical Interface</b>   |   |                              |
| <b>Physical interface</b>   | Name of the physical interface.   | All levels                   |
| <b>Enabled</b>  | State of the interface.   | All levels                   |
| <b>Interface index</b>  | Index number of the physical interface, which reflects its initialization sequence.   | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>   | SNMP index number for the physical interface.   | <b>detail extensive none</b> |
| <b>Generation</b>   | Unique number for use by Juniper Networks technical support only.   | <b>detail extensive</b>      |
| <b>Link-level type</b>  | Encapsulation being used on the physical interface.   | All levels                   |
| <b>MTU</b>  | Maximum transmission unit size on the physical interface.   | All levels                   |
| <b>Speed</b>  | Speed at which the interface is running.  | All levels                   |
| <b>Duplex</b>   | Duplex mode of the interface, either <b>Full-Duplex</b> or <b>Half-Duplex</b> .   | All levels                   |
| <b>Loopback</b>   | Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .  | All levels                   |
| <b>Source filtering</b>   | Source filtering status: <b>Enabled</b> or <b>Disabled</b> .  | All levels                   |
| <b>LAN-PHY mode</b>   | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels                   |
| <b>Unidirectional</b>   | Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.             | All levels                   |
| <b>Flow control</b>   | Flow control status: <b>Enabled</b> or <b>Disabled</b> .  | All levels                   |
| <b>NOTE:</b> This field is only displayed if asymmetric flow control is not configured. |   |                              |

Table 260: show interfaces xe Output Fields (*continued*)

| Field Name                     | Field Description   | Level of Output              |
|--------------------------------|---|------------------------------|
| <b>Configured-flow-control</b> | Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ): <ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.<br/><b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.<br/><b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul> <p><b>NOTE:</b> This field is only displayed if asymmetric flow control is configured.</p> | All levels                   |
| <b>Auto-negotiation</b>        | Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .   | All levels                   |
| <b>Remote-fault</b>            | Remote fault status: <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>   | All levels                   |
| <b>Device flags</b>            | Information about the physical device.  | All levels                   |
| <b>Interface flags</b>         | Information about the interface.  | All levels                   |
| <b>Link flags</b>              | Information about the link.   | All levels                   |
| <b>Wavelength</b>              | Configured wavelength, in nanometers (nm).  | All levels                   |
| <b>Frequency</b>               | Frequency associated with the configured wavelength, in terahertz (THz).  | All levels                   |
| <b>CoS queues</b>              | Number of CoS queues configured.  | <b>detail extensive none</b> |
| <b>Schedulers</b>              | Number of CoS schedulers configured.  | <b>extensive</b>             |
| <b>Hold-times</b>              | Current interface hold-time up and hold-time down, in milliseconds.   | <b>detail extensive</b>      |
| <b>Current address</b>         | Configured MAC address.   | <b>detail extensive none</b> |
| <b>Hardware address</b>        | Hardware MAC address.   | <b>detail extensive none</b> |
| <b>Last flapped</b>            | Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .  | <b>detail extensive none</b> |
| <b>Input Rate</b>              | Input rate in bits per second (bps) and packets per second (pps).   | None specified               |
| <b>Output Rate</b>             | Output rate in bps and pps.   | None specified               |
| <b>Statistics last cleared</b> | Time when the statistics for the interface were last set to zero.   | <b>detail extensive</b>      |



Table 260: show interfaces xe Output Fields (*continued*)

| Field Name                | Field Description   | Level of Output         |
|---------------------------|---|-------------------------|
| <b>Traffic statistics</b> | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>  | <b>detail extensive</b> |
| <b>Input errors</b>       | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |

Table 260: show interfaces xe Output Fields (*continued*)

| Field Name                      | Field Description  | Level of Output         |
|---------------------------------|--|-------------------------|
| <b>Output errors</b>            | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul> | <b>extensive</b>        |
| <b>Egress queues</b>            | Total number of egress queues supported on the specified interface.  | <b>detail extensive</b> |
| <b>Queue counters (Egress)</b>  | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>  | <b>detail extensive</b> |
| <b>Queue Number</b>             | The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.   | <b>detail extensive</b> |
| <b>Ingress queues</b>           | Total number of ingress queues supported on the specified interface.   | <b>extensive</b>        |
| <b>Queue counters (Ingress)</b> | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>  | <b>extensive</b>        |

Table 260: show interfaces xe Output Fields (*continued*)

| Field Name                              | Field Description  | Level of Output              |
|---|--|------------------------------|
| <b>Active alarms and Active defects</b> | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>  | <b>detail extensive none</b> |
| <b>PCS statistics</b>                   | Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.   | <b>detail extensive</b>      |
| <b>MAC statistics</b>                   | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul> | <b>extensive</b>             |
| <b>Filter statistics</b>                | Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.  | <b>extensive</b>             |

Table 260: show interfaces xe Output Fields (*continued*)

| Field Name                  | Field Description  | Level of Output |
|-----------------------------|--|-----------------|
| Autonegotiation information | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive). For asymmetric <b>PAUSE</b>, shows if the <b>PAUSE</b> transmit and <b>PAUSE</b> receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul> | extensive       |

Table 260: show interfaces xe Output Fields (*continued*)

| Field Name                                    | Field Description  | Level of Output              |
|---|--|------------------------------|
| <b>Packet Forwarding Engine configuration</b> | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul> | <b>extensive</b>             |
| <b>Logical Interface</b>                      |  |                              |
| <b>Logical interface</b>                      | Name of the logical interface.   | All levels                   |
| <b>Index</b>                                  | Index number of the logical interface, which reflects its initialization sequence.   | <b>detail extensive none</b> |
| <b>SNMP ifIndex</b>                           | SNMP interface index number for the logical interface.   | <b>detail extensive none</b> |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.  | <b>detail extensive</b>      |
| <b>Flags</b>                                  | Information about the logical interface.   | All levels                   |
| <b>Encapsulation</b>                          | Encapsulation on the logical interface.  | All levels                   |
| <b>Protocol</b>                               | Protocol family.   | <b>detail extensive none</b> |
| <b>Traffic statistics</b>                     | Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.   | <b>detail extensive</b>      |
| <b>IPv6 transit statistics</b>                | If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.   | <b>extensive</b>             |
| <b>Local statistics</b>                       | Number and rate of bytes and packets destined to and from the switch.  | <b>extensive</b>             |
| <b>Transit statistics</b>                     | Number and rate of bytes and packets transiting the switch.  | <b>extensive</b>             |
| <b>Generation</b>                             | Unique number for use by Juniper Networks technical support only.  | <b>detail extensive</b>      |
| <b>Route Table</b>                            | Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.   | <b>detail extensive none</b> |

Table 260: show interfaces xe Output Fields (*continued*)

| Field Name              | Field Description  | Level of Output              |
|-------------------------|--|------------------------------|
| <b>Input Filters</b>    | Names of any input filters applied to this interface.  | <b>detail extensive</b>      |
| <b>Output Filters</b>   | Names of any output filters applied to this interface.   | <b>detail extensive</b>      |
| <b>Flags</b>            | Information about protocol family flags.<br><br>If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled. | <b>detail extensive</b>      |
| <b>Addresses, Flags</b> | Information about the address flags.   | <b>detail extensive none</b> |
| <i>protocol-family</i>  | Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.   | <b>brief</b>                 |
| <b>Flags</b>            | Information about the address flag.  | <b>detail extensive none</b> |
| <b>Destination</b>      | IP address of the remote side of the connection.   | <b>detail extensive none</b> |
| <b>Local</b>            | IP address of the logical interface.   | <b>detail extensive none</b> |
| <b>Broadcast</b>        | Broadcast address of the logical interlace.  | <b>detail extensive none</b> |
| <b>Generation</b>       | Unique number for use by Juniper Networks technical support only.  | <b>detail extensive</b>      |

## Sample Output

### show interfaces

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped  : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0

```

```

Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode

```

### show interfaces (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped  : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode

```

### show interfaces brief

```

user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface xe-0/0/1.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch

```

### show interfaces detail

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1

```

Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)  
 Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)

## Traffic statistics:

|                 |   |       |
|-----------------|---|-------|
| Input bytes :   | 0 | 0 bps |
| Output bytes :  | 0 | 0 bps |
| Input packets:  | 0 | 0 pps |
| Output packets: | 0 | 0 pps |

## IPv6 transit statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

Egress queues: 12 supported, 9 in use

| Queue counters: | Queued packets | Transmitted packets | Dropped packets |
|-----------------|----------------|---------------------|-----------------|
| 0 best-effort   | 0              | 0                   | 0               |
| 1 fc7           | 0              | 0                   | 0               |
| 2 no-loss       | 0              | 0                   | 0               |
| 3 fcoe          | 0              | 0                   | 0               |
| 4 fc4           | 0              | 0                   | 0               |
| 5 fc5           | 0              | 0                   | 0               |
| 6 fc6           | 0              | 0                   | 0               |
| 7 network-cont  | 0              | 0                   | 0               |
| 8 mcast         | 0              | 0                   | 0               |

| Queue number: | Mapped forwarding classes |
|---------------|---------------------------|
| 0             | best-effort               |
| 1             | fc7                       |
| 2             | no-loss                   |
| 3             | fcoe                      |
| 4             | fc4                       |
| 5             | fc5                       |
| 6             | fc6                       |
| 7             | network-control           |
| 8             | mcast                     |

Active alarms : None

Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

## Traffic statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

## Local statistics:

|                 |   |
|-----------------|---|
| Input bytes :   | 0 |
| Output bytes :  | 0 |
| Input packets:  | 0 |
| Output packets: | 0 |

## Transit statistics:

|                |   |       |
|----------------|---|-------|
| Input bytes :  | 0 | 0 bps |
| Output bytes : | 0 | 0 bps |



```

Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces detail (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped  : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
    Input bytes :          0          0 bps
    Output bytes :          0          0 bps
    Input packets:          0          0 pps
    Output packets:          0          0 pps
  IPv6 transit statistics:
    Input bytes :          0
    Output bytes :          0
    Input packets:          0
    Output packets:          0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort          0              0              0
    1 fc7                 0              0              0
    2 no-loss              0              0              0
    3 fcoe                 0              0              0
    4 fc4                  0              0              0
    5 fc5                  0              0              0
    6 fc6                  0              0              0
    7 network-cont        0              0              0
    8 mcast                0              0              0

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                fc7
    2                no-loss
    3                fcoe
    4                fc4
    5                fc5
    6                fc6

```

```

7          network-control
8          mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

#### show interfaces extensive

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

|                |   |   |   |
|----------------|---|---|---|
| 0 best-effort  | 0 | 0 | 0 |
| 1 fc7          | 0 | 0 | 0 |
| 2 no-loss      | 0 | 0 | 0 |
| 3 fcoe         | 0 | 0 | 0 |
| 4 fc4          | 0 | 0 | 0 |
| 5 fc5          | 0 | 0 | 0 |
| 6 fc6          | 0 | 0 | 0 |
| 7 network-cont | 0 | 0 | 0 |
| 8 mcast        | 0 | 0 | 0 |

Queue number:            Mapped forwarding classes

|   |                 |
|---|-----------------|
| 0 | best-effort     |
| 1 | fc7             |
| 2 | no-loss         |
| 3 | fcoe            |
| 4 | fc4             |
| 5 | fc5             |
| 6 | fc6             |
| 7 | network-control |
| 8 | mcast           |

Active alarms : None

Active defects : None

MAC statistics:

|                    | Receive | Transmit |
|--------------------|---------|----------|
| Total octets       | 0       | 0        |
| Total packets      | 0       | 0        |
| Unicast packets    | 0       | 0        |
| Broadcast packets  | 0       | 0        |
| Multicast packets  | 0       | 0        |
| CRC/Align errors   | 0       | 0        |
| FIFO errors        | 0       | 0        |
| MAC control frames | 0       | 0        |
| MAC pause frames   | 0       | 0        |
| Oversized frames   | 0       |          |
| Jabber frames      | 0       |          |
| Fragment frames    | 0       |          |
| VLAN tagged frames | 0       |          |
| Code violations    | 0       |          |

MAC Priority Flow Control Statistics:

|              |   |   |
|--------------|---|---|
| Priority : 0 | 0 | 0 |
| Priority : 1 | 0 | 0 |
| Priority : 2 | 0 | 0 |
| Priority : 3 | 0 | 0 |
| Priority : 4 | 0 | 0 |
| Priority : 5 | 0 | 0 |
| Priority : 6 | 0 | 0 |
| Priority : 7 | 0 | 0 |

Filter statistics:

|                      |   |   |
|----------------------|---|---|
| Input packet count   | 0 |   |
| Input packet rejects | 0 |   |
| Input DA rejects     | 0 |   |
| Input SA rejects     | 0 |   |
| Output packet count  |   | 0 |

```

Output packet pad count          0
Output packet error count        0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      %      bps      %      usec
0 best-effort      75      7500000000      75      0      low
none
7 network-control      5      500000000      5      0      low
none
8 mcast      20      2000000000      20      0      low
none

```

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)

Flags: SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Local statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Transit statistics:

```

Input bytes :      0      0 bps
Output bytes :      0      0 bps
Input packets:      0      0 pps
Output packets:      0      0 pps

```

Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0

Flags: Trunk-Mode

### show interfaces extensive (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1 extensive
```

Physical interface: xe-0/0/1, Enabled, Physical link is Up

Interface index: 49195, SNMP ifIndex: 591, Generation: 169

Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU

Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

Disabled,

Configured-flow-control tx-buffers: off rx-buffers: on

Device flags : Present Running

Interface flags: SNMP-Traps Internal: 0x0

Link flags : None

CoS queues : 12 supported, 12 maximum usable queues

Hold-times : Up 0 ms, Down 0 ms

Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1

Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)

Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)

Traffic statistics:

```

Input bytes :      0      0 bps
Output bytes :      0      0 bps
Input packets:      0      0 pps
Output packets:      0      0 pps

```

IPv6 transit statistics:

```

Input bytes :      0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          0              0              0
1 fc7                  0              0              0
2 no-loss              0              0              0
3 fcoe                 0              0              0
4 fc4                  0              0              0
5 fc5                  0              0              0
6 fc6                  0              0              0
7 network-cont         0              0              0
8 mcast                0              0              0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive      Transmit
Total packets         0            0
Unicast packets       0            0
Broadcast packets     0            0
Multicast packets     0            0
CRC/Align errors      0            0
FIFO errors           0            0
MAC control frames    0            0
MAC pause frames      0            0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations        0
MAC Priority Flow Control Statistics:
Priority : 0           0            0
Priority : 1           0            0

```

```

Priority : 2          0          0
Priority : 3          0          0
Priority : 4          0          0
Priority : 5          0          0
Priority : 6          0          0
Priority : 7          0          0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort         75    7500000000    75      0    low   none
7 network-control     5     500000000     5      0    low   none
8 mcast              20    2000000000    20      0    low   none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0          0 bps
Output bytes : 0          0 bps
Input packets: 0          0 pps
Output packets: 0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface      Admin Link Proto  Local      Remote

xe-0/0/1       up    up
xe-0/0/1.0     up    up    eth-switch

```

### show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

```
Interface flags: Internal: 0x4000
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped    : Never
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

## show lacp interfaces

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show lacp interfaces</code><br><code>&lt;interface-name&gt;</code>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.   |
| <b>Options</b>                  | <p><code>none</code>—Display LACP information for all interfaces.</p> <p><code>interface-name</code>—(Optional) Display LACP information for the specified interface:</p> <ul style="list-style-type: none"><li>• Aggregated Ethernet—<code>aex</code></li><li>• Gigabit Ethernet—<code>ge-fpc/pic/port</code></li><li>• 10-Gigabit Ethernet—<code>xe-fpc/pic/port</code></li></ul>   |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li><li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462</a></li><li>• <a href="#">Configuring Aggregated Ethernet Links (CLI Procedure)</a></li><li>• <a href="#">Configuring Link Aggregation on page 2593</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP (CLI Procedure)</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 2589</a></li><li>• <a href="#">Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 2393</a></li><li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp interfaces (EX Series Switches) on page 2872</a><br><a href="#">show lacp interfaces (QFX Series) on page 2873</a>  |
| <b>Output Fields</b>            | <a href="#">Table 251 on page 2766</a> lists the output fields for the <code>show lacp interfaces</code> command. Output fields are listed in the approximate order in which they appear.   |



Table 261: show lacp interfaces Output Fields

| Field Name           | Field Description  |
|----------------------|--|
| Aggregated interface | Aggregated Ethernet interface name.  |
| LACP State           | <p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> <li>For a child interface configured with the <b>force-up</b> statement, LACP state displays <b>FUP</b> along with the interface name.</li> <li><b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li><b>Actor</b>—Local device participating in the LACP negotiation.</li> <li><b>Partner</b>—Remote device participating in the LACP negotiation.</li> </ul> </li> <li><b>Exp</b>—Expired state. <b>Yes</b> indicates that the actor or partner is in an expired state. <b>No</b> indicates that the actor or partner is not in an expired state.</li> <li><b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. <b>No</b> indicates that the operational partner information in use has been received in an LACP PDU.</li> <li><b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li><b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li><b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered to be synchronized. The link has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. The link is currently not in the right aggregation.</li> <li><b>Aggr</b>—Ability of the aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li><b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or a fast transmission rate, depending upon the expressed LACP timeout preference (<b>Long Timeout</b> or <b>Short Timeout</b>).</li> <li><b>Activity</b>—Actor's or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> </ul> |

Table 261: show lacp interfaces Output Fields (*continued*)

| Field Name    | Field Description  |
|---------------|--|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li><b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li><b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li><b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li><b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li><b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li><b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li><b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li><b>Transmit State</b>—Transmit state of the state machine. The transmit state is one of the following values: <ul style="list-style-type: none"> <li><b>Fast periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li><b>No periodic</b>—Periodic transmissions are disabled.</li> <li><b>Periodic timer</b>—Transitory state entered when the periodic timer expires.</li> <li><b>Slow periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li><b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> <li><b>Attached</b>—The multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li><b>Collecting—Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li><b>Collecting distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li><b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li><b>Distributing—Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames can be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li><b>Waiting</b>—The multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> </li> </ul> |

## Sample Output

### show lacp interfaces (EX Series Switches)

```

user@switch> show lacp interfaces ae5
Aggregated interface: ae5
LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-2/0/7        Actor No   No   Yes  Yes  Yes  Yes   Fast    Active
xe-2/0/7        Partner No   No   Yes  Yes  Yes  Yes   Fast    Passive

```

|          |         |    |    |    |     |     |     |      |         |
|----------|---------|----|----|----|-----|-----|-----|------|---------|
| xe-4/0/7 | Actor   | No | No | No | No  | No  | Yes | Fast | Active  |
| xe-4/0/7 | Partner | No | No | No | Yes | Yes | Yes | Fast | Passive |

| LACP protocol:     | Receive State | Transmit State | Mux State               |
|--------------------|---------------|----------------|-------------------------|
| xe-2/0/7(Active)   | Current       | Fast periodic  | Collecting distributing |
| xe-34/0/7(Standby) | Current       | Fast periodic  | Waiting                 |

### show lacp interfaces (QFX Series)

```
user@switch> show lacp interfaces nodegroup1:ae0 extensive
```

```
Aggregated interface: nodegroup1:ae0
```

| LACP state:       | Role    | Exp | Def | Dist | CoI | Syn | Aggr | Timeout | Activity |
|-------------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| node1:xe-0/0/1FUP | Actor   |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Active            |         |     |     |      |     |     |      |         |          |
| node1xe-0/0/1FUP  | Partner |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Passive           |         |     |     |      |     |     |      |         |          |
| node2:xe-0/0/2    | Actor   |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Active            |         |     |     |      |     |     |      |         |          |
| node2:xe-0/0/2    | Partner |     | No  | Yes  | No  | No  | No   | Yes     | Fast     |
| Passive           |         |     |     |      |     |     |      |         |          |

|              | LACP protocol:           | Receive State | Transmit State | Mux State  |
|--------------|--------------------------|---------------|----------------|------------|
|              | node1:xe-0/0/1FUP        | Current       | Fast periodic  | Collecting |
| distributing | node2:xe-0/0/2           | Current       | Fast periodic  | Collecting |
| distributing | node1:xe-0/0/1 (active)  | Current       | Fast periodic  | Collecting |
| distributing | node2:xe-0/0/2 (standby) | Current       | Fast periodic  | WAITING    |

## show lacp statistics interfaces (View)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show lacp statistics interfaces</b> <i>interface-name</i>  |
| <b>Release Information</b>      | Command modified in Release 10.2 of Junos OS.<br>Command introduced in Release 11.1 of Junos OS for the QFX Series.   |
| <b>Description</b>              | Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.   |
| <b>Options</b>                  | <i>interface-name</i> —(Optional) Name of an interface.   |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 2593</a></li> <li>• <a href="#">Verifying the Status of a LAG Interface on page 2750</a></li> <li>• <a href="#">Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 2751</a></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462</a></li> <li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 2466</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show lacp statistics interfaces on page 2876</a><br><a href="#">show lacp statistics interfaces (QFX Series) on page 2876</a><br><a href="#">show lacp statistics interfaces (QFabric Systems) on page 2876</a>   |
| <b>Output Fields</b>            | <a href="#">Table 262 on page 2875</a> lists the output fields for the <b>show lacp statistics interfaces</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 262: show lacp statistics interfaces Output Fields**

| Field Name                  | Field Description   |
|-----------------------------|---|
| <b>Aggregated interface</b> | Aggregated interface value.   |
| <b>LACP Statistics</b>      | <p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—LACP transmit counter that increments for each normal hello.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul> |

## Sample Output

### show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-2/0/0              1352        2035          0                0
ge-2/0/1              1352        2056          0                0
ge-2/2/0              1352        2045          0                0
ge-2/2/1              1352        2043          0                0
```

### show lacp statistics interfaces (QFX Series)

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
xe-0/0/2              1352        2035          0                0
xe-0/0/3              1352        2056          0                0
```

### show lacp statistics interfaces (QFabric Systems)

```
user@host> show lacp statistics interfaces nodegroup1:ae0
Aggregated interface: nodegroup1:ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
node1:xe-0/0/2        1352        2035          0                0
node2:xe-0/0/3        1352        2056          0                0
```

## show oam ethernet link-fault-management

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | show oam ethernet link-fault-management<br><brief   detail><br><interface-name>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.4 for EX Series switches.  |
| <b>Description</b>              | Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.  |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —(Optional) Display link fault management information for the specified Ethernet interface only.  |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 2583</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 2604</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet link-fault-management brief on page 2881</a><br><a href="#">show oam ethernet link-fault-management detail on page 2881</a>   |
| <b>Output Fields</b>            | Table 263 on page 2877 lists the output fields for the <b>show oam ethernet link-fault-management</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 263: show oam ethernet link-fault-management Output Fields

| Field Name             | Field Description   | Level of Output |
|------------------------|---|-----------------|
| <b>Status</b>          | Indicates the status of the established link. <ul style="list-style-type: none"> <li>• <b>Fail</b>—A link fault condition exists.</li> <li>• <b>Running</b>—A link fault condition does not exist.</li> </ul>   | All levels      |
| <b>Discovery state</b> | State of the discovery mechanism: <ul style="list-style-type: none"> <li>• <b>Passive Wait</b></li> <li>• <b>Send Any</b></li> <li>• <b>Send Local Remote</b></li> <li>• <b>Send Local Remote Ok</b></li> </ul> | All levels      |
| <b>Peer address</b>    | Address of the OAM peer.  | All levels      |

Table 263: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                       | Field Description  | Level of Output |
|----------------------------------|--|-----------------|
| <b>Flags</b>                     | Information about the interface. <ul style="list-style-type: none"> <li><b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE has either not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li><b>Local-Stable</b>—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li><b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul>  | All levels      |
| <b>Remote loopback status</b>    | Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).  | All levels      |
| <b>Remote entity information</b> | Remote entity information. <ul style="list-style-type: none"> <li><b>Remote MUX action</b>—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs.</li> <li><b>Remote parser action</b>—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs.</li> <li><b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li><b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li><b>Remote loopback mode</b>—Indicates whether remote loopback is supported or not supported.</li> <li><b>Link events</b>—Indicates whether interpreting link events is supported or not supported on the remote peer.</li> <li><b>Variable requests</b>—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.</li> </ul> | All levels      |
| <b>OAM Receive Statistics</b>    |  |                 |
| <b>Information</b>               | The number of information PDUs received.   | <b>detail</b>   |
| <b>Event</b>                     | The number of loopback control PDUs received.  | <b>detail</b>   |
| <b>Variable request</b>          | The number of variable request PDUs received.  | <b>detail</b>   |
| <b>Variable response</b>         | The number of variable response PDUs received.   | <b>detail</b>   |
| <b>Loopback control</b>          | The number of loopback control PDUs received.  | <b>detail</b>   |



Table 263: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name   | Field Description  | Level of Output |
|--|--|-----------------|
| <b>Organization specific</b>                       | The number of vendor organization specific PDUs received.  | <b>detail</b>   |
| <b>OAM Transmit Statistics</b>                     |  |                 |
| <b>Information</b>                                 | The number of information PDUs transmitted.  | <b>detail</b>   |
| <b>Event</b>                                       | The number of event notification PDUs transmitted.   | <b>detail</b>   |
| <b>Variable request</b>                            | The number of variable request PDUs transmitted.   | <b>detail</b>   |
| <b>Variable response</b>                           | The number of variable response PDUs transmitted.  | <b>detail</b>   |
| <b>Loopback control</b>                            | The number of loopback control PDUs transmitted.   | <b>detail</b>   |
| <b>Organization specific</b>                       | The number of vendor organization specific PDUs transmitted.   | <b>detail</b>   |
| <b>OAM Received Symbol Error Event information</b> |  |                 |
| <b>Events</b>                                      | The number of symbol error event TLVs that have been received after the OAM sublayer was reset.  | <b>detail</b>   |
| <b>Window</b>                                      | The symbol error event window in the received PDU.<br><br>The protocol default value is the number of symbols that can be received in one second on the underlying physical layer. | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of errored symbols in the period required for the event to be generated.  | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of symbol errors in the period reported in the received event PDU.  | <b>detail</b>   |
| <b>Total errors</b>                                | The number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>Symbol errors are coding symbol errors.                      | <b>detail</b>   |
| <b>OAM Received Frame Error Event Information</b>  |  |                 |
| <b>Events</b>                                      | The number of errored frame event TLVs that have been received after the OAM sublayer was reset.   | <b>detail</b>   |
| <b>Window</b>                                      | The duration of the window in terms of the number of 100 ms period intervals.  | <b>detail</b>   |
| <b>Threshold</b>                                   | The number of detected errored frames required for the event to be generated.  | <b>detail</b>   |
| <b>Errors in period</b>                            | The number of detected errored frames in the period.   | <b>detail</b>   |

Table 263: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name   | Field Description   | Level of Output |
|--|---|-----------------|
| <b>Total errors</b>                                      | The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>A frame error is any frame error on the underlying physical layer. | <b>detail</b>   |
| <b>OAM Received Frame Period Error Event Information</b> |   |                 |
| <b>Events</b>  | The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.   | <b>detail</b>   |
| <b>Window</b>  | The duration of the frame seconds window.   | <b>detail</b>   |
| <b>Threshold</b>   | The number of frame seconds errors in the period.   | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of frame seconds errors in the period.   | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.   | <b>detail</b>   |
| <b>OAM Transmitted Symbol Error Event Information</b>    |   |                 |
| <b>Events</b>  | The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.  | <b>detail</b>   |
| <b>Window</b>  | The symbol error event window in the transmitted PDU.   | <b>detail</b>   |
| <b>Threshold</b>   | The number of errored symbols in the period required for the event to be generated.   | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of symbol errors in the period reported in the transmitted event PDU.  | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.   | <b>detail</b>   |
| <b>OAM Transmitted Frame Error Event Information</b>     |   |                 |
| <b>Events</b>  | The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.   | <b>detail</b>   |
| <b>Window</b>  | The duration of the window in terms of the number of 100 ms period intervals.   | <b>detail</b>   |
| <b>Threshold</b>   | The number of detected errored frames required for the event to be generated.   | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of detected errored frames in the period.  | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored frames that have been detected after the OAM sublayer was reset.  | <b>detail</b>   |

## Sample Output

### show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
  Remote MUX action: discarding, Remote parser action: loopback
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

### show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
  Information: 186365, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM transmit statistics:
  Information: 186347, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame period error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

## show redundant-trunk-group

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show redundant-trunk-group &lt;group-name group-name&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  |
| <b>Description</b>              | Display information about redundant trunk groups.   |
| <b>Options</b>                  | <code>group-name group-name</code> —Display information about the specified redundant trunk group.  |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery</a></li> <li>• <a href="#">Example: Configuring Redundant Trunk Links for Faster Recovery on page 2578</a></li> <li>• <a href="#">Understanding Redundant Trunk Links on page 2447</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show redundant-trunk-group group-name Group1 on page 2882</a>   |
| <b>Output Fields</b>            | <a href="#">Table 264 on page 2882</a> lists the output fields for the <code>show redundant-trunk-group</code> command. Output fields are listed in the approximate order in which they appear.   |

Table 264: show redundant-trunk-group Output Fields

| Field Name        | Field Description  |
|-------------------|--|
| Group name        | Name of the redundant trunk port group.  |
| Interface         | Name of an interface belonging to the trunk port group.  |
| State             | Operating state of the interface. <ul style="list-style-type: none"> <li>• <b>Up</b> denotes the interface is up.</li> <li>• <b>Down</b> denotes the interface is down.</li> <li>• <b>Pri</b> denotes a primary interface.</li> <li>• <b>Act</b> denotes an active interface.</li> </ul> |
| Time of last flap | Date and time at which the advertised link became unavailable, and then, available again.  |
| Flap count        | Total number of flaps since the last switch reboot.  |

## Sample Output

### show redundant-trunk-group group-name Group1

```
user@switch> show redundant-trunk-group group-name Group1
```

| Group name | Interface | State | Time of last flap | Flap Count |
|------------|-----------|-------|-------------------|------------|
|------------|-----------|-------|-------------------|------------|

|        |             |            |       |   |
|--------|-------------|------------|-------|---|
| Group1 | ge-0/0/45.0 | UP/Pri/Act | Never | 0 |
|        | ge-0/0/47.0 | UP         | Never | 0 |

## show uplink-failure-detection

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>show uplink-failure-detection</code><br><code>&lt;group group-name&gt;</code>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for EX Series switches.  |
| <b>Description</b>              | Display information about the uplink-failure-detection group, the member interfaces, and their status.   |
| <b>Options</b>                  | <b>none</b> —Display information about all groups configured for uplink failure detection.<br><b>group group-name</b> —(Optional) Display information about the specified group only.  |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Uplink Failure Detection on page 2392</a></li> <li>• <a href="#">Configuring Interfaces for Uplink Failure Detection on page 2592</a></li> <li>• <a href="#">Example: Configuring Interfaces for Uplink Failure Detection on page 2457</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show uplink-failure-detection on page 2884</a>   |
| <b>Output Fields</b>            | <a href="#">Table 265 on page 2884</a> lists the output fields for the <b>show uplink-failure-detection</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 265: show uplink-failure-detection Output Fields**

| Field Name     | Field Description   |
|----------------|---|
| Group          | Name of the group.  |
| Uplink         | The uplink interface or interfaces configured as link-to-monitor.<br><b>NOTE:</b> The asterisk (*) indicates that the link is up.   |
| Downlink       | The downlink interface or interfaces configured as link-to-disable.<br><b>NOTE:</b> The asterisk (*) indicates that the link is up.   |
| Failure Action | Status of uplink failure detection: <ul style="list-style-type: none"> <li>• Active—The switch has detected an uplink failure and has brought the downlink down.</li> <li>• Inactive—The uplink or uplinks are up.</li> </ul> |

## Sample Output

### show uplink-failure-detection

```
user@switch> show uplink-failure-detection
```

Group : group1  
Uplink : ge-0/0/0\*  
Downlink : ge-0/0/1\*  
Failure Action : Inactive

Group : group2  
Uplink : ge-0/0/3.0  
Downlink : ge-0/0/4.0  
Failure Action : Active





## CHAPTER 36

# Troubleshooting

- [Troubleshooting Procedures on page 2887](#)

## Troubleshooting Procedures

---

- [Troubleshooting an Aggregated Ethernet Interface on page 2887](#)
- [Troubleshooting Multichassis Link Aggregation on page 2887](#)
- [Troubleshooting Network Interfaces on page 2893](#)

## Troubleshooting an Aggregated Ethernet Interface

**Problem**    **Description:** The `show interfaces terse` command shows that the LAG is down.

**Solution**    Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related Documentation**

- [Verifying the Status of a LAG Interface on page 2750](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)

## Troubleshooting Multichassis Link Aggregation

Use the following information to troubleshoot multichassis link aggregation configuration.

- [MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table on page 2888](#)
- [MC-LAG Peer Does Not Go into Standby Mode on page 2889](#)

- [Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive on page 2889](#)
- [Redirect Filters Take Priority over User-Defined Filters on page 2889](#)
- [Operational Command Output Is Wrong on page 2889](#)
- [ICCP Connection Might Take Up to 60 Seconds to Become Active on page 2890](#)
- [MAC Address Age Learned on an MC-AE Interface Is Reset to Zero on page 2890](#)
- [MAC Address Is Not Learned Remotely in a Default VLAN on page 2890](#)
- [Snooping Entries Learned on MC-AE Interfaces Are Not Removed on page 2890](#)
- [ICCP Does Not Come Up After You Add or Delete an Authentication Key on page 2891](#)
- [Local Status Is Standby When It Should Be Active on page 2891](#)
- [Packets Loop on the Server When ICCP Fails on page 2891](#)
- [Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change on page 2891](#)
- [No Commit Checks Are Done for ICL-PL Interfaces on page 2892](#)
- [Double Failover Scenario on page 2892](#)
- [Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up on page 2892](#)
- [Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer on page 2892](#)
- [AE Interfaces Go Down on page 2892](#)
- [Flooding of Upstream Traffic on page 2893](#)

### MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table

**Problem Description:** When both of the multichassis aggregated Ethernet (MC-AE) interfaces on both connected multichassis link aggregation group (MC-LAG) peers are down, the MAC addresses learned on the MC-AE interfaces are not removed from the MAC address table. For example, if you disable the MC-AE interface (ae0) on both MC-LAG peers by issuing the **set interfaces ae0 disable** command and commit the configuration, the MAC table still shows the MAC addresses as being learned on the MC-AE interfaces of both MC-LAG peers:

```
user@switchA> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
VLAN      MAC address      Type      Age Interfaces
v10        *                Flood     - All-members
v10        00:10:94:00:00:01 Learn(L)    3:55 ae0.0 (MCAE)
v10        00:10:94:00:00:02 Learn(R)    0 xe-0/0/9.0
v20        *                Flood     - All-members
v30        *                Flood     - All-members
v30        84:18:88:de:b1:2e Static      - Router
```

```
user@switchB> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
VLAN      MAC address      Type      Age Interfaces
v10        *                Flood     - All-members
```

|     |                   |          |                |
|-----|-------------------|----------|----------------|
| v10 | 00:10:94:00:00:01 | Learn(R) | 0 ae0.0 (MCAE) |
| v10 | 00:10:94:00:00:02 | Learn    | 40 xe-0/0/10.0 |
| v20 | *                 | Flood    | - All-members  |
| v30 | *                 | Flood    | - All-members  |
| v30 | 84:18:88:df:83:0a | Static   | - Router       |

**Solution** This is expected behavior.

### MC-LAG Peer Does Not Go into Standby Mode

**Problem** **Description:** A multichassis link aggregation group (MC-LAG) peer does not go into standby mode if the MC-LAG peer IP address specified in the Interchassis Control Protocol (ICCP) configuration and the IP address specified in the multichassis protection configuration are different.

**Solution** To prevent failure to enter standby mode, make sure the peer IP address in the ICCP configurations and the IP address in multichassis protection configurations are the same.

### Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive

**Problem** **Description:** When the interchassis control link-protection link (ICL-PL) and multichassis aggregated Ethernet (MC-AE) interfaces go down on the primary multichassis link aggregation group (MC-LAG) peer, the secondary MC-LAG peer's MC-AE interfaces with status control set to standby become inactive instead of active.

**Solution** This is expected behavior.

### Redirect Filters Take Priority over User-Defined Filters

**Problem** **Description:** Multichassis link aggregation group (MC-LAG) implicit failover redirection filters take precedence over user-configured explicit filters. This is expected behavior.

**Solution** This is expected behavior.

### Operational Command Output Is Wrong

**Problem** **Description:** After you deactivate the Interchassis Control Protocol (ICCP), the **show iccp** operational command output still shows registered client daemons, such as mcsnoopd, lacpd, and eswd.

For example:

```
user@switch> show iccp
Client Application: MCSNOOPD
Redundancy Group IDs Joined: None
```

```
Client Application: lacpd
Redundancy Group IDs Joined: 1
```

```
Client Application: eswd
Redundancy Group IDs Joined: 1
```

The **show iccp** command output always shows registered modules regardless of whether or not ICCP peers are configured.

**Solution** This is expected behavior.

---

### ICCP Connection Might Take Up to 60 Seconds to Become Active

---

**Problem** **Description:** When the Interchassis Control Protocol (ICCP) configuration and the routed VLAN interface (RVI) configuration are committed together, the ICCP connection might take up to 60 seconds to become active.

**Solution** This is expected behavior.

---

### MAC Address Age Learned on an MC-AE Interface Is Reset to Zero

---

**Problem** **Description:** When you activate and then deactivate an interchassis control link-protection link (ICL-PL), the MAC address age learned on the multichassis aggregated Ethernet (MC-AE) interface is reset to zero. The next-hop interface changes trigger MAC address updates in the hardware, which then triggers aging updates in the Packet Forwarding Engine (PFE). The result is that the MAC address age is updated to zero.

For example, the ICL-PL has been deactivated, and the **show ethernet-switching table** command output shows that the MAC addresses have an age of 0.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 2 learned, 0 persistent entries
  VLAN      MAC address      Type      Age Interfaces
  ---      -
v100        *                Flood     - All-members
v100        00:10:00:00:00:01 Learn(L)   0 ae0.0 (MCAE)
v100        00:10:00:00:00:02 Learn(L)   0 ae0.0 (MCAE)
```

**Solution** This is expected behavior.

---

### MAC Address Is Not Learned Remotely in a Default VLAN

---

**Problem** **Description:** If a multichassis link aggregation group (MC-LAG) peer learns a MAC address in the default VLAN, the Interchassis Control Protocol (ICCP) does not synchronize the MAC address with the MAC address of the other MC-LAG peer.

**Solution** This is expected behavior.

---

### Snooping Entries Learned on MC-AE Interfaces Are Not Removed

---

**Problem** **Description:** When multichassis aggregated Ethernet (MC-AE) interfaces are configured on a VLAN that is enabled for multicast snooping, the membership entries learned on the MC-AE interfaces on the VLAN are not cleared when the MC-AE interfaces go down.

This is done to speed up convergence time when the interfaces come up, or come up and go down.

**Solution** This is expected behavior.

---

#### ICCP Does Not Come Up After You Add or Delete an Authentication Key

---

**Problem** **Description:** The Interchassis Control Protocol (ICCP) connection is not established when you add an authentication key and then delete it only at the global ICCP level. However, authentication works correctly at the ICCP peer level.

**Solution** Delete the ICCP configuration , and then add the ICCP configuration.

---

#### Local Status Is Standby When It Should Be Active

---

**Problem** **Description:** If the multichassis aggregated Ethernet (MC-AE) interface is down when the state machine is in a synchronized state, the multichassis link aggregation group (MC-LAG) peer local status is standby. If the MC-AE interface goes down after the state machine is in an active state, then the local status remains active, and the local state indicates that the interface is down.

**Solution** This is expected behavior.

---

#### Packets Loop on the Server When ICCP Fails

---

**Problem** **Description:** When you enable backup liveness detection for a multichassis link aggregation group (MC-LAG), and the backup liveness detection packets are lost because of a temporary failure on the MC-LAG, then both of the peers in the MC-LAG remain active. If this happens, both of the MC-LAG peers send packets to the connected server.

**Solution** This is expected behavior.

---

#### Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change

---

**Problem** **Description:** After a reboot or after a new Interchassis Control Protocol (ICCP) configuration has been committed, and the ICCP connection does not become active, the Link Aggregation Control Protocol (LACP) messages transmitted over the multichassis aggregated Ethernet (MC-AE) interfaces use the default system ID. The configured system ID is used instead of the default system ID only after the MC-LAG peers synchronize with each other.

**Solution** This is expected behavior.

### No Commit Checks Are Done for ICL-PL Interfaces

---

**Problem** **Description:** There are no commit checks on the interface being configured as an interchassis control link-protection link (ICL-PL), so you must provide a valid interface name for the ICL-PL.

**Solution** This is expected behavior.

### Double Failover Scenario

---

**Problem** **Description:** If the following events happen in this exact order—the Interchassis Control Protocol (ICCP) goes down, and the multichassis aggregated Ethernet (MC-AE) interface on the multichassis link aggregation group (MC-LAG) peer in active mode goes down—a double failover occurs. In this scenario, the MC-LAG peer in standby mode does not detect what happens on the active MC-LAG peer. The MC-LAG peer in standby mode operates as if the MC-AE interface on the MC-LAG in active mode were up and blocks the interchassis control protocol-protection link (ICL-PL) traffic. The ICL-PL traffic is not forwarded.

**Solution** This is expected behavior.

### Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up

---

**Problem** **Description:** When the interchassis control link-protection link (ICL-PL) goes down and up, multicast traffic is flooded to all of the interfaces in the VLAN. The Packet Forwarding Engine (PFE) flag Ip4McastFloodMode for the VLAN is changed to MCAST\_FLOOD\_ALL. This problem only occurs when a multichassis link aggregation group (MC-LAG) is configured for Layer 2.

**Solution** This is expected behavior.

### Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer

---

**Problem** **Description:** When the Interchassis Control Protocol (ICCP) is down, the status of a remote MC-LAG peer is unknown. Even if the MC-LAG peer is configured as standby, the traffic is not redirected to this peer because it is assumed that this peer is down.

**Solution** This is expected behavior.

### AE Interfaces Go Down

---

**Problem** **Description:** When a multichassis aggregated Ethernet (MC-AE) interface is converted to an aggregated Ethernet (AE) interface, it retains some MC-AE properties. For example, the AE interface might retain the administrative key of the MC-AE. When this happens, the AE interface goes down.

**Solution** Restart the Link Aggregation Control Protocol (LACP) on the multichassis link aggregation group (MC-LAG) peer hosting the AE interface to bring up the AE interface. Restarting LACP removes the MC-AE properties of the AE interface.

### Flooding of Upstream Traffic

**Problem** **Description:** When MAC synchronization is enabled, the multichassis link aggregation group (MC-LAG) peer can resolve Address Resolution Protocol (ARP) entries for the MC-LAG routed VLAN interface (RVI) with either of the MC-LAG peer MAC addresses. If the downstream traffic is sent with one MAC address (MAC1) but the peer has resolved the MAC address with a different MAC address (MAC2), the MAC2 address might not be learned by any of the access layer switches. Flooding of the upstream traffic for the MAC2 address might then occur.

**Solution** Make sure that downstream traffic is sent from the MC-LAG peers periodically to prevent the MAC addresses from aging out.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)

## Troubleshooting Network Interfaces

### The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

**Problem** **Description:** The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command **show interfaces *interface-name***, the disabled port is not listed.

**Cause** By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution** Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.





## PART 10

# Routing Options

- [Overview on page 2897](#)
- [Configuration on page 2903](#)
- [Administration on page 3055](#)
- [Troubleshooting on page 3249](#)



## CHAPTER 37

# Overview

- [Routing Options Overview on page 2897](#)

## Routing Options Overview

---

- [Overview of Routing Options on page 2897](#)
- [Understanding Virtual Router Routing Instances on page 2898](#)
- [Understanding Distributed Periodic Packet Management on page 2898](#)
- [Understanding Bidirectional Forwarding Detection \(BFD\) on page 2899](#)
- [Understanding the Unified Forwarding Table on page 2899](#)

## Overview of Routing Options

In addition to dynamic routing protocols, you can configure static routing on QFX Series switches. You can also configure a variety of protocol-independent routing properties, such as

- Per-packet load balancing (equal cost multipath routing)
- Autonomous system numbers
- Autonomous system confederation members
- Router identifiers
- Routing table groups
- Multicast scoping

### Related Documentation

- [Understanding Distributed Periodic Packet Management on page 2898](#)
- [Understanding Virtual Router Routing Instances on page 2898](#)

## Understanding Virtual Router Routing Instances

Virtual router routing instances allow administrators to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

Each virtual router routing instance consists of sets of the following:

- Routing tables
- Interfaces that belong to these routing tables
- Routing protocol configurations
- Routing option configurations

You can use virtual router routing instances to isolate customer traffic on a network and to bind customer-specific routing instances to customer-owned interfaces. Each interface can belong to only one routing instance. QFX 3500 and QFX3600 switches and QFabric systems support as many as 256 virtual router routing instances. QFX 5100 switches support as many as 512 virtual router routing instances.

### Related Documentation

- [Configuring Virtual Router Routing Instances on page 2908](#)

## Understanding Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocol, and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.



**BEST PRACTICE:** We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

---

- Related Documentation**
- [Configuring Distributed Periodic Packet Management on page 2906](#)

## Understanding Bidirectional Forwarding Detection (BFD)

The Bidirectional Forwarding Detection (BFD) protocol is a simple mechanism that detects failures in a network and works in a wide variety of network environments and topologies. In BFD operation, switches exchange BFD hello packets at a specified interval and detect a neighbor failure if they do not receive a reply after a specified interval. The BFD failure detection timers support shorter time limits than the static route failure detection mechanisms, so they can provide faster detection of failures.

To configure faster failure detection, use lower BFD timer values. The timers can automatically adapt to a higher value if an adjacency fails, and they also adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. In this case, a back-off algorithm increases the receive interval by two if the local BFD instance is the reason for the session flap and increases the transmission interval by two if the remote BFD instance is the reason for the session flap.

You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. This command is hitless, meaning that it does not affect traffic flow.



**NOTE:** QFX and EX4600 switches do not support BFD timer values of less than 1 second.

- Related Documentation**
- [Examples: Configuring BFD for Static Routes on page 2914](#)
  - [Example: Configuring BFD Authentication for Static Routes on page 2929](#)

## Understanding the Unified Forwarding Table

- [Using the Unified Forwarding Table to Optimize Address Storage on page 2899](#)
- [MAC Address and Host Address Memory Allocation on page 2900](#)
- [LPM Table Memory Allocation on page 2901](#)

### Using the Unified Forwarding Table to Optimize Address Storage

On QFX5100 and EX4600 switches, you can control the allocation of forwarding table memory available to store the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match (LPM) table entries.



**NOTE:** Starting with Junos OS 13.2X51-D15, you can allocate more memory to store prefixes in the range /65 to /127 range.

This feature gives you the flexibility to configure your switch to match the needs of your particular network environment.

### MAC Address and Host Address Memory Allocation

There are several profiles that allocate memory differently for MAC addresses and host addresses. You configure the mix that best meets your needs by choosing the appropriate profile. [Table 96 on page 1546](#) lists the profiles you can choose and the associated maximum values for the MAC address and host table entries.

**Table 266: Unified Forwarding Table Profiles**

| Profile Name                      | MAC Table     | Host Table (unicast and multicast addresses) |              |             |             |             |             |
|-----------------------------------|---------------|--|--------------|-------------|-------------|-------------|-------------|
|                                   | MAC Addresses | IPv4 unicast                                 | IPv6 unicast | IPv4 (*, G) | IPv4 (S, G) | IPv6 (*, G) | IPv6 (S, G) |
| <b>l2-profile-one</b>             | 288K          | 16K  | 8K           | 8K          | 8K          | 4K          | 4K          |
| <b>l2-profile-two</b>             | 224K          | 80K  | 40K          | 40K         | 40K         | 20K         | 20K         |
| <b>l2-profile-three (default)</b> | 160K          | 144K   | 72K          | 72K         | 72K         | 36K         | 36K         |
| <b>l3-profile</b>                 | 96K           | 208K   | 104K         | 104K        | 104K        | 52K         | 52K         |
| <b>lpm-profile</b>                | 32K           | 16K  | 8K           | 8K          | 8K          | 4K          | 4K          |

Note that all entries in the host table share the same memory space. If the host table stores the maximum number of entries for any given type, the entire shared table is full and is unable to accommodate *any* entries of any other type. As you can see, different entry types occupy different amounts of memory. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address.

[Table 97 on page 1546](#) lists various valid combinations that the host table can store if you use the **l2-profile-one** profile. Each row in the table represents a case in which the host table is full and cannot accommodate any more entries. .

**Table 267: Example Host Table Combinations Using l2-profile-one**

| IPv4 unicast | IPv6 unicast | IPv4 multicast (*, G) | IPv4 multicast (S, G) | IPv6 multicast (*, G) | IPv6 multicast (S, G) |
|--------------|--------------|-----------------------|-----------------------|-----------------------|-----------------------|
| 16K          | 0            | 0                     | 0                     | 0                     | 0                     |
| 12K          | 2K           | 0                     | 0                     | 0                     | 0                     |
| 12K          | 0            | 2K                    | 2K                    | 0                     | 0                     |
| 8K           | 4K           | 0                     | 0                     | 0                     | 0                     |
| 4K           | 2K           | 2K                    | 2K                    | 0                     | 0                     |

Table 267: Example Host Table Combinations Using l2-profile-one (continued)

| IPv4 unicast | IPv6 unicast | IPv4 multicast (*, G) | IPv4 multicast (S, G) | IPv6 multicast (*, G) | IPv6 multicast (S, G) |
|--------------|--------------|-----------------------|-----------------------|-----------------------|-----------------------|
| 0            | 4K           | 0                     | 0                     | 1K                    | 1K                    |

LPM Table Memory Allocation

You configure the memory allocation for LPM table entries differently depending on which version of Junos OS you use. To learn how to configure memory allocation for LPM table entries see [“Configuring the Unified Forwarding Table” on page 1697](#).

Related Documentation

- [Configuring the Unified Forwarding Table on page 1697](#)





## CHAPTER 38

# Configuration

- [Configuration Tasks on page 2903](#)
- [Configuration Examples on page 2913](#)
- [Configuration Statements on page 2937](#)

### Configuration Tasks

---

- [Configuring Static Routing on page 2904](#)
- [Configuring Per-Packet Load Balancing on page 2904](#)
- [Configuring Distributed Periodic Packet Management on page 2906](#)
- [Configuring Virtual Router Routing Instances on page 2908](#)
- [Configuring the Unified Forwarding Table on page 2909](#)

## Configuring Static Routing

Static routes are routes that are manually configured and entered into the routing table.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:  

```
[edit]
user@switch# set routing-options static route 0.0.0.0/0 next-hop 10.0.1.1
```
- To configure a static route and specify the next address to be used when routing traffic to the static route:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 next-hop 10.0.0.2.1
```
- To always keep the static route in the forwarding table:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 retain
```
- To prevent the static route from being readvertised:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 no-readvertise
```
- To remove inactive routes from the forwarding table:  

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 active
```

### Related Documentation

- [Monitoring Routing Information on page 3055](#)

## Configuring Per-Packet Load Balancing

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS chooses one of the next-hop addresses to install into the forwarding table in a random fashion. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use this feature to spread traffic across multiple paths.

On a QFX3500 standalone switch, with static routing configured, whenever a route pointing to an ECMP next-hop changes to a new ECMP next-hop with a different member list but contains the exact member count as before, the location of the retained members in the new member list is the same as in the old member list.

For example, if you have the following configuration on the switch:

```
set routing-options static route 0.0.0.0/0 next-hop 11.8.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.9.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.10.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.11.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.12.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.13.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.14.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.15.12.2
```

and want to change the first and eighth route to point to another location, you can issue the following commands:

```
delete routing-options static route 0.0.0.0/0 next-hop 11.8.12.2
delete routing-options static route 0.0.0.0/0 next-hop 11.15.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.16.12.2
set routing-options static route 0.0.0.0/0 next-hop 11.17.12.2
```

This configuration does not affect the second next-hop through the seventh next-hop.

When per-packet load balancing is configured, traffic is divided into individual flows (up to a maximum of 16). Packets for an individual flow are sent out a single interface. To determine flows, the switch examines each of the following packet fields:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

The switch recognizes packets in which all of these parameters are identical and ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.



**NOTE:** Load balancing is not supported on management interfaces.

The following steps show how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {
```

```
from {  
    match-conditions;  
    route-filter destination-prefix match-type <actions>;  
    prefix-list name;  
}  
then {  
    load-balance per-packet;  
}  
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {  
    export policy-name;  
}
```

When you enable per-packet load balancing on a QFabric system, packets might be switched across the fabric even though there is a local port with a same-cost route to the destination. For example, if per-packet load balancing is enabled and a packet arrives at network Node device A, it might be switched to network Node device B and forwarded from there even if there is a same-cost route through a port on Node device A to the destination. In this case, traffic transits the fabric needlessly. You can configure a QFabric system to choose a locally switched route if one is available. To enable this feature, include the **ecmp-do-local-lookup** statement at the **[edit forwarding-options]** hierarchy level.

#### Related Documentation

- [Examples: Configuring Per-Packet Load Balancing on page 2913](#)

## Configuring Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for Link Aggregation Control Protocol (LACP) packets only.



**BEST PRACTICE:** We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- [Disabling or Enabling Distributed Periodic Packet Management Globally on page 2907](#)
- [Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 2907](#)

### Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
user@switch# delete ppm no-delegate-processing
```

### Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

#### Related Documentation

- [Understanding Distributed Periodic Packet Management on page 2898](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 2393](#)

## Configuring Virtual Router Routing Instances

Use virtual router routing instances to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

To configure virtual router routing instances:

1. Create a routing instance:

```
[edit routing-instances]user@switch# set routing-instance-name instance-type virtual-router
```



**NOTE:** The default routing instance, master, refers to the main inet.0 routing table. The master routing instance is reserved and cannot be specified as a routing instance.

2. Bind each routing instance to the corresponding interfaces:

```
[edit routing-instances]user@switch# set routing-instance-name interface
device-name:type-fpc/pic/port.logical-unit-number
```



**NOTE:**

- You must bind routing instances to interfaces from the Node devices assigned to the network Node group only. If you try to bind routing instances to interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
- You can bind an interface to one routing instance only.

3. Create each of the logical interfaces bound to each routing instance:

```
[edit interfaces]user@switch# set device-name:type-fpc/pic/port unit logical-unit-number
family inet address ip-address
```



**NOTE:** Do not create a logical interface using the family ethernet-switching option in this step. Binding an interface using the family ethernet-switching option to a routing instance can cause the interface to shut down.

4. (Optional) Configure routing protocols for the routing instance at the `[edit routing-instances routing-instance-name protocols]` hierarchy level.
5. (Optional) Configure routing options for the routing instance at the `[edit routing-instances routing-instance-name routing-options]` hierarchy level.

### Related Documentation

- [Understanding Virtual Router Routing Instances on page 2898](#)
- [Understanding Interfaces on the QFabric System](#)
- [Understanding Node Groups](#)

- [Verifying That Virtual Router Routing Instances Are Working on page 3056](#)

## Configuring the Unified Forwarding Table

To optimize the way your switch allocates memory for different types of addresses, you can choose a unified forwarding table profile. In addition to choosing this profile, you can also decide how you want memory allocated for longest prefix match (LPM) entries.

- [Configuring an Address-Storage Profile on page 2909](#)
- [Configuring the LPM Allocation on page 2910](#)

### Configuring an Address-Storage Profile

On QFX5100 and EX4600 switches, you can control the allocation of memory available to store the following:

- MAC addresses
- Layer 3 host entries
- Longest prefix match (LPM) table entries

You configure the mix that best meets your needs by choosing the appropriate profile. [Table 109 on page 1698](#) lists the profiles you can choose and the maximum values for the MAC address and host table entries.

**Table 268: Unified Forwarding Table Profiles**

| Profile Name                      | MAC Table | Host Table (unicast and multicast addresses) |              |             |             |             |             |
|-----------------------------------|-----------|--|--------------|-------------|-------------|-------------|-------------|
|                                   |           | IPv4 unicast                                 | IPv6 unicast | IPv4 (*, G) | IPv4 (S, G) | IPv6 (*, G) | IPv6 (S, G) |
| <b>l2-profile-one</b>             | 288K      | 16K  | 8K           | 8K          | 8K          | 4K          | 4K          |
| <b>l2-profile-two</b>             | 224K      | 80K  | 40K          | 40K         | 40K         | 20K         | 20K         |
| <b>l2-profile-three (default)</b> | 160K      | 144K   | 72K          | 72K         | 72K         | 36K         | 36K         |
| <b>l3-profile</b>                 | 96K       | 208K   | 104K         | 104K        | 104K        | 52K         | 52K         |
| <b>lpm-profile*</b>               | 32K       | 16K  | 8K           | 8K          | 8K          | 4K          | 4K          |

Note that if the host table stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address. For more information about valid combinations of table entries see *Understanding the Unified Forwarding Table*.

To configure the profile that you want, enter and commit the following statement:

```
[edit]
user@switch# set chassis forwarding-options profile-name
```



**NOTE:** When you configure and commit a profile, the PFE process restarts and all the data interfaces on the switch go down and come back up.

The settings for **l2-profile-three** are configured by default. That is, if you do not enter a **set forwarding-options chassis profile-name** statement, these settings are configured.

### Configuring the LPM Allocation

---

In addition to choosing a profile, you can further optimize memory allocation for LPM table entries by configuring how many IPv6 prefixes in the range /65 through /127 you want the switch to store. The switch uses LPM entries during address lookup to match addresses to the most-specific (longest) applicable prefix. The procedures for configuring the LPM table are different depending on which version of Junos OS you are using.

- [Configuring the LPM Table With Junos OS 13.2X51-D10 and 13.2X52-D10 on page 2910](#)
- [Configuring the LPM Table With Junos OS 13.2x51-D15 on page 2911](#)

#### **Configuring the LPM Table With Junos OS 13.2X51-D10 and 13.2X52-D10**

With Junos OS 13.2x51-D10 and 13.2X52-D10, the switch allocates memory for 16 IPv6 prefixes in the range /65 through /127 by default. If you want to use more than 16 IPv6 prefixes in this range, you must enter and commit the following statement:

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix [1-128]
```

Each increment adds support for 16 IPv6 prefixes between /65 and /127, for a maximum of 2048 such prefixes (16 x 128 = 2048). The system supports 16 of these prefixes by default, so to increase the number of supported prefixes, you must enter a value of 2 or greater. For example, if you enter **2**, the system will support 32 IPv6 prefixes in the range /65 through /127.



**NOTE:** When you configure and commit the **num-65-127-prefix** value, all the data interfaces on the switch restart. The management interfaces are unaffected.

The LPM table is shared, and each increment that you add for IPv6 prefixes in the range /65 through /127 reduces the number of table entries that are available for IPv4 prefixes and IPv6 prefixes shorter than /65. Note that IPv6 prefixes /65 and longer consume twice as much memory as shorter IPv6 prefixes and four times as much memory as IPv4 prefixes. So, for example, entering the following statement

```
user@switch# set chassis forwarding-options l2-profile-one num-65-127-prefix 2
```

provides for 16 additional IPv6 prefixes /65 or longer (for a total of 32 such prefixes) and reduces the numbers of other prefixes that can be stored, as indicated:



- 32 fewer IPv6 prefixes shorter than /65 (16 IPv6 prefixes /65 or longer consume the same amount of memory as 32 IPv6 prefixes shorter than /65), or
- 64 fewer IPv4 prefixes (16 IPv6 prefixes /65 or longer consume the same amount of memory as 64 IPv4 prefixes)

Table 110 on page 1700 provides examples of valid combinations that the LPM table can store using the **l2** and **l3** profiles. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

**Table 269: Example LPM Table Combinations Using l2-and l3 Profiles With Junos OS 13.2X51-D10 and 13.2X52-D10**

| num-65-127-prefix Value | IPv4 Entries | IPv6 Entries (Prefix <= 64) | IPv6 Entries (Prefix >= 65) |
|-------------------------|--------------|-----------------------------|-----------------------------|
| 1 (default)             | 16K-16       | 0K                          | 16                          |
| 1 (default)             | 0K           | 8K-16                       | 16                          |
| 1 (default)             | 8K-16        | 4K                          | 16                          |
| 64                      | 4K           | 4K                          | 1K                          |
| 64                      | 2K           | 5K                          | 1K                          |
| 64                      | 0K           | 6K                          | 1K                          |
| 128                     | 4K           | 2K                          | 2K                          |
| 128                     | 2K           | 3K                          | 2K                          |
| 128                     | 0K           | 4K                          | 2K                          |



**NOTE:** With Junos OS 13.2X51-D10 and 13.2X52-D10, the **lpm-profile** does not support IPv6 prefixes. If you use this version of Junos OS and also use the **lpm-profile**, do not configure the **num-65-127-prefix** statement. That is, leave it at its default value of 1, which allows for as many as 128K IPv4 prefixes (the maximum possible).

#### *Configuring the LPM Table With Junos OS 13.2x51-D15*

With Junos OS 13.2X51-D15, you can configure the memory allocation for the LPM table for the **lpm-profile** profile independently of the other profiles. In addition, Junos OS 13.2x51-D15 offers twice as much storage for IPv6 prefixes /65 through /127 (4K instead of 2K) for the **l2** and **l3** profiles.

- [Configuring the l2 and l3 profiles With Junos OS 13.2x51-D15 on page 2912](#)
- [Configuring The lpm-profile With Junos OS 13.2x51-D15 on page 2912](#)

### Configuring the l2 and l3 profiles With Junos OS 13.2x51-D15

With Junos OS 13.2x51-D15, you can configure the switch to support as many as 4K IPv6 prefixes /65 through /127 if you are using any profile other than the **lpm-profile** profile. To do so, enter and commit the following statement:

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix [0-4]
```

Each increment adds support for 1K IPv6 prefixes between /65 and /127, for a maximum of 4K such prefixes. The default value is 1, which allocates memory for 1K of IPv6 prefixes in this range. Each increment that you add for IPv6 prefixes in the range /65 through /127 reduces the number of table entries that are available for IPv6 prefixes shorter than /65 and IPv4 prefixes. [Table 111 on page 1701](#) shows the numbers of entries that you can allocate by using the **num-65-127-prefix** statement with Junos OS 13.2X51-D15. Once again, each row represents a case in which the table is full and cannot accommodate any more entries.

**Table 270: LPM Table Combinations for l2 and l3 profiles With Junos OS 13.2X51-D15**

| num-65-127-prefix Value | IPv4 Entries | IPv6 Entries (Prefix <= 64) | IPv6 Entries (Prefix >= 65) |
|-------------------------|--------------|-----------------------------|-----------------------------|
| 0                       | 16K          | 8K                          | 0K                          |
| 1 (default)             | 12K          | 6K                          | 1K                          |
| 2                       | 8K           | 4K                          | 2K                          |
| 3                       | 4K           | 2K                          | 3K                          |
| 4                       | 0K           | 0K                          | 4K                          |



**NOTE:** When you configure the **num-65-127-prefix** value, the PFE process restarts and all the data interfaces on the switch go down and come back up. The management interfaces are unaffected.

### Configuring The lpm-profile With Junos OS 13.2x51-D15

If you use the **lpm-profile** profile with Junos OS 13.2x51-D15, you can control whether the switch allocates any memory for IPv6 prefixes /65 through /127. By default, the switch supports the following with this profile:

- 128K IPv4 prefixes
- 16K IPv6 prefixes (all lengths)

You can disable support for IPv6 prefixes /65 through /127 with the **lpm-profile** profile so that there is more memory for IPv6 prefixes shorter than /65. To do so, enter and commit the following statement:

```
user@switch# set chassis forwarding-options profile-name prefix-65-127-disable
```

If you enter this statement, the switch allocates memory for the following:

- 128K IPv4 and IPv6 prefixes shorter than /65
- 0K IPv6 prefixes /65 through /127

For example, if you use the **prefix-65-127-disable** statement, each of the following combinations are valid:

- 100K IPv4 and 28K IPv6 /64 prefixes
- 64K IPv4 and 64K IPv6 /64 prefixes
- 128K IPv4 and 0K IPv6 /64 prefixes
- 0K IPv4 and 128K IPv6 /64 prefixes

**Related  
Documentation**

- [Understanding the Unified Forwarding Table on page 1545](#)

## Configuration Examples

- [Examples: Configuring Per-Packet Load Balancing on page 2913](#)
- [Examples: Configuring BFD for Static Routes on page 2914](#)
- [Example: Configuring BFD Authentication for Static Routes on page 2929](#)

### Examples: Configuring Per-Packet Load Balancing

Perform per-packet load balancing for all routes:

```
[edit]
policy-options {
  policy-statement load-balancing-policy {
    then {
      load-balance per-packet;
    }
  }
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}
```

Perform per-packet load balancing for a limited set of routes:

```
[edit]
policy-options {
  policy-statement load-balancing-policy {
    from {
      route-filter 192.168.10/24 orlonger;
      route-filter 9.114/16 orlonger;
    }
    then {
      load-balance per-packet;
    }
  }
}
```

```
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}
```

**Related Documentation**

- [Configuring Per-Packet Load Balancing on page 2904](#)

## Examples: Configuring BFD for Static Routes

- [Understanding BFD for Static Routes on page 2914](#)
- [Example: Configuring BFD for Static Routes on page 2918](#)
- [Example: Enabling BFD on Qualified Next Hops in Static Routes on page 2923](#)

### Understanding BFD for Static Routes

---

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes. In Junos OS Release 8.2 and later, BFD also supports multihop static routes.

To enable failure detection, include the **bfd-liveness-detection** statement in the static route configuration.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is not supported for any other protocol.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route destination-prefix]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.



**NOTE:** If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.



**NOTE:** SRX Series devices do not support distributed BFD.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

To specify the minimum transmit interval for failure detection, include the **transmit-interval** **minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.



**NOTE:** You must configure the **neighbor** statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.



**NOTE:** We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.



**NOTE:** If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

Junos OS also supports BFD over multihop static routes. For example, you can configure BFD over a Layer 3 path to provide path integrity over that path. You can limit the number of hops by specifying the time to live (TTL).

To configure BFD over multihop static routes, include the following statements:

```
static route destination-prefix {
  bfd-liveness-detection {
    local-address ip-address;
    minimum-receive-ttl number;
  }
}
```

To specify the source address for the multihop static route and to enable multihop BFD support, include the **local-address** statement.

To specify the number of hops, include the **minimum-receive-ttl** statement. You must configure this statement for a multihop BFD session. You can configure a value in the range from 1 through 255. It is optional for a single-hop BFD session. If you configure the **minimum-receive-ttl** statement for a single-hop session, the value must be 255.

On M Series and T Series platforms only, starting in Junos OS Release 12.3, multihop BFD runs on the CPU in the FPC, DPC, or MPC. This is referred to as *distributed BFD*. Previously, multihop BFD ran from the Routing Engine.

### Example: Configuring BFD for Static Routes

---

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- [Requirements on page 2918](#)
- [Overview on page 2918](#)
- [Configuration on page 2919](#)
- [Verification on page 2922](#)

#### **Requirements**

In this example, no special configuration beyond device initialization is required.

#### **Overview**

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.


In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

[Figure 56 on page 2919](#) shows the sample network.



Figure 56: Customer Routes Connected to a Service Provider

**Configuration**

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device B**

```

set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Device D**

```

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```
[edit interfaces]
```

```
user@B# set ge-1/2/0 unit 0 description B->D
```

```
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
```

```
user@B# set lo0 unit 57 family inet address 10.0.0.1/32
```

```
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```

2. On Device B, create a static route and set the next-hop address.

```
[edit routing-options]
```

```
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```

3. On Device B, configure BFD for the static route.

```
[edit routing-options]
```

```
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval  
1000
```

4. On Device B, configure tracing operations for BFD.

```
[edit protocols]
```

```
user@B# set bfd traceoptions file bfd-trace
```

```
user@B# set bfd traceoptions flag all
```

5. If you are done configuring Device B, commit the configuration.

```
[edit]
```

```
user@B# commit
```

6. On Device D, configure the interfaces.

```
[edit interfaces]
```

```
user@D# set ge-1/2/0 unit 1 description D->B
```

```
user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
```

```
user@D# set lo0 unit 2 family inet address 192.168.47.5/32
```

```
user@D# set lo0 unit 2 family inet address 192.168.47.6/32
```

7. On Device D, create a static route and set the next-hop address.

```
[edit routing-options]
```

```
user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
```

8. On Device D, configure BFD for the static route.

```
[edit routing-options]
```

```
user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
```

9. On Device D, configure tracing operations for BFD.

```
[edit protocols]
```

```
user@D# set bfd traceoptions file bfd-trace
```

```
user@D# set bfd traceoptions flag all
```

10. If you are done configuring Device D, commit the configuration.

```
[edit]
```

```
user@D# commit
```

## Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

Device B  user@B# show interfaces
          ge-1/2/0 {
            unit 0 {
              description B->D;
              family inet {
                address 172.16.1.1/24;
              }
            }
          }
          lo0 {
            unit 57 {
              family inet {
                address 10.0.0.1/32;
                address 10.0.0.2/32;
              }
            }
          }

          user@D# show protocols
          bfd {
            traceoptions {
              file bfd-trace;
              flag all;
            }
          }

          user@B# show routing-options
          static {
            route 192.168.47.0/24 {
              next-hop 172.16.1.2;
              bfd-liveness-detection {
                minimum-interval 1000;
              }
            }
          }

Device D  user@D# show interfaces
          ge-1/2/0 {
            unit 1 {
              description D->B;
              family inet {
                address 172.16.1.2/24;
              }
            }
          }
          lo0 {
            unit 2 {
              family inet {
                address 192.168.47.5/32;
                address 192.168.47.6/32;
              }
            }
          }

          user@D# show routing-options
          static {

```

```

route 0.0.0.0/0 {
  next-hop 172.16.1.1;
  bfd-liveness-detection {
    minimum-interval 1000;
  }
}

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 2922](#)
- [Viewing Detailed BFD Events on page 2923](#)

### Verifying That BFD Sessions Are Up

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2 | Up    | lt-1/2/0.0 | 3.000       | 1.000             | 3          |

Client Static, TX interval 1.000, RX interval 1.000  
 Session up time 00:14:30  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Replicated, routing table index 172  
 Min async interval 1.000, min slow interval 1.000  
 Adaptive async TX interval 1.000, RX interval 1.000  
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
 Local discriminator 2, remote discriminator 1  
 Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

```
user@D> show bfd session extensive
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.1 | Up    | lt-1/2/0.1 | 3.000       | 1.000             | 3          |

Client Static, TX interval 1.000, RX interval 1.000  
 Session up time 00:14:35  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Replicated, routing table index 170  
 Min async interval 1.000, min slow interval 1.000  
 Adaptive async TX interval 1.000, RX interval 1.000  
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
 Local discriminator 1, remote discriminator 2  
 Echo mode disabled/inactive

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

#### *Viewing Detailed BFD Events*

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
Nov 23 14:26:55 Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: ppm_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55 IfIndex (3) len 4: 0
Nov 23 14:26:55 Protocol (1) len 1: BFD
Nov 23 14:26:55 Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 74
```

**Meaning** BFD messages are being written to the trace file.

#### Example: Enabling BFD on Qualified Next Hops in Static Routes

This example shows how to configure a static route with multiple possible next hops. Each next hop has Bidirectional Forwarding Detection (BFD) enabled.

- [Requirements on page 2923](#)
- [Overview on page 2923](#)
- [Configuration on page 2924](#)
- [Verification on page 2927](#)

#### **Requirements**

In this example, no special configuration beyond device initialization is required.

#### **Overview**


In this example, Device B has the static route **192.168.47.0/24** with two possible next hops. The two next hops are defined using two **qualified-next-hop** statements. Each next hop has BFD enabled.

BFD is also enabled on Device D because BFD must be enabled on both ends of the connection.

A next hop is included in the routing table if the BFD session is up. The next hop is removed from the routing table if the BFD session is down.

See [Figure 57 on page 2924](#).

**Figure 57: BFD Enabled on Qualified Next Hops**



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device B
set interfaces fe-0/1/0 unit 2 description secondary-B->D
set interfaces fe-0/1/0 unit 2 family inet address 192.168.2.1/24
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2
  bfd-liveness-detection minimum-interval 60
set routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2
  bfd-liveness-detection minimum-interval 60

Device D
set interfaces fe-0/1/0 unit 3 description secondary-D->B
set interfaces fe-0/1/0 unit 3 family inet address 192.168.2.2/24
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 60
  
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a static route with two possible next hops, both with BFD enabled:

1. On Device B, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@B# set unit 2 description secondary-B->D
user@B# set unit 2 family inet address 192.168.2.1/24
```

```
[edit interfaces ge-1/2/0]
user@B# set unit 0 description B->D
user@B# set unit 0 family inet address 172.16.1.1/24
```

2. On Device B, configure the static route with two next hops, both with BFD enabled.

```
[edit routing-options static route 192.168.47.0/24]
user@B# set qualified-next-hop 192.168.2.2 bfd-liveness-detection minimum-interval
60
user@B# set qualified-next-hop 172.16.1.2 bfd-liveness-detection minimum-interval
60
```

3. On Device D, configure the interfaces.

```
[edit interfaces fe-0/1/0]
user@D# set unit 3 description secondary-D->B
user@D# set unit 3 family inet address 192.168.2.2/24
```

```
[edit interfaces ge-1/2/0]
user@D# set unit 1 description D->B
user@D# set unit 1 family inet address 172.16.1.2/24
```

4. On Device D, configure a BFD-enabled default static route with two next hops to the provider network.

In this case, BFD is enabled on the route, not on the next hops.

```
[edit routing-options static route 0.0.0.0/0]
user@D# set qualified-next-hop 192.168.2.1
user@D# set qualified-next-hop 172.16.1.1
user@D# set bfd-liveness-detection minimum-interval 60
```

**Results** Confirm your configuration by issuing the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/1/0 {
  unit 2 {
    description secondary-B->D;
    family inet {
      address 192.168.2.1/24;
    }
  }
}
```

```
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}

user@B# show routing-options
static {
  route 192.168.47.0/24 {
    qualified-next-hop 192.168.2.2 {
      bfd-liveness-detection {
        minimum-interval 60;
      }
    }
    qualified-next-hop 172.16.1.2 {
      bfd-liveness-detection {
        minimum-interval 60;
      }
    }
  }
}

user@D# show interfaces
fe-0/1/0 {
  unit 3 {
    description secondary-D->B;
    family inet {
      address 192.168.2.2/24;
    }
  }
}
ge-1/2/0 {
  unit 1 {
    description D->B;
    family inet {
      address 172.16.1.2/24;
    }
  }
}

user@D# show routing-options
static {
  route 0.0.0.0/0 {
    qualified-next-hop 192.168.2.1;
    qualified-next-hop 172.16.1.1;
    bfd-liveness-detection {
      minimum-interval 60;
    }
  }
}
```

If you are done configuring the devices, enter **commit** from configuration mode.



**Verification**

Confirm that the configuration is working properly.

- [Checking the Routing Tables on page 2927](#)
- [Verifying the BFD Sessions on page 2927](#)
- [Removing BFD from Device D on page 2927](#)
- [Removing BFD from One Next Hop on page 2928](#)

**Checking the Routing Tables**

**Purpose** Make sure that the static route appears in the routing table on Device B with two possible next hops.

**Action** user@B> show route 192.168.47.0 extensive  
 inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
 192.168.47.0/24 (1 entry, 1 announced)  
 TSI:  
 KRT in-kernel 192.168.47.0/24 -> {192.168.2.2}  
     \*Static Preference: 5  
         Next hop type: Router  
         Address: 0x9334010  
         Next-hop reference count: 1  
         Next hop: 172.16.1.2 via ge-1/2/0.0  
         Next hop: 192.168.2.2 via fe-0/1/0.2, selected  
         State: <Active Int Ext>  
         Age: 9  
         Task: RT  
         Announcement bits (1): 3-KRT  
         AS path: I

**Meaning** Both next hops are listed. The next hop 192.168.2.2 is the selected route.

**Verifying the BFD Sessions**

**Purpose** Make sure that the BFD sessions are up.

**Action** user@B> show bfd session

| Address     | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2  | Up    | ge-1/2/0.0 | 0.720       | 0.240             | 3          |
| 192.168.2.2 | Up    | fe-0/1/0.2 | 0.720       | 0.240             | 3          |

2 sessions, 2 clients  
 Cumulative transmit rate 8.3 pps, cumulative receive rate 8.3 pps

**Meaning** The output shows that the BFD sessions are up.

**Removing BFD from Device D**

**Purpose** Demonstrate what happens when the BFD session is down for both next hops.

- Action** 1. Deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

| Address     | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2  | Down  | ge-1/2/0.0 | 3.000       | 1.000             | 3          |
| 192.168.2.2 | Down  | fe-0/1/0.2 | 3.000       | 1.000             | 3          |

```
2 sessions, 2 clients
```

```
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps
```

3. Rerun the **show route 192.168.47.0** command on Device B.

```
user@B> show route 192.168.47.0
```

**Meaning** As expected, when the BFD sessions are down, the static route is removed from the routing table.

### *Removing BFD from One Next Hop*

**Purpose** Demonstrate what happens when only one next hop has BFD enabled.

- Action** 1. If it is not already deactivated, deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Deactivate BFD on one of the next hops on Device B.

```
[edit routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2]
user@B# deactivate bfd-liveness-detection
user@B# commit
```

3. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session
```

| Address     | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|------------|-------------|-------------------|------------|
| 192.168.2.2 | Down  | fe-0/1/0.2 | 3.000       | 1.000             | 3          |

4. Rerun the **show route 192.168.47.0 extensive** command on Device B.

```
user@B> show route 192.168.47.0 extensive
```

```
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

```
192.168.47.0/24 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.47.0/24 -> {172.16.1.2}
```

```
*Static Preference: 5
```

```
Next hop type: Router, Next hop index: 624
```

```
Address: 0x92f0178
```

```
Next-hop reference count: 3
```

```

Next hop: 172.16.1.2 via ge-1/2/0.0, selected
State: <Active Int Ext>
Age: 2:36
Task: RT
Announcement bits (1): 3-KRT
AS path: I

```

**Meaning** As expected, the BFD session is down for the 192.168.2.2 next hop. The 172.16.1.2 next hop remains in the routing table, and the route remains active, because BFD is not a condition for this next hop to remain valid.

- Related Documentation**
- [Example: Configuring BFD Authentication for Static Routes on page 2929](#)
  - [Example: Configuring BFD for OSPF on page 4123](#)
  - [Example: Configuring BFD for BGP on page 3462](#)
  - [Example: Configuring BFD for IS-IS](#)
  - [Configuring PIM and the Bidirectional Forwarding Detection \(BFD\) Protocol](#)

## Example: Configuring BFD Authentication for Static Routes

- [Understanding BFD Authentication for Static Routes on page 2929](#)
- [Example: Configuring BFD Authentication for Static Routes on page 2931](#)

### Understanding BFD Authentication for Static Routes

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.



**NOTE:** We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IPv4 and IPv6 static routes. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 2930](#)
- [Security Authentication Keychains on page 2930](#)
- [Strict Versus Loose Authentication on page 2931](#)

### ***BFD Authentication Algorithms***

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

### ***Security Authentication Keychains***

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and

associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### ***Strict Versus Loose Authentication***

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Example: Configuring BFD Authentication for Static Routes**

---

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for static routes.

- [Requirements on page 2931](#)
- [Overview on page 2931](#)
- [Configuration on page 2932](#)
- [Verification on page 2935](#)

#### ***Requirements***

Junos OS Release 9.6 or later (Canda and United States version).

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

#### ***Overview***

You can configure authentication for BFD sessions running over IPv4 and IPv6 static routes. Routing instances and logical systems are also supported.

The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the static route.
2. Associate the authentication keychain with the static route.
3. Configure the related security authentication keychain. This must be configured on the main router.




**TIP:** We recommend that you specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

[edit]

```
user@host> set routing-options static route ipv4 bfd-liveness-detection
authentication loose-check
```

Figure 58 on page 2932 shows the sample network.

**Figure 58: Customer Routes Connected to a Service Provider**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

#### Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
key-chain bfd-kc4
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
algorithm keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
"$9$JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
DkP5Ft0IQFcleV7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
"2011-1-12:00:00 -0800"
```

**Device D**

```

set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication key-chain
  bfd-kc4
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication algorithm
  keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
  "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
  DkP5FtOIQFclev7N"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
  "2011-1-1.12:00:00 -0800"

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```

[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24

```

```

user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32

```

2. On Device B, create a static route and set the next-hop address.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2

```

3. On Device B, configure BFD for the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000

```

4. On Device B, specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on the static route.

```

[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
  algorithm keyed-sha-1

```



**NOTE:** Nonstop active routing (NSR) is not supported with the **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

5. On Device B, specify the keychain to be used to associate BFD sessions on the specified route with the unique security authentication keychain attributes.

This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
key-chain bfd-kc4
```

6. On Device B, specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 5.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-kc4]
user@B# set key 5 secret
"$9$JhZHmn6Ap0In/9ApOcSs24oaZikPft3wY24ZG.mz36AtOIeyMWxSrlKvM-dbs2a
DkP5Ft0IQFclev7N"
user@B# set key 5 start-time "2011-1-12:00:00 -0800"
```

7. If you are done configuring Device B, commit the configuration.

```
[edit]
user@B# commit
```

8. Repeat the configuration on Device D.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

### Results

Confirm your configuration by issuing the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device B user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
  }
}
lo0 {
  unit 57 {
    family inet {
      address 10.0.0.1/32;
    }
  }
}
```



```

        address 10.0.0.2/32;
    }
}

user@B# show routing-options
static {
    route 192.168.47.0/24 {
        next-hop 172.16.1.2;
        bfd-liveness-detection {
            minimum-interval 1000;
            authentication {
                key-chain bfd-kc4;
                algorithm keyed-sha-1;
            }
        }
    }
}

user@B# show security
authentication-key-chains {
    key-chain bfd-kc4 {
        key 5 {
            secret
            "$9$JhZHmn6Ap0In/9ApOcSs24oaZikPfT3wY24ZG.mz36AtOIEyMWxSrlKvM-dbs2a
            DkP5FtOIQFclev7N"; ## SECRET-DATA
            start-time "2011-1-1.12:00:00 -0800";
        }
    }
}

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That BFD Sessions Are Up on page 2935](#)
- [Viewing Details About the BFD Session on page 2936](#)
- [Viewing Extensive BFD Session Information on page 2936](#)

### Verifying That BFD Sessions Are Up

**Purpose** Verify that the BFD sessions are up.

**Action** From operational mode, enter the **show bfd session** command.

```
user@B> show bfd session
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2 | Up    | ge-1/2/0.0 | 3.000       | 1.000             | 3          |

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** The command output shows that the BFD session is up.

**Viewing Details About the BFD Session**

**Purpose** View details about the BFD sessions and make sure that authentication is configured.

**Action** From operational mode, enter the **show bfd session detail** command.

```
user@B> show bfd session detail
```

| Address    | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|------------|-------|------------|-------------|-------------------|------------|
| 172.16.1.2 | Up    | ge-1/2/0.0 | 3.000       | 1.000             | 3          |

Client Static, TX interval 1.000, RX interval 1.000, **Authenticate**  
 Session up time 00:53:58  
 Local diagnostic NbrSignal, remote diagnostic None  
 Remote state Up, version 1  
 Logical system 9, routing table index 22

1 sessions, 1 clients  
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured.

**Viewing Extensive BFD Session Information**

**Purpose** View more detailed information about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
```

| Address    | State | Interface  | Time  | Interval | Multiplier |
|------------|-------|------------|-------|----------|------------|
| 172.16.1.2 | Up    | ge-1/2/0.0 | 3.000 | 1.000    | 3          |

Client Static, TX interval 1.000, RX interval 1.000, **Authenticate**  
 keychain bfd-kc4, algo keyed-sha-1, mode strict  
 Session up time 01:39:45  
 Local diagnostic NbrSignal, remote diagnostic None  
 Remote state Up, version 1  
 Logical system 9, routing table index 22  
 Min async interval 1.000, min slow interval 1.000  
 Adaptive async TX interval 1.000, RX interval 1.000  
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
 Local discriminator 3, remote discriminator 4  
 Echo mode disabled/inactive  
 Authentication enabled/active, keychain bfd-kc4, algo keyed-sha-1, mode strict

1 sessions, 1 clients  
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured. The output for the **extensive** command provides the keychain name, the authentication algorithm, and the mode for each client in the session.

**Related Documentation**

- [Examples: Configuring BFD for Static Routes on page 2914](#)

---

## Configuration Statements

---

- [active](#) on page 2940
- [aggregate \(Routing\)](#) on page 2941
- [as-path \(Routing Options\)](#) on page 2943
- [autonomous-system](#) on page 2945
- [backup-pe-group](#) on page 2947
- [backups](#) on page 2948
- [bandwidth \(Multicast Flow Map\)](#) on page 2949
- [bfd-liveness-detection \(Routing Options Static Route\)](#) on page 2950
- [bgp-orf-cisco-mode](#) on page 2954
- [bmp](#) on page 2956
- [brief](#) on page 2958
- [centralized](#) on page 2959
- [community \(Routing Options\)](#) on page 2960
- [confederation](#) on page 2962
- [disable \(Routing Options\)](#) on page 2963
- [description \(Routing Instances\)](#) on page 2963
- [discard](#) on page 2964
- [export \(Routing Options\)](#) on page 2965
- [export-rib](#) on page 2966
- [fate-sharing](#) on page 2968
- [flow](#) on page 2969
- [flow-map](#) on page 2970
- [forwarding-cache \(Flow Maps\)](#) on page 2971
- [forwarding-cache \(Multicast\)](#) on page 2972
- [forwarding-options \(chassis\)](#) on page 2974
- [forwarding-table](#) on page 2975
- [generate](#) on page 2976
- [graceful-restart \(Enabling Globally\)](#) on page 2978
- [import \(Routing Options\)](#) on page 2979
- [import-policy](#) on page 2980
- [import-rib](#) on page 2981
- [indirect-next-hop](#) on page 2982
- [install \(Routing Options\)](#) on page 2983
- [instance-export](#) on page 2984
- [instance-import](#) on page 2984

- [instance-type](#) on page 2985
- [interface \(Multicast Static Routes\)](#) on page 2986
- [interface \(Routing Instances\)](#) on page 2987
- [interface \(Routing Options\)](#) on page 2988
- [interface-routes](#) on page 2989
- [local-address \(Routing Options\)](#) on page 2990
- [martians](#) on page 2991
- [maximum-bandwidth \(Routing Options\)](#) on page 2992
- [maximum-paths](#) on page 2993
- [maximum-prefixes](#) on page 2995
- [med-igp-update-interval](#) on page 2996
- [metric \(Aggregate, Generated, or Static Route\)](#) on page 2997
- [multicast \(Routing Options\)](#) on page 2998
- [no-qos-adjust](#) on page 2999
- [num-65-127-prefix](#) on page 3000
- [options \(Routing Options\)](#) on page 3001
- [pim-to-igmp-proxy](#) on page 3002
- [pim-to-mld-proxy](#) on page 3003
- [policy \(Aggregate and Generated Routes\)](#) on page 3004
- [policy \(Flow Maps\)](#) on page 3005
- [policy-options](#) on page 3006
- [policy-statement](#) on page 3007
- [ppm](#) on page 3011
- [ppm \(Ethernet Switching\)](#) on page 3012
- [preference \(Routing Options\)](#) on page 3013
- [prefix](#) on page 3014
- [prefix-65-127-disable](#) on page 3014
- [protocols](#) on page 3015
- [qualified-next-hop \(Static Routes\)](#) on page 3017
- [readvertise](#) on page 3019
- [redundant-sources](#) on page 3020
- [resolution](#) on page 3021
- [resolution-ribs](#) on page 3022
- [resolve](#) on page 3023
- [restart-duration \(Routing Options\)](#) on page 3024
- [retain](#) on page 3025
- [reverse-oif-mapping](#) on page 3026

- [rpf-check-policy \(Routing Options RPF\) on page 3027](#)
- [rib \(General\) on page 3028](#)
- [rib \(Route Resolution\) on page 3030](#)
- [rib-group \(Routing Options\) on page 3031](#)
- [rib-groups on page 3032](#)
- [route-distinguisher-id on page 3034](#)
- [route-record on page 3035](#)
- [router-id on page 3036](#)
- [routing-instances on page 3037](#)
- [routing-options on page 3038](#)
- [scope on page 3038](#)
- [scope-policy on page 3039](#)
- [source \(Source-Specific Multicast\) on page 3040](#)
- [source-routing on page 3041](#)
- [ssm-groups on page 3042](#)
- [ssm-map \(Routing Options Multicast\) on page 3043](#)
- [static \(Routes\) on page 3044](#)
- [subscriber-leave-timer on page 3046](#)
- [tag \(Routing Options\) on page 3047](#)
- [threshold \(Multicast Forwarding Cache\) on page 3048](#)
- [timeout \(Flow Maps\) on page 3049](#)
- [timeout \(Multicast\) on page 3050](#)
- [traceoptions \(Routing Options\) on page 3051](#)
- [upstream-interface on page 3054](#)

## active

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | (active   passive);  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with <b>reject</b> next hops when they are inactive.</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Remove a route from the routing and forwarding tables when it becomes inactive.</li> <li>• <b>passive</b>—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.</li> </ul> <p>Include the <b>active</b> statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>passive</b> option specified in the <b>defaults</b> portion of the statement.</p>  |
| <b>Default</b>                  | active   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> </ul>   |

## aggregate (Routing)

|                     |   |
|---------------------|---|
| Syntax              | <pre> aggregate {   defaults {     ... aggregate-options ...   }   route destination-prefix {     policy policy-name;     ... aggregate-options ...   } } </pre>  |
| Hierarchy Level     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options <b>rib</b> <i>routing-table-name</i>]</p>  |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>  |
| Description         | Configure aggregate routes.   |
| Options             | <p><b>aggregate-options</b>—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in <b>aggregate-options</b>. Each option is explained separately.</p> <ul style="list-style-type: none"> <li>• (<b>active</b>   <b>passive</b>);</li> <li>• <b>as-path</b> &lt;<i>as-path</i>&gt; &lt;origin (egp   igp   incomplete)&gt; &lt;atomic-aggregate&gt; &lt;aggregator <i>as-number</i> <i>ip-address</i>&gt;;</li> <li>• (<b>brief</b>   <b>full</b>);</li> <li>• <b>community</b> [ <i>community-ids</i> ];</li> <li>• <b>discard</b>;</li> <li>• (<b>metric</b>   <b>metric2</b>   <b>metric3</b>   <b>metric4</b>) <i>value</i> &lt;type <i>type</i>&gt;;</li> <li>• (<b>preference</b>   <b>preference2</b>   <b>color</b>   <b>color2</b>) <i>preference</i> &lt;type <i>type</i>&gt;;</li> <li>• <b>tag</b> <i>metric type number</i>;</li> </ul> <p><b>defaults</b>—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults</p> |

and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

**route *destination-prefix***—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.   |
|                                 | routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li></ul> |



## as-path (Routing Options)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>as-path &lt;as-path&gt; &lt;aggregator as-number ip-address&gt; &lt;atomic-aggregate&gt; &lt;origin (egp   igp   incomplete)&gt;;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>         | <p>Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.</p> <p>In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i>&lt;16-bit high-order value in decimal&gt;.&lt;16-bit low-order value in decimal&gt;</i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.</p>  |
| <b>Default</b>             | No AS path information is associated with static routes.   |
| <b>Options</b>             | <p><b>aggregator</b>—(Optional) Attach the BGP <b>aggregator</b> path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for <i>as-number</i>, followed by the IP address of the BGP system that formed the aggregate route for <i>ip-address</i>.</p>  |

**as-path**—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ( [ ] ). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for **as-path**. You must use a complete, valid AS path.

**atomic-aggregate**—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

**origin egp**—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

**origin igp**—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

**origin incomplete**—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Examples: Configuring Static Routes</i></li><li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li><li>• <i>Example: Conditionally Generating Static Routes</i></li><li>• <i>Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview</i></li></ul> |
|------------------------------|--|

## autonomous-system

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>autonomous-system <i>autonomous-system</i> &lt;asdot-notation&gt; &lt;loops <i>number</i>&gt; {     independent-domain &lt;no-attrset&gt;; }</pre>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>asdot-notation</b> option introduced in Junos OS Release 9.3.</p> <p><b>asdot-notation</b> option introduced in Junos OS Release 9.3 for EX Series switches.</p> <p><b>no-attrset</b> option introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>  |
| <b>Description</b>         | <p>Specify the routing device's AS number.</p> <p>An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (<a href="http://www.isi.edu">http://www.isi.edu</a>).</p> <p>If you are using BGP on the routing device, you must configure an AS number.</p> <p>The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.</p> <p>In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i>&lt;16-bit high-order value in decimal&gt;.&lt;16-bit low-order value in decimal&gt;</i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> |
| <b>Options</b>             | <p><b><i>autonomous-system</i></b>—AS number. Use a number assigned to you by the NIC.</p>  |

**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

**Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

**Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

**asdot-notation**—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

**Default:** Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

**loops number**—(Optional) Specify the number of times detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

**Range:** 1 through 10

**Default:** 1



**NOTE:** When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the loops statement must be enabled only on a subset of routing instances.

---

The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3261](#)
- [Examples: Configuring Internal BGP Peering on page 3284](#)

## backup-pe-group

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>backup-pe-group <i>group-name</i> {<br/>    <b>backups</b> [ <i>addresses</i> ];<br/>    <b>local-address</b> <i>address</i>;<br/>}</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>              | Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.  |
| <b>Options</b>                  | <b>backups <i>addresses</i></b> —Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.<br><br><b>local-address <i>address</i></b> —Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.<br><br><b>pe-group-name</b> —Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ingress PE Redundancy</a></li> <li>• <a href="#">Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs</a></li> </ul>  |

## backups

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>backups [ <i>addresses</i> ];</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>backup-pe-group</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>backup-pe-group</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>backup-pe-group</b> <i>group-name</i> ],<br>[edit routing-options multicast <b>backup-pe-group</b> <i>group-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.   |
| <b>Options</b>                  | <b><i>addresses</i></b> —Addresses of other PEs in the backup group.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Ingress PE Redundancy</i></li></ul>   |

## bandwidth (Multicast Flow Map)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>bandwidth ( <i>bps</i>   adaptive );</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">flow-map</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">flow-map</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">flow-map</a>],</p> <p>[edit routing-options multicast <a href="#">flow-map</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | Configure the bandwidth property for multicast flow maps.  |
| <b>Options</b>                  | <p><b>adaptive</b>—Specify that the bandwidth is measured for the flows that are matched by the flow map.</p> <p><b><i>bps</i></b>—Bandwidth, in bits per second, for the flow map.</p> <p><b>Range:</b> 0 through any amount of bandwidth</p> <p><b>Default:</b> 2 Mbps</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring a Multicast Flow Map</i></li> </ul>   |

## bfd-liveness-detection (Routing Options Static Route)

**Syntax** `bfd-liveness-detection {`  
     `authentication {`  
         `algorithm` *algorithm-name*;  
         `key-chain` *key-chain-name*;  
         `loose-check`;  
     `}`  
     `detection-time {`  
         `threshold` *milliseconds*;  
     `}`  
     `holddown-interval` *milliseconds*;  
     `local-address` *ip-address*;  
     `minimum-interval` *milliseconds*;  
     `minimum-receive-interval` *milliseconds*;  
     `minimum-receive-ttl` *number*;  
     `multiplier` *number*;  
     `neighbor` *address*;  
     `no-adaptation`;  
     `transmit-interval {`  
         `minimum-interval` *milliseconds*;  
         `threshold` *milliseconds*;  
     `}`  
     `version` (1 | automatic);  
`}`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix*],  
 [edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],  
 [edit logical-systems *logical-system-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],  
 [edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],  
 [edit routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit routing-options rib *routing-table-name* static route *destination-prefix*],  
 [edit routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],  
 [edit routing-options static route *destination-prefix*],



[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

|                            |   |
|----------------------------|---|
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>detection-time threshold</b> and <b>transmit-interval threshold</b> options introduced in Junos OS Release 8.2.</p> <p><b>local-address</b> statement introduced in Junos OS Release 8.2.</p> <p><b>minimum-receive-ttl</b> statement introduced in Junos OS Release 8.2.</p> <p>Support for logical routers introduced in Junos OS Release 8.3.</p> <p><b>holddown-interval</b> statement introduced in Junos OS Release 8.5.</p> <p><b>no-adaptation</b> statement introduced in Junos OS Release 9.0.</p> <p>Support for IPv6 static routes introduced in Junos OS Release 9.1.</p> <p><b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> statements introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> |
| <b>Description</b>         | <p>Configure bidirectional failure detection timers and authentication criteria for static routes.</p>  |

**Options**    **authentication algorithm** *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

**authentication key-chain** *key-chain-name* —Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold** *milliseconds*—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**holddown-interval** *milliseconds*—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

**Range:** 0 through 255,000

**Default:** 0

**local-address** *ip-address*—Enable a multihop BFD session and configure the source address for the BFD session.

**minimum-interval** *milliseconds*—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000

**minimum-receive-interval** *milliseconds*—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route destination-prefix bfd-liveness-detection]** hierarchy level.

**Range:** 1 through 255,000

**minimum-receive-ttl** *number*—Configure the time to live (TTL) for the multihop BFD session.

**Range:** 1 through 255

**Default:** 255

**multiplier** *number*—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**neighbor *address***—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

**no-adaptation**—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295

**transmit-interval minimum-interval *milliseconds***—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

**Range:** 1 through 255,000


**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li> </ul> |
|------------------------------|---|

## bgp-orf-cisco-mode

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | bgp-orf-cisco-mode;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit routing-options <b>outbound-route-filter</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>  |
| <b>Description</b>              | Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.  |
|                                 | <p> <b>NOTE:</b> To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>  |
| <b>Default</b>                  | Disabled   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3429](#)

## bmp

---

```
Syntax  bmp {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
station station-name {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
        seconds;
        flaps flaps;
        period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
station-address (ip-address | name);
station-port port-number;
statistics-timeout seconds;
traceoptions {
```

```

        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
    }
}
station-address (ip-address | name);
station-port port-number;
statistics-timeout seconds;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier>;
}
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [bgp](#)],  
 [edit logical-systems *logical-system-name* protocols bgp [group](#) *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* [neighbor](#) *address*],  
 [edit logical-systems *logical-system-name* routing-options],  
 [edit protocols [bgp](#)],  
 [edit protocols bgp [group](#) *group-name*],  
 [edit protocols bgp group *group-name* [neighbor](#) *address*],  
 [edit routing-options]

**Release Information** Statement introduced in Junos OS Release 9.5.  
 Statement introduced in Junos OS Release 9.5 for EX Series switches.  
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.  
 Support for BMP version 3 introduced in Junos OS Release 13.3.

**Description** Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station.

**Options** The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**


- *Example: Configuring the BGP Monitoring Protocol*

## brief

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | (brief   full);  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.</p> <ul style="list-style-type: none"> <li>• <b>brief</b>—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP <b>ATOMIC_ATTRIBUTE</b> path attribute is included with the aggregate route.</li> <li>• <b>full</b>—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the <b>route</b> portion of the <b>generate</b> statement to override a <b>retain</b> option specified in the <b>defaults</b> portion of the statement.</li> </ul>   |
| <b>Default</b>                  | full   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 2941</a></li> <li>• <a href="#">generate on page 2976</a></li> </ul>  |



## centralized

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | centralized;   |
| <b>Hierarchy Level</b>          | [edit protocols lacp ppm]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | <p>Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.</p> <p>This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the <b>no-delegate-processing</b> statement in the [edit routing-options ppm] hierarchy.</p> |
|                                 | <div>  <p><b>BEST PRACTICE:</b> We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.</p> </div>   |
| <b>Default</b>                  | Distributed PPM processing is enabled for all packets that use PPM.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i></li> <li>• <i>Configuring Aggregated Ethernet LACP (CLI Procedure)</i></li> <li>• <a href="#">Configuring Distributed Periodic Packet Management on page 2906</a></li> <li>• <a href="#">Configuring Link Aggregation on page 2593</a></li> </ul>  |

## community (Routing Options)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>community ([ <i>community-ids</i> ]   no-advertise   no-export   no-export-subconfed   none);</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>   static) (defaults   route)]</p>         |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>         | Associate BGP community information with a static, aggregate, or generated route.  |
| <b>Default</b>             | No BGP community information is associated with static routes.   |
| <b>Options</b>             | <p><b><i>community-ids</i></b>—One or more community identifiers. The <b><i>community-ids</i></b> format varies according to the type of attribute that you use.</p> <p>The BGP community attribute format is <b><i>as-number:community-value</i></b>:</p> <ul style="list-style-type: none"> <li>• <b><i>as-number</i></b>—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value.</li> <li>• <b><i>community-value</i></b>—Identifier of the community member. It can be a number from 0 through 65,535.</li> </ul> <p>For more information about BGP community attributes, see the “Configuring the Extended Communities Attribute” section in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>For specifying the BGP community attribute only, you also can specify <b><i>community-ids</i></b> as one of the following well-known community names defined in RFC 1997:</p> <ul style="list-style-type: none"> <li>• <b>no-advertise</b>—Routes containing this community name are not advertised to other BGP peers.</li> <li>• <b>no-export</b>—Routes containing this community name are not advertised outside a BGP confederation boundary.</li> </ul> |

- **no-export-subconfed**—Routes containing this community are advertised to IBGP peers with the same AS number, but not to members of other confederations.



**NOTE:** Extended community attributes are not supported at the [edit routing-options] hierarchy level. You must configure extended communities at the [edit policy-options] hierarchy level. For information about configuring extended communities, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 2941</a></li> <li>• <a href="#">generate on page 2976</a></li> <li>• <i>static</i></li> </ul> |

## confederation

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>confederation <i>confederation-autonomous-system</i> members [ <i>autonomous-systems</i> ];</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | <p>Specify the routing device's confederation AS number.</p> <p>If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more <i>confederations</i>. Each confederation is identified by its own AS number, which is called a <i>confederation AS number</i>. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.</p> <p>The BGP path attributes <b>NEXT_HOP</b>, <b>LOCAL_PREF</b>, and <b>MULTI_EXIT_DISC</b>, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.</p> <p>Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.</p> <p>Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.</p> <p>If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the <b>confederation</b> statement.</p> <p>Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.</p> <p>Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.</p> |
| <b>Options</b>                  | <p><b><i>autonomous-systems</i></b>—AS numbers of the confederation members.<br/><b>Range:</b> 1 through 65,535</p> <p><b><i>confederation-autonomous-system</i></b>—Confederation AS number. Use one of the numbers assigned to you by the NIC.<br/><b>Range:</b> 1 through 65,535</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |

- Related Documentation**
- [Example: Configuring BGP Confederations on page 3564](#)

## disable (Routing Options)

---


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | disable;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-options graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],<br>[edit routing-options graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>              | Disable graceful restart.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Junos OS High Availability Library for Routing Devices</a></li> </ul>  |

## description (Routing Instances)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | description <i>text</i> ;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 11.1 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                                    |
| <b>Description</b>              | Provide a text description for the routing instance. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the <b>show route instance detail</b> command and has no effect on the operation of the routing instance. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Routing Instances on PE Routers in VPNs</a></li> <li>• <a href="#">show route instance on page 3163</a></li> </ul>   |

## discard

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | discard;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.</p> <p>To propagate static routes into the routing protocols, include the <b>discard</b> statement when you define the route, along with a routing policy.</p>  |
|                                 | <p> <b>NOTE:</b> In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as <b>route <i>address/mask-length</i> next-hop 127.0.0.1</b>) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.</p>  |
| <b>Default</b>                  | When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> </ul>   |

- [aggregate on page 2941](#)
- [generate on page 2976](#)

## export (Routing Options)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>export [ <i>policy-name</i> ];</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options forwarding-table],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options forwarding-table],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options forwarding-table],</p> <p>[edit routing-options forwarding-table]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Apply one or more policies to routes being exported from the routing table into the forwarding table.</p> <p>In the <b>export</b> statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.</p> <p>You can reference the same routing policy one or more times in the same or a different <b>export</b> statement.</p> <p>You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:</p> <ul style="list-style-type: none"> <li>• Per-packet load balancing</li> <li>• Class of service (CoS)</li> </ul> |
| <b>Options</b>                  | <i>policy-name</i> —Name of one or more policies.  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Load Balancing BGP Traffic on page 3477</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> <li>• <i>How a Routing Policy Is Evaluated</i></li> </ul>  |

## export-rib

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>export-rib <i>routing-table-name</i>;</code>  |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i></code><br><code>routing-options <b>rib-groups</b> <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</code><br><code>[edit routing-options <b>rib-groups</b> <i>group-name</i>]</code> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>  |
| <b>Description</b>         | <p>Specify the name of the routing table from which Junos OS should export routing information. For any individual RIB group, only one table can be specified in the <b>export-rib</b> statement.</p>   |

The **export-rib** statement specifies the source table from which routing information is advertised.

One common use of the **export-rib** statement is interdomain routing. The export RIB is the table used when BGP extracts routes to advertise to peers. In multicast interdomain routing, for example, the export RIB is likely to be inet.2.

Another use of **export-rib** is dynamic route leaking between the global routing table (inet.0) and a VRF routing table (*instance.inet.0*). For example, you can use a RIB group to copy routes learned in the VRF into the global routing table, inet.0, or copy routes learned in inet.0 into a VRF. You define the use of this RIB group in the VRF's BGP configuration. In a routing policy you can do dynamic filtering of routes. For instance, you can use an import policy to only copy routes with certain communities into the global routing table.

For example:

```
rib-groups {
  rib-interface-routes-v4 {
    import-rib [ inet.0 VRF.inet.0 ];
  }
  rib-import-VRF-routes-to-inet0-v4 {
    export-rib VRF.inet.0;
    import-rib [ VRF.inet.0 inet.0 ];
    import-policy rib-import-VRF-routes-to-inet0-v4;
  }
  rib-import-inet0-routes-to-VRF-v4 {
    export-rib inet.0;
    import-rib [ inet.0 VRF.inet.0 ];
    import-policy rib-import-inet0-routes-to-VRF-v4;
  }
}
routing-options {
  interface-routes {
    rib-group {
```



```

        inet rib-interface-routes-v4;
    }
}
protocols {
    bgp {
        group iBGP-peers {
            type internal;
            family inet {
                unicast {
                    rib-group rib-import-inet0-routes-to-VRF-v4;
                }
            }
        }
    }
}
routing-instances {
    VRF {
        routing-options {
            interface-routes {
                rib-group {
                    inet rib-interface-routes-v4;
                }
            }
        }
        protocols {
            bgp {
                group peersin-VRF {
                    family inet {
                        unicast {
                            rib-group rib-import-VRF-routes-to-inet0-v4;
                        }
                    }
                }
            }
        }
    }
}
}

```

**Options** *routing-table-name*—Routing table group name.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- *Example: Configuring a PIM RPF Routing Table*
- *Example: Configuring DVMRP to Announce Unicast Routes*
- *Example: Configuring a Dedicated PIM RPF Routing Table*
- *Example: Configuring Any-Source Multicast for Draft-Rosen VPNs*
- [import-rib on page 2981](#)
- *passive*

## fate-sharing

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>fate-sharing {<br/>    group <i>group-name</i> {<br/>        cost <i>value</i>;<br/>        from <i>address</i> &lt;to <i>address</i>&gt;;<br/>    }<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | <p>Specify a backup path in case the primary path becomes unusable.</p> <p>You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.</p> <p>Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.</p>   |
| <b>Options</b>                  | <p><b>cost <i>value</i></b>—Cost assigned to the group.<br/><b>Range:</b> 1 through 65,535<br/><b>Default:</b> 1</p> <p><b>from <i>address</i></b>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate <b>from</b> addresses in the group.</p> <p><b>group <i>group-name</i></b>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p><b>to <i>address</i></b>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a <b>from</b> and a <b>to</b> address.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Ingress Router for MPLS-Signaled LSPs</i></li><li>• <i>Junos OS MPLS Applications Library for Routing Devices</i></li></ul>   |

## flow

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> flow {     route <i>name</i> {         match {             <i>match-conditions</i>;         }         term-order (legacy   standard);         then {             <i>actions</i>;         }     }     firewall-install-disable;     term-order (legacy   standard);     validation {         traceoptions {             file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;             flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;;         }     } } </pre>  |
| <b>Hierarchy Level</b>     | <p>[edit routing-options],<br/> [edit logical-systems <i>logical-system-name</i> routing-options],<br/> [edit routing-instances <i>routing-instance-name</i> routing-options],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options]</p>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> <b>term-order</b> statement introduced in Junos OS Release 10.0<br/> Statement introduced in Junos OS Release 11.3 for the QFX Series.<br/> <b>firewall-install-disable</b> statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.</p>   |
| <b>Description</b>         | Configure a flow route.  |
| <b>Default</b>             | legacy   |
| <b>Options</b>             | <p><b>actions</b>—An action to take if conditions match.</p> <p><b>firewall-install-disable</b>—(PTX Series routers only) Disable installing flow-specification firewall filters in the firewall process (dfwd).</p> <p><b>Default:</b> If you omit the <b>firewall-install-disable</b> statement, the default behavior is <b>firewall-install-disable</b> mode.</p> <p><b>match-conditions</b>—Match packets to these conditions.</p> <p><b>route <i>name</i></b>—Name of the flow route.</p> <p><b>standard</b>—Specify to use version 7 or later of the flow-specification algorithm.</p> |

**term-order (legacy | standard)**—Specify the version of the flow-specification algorithm.

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

**then**—Actions to take on matching packets.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Flow Routes*

---

## flow-map

---

**Syntax**

```
flow-map flow-map-name {  
    bandwidth (bps | adaptive);  
    forwarding-cache {  
        timeout (never non-discard-entry-only | minutes);  
    }  
    policy [ policy-names ];  
    redundant-sources [ addresses ];  
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],  
[edit logical-systems *logical-system-name* routing-options multicast],  
[edit routing-instances *routing-instance-name* routing-options multicast],  
[edit routing-options multicast]

**Release Information** Statement introduced in Junos OS Release 8.2.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure multicast flow maps.

**Options** *flow-map-name*—Name of the flow-map.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring a Multicast Flow Map*

---

## forwarding-cache (Flow Maps)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | forwarding-cache {<br>timeout (minutes   never non-discard-entry-only );<br>}   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Configure multicast forwarding cache properties for the flow map.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |

## forwarding-cache (Multicast)

---

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>forwarding-cache {   allow-maximum;   family (inet   inet6) {     threshold {       log-warning value;       suppress value &lt;reuse value&gt;;     }   }   threshold {     log-warning value;     suppress value &lt;reuse value&gt;;   }   timeout minutes; }</pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>         | <p>Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits, the threshold at which a warning message is logged, and timeout values.</p> <p>Specify a value for the threshold at which to suppress new multicast forwarding cache entries and an optional reuse value for the threshold at which the router begins to create new multicast forwarding cache entries. The range for both is from 1 through 200,000. If configured, the reuse value should be less than the suppression threshold value. The suppression value is mandatory. If you do not specify the optional reuse value, then the number of multicast forwarding cache entries is limited to the suppression value. A new entry is created as soon as the number of multicast forwarding cache entries falls below the suppression value.</p> <p>You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the <b>threshold</b> statement globally for the multicast forwarding cache or including the <b>family</b> statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> |
| <b>Default</b>             | By default, there are no limits on the number of multicast forwarding cache entries.   |
| <b>Options</b>             | <p><b>family (inet   inet6)</b>—(Optional) Apply the configured thresholds to either IPv4 or IPv6 multicast forwarding cache entries.</p> <p><b>Default:</b> By default, the configured thresholds are applied to both IPv4 and IPv6 multicast forwarding cache entries.</p>   |

The remaining statements are explained separately.

|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Multicast Forwarding Cache</i></li></ul> |
|------------------------------|--|

## forwarding-options (chassis)

**Syntax** forwarding options *profile-name* {  
     num-65-127-prefix *value*  
     lpm-profile *prefix-65-127-disable*  
 }

**Hierarchy Level** [edit *chassis*]

**Release Information** Statement introduced in Junos 13.2 for the QFX Series.

**Description** Configure a unified forwarding table profile to allocate the amount a memory available for the following:

- MAC addresses
- Layer 3 host entries
- Longest prefix match table entries

**Options** *profile-name*—name of the profile to use for memory allocation in the unified forwarding table. [Table 112 on page 1724](#) lists the profiles you can choose and the associated values for each type of entry.

**Table 271: Unified Forwarding Table Profiles**

| Profile Name               | MAC Table     | Host Table (unicast and multicast addresses) |              |             |             |             |             |
|----------------------------|---------------|--|--------------|-------------|-------------|-------------|-------------|
|                            | MAC Addresses | IPv4 unicast                                 | IPv6 unicast | IPv4 (*, G) | IPv4 (S, G) | IPv6 (*, G) | IPv6 (S, G) |
| l2-profile-one             | 288K          | 16K  | 8K           | 8K          | 8K          | 4K          | 4K          |
| l2-profile-two             | 224K          | 80K  | 40K          | 40K         | 40K         | 20K         | 20K         |
| l2-profile-three (default) | 160K          | 144K   | 72K          | 72K         | 72K         | 36K         | 36K         |
| l3-profile                 | 96K           | 208K   | 104K         | 104K        | 104K        | 52K         | 52K         |
| lpm-profile*               | 32K           | 16K  | 8K           | 8K          | 8K          | 4K          | 4K          |

\* This profile supports only IPv4 in Junos OS 13.2X51-D10. With Junos OS 13.2X51-D15 it supports IPv4 and IPv6.

Note that if the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see [“Understanding the Unified Forwarding Table” on page 1545](#).

You configure the memory allocation for LPM table entries differently depending on whether you use Junos OS 13.2X51-D10 or Junos OS 13.2X51-D15 and later. To learn



how to configure memory allocation for LPM table entries see [“Configuring the Unified Forwarding Table” on page 1697](#).

**Required Privilege  
Level**

- Related  
Documentation**
- [Understanding the Unified Forwarding Table on page 1545](#)
  - [Configuring the Unified Forwarding Table on page 1697](#)

## forwarding-table

---

**Syntax** forwarding-table {  
     [export](#) [ *policy--names* ];  
     ([indirect-next-hop](#) | no-indirect-next-hop);  
 }

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options],  
 [edit routing-options]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure information about the routing device's forwarding table.

The remaining statements are explained separately.

**Required Privilege  
Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

- Related  
Documentation**
- [Configuring Per-Packet Load Balancing on page 2904](#)

## generate

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> generate {   defaults {     generate-options;   }   route destination-prefix {     policy policy-name;     generate-options;   } } </pre>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options <b>rib</b> <i>routing-table-name</i>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>         | Configure generated routes, which are used as routes of last resort.  |
| <b>Options</b>             | <p><b>defaults</b>—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the <b>generate</b> statement.</p> <p><b>generate-options</b>—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in <b>generate-options</b>. Each option is explained separately.</p> <ul style="list-style-type: none"> <li>• (<b>active</b>   <b>passive</b>);</li> <li>• <b>as-path</b> <i>&lt;as-path&gt;</i> <i>&lt;origin (egp   igp   incomplete)&gt;</i> <i>&lt;atomic-aggregate&gt;</i> <i>&lt;aggregator as-number in-address&gt;</i>;</li> <li>• (<b>brief</b>   <b>full</b>);</li> <li>• <b>community</b> [ <i>community-ids</i> ];</li> <li>• <b>discard</b>;</li> <li>• (<b>metric</b>   <i>metric2</i>   <i>metric3</i>   <i>metric4</i>) <i>value</i> <i>&lt;type type&gt;</i>;</li> <li>• (<b>preference</b>   <i>preference2</i>   <b>color</b>   <i>color2</i>) <i>preference</i> <i>&lt;type type&gt;</i>;</li> <li>• <b>tag</b> <i>metric type number</i>;</li> </ul> <p><b>route destination-prefix</b>—Configure a non-default generated route:</p> <ul style="list-style-type: none"> <li>• <b>default</b>—For the default route to the destination. This is equivalent to specifying an IP address of 0.0.0.0/0.</li> </ul> |

- *destination-prefix/prefix-length—/destination-prefix* is the network portion of the IP address, and *prefix-length* is the destination prefix length.

The [policy](#) statement is explained separately.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.   |
|                                 | routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Conditionally Generating Static Routes</i></li></ul> |

## graceful-restart (Enabling Globally)

---

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>graceful-restart {<br/>    disable;<br/>    helper-disable;<br/>    maximum-helper-recovery-time <i>seconds</i>;<br/>    maximum-helper-restart-time <i>seconds</i>;<br/>    notify-duration <i>seconds</i>;<br/>    recovery-time <i>seconds</i>;<br/>    restart-duration <i>seconds</i>;<br/>    stale-routes-time <i>seconds</i>;<br/>}</pre> |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options]  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>         | Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.   |



### NOTE:

- For VPNs, the `graceful-restart` statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
  - For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
  - LDP sessions flap when `graceful-restart` configurations change.
- 

|                                 |   |
|---------------------------------|---|
| <b>Default</b>                  | Graceful restart is disabled by default.  |
| <b>Options</b>                  | The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling Graceful Restart</a></li><li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2261</a></li></ul> |


- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

## import (Routing Options)


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution <a href="#">rib</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options resolution <a href="#">rib</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options resolution <a href="#">rib</a> ],<br>[edit routing-options resolution <a href="#">rib</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Specify one or more import policies to use for route resolution.  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Route Resolution on PE Routers</i></li> </ul>  |

## import-policy

---


|   |   |
|---|---|
| Syntax  | import-policy [ <i>policy-names</i> ];  |
| Hierarchy Level   | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib-groups</b> <i>group-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i> ],<br>[edit routing-options <b>rib-groups</b> <i>group-name</i> ] |
| Release Information   | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| Description   | Apply one or more policies to routes imported into the routing table group. The <b>import-policy</b> statement complements the <b>import-rib</b> statement and cannot be used unless you first specify the routing tables to which routes are being imported.   |
| <div> <b>NOTE:</b> On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.</div> |   |
| Options   | <i>policy-names</i> —Name of one or more policies.  |
| Required Privilege Level  | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| Related Documentation   | <ul style="list-style-type: none"><li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li><li>• <a href="#">export-rib on page 2966</a></li><li>• <i>passive</i></li></ul>  |

## import-rib

|   |  |
|---|--|
| <b>Syntax</b>   | <code>import-rib [ <i>routing-table-names</i> ];</code>  |
| <b>Hierarchy Level</b>  | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib-groups</b> <i>group-name</i>],</p> <p>[edit routing-options <b>rib-groups</b> <i>group-name</i>]</p>   |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>  | <p>Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be <b>inet.0</b> or <b>routing-instance-name.inet.0</b>. For IPv6 import routing tables, the primary routing table must be <b>inet6.0</b>.</p> <p>In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.</p> |
| <div>  <b>NOTE:</b> On EX Series switches, only dynamically learned routes can be imported from one routing table group to another. </div> |  |
| <b>Options</b>  | <b><i>routing-table-names</i></b> —Name of one or more routing tables.   |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <a href="#">export-rib on page 2966</a></li> <li>• <i>passive</i></li> </ul>   |

## indirect-next-hop

---

|  |  |
|--|--|
| <b>Syntax</b>  | (indirect-next-hop   no-indirect-next-hop);  |
| <b>Hierarchy Level</b>   | [edit logical-systems <i>logical-system-name</i> routing-options forwarding-table],<br>[edit routing-options forwarding-table]   |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>   | Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable. |
| <div> <b>NOTE:</b><ul style="list-style-type: none"><li>• When virtual private LAN service (VPLS) is configured on the routing device, the <b>indirect-next-hop</b> statement is configurable at the [edit routing-options forwarding-table] hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances.</li><li>• By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the <b>no-indirect-next-hop</b> statement.</li><li>• By default, indirectly connected next hops are enabled on PTX Series routers.</li></ul></div> |  |
| <b>Default</b>   | Disabled.  |
| <b>Options</b>   | <b>indirect-next-hop</b> —Enable indirectly connected next hops.<br><b>no-indirect-next-hop</b> —Explicitly disable indirect next hops.  |
| <b>Required Privilege Level</b>  | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <i>Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine</i></li></ul>  |



## install (Routing Options)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | (install   no-install);   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options static (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> static (defaults   route)]</p> <p>[edit routing-options static (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.  |
| <b>Options</b>                  | <p><b>install</b>—Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>no-install</b> option specified in the <b>defaults</b> portion of the statement.</p> <p><b>no-install</b>—Do not install the route into the forwarding table, even if it is the route with the lowest preference.</p> <p><b>Default:</b> install</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>static</i></li> </ul>   |

## instance-export

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>instance-export [ <i>policy-names</i> ];</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Apply one or more policies to routes being exported from a routing instance.  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more export policies.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li></ul>   |

## instance-import

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>instance-import [ <i>policy-names</i> ];</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Apply one or more policies to routes being imported into a routing instance.  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more import policies.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li></ul>   |

---

## instance-type

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | instance-type virtual-router  |
| <b>Hierarchy Level</b>          | [edit <a href="#">routing-instances</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | Specify the type of routing instance.   |
| <b>Options</b>                  | <b>virtual-router</b> —Virtual router routing instance.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches</i></li><li>• <i>Configuring Virtual Routing Instances (CLI Procedure)</i></li><li>• <a href="#">Configuring Virtual Router Routing Instances on page 2908</a></li></ul> |

## interface (Multicast Static Routes)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>interface <i>interface-names</i> {<br/>    disable;<br/>    maximum-bandwidth <i>bps</i>;<br/>    no-qos-adjust;<br/>    reverse-oif-mapping {<br/>        no-qos-adjust;<br/>    }<br/>    subscriber-leave-timer <i>seconds</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>    routing-options <a href="#">multicast</a>],<br/>[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">multicast</a>],<br/>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">multicast</a>],<br/>[edit routing-options <a href="#">multicast</a>]</pre>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Enable multicast traffic on an interface.</p> <p>By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.</p> <p>You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.</p> <p>You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.</p> <p>Static routes must be configured before you can enable multicast on an interface. Configuring the <b>interface</b> statement alone does not install any routes into the routing table. This feature relies on the static route configuration.</p> |
| <b>Options</b>                  | <p><b><i>interface-names</i></b>—Name of one or more interfaces on which to enable multicast traffic.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Defining Interface Bandwidth Maximums</i></li><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>   |

## interface (Routing Instances)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>interface <i>interface-name</i>;</code>                     |
| <b>Hierarchy Level</b>     | [edit <a href="#">routing-instances</a> ]                         |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.3 for the QFX Series. |
| <b>Description</b>         | For virtual router routing instances, configure an interface.     |



### NOTE:

- You must configure only interfaces from the Node devices assigned to the network Node group. If you try to configure interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
- You can configure an interface for one routing instance only.

|                                 |   |
|---------------------------------|---|
| <b>Options</b>                  | <i>interface-name</i> —Name of an interface.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Virtual Router Routing Instances on page 2908</a></li> </ul> |

## interface (Routing Options)

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>interface <i>interface-names</i> {<br/>    <b>maximum-bandwidth</b> <i>bps</i>;<br/>    <b>no-qos-adjust</b>;<br/>    <b>reverse-oif-mapping</b> {<br/>        <b>no-qos-adjust</b>;<br/>    }<br/>    <b>subscriber-leave-timer</b> <i>seconds</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>         | Enable multicast traffic on an interface.   |



**TIP:** You cannot enable multicast traffic on an interface by using the **routing-options multicast interface** statement and configure PIM on the interface.

---

|                                 |   |
|---------------------------------|---|
| <b>Options</b>                  | <b><i>interface-name</i></b> —Names of the physical or logical interface.<br><br>The remaining statements are explained separately.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Defining Interface Bandwidth Maximums</i></li><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul> |

## interface-routes

**Syntax**

```
interface-routes {
    family (inet | inet6) {
        export {
            lan;
            point-to-point;
        }
    }
    rib-group group-name;
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.



**NOTE:** On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

**Description** Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0**, and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

**Options** **inet**—Specify the IPv4 address family.

**inet6**—Specify the IPv6 address family.

**lan**—Export LAN routes.

**point-to-point**—Export point-to-point routes.

The remaining statements are explained separately.

**Required Privilege** **routing**—To view this statement in the configuration.

**Level** **routing-control**—To add this statement to the configuration.

- Related Documentation**
- *Example: Importing Direct and Static Routes Into a Routing Instance*
  - *Example: Configuring Multiple Routing Instances of OSPF*
  - *passive*

---

## local-address (Routing Options)

---

**Syntax** **local-address** *address*;

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],  
[edit logical-systems *logical-system-name* routing-options multicast **backup-pe-group** *group-name*],  
[edit routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],  
[edit routing-options multicast **backup-pe-group** *group-name*]

**Release Information** Statement introduced in Junos OS Release 9.0.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

**Options** **address**—Address of local PEs in the backup group.

**Required Privilege** **routing**—To view this statement in the configuration.

**Level** **routing-control**—To add this statement to the configuration.

- Related Documentation**
- *Example: Configuring Ingress PE Redundancy*



## martians

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>martians {     destination-prefix match-type &lt;allow&gt;; }</pre>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>rib</b> <i>routing-table-name</i>],</p> <p>[edit routing-options],</p> <p>[edit routing-options <b>rib</b> <i>routing-table-name</i>]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | Configure martian addresses.   |
| <b>Options</b>                  | <p><b>allow</b>—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The <b>allow</b> option is the only supported action.</p> <p><b>destination-prefix</b>—Destination route you are configuring:</p> <ul style="list-style-type: none"> <li>• <b>destination-prefix/prefix-length—destination-prefix</b> is the network portion of the IP address, and <b>prefix-length</b> is the destination prefix length.</li> <li>• <b>default</b>—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address <b>0.0.0.0/0</b>.</li> </ul> <p><b>match-type</b>—Criteria that the destination must match:</p> <ul style="list-style-type: none"> <li>• <b>exact</b>—Exactly match the route's mask length.</li> <li>• <b>longer</b>—The route's mask length is greater than the specified mask length.</li> <li>• <b>orlonger</b>—The route's mask length is equal to or greater than the specified mask length.</li> <li>• <b>through destination-prefix</b>—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.</li> <li>• <b>upto prefix-length</b>—The route's mask length falls between the two destination prefix lengths, inclusive.</li> </ul> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

**Related Documentation**    • *Example: Configuring Martian Addresses*

## maximum-bandwidth (Routing Options)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>maximum-bandwidth <i>bps</i>;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit dynamic-profiles <i>profile-name</i> routing-instances <i>instance-name</i> routing-options multicast interface <i>interface-name</i>],</code><br><code>[edit dynamic-profiles <i>profile-name</i> routing-options multicast interface <i>interface-name</i>]</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast interface <i>interface-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>],</code><br><code>[edit routing-options multicast interface <i>interface-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br><b>dynamic-profiles</b> hierarchy level added in Junos OS Release 11.2.   |
| <b>Description</b>              | Configure the multicast bandwidth for the interface.  |
| <b>Options</b>                  | <b><i>bps</i></b> —Bandwidth rate, in bits per second, for the multicast interface.<br><b>Range:</b> 0 through any amount of bandwidth  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | • <i>Example: Defining Interface Bandwidth Maximums</i>   |

## maximum-paths

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>maximum-paths <i>path-limit</i> &lt;log-interval <i>seconds</i>&gt; &lt;log-only   threshold <i>value</i>&gt;;</code>   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options] |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>         | Configure a limit for the number of routes installed in a routing table based upon the route path.  |



**NOTE:** The `maximum-paths` statement is similar to the `maximum-prefixes` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

|                |   |
|----------------|---|
| <b>Options</b> | <p><code>log-interval <i>seconds</i></code>—(Optional) Minimum time interval (in seconds) between log messages.<br/> <b>Range:</b> 5 through 86,400</p> <p><code>log-only</code>—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><code><i>path-limit</i></code>—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.<br/> <b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>)<br/> <b>Default:</b> No default</p> <p><code>threshold <i>value</i></code>—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the <code><i>path-limit</i></code> value that starts triggering the warnings.<br/> <b>Range:</b> 1 through 100</p> |
|----------------|---|



NOTE: When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *path-limit* value, then additional routes are rejected.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

## maximum-prefixes

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>maximum-prefixes <i>prefix-limit</i> &lt;log-interval <i>seconds</i>&gt; &lt;log-only   threshold <i>percentage</i>&gt;;</code>  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options]                              |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>         | Configure a limit for the number of routes installed in a routing table based upon the route prefix.<br><br>Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes. |



**NOTE:** The `maximum-prefixes` statement is similar to the `maximum-paths` statement. The `maximum-prefixes` statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The `maximum-paths` limit applies the total number of routes (two). The `maximum-prefixes` limit applies to the total number of unique prefixes (one).

|                |   |
|----------------|---|
| <b>Options</b> | <p><code>log-interval <i>seconds</i></code>—(Optional) Minimum time interval (in seconds) between log messages.</p> <p><b>Range:</b> 5 through 86,400</p> <p><code>log-only</code>—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.</p> <p><code>prefix-limit</code>—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.</p> <p><b>Range:</b> 1 through 4,294,967,295</p> <p><b>Default:</b> No default</p> <p><code>threshold <i>value</i></code>—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the <code>prefix-limit</code> value that starts triggering the warnings.</p> |
|----------------|---|

**Range:** 1 through 100



**NOTE:** When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *prefix-limit* value, then additional routes are rejected.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs](#)

---

## med-igp-update-interval

---

**Syntax** med-igp-update-interval *minutes*;

**Hierarchy Level** [edit routing-options]

**Release Information** Statement introduced in Junos OS Release 9.0  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp offset delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

**Options** *minutes*—Interval to delay MED updates.  
**Range:** 10 through 600  
**Default:** 10 minutes

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3352](#)
- [metric-out on page 3693](#)

## metric (Aggregate, Generated, or Static Route)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | (metric   metric2   metric3   metric4) <i>metric</i> <type type>;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options ( <a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],<br>[edit routing-options ( <a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>              | Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with <b>metric</b> (for the first metric value) and continuing with <b>metric2</b> , <b>metric3</b> , and <b>metric4</b> .   |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )<br><br><i>type type</i> —(Optional) Type of route.<br><br>When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.<br><b>Range:</b> 1 through 16 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Static Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 2941</a></li> <li>• <a href="#">generate on page 2976</a></li> <li>• <i>static</i></li> </ul>  |

## multicast (Routing Options)

```
Syntax  multicast {
        forwarding-cache {
            threshold suppress value <reuse value>;
        }
        interface interface-name {
            enable;
        }
        scope scope-name {
            interface [ interface-names ];
            prefix destination-prefix;
        }
        ssm-groups {
            address;
        }
    }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Configure generic multicast properties.



**NOTE:** You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Examples: Configuring Administrative Scoping*
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895](#)
- *Examples: Configuring the Multicast Forwarding Cache*
- *Multicast Protocols Feature Guide for Routing Devices*
- ([indirect-next-hop on page 2982](#) | no-indirect-next-hop)



## no-qos-adjust

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-qos-adjust;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i> <a href="#">reverse-oif-mapping</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i> <a href="#">reverse-oif-mapping</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i> <a href="#">reverse-oif-mapping</a>],</p> <p>[edit routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-options multicast <a href="#">interface</a> <i>interface-name</i> <a href="#">reverse-oif-mapping</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement added to [edit routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast interface <i>interface-name</i>], and [edit routing-options multicast interface <i>interface-name</i>] hierarchy levels in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li> </ul>  |

## num-65-127-prefix

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | num-65-127-prefix <i>value</i>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis forwarding-options</a> <i>profile-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos 13.2 for the QFX Series.  |
| <b>Description</b>              | Configure the number of supported IPv6 prefixes in the range /65 through /127.  |
| <b>Options</b>                  | <p><b>value</b>—With Junos OS 13.2X51D10: Value in the range 1 through 128. Each increment adds support for 16 IPv6 addresses with prefixes between /65 and /127, for a maximum of 2048 such addresses (16 x 128 = 2048).</p> <p><b>value</b>—With Junos OS 13.2X51D15: Value in the range 0 through 4. Each increment adds support for 1K IPv6 addresses with prefixes between /65 and /127, for a maximum of 4K such addresses.</p> |
| <b>Required Privilege Level</b> |   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Unified Forwarding Table on page 1697</a></li></ul>   |

## options (Routing Options)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>options {   syslog (level <i>level</i>   upto level <i>level</i>); }</pre>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.</p>   |
| <b>Options</b>                  | <p><b>level <i>level</i></b>—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:</p> <ul style="list-style-type: none"> <li>• <b>alert</b>—Conditions that should be corrected immediately, such as a corrupted system database.</li> <li>• <b>critical</b>—Critical conditions, such as hard drive errors.</li> <li>• <b>debug</b>—Software debugging messages.</li> <li>• <b>emergency</b>—Panic or other conditions that cause the system to become unusable.</li> <li>• <b>error</b>—Standard error conditions.</li> <li>• <b>info</b>—Informational messages.</li> <li>• <b>notice</b>—Conditions that are not error conditions, but might warrant special handling.</li> <li>• <b>warning</b>—System warning messages.</li> </ul> <p><b>upto level <i>level</i></b>—Log all messages up to a particular level.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <b>syslog</b> in the <i>Junos OS Administration Library for Routing Devices</i></li> </ul>  |

## pim-to-igmp-proxy

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>pim-to-igmp-proxy {<br/>  upstream-interface [ interface-names ];<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | <p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the <b>pim-to-igmp-proxy</b> statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>Configuring PIM-to-IGMP Message Translation</li></ul>   |

## pim-to-mld-proxy

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>pim-to-mld-proxy {<br/>    upstream-interface [ interface-names ];<br/>}</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.6.<br>Statement introduced in Junos OS Release 9.6 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | <p>Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the <b>pim-to-mld-proxy</b> statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring PIM-to-MLD Message Translation</li> </ul>   |

## policy (Aggregate and Generated Routes)

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>policy <i>policy-name</i>;</code>   |
| <b>Hierarchy Level</b>     | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit routing-options (<b>aggregate</b>   <b>generate</b>) (defaults   route)],</code><br><code>[edit routing-options rib <i>routing-table-name</i> (<b>aggregate</b>   <b>generate</b>) (defaults   route)]</code>   |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>         | <p>Associate a routing policy when configuring an aggregate or generated route's destination prefix in the <b>routes</b> part of the <b>aggregate</b> or <b>generate</b> statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.</p> <p>If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or <i>primary</i>, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.</p> <p>The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:</p> <ol style="list-style-type: none"><li>1. Compare the protocol's <b>preference</b> of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table.</li><li>2. Compare the protocol's <b>preference2</b> of the contributing routes. The lower <b>preference2</b> value is better. If only one route has <b>preference2</b>, then this route is preferred.</li><li>3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values.<ol style="list-style-type: none"><li>a. The primary contributor is the numerically smallest prefix value.</li><li>b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.</li></ol></li></ol> |

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

|                                 |   |
|---------------------------------|---|
| <b>Options</b>                  | <i>policy-name</i> —Name of a routing policy.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 2941</a></li> <li>• <a href="#">generate on page 2976</a></li> </ul> |

## policy (Flow Maps)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>policy [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | Configure a flow map policy.  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies for flow mapping.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.  |

## policy-options

---

**Syntax**    `policy-options`

```
  application-maps application-map-name {
    application application-name {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name match-type <actions>;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
    }
  }
```

**Hierarchy Level**    [edit]

**Release Information**    Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.1 for the EX Series.

**Description**    Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level**    storage—To view this statement in the configuration.  
storage-control—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)



## policy-statement

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> policy-statement <i>policy-name</i> {   term <i>term-name</i> {     from {       family <i>family-name</i>;       match-conditions;       policy <i>subroutine-policy-name</i>;       prefix-list <i>prefix-list-name</i>;       prefix-list-filter <i>prefix-list-name</i> match-type &lt;actions&gt;;       route-filter <i>destination-prefix</i> match-type &lt;actions&gt;;       source-address-filter <i>source-prefix</i> match-type &lt;actions&gt;;     }     to {       match-conditions;       policy <i>subroutine-policy-name</i>;     }     then <i>actions</i>;   } } </pre>   |
| <b>Hierarchy Level</b>     | <p>[edit dynamic policy-options],</p> <p>[edit logical-systems <i>logical-system-name</i> policy-options],</p> <p>[edit policy-options]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>inet-mdt</b> option introduced in Junos OS Release 10.0R2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>route-target</b> option introduced in Junos OS Release 12.2.</p>   |
| <b>Description</b>         | <p>Define a routing policy, including subroutine policies.</p> <p>A <i>term</i> is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.</p> <p>Each term contains a set of match conditions and a set of actions:</p> <ul style="list-style-type: none"> <li>• Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.</li> <li>• Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.</li> </ul> <p>Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of</p> |

**accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement *policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

**Options** *actions*—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

**family** *family-name*—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**.



**NOTE:** When *family* is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

**from**—(Optional) Match a route based on its source address.

**match-conditions**—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

**policy** *subroutine-policy-name*—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **\_\_.\*-internal\_\_**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

**policy-name**—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

**prefix-list** *prefix-list-name*—Name of a list of IPv4 or IPv6 prefixes.

**prefix-list-filter** *prefix-list-name*—Name of a prefix list to evaluate using qualifiers; *match-type* is the type of match (see *Configuring Prefix List Filters*), and *actions* is the action to take if the prefixes match.

**route-filter** *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *destination-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *destination-prefix* matches.

**source-address-filter** *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is

the IPv4 or IPv6 route prefix to match, **match-type** is the type of match (see *Configuring Route Lists*), and **actions** is the action to take if the **source-prefix** matches.

**term term-name**—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

**to**—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

**then**—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.                |
|                                 | routing-control—To add this statement to the configuration.         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>dynamic-db</i></li></ul> |

## ppm

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>ppm {     no-delegate-processing; }</pre>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.4.</p> <p>Statement introduced in Junos OS Release 10.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>(M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).</p> <p>After you disable PPM, PPM processing continues to run on the Routing Engine.</p> <p>In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The <b>no-delegate-processing</b> statement disables the default behavior and restores the legacy behavior.</p> |
| <b>Default</b>                  | Distributed PPM processing is enabled for all protocols that use PPM.   |
| <b>Options</b>                  | <b>no-delegate-processing</b> —Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i></li> <li><i>Ensuring That Distributed ppm Is Not Disabled</i></li> </ul>  |

## ppm (Ethernet Switching)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ppm {<br>centralized;<br>}  |
| <b>Hierarchy Level</b>          | [edit protocols lacp]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 10.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for T Series devices.   |
| <b>Description</b>              | <p>Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.</p> <p>This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the <b>no-delegate-processing</b> configuration statement in the [edit routing-options ppm] statement hierarchy.</p> |
| <b>Default</b>                  | Distributed PPM processing is enabled for all packets that use PPM.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)</i></li><li>• <a href="#">Configuring Distributed Periodic Packet Management on page 2906</a></li></ul>  |

## preference (Routing Options)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>(preference   preference2   color   color2) preference &lt;type type&gt;;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value (<b>preference2</b>), as well as colors, which are even finer-grained preference values (<b>color</b> and <b>color2</b>).</p> <p>If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.</p>   |
| <b>Options</b>                  | <p><b>preference</b>—Preference value. A lower number indicates a more preferred route.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 5 (for static routes), 130 (for aggregate and generated routes)</p> <p><b>type type</b>—(Optional) Type of route.</p> <p><b>Range:</b> 1 through 16</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 2941</a></li> <li>• <a href="#">generate on page 2976</a></li> <li>• <i>static</i></li> </ul>   |

## prefix

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>prefix destination-prefix;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">scope scope-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">scope scope-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">scope scope-name</a> ],<br>[edit routing-options multicast <a href="#">scope scope-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | Configure the prefix for multicast scopes.  |
| <b>Options</b>                  | <i>destination-prefix</i> —Address range for the multicast scope.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Examples: Configuring Administrative Scoping</i></li><li>• <i>Example: Creating a Named Scope for Multicast Scoping</i></li><li>• <i>multicast</i></li></ul>   |

## prefix-65-127-disable

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>prefix-65-127-disable</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">chassis forwarding-options</a> lpm-profile]   |
| <b>Release Information</b>      | Statement introduced in Junos 13.2X51-D15 for the QFX Series.   |
| <b>Description</b>              | Disable support in the longest prefix match (LPM) table for IPv6 prefixes in the range /65 through /127.                |
| <b>Required Privilege Level</b> |   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Unified Forwarding Table on page 1697</a></li></ul> |



## protocols

```

Syntax protocols {
    bgp {
        ... bgp-configuration ...
    }
    isis {
        ... isis-configuration ...
    }
    ldp {
        ... ldp-configuration ...
    }
    msdp {
        ... msdp-configuration ...
    }
    mstp {
        ... mstp-configuration ...
    }
    ospf {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        ... ospf-configuration ...
    }
    ospf3 {
        domain-id domain-id;
        domain-vpn-tag number;
        route-type-community (iana | vendor);
        ... ospf3-configuration ...
    }
    pim {
        ... pim-configuration ...
    }
    rip {
        ... rip-configuration ...
    }
    ripng {
        ... ripng-configuration ...
    }
    rstp {
        rstp-configuration;
    }
    vstp {
        vstp configuration;
    }
    vpls {
        vpls configuration;
    }
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],  
[edit routing-instances *routing-instance-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.

Support for RIPng introduced in Junos OS Release 9.0.  
Statement introduced in Junos OS Release 11.1 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

**Options** **bgp**—Specify BGP as the protocol for a routing instance.  
**isis**—Specify IS-IS as the protocol for a routing instance.  
**ldp**—Specify LDP as the protocol for a routing instance.  
**l2vpn**—Specify Layer 2 VPN as the protocol for a routing instance.  
**msdp**—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.  
**mstp**—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.  
**ospf**—Specify OSPF as the protocol for a routing instance.  
**ospf3**—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.



**NOTE:** OSPFv3 supports the **no-forwarding**, **virtual-router**, and **vrf** routing instance types only.

---

**pim**—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.  
**rip**—Specify RIP as the protocol for a routing instance.  
**ripng**—Specify RIP next generation (RIPng) as the protocol for a routing instance.  
**rstp**—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.  
**vstp**—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.  
**vpls**—Specify VPLS as the protocol for a routing instance.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Multiple Routing Instances of OSPF*

## qualified-next-hop (Static Routes)

**Syntax** `qualified-next-hop (address | interface-name) {  
     bfd-liveness-detection {  
         authentication {  
             algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |  
                 simple-password);  
             key-chain key-chain-name;  
             loose-check;  
         }  
         detection-time {  
             threshold milliseconds;  
         }  
         holddown-interval milliseconds;  
         minimum-interval milliseconds;  
         minimum-receive-interval milliseconds;  
         multiplier number;  
         no-adaptation;  
         transmit-interval {  
             minimum-interval milliseconds;  
             threshold milliseconds;  
         }  
         version (1 | automatic);  
     }  
     interface interface-name;  
     metric metric;  
     preference preference;  
}`

**Hierarchy Level** `[edit logical-systems logical-system-name routing-instances routing-instance-name  
     routing-options static route destination-prefix],  
     [edit logical-systems logical-system-name routing-options rib inet6.0 static route  
         destination-prefix],  
     [edit logical-systems logical-system-name routing-options static route destination-prefix],  
     [edit routing-instances routing-instance-name routing-options static route destination-prefix],  
     [edit routing-options rib inet6.0 static route destination-prefix],  
     [edit routing-options static route destination-prefix]`

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 11.3 for the QFX Series.  
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

**Options** *address*—IPv4, IPv6, or ISO network address of the next hop.  
*interface-name*—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop

interface for a static route, specify **qualified-next-hop *interface-name***, where *interface-name* is the name of the IPv4 or IPv6 unnumbered interface.



**NOTE:** For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the `unnumbered-address <interface-name>` statement at the `[edit interfaces <interface-name> unit <logical-unit-number> family <family-name>]` hierarchy level as described in *Configuring an Unnumbered Interface*.

---

The remaining statements are explained separately.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Static Route Preferences and Qualified Next Hops</a></li><li>• <a href="#">Example: Enabling BFD on Qualified Next Hops in Static Routes on page 2923</a></li></ul> |
|------------------------------|--|

## readvertise


|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | (readvertise   no-readvertise);  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options static (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> static (defaults   route)],</p> <p>[edit routing-options static (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | Configure whether static routes are eligible to be readvertised by routing protocols:  |
| <b>Default</b>                  | Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the <b>no-readvertise</b> statement.   |
| <b>Options</b>                  | <p><b>readvertise</b>—Readvertise static routes. Include the <b>readvertise</b> statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>no-readvertise</b> option specified in the <b>defaults</b> portion of the statement.</p> <p><b>no-readvertise</b>—Mark a static route as being ineligible for readvertisement. Include the <b>no-readvertise</b> option when configuring the route.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Controlling Static Routes in Routing and Forwarding Tables</i></li> <li>• <i>static</i></li> </ul>  |

## redundant-sources

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>redundant-sources [ <i>addresses</i> ];</code>  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i>],</code><br><code>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Configure a list of redundant sources for multicast flows defined by a flow map.  |
| <b>Options</b>                  | <b><i>addresses</i></b> —List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring a Multicast Flow Map</i></li></ul>  |

## resolution

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> resolution {     rib <i>routing-table-name</i> {         import [ <i>policy-names</i> ];         resolution-ribs [ <i>routing-table-names</i> ];     } } </pre>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-options]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.</p> <p>For example, you might want to direct <b>inet.2</b> route resolution to use topology routing tables <b>:red.inet.0</b> and <b>:blue.inet.0</b> for protocol next-hop IP address lookups. Or you might want to direct <b>bgp.l3vpn.0</b> to use the information in <b>inet.0</b> to resolve routes, thus overriding the default behavior, which is to use <b>inet.3</b>.</p> <p>You can specify up to two routing tables in the <b>resolution-ribs</b> statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure <b>resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3]</b>, <b>inet.0</b> is checked first and then <b>inet.3</b> is checked.</p> |
|                                 | <p> <b>NOTE:</b> Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the <b>show system processes</b> and the <b>show task memory</b> commands. For more information, see <i>Routing Protocol Process Memory FAQs</i>.</p>  |
|                                 | <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- *Example: Configuring Route Resolution on PE Routers*
  - *Example: Configuring Route Resolution on Route Reflectors*
  - *Example: Configuring Multitopology Routing Based on a Multicast Source*

## resolution-ribs

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>resolution-ribs [ <i>routing-table-names</i> ];</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options resolution <a href="#">rib</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options resolution <a href="#">rib</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options resolution <a href="#">rib</a>],</code><br><code>[edit routing-options resolution <a href="#">rib</a>]</code>  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | <p>Specify one or more routing tables to use for route resolution.</p> <p>This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is <b>inet.3</b>, but you want to allow fallback resolution through <b>inet.0</b>. One example use case is overriding the <b>bgp.rtarget.0 (family route-target)</b> routing table resolution from using only <b>inet.3</b> to using both <b>inet.3</b> and <b>inet.0</b>.</p> |
| <b>Options</b>                  | <b><i>routing-table-names</i></b> —Name of one or more routing tables.   |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Route Resolution on PE Routers</i></li><li>• <i>Example: Configuring Multitopology Routing Based on a Multicast Source</i></li></ul>   |



## resolve

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | resolve;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)],<br>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],<br>[edit logical-systems <i>logical-system-name</i> routing-options static (defaults   route)],<br>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static (defaults   route)],<br>[edit routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)],<br>[edit routing-options rib <i>routing-table-name</i> static (defaults   route)],<br>[edit routing-options static (defaults   route)] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the <b>inet.0</b> and <b>inet.3</b> routing tables.  |
| <b>Default</b>                  | Static routes can point only to a directly connected next hop.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>static</i></li> </ul>   |

## restart-duration (Routing Options)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>restart-duration <i>seconds</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> routing-options graceful-restart],<br>[edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart],<br>[edit routing-options graceful-restart] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Configure the restart timer for graceful restart.   |
| <b>Options</b>                  | <b><i>seconds</i></b> —Configure the time period for the restart to last.<br><b>Range:</b> 120 through 900 seconds<br><b>Default:</b> 300 seconds   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Junos OS High Availability Library for Routing Devices</i></li></ul>   |

## retain

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | (no-retain   retain);   |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   routing-options rib <i>routing-table-name</i> static (defaults   route)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>   routing-options static (defaults   route)], [edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static   (defaults   route)], [edit logical-systems <i>logical-system-name</i> routing-options static (defaults   route)], [edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> static   (defaults   route)], [edit routing-instances <i>routing-instance-name</i> routing-options static (defaults   route)], [edit routing-options rib <i>routing-table-name</i> static (defaults   route)], [edit routing-options static (defaults   route)]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:  |
| <b>Default</b>                  | Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.  |
| <b>Options</b>                  | <p><b>no-retain</b>—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the <b>no-retain</b> statement. Include this statement when configuring an individual route in the <b>route</b> portion of the <b>static</b> statement to override a <b>retain</b> option specified in the <b>defaults</b> portion of the statement.</p> <p><b>retain</b>—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>static</i></li> </ul>   |

## reverse-oif-mapping

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>reverse-oif-mapping {<br/>    no-qos-adjust;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-options multicast <a href="#">interface</a> <i>interface-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>The <b>no-qos-adjust</b> statement added in Junos OS Release 9.5.</p> <p>The <b>no-qos-adjust</b> statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.</p> <p>The remaining statement is explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>  |

## rpf-check-policy (Routing Options RPF)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>rpf-check-policy [ <i>policy-names</i> ];</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more multicast RPF check policies.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring RPF Policies</i></li> </ul>   |

## rib (General)

---

```
Syntax  rib routing-table-name {  
        aggregate {  
            defaults {  
                ... aggregate-options ...  
            }  
            route destination-prefix {  
                policy policy-name;  
                ... aggregate-options ...  
            }  
        }  
        generate {  
            defaults {  
                generate-options;  
            }  
            route destination-prefix {  
                policy policy-name;  
                generate-options;  
            }  
        }  
        martians {  
            destination-prefix match-type <allow>;  
        }  
    }  
    static {  
        defaults {  
            static-options;  
        }  
        rib-group group-name;  
        route destination-prefix {  
            next-hop;  
            static-options;  
        }  
    }  
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],  
[edit logical-systems *logical-system-name* routing-options],  
[edit routing-instances *routing-instance-name* routing-options],  
[edit routing-options]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.



**NOTE:** The IPv4 multicast routing table (`inet.1`) and the IPv6 multicast routing table (`inet6.1`) are not supported for this statement.

**Default** If you do not specify a routing table name with the *routing-table-name* option, the software uses the default routing tables, which are `inet.0` for unicast routes and `inet.1` for the multicast cache.

**Options** *routing-table-name*—Name of the routing table, in the following format:  
*protocol [.identifier]*.

In a routing instance, the routing table name must include the routing instance name.

For example, if the routing instance name is `link0`, the routing table name might be `link0.inet6.0`.

- *protocol* is the protocol family. It can be `inet6` for the IPv6 family, `inet` for the IPv4 family, `iso` for the ISO protocol family, or *instance-name.iso.0* for an ISO routing instance.
- *identifier* is a positive integer that specifies the instance of the routing table.

**Default:** `inet.0`

The remaining statements are explained separately.

**Required Privilege Level** `routing`—To view this statement in the configuration.  
`routing-control`—To add this statement to the configuration.

**Related Documentation**

- *Example: Creating Routing Tables*
- *passive*

## rib (Route Resolution)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>rib <i>routing-table-name</i> {<br/>    <b>import</b> [ <i>policy-names</i> ];<br/>    <b>resolution-ribs</b> [ <i>routing-table-names</i> ];<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>resolution</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options <b>resolution</b> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options <b>resolution</b> ],<br>[edit routing-options <b>resolution</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Specify a routing table name for route resolution.<br><br>The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Route Resolution on PE Routers</i></li></ul>  |



## rib-group (Routing Options)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <a href="#">interface-routes</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">interface-routes</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> static],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options static],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">interface-routes</a>],</p> <p>[edit routing-options <a href="#">interface-routes</a>],</p> <p>[edit routing-options rib <i>routing-table-name</i> static],</p> <p>[edit routing-options static]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | Configure which routing table groups interface routes are imported into.   |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <a href="#">interface-routes on page 2989</a></li> <li>• <a href="#">rib-groups on page 3032</a></li> </ul>  |

## rib-groups

---

|                     |  |
|---------------------|--|
| Syntax              | <pre>rib-groups {<br/>    group-name {<br/>        export-rib group-name;<br/>        import-policy [ policy-names ];<br/>        import-rib [ group-names ];<br/>    }<br/>}</pre>  |
| Hierarchy Level     | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]   |
| Release Information | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| Description         | <p>Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.</p> <p>Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the <b>import-rib</b> statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the <b>export-rib</b> statement).</p> <p>The first routing table you specify is the <i>primary routing table</i>, and any additional routing tables are the <i>secondary routing tables</i>.</p> <p>The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify <b>inet.0</b> as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify <b>inet6.0</b> as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (<b>inet6.x</b>).</p> <p>In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the <b>import-rib</b> statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same <b>import-rib</b> statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the <b>inet6.3</b> routing table with IPv6 addresses that are compatible with IPv4. Specify <b>inet.0</b> as the primary routing table, and specify <b>inet6.3</b> as a secondary routing table.</p> |



**NOTE:** On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

---



**NOTE:** If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

**Options** *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.


**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- [rib-group on page 3031](#)

## route-distinguisher-id

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>route-distinguisher-id <i>ip-address</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | <p>Automatically assign a route distinguisher to the routing instance.</p> <p>If you configure the <b>route-distinguisher</b> statement in addition to the <b>route-distinguisher-id</b> statement, the value configured for <b>route-distinguisher</b> supersedes the value generated from <b>route-distinguisher-id</b>.</p> <div><p><b>NOTE:</b> To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the <b>route-distinguisher</b> statement is different from the first half of the route distinguisher obtained by configuring the <b>route-distinguisher-id</b> statement.</p></div> |
| <b>Options</b>                  | <i>ip-address</i> —Address for routing instance.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring BGP Route Target Filtering for VPNs</i></li><li>• <i>Configuring Routing Instances on PE Routers in VPNs</i></li></ul>   |

---

## route-record

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | route-record;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-options]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.                                    |
| <b>Description</b>              | Export the AS path and routing information to the traffic sampling process.<br><br>Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling Flow Aggregation</i></li><li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li></ul>  |

## router-id

---

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>router-id address;</code>  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],<br>[edit logical-systems <i>logical-system-name</i> routing-options],<br>[edit routing-instances <i>routing-instance-name</i> routing-options],<br>[edit routing-options]  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>         | Specify the routing device's IP address.<br><br>The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used. |



**NOTE:** We strongly recommend that you configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

For more information about the router identifier in OSPF, see [“Example: Configuring an OSPF Router Identifier” on page 4048](#).

You must configure a router-id in order for BGP and OSPF to function in a routing instance. Use the **show route instance detail** command to display the router-id value for a routing instance. If the router-id is **0.0.0.0**, then the routing instance has no router-id.

For more information about the router identifier in OSPF, see [“Example: Configuring an OSPF Router Identifier” on page 4048](#).



**NOTE:** If you run OSPF for IPv6 or BGP for IPv6 in a routing instance, you must configure an IPv4 router identifier (router-id) in the routing instance itself. In other words, the IPv4 router-id in the main routing instance is not inherited by other routing instances. Even if you run *only* IPv6 OSPF or BGP in a routing instance, the IPv4 router-id must be configured because OSPF and BGP, even when used exclusively with IPv6, use the IPv4 router-id for handshaking. If you do not configure the IPv4 router-id in the IPv6 OSPF or BGP routing instance, then the IPv6 protocols will use invalid IPv4 address 0.0.0.0 and the adjacencies and connections will fail.

|                                 |  |
|---------------------------------|--|
| <b>Options</b>                  | <p><b>address</b>—IP address of the routing device.</p> <p><b>Default:</b> Address of the first interface encountered by Junos OS</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Examples: Configuring External BGP Peering on page 3261</a></li> <li>• <a href="#">Examples: Configuring Internal BGP Peering on page 3284</a></li> </ul> |

## routing-instances

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>routing-instances <i>routing-instance-name</i> {   description;   instance-type virtual-router;   interface <i>interface-name</i>;   protocols;   routing-options }</pre> |
| <b>Hierarchy Level</b>          | [edit]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | (QFabric switches only) Configure a virtual router routing instance.   |
| <b>Options</b>                  | <p><b><i>routing-instance-name</i></b>—Name of this routing instance.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Virtual Router Routing Instances on page 2908</a></li> </ul>  |

## routing-options

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | routing-options { ... }   |
| <b>Hierarchy Level</b>          | [edit],<br>[edit routing-instances <i>routing-instance-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | Configure protocol-independent routing properties.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Routing Options on page 2897</a></li><li>• <a href="#">Understanding Distributed Periodic Packet Management on page 2898</a></li><li>• <a href="#">Example: Configuring SSM Maps for Different Groups to Different Sources on page 4899</a></li></ul> |

## scope

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | scope <i>scope-name</i> {<br>interface [ <i>interface-names</i> ];<br>prefix <i>destination-prefix</i> ;<br>}   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],<br>[edit routing-options multicast] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | Configure multicast scoping.  |
| <b>Options</b>                  | <i>scope-name</i> —Name of the multicast scope.<br><br>The remaining statements are explained separately.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Creating a Named Scope for Multicast Scoping</a></li></ul>   |



## scope-policy

**Syntax** `scope-policy [ policy-names ];`

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options multicast],  
[edit routing-options multicast]



**NOTE:** You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the `scope` statement to a specific routing instance at the [edit routing-instances *routing-instance-name* routing-options multicast] or [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast] hierarchy level.

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

**Description** Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

**Options** *policy-names*—Name of one or more multicast scope policies.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**


- [scope on page 3038](#)
- *Example: Using a Scope Policy for Multicast Scoping*

## source (Source-Specific Multicast)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>source [ <i>addresses</i> ];</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>],</code><br><code>[edit routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>              | Specify IPv4 or IPv6 source addresses for an SSM map.   |
| <b>Options</b>                  | <i>addresses</i> —IPv4 or IPv6 source addresses.  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To view this statement in the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 4893</a></li></ul>   |

## source-routing

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | source-routing {<br>(ip   ipv6)<br>}  |
| <b>Hierarchy Level</b>          | [edit routing-options]  |
| <b>Release Information</b>      | Statement for IPv6 introduced in Junos OS Release 8.2.<br>Statement for IPv4 introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.                               |
| <b>Description</b>              | <p>Enable source routing.</p> <p>Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.</p>   |
|                                 | <div>  <p><b>NOTE:</b> We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.</p> </div> |
| <b>Default</b>                  | Disabled  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring Filter-Based Forwarding on the Source Address</i></li> </ul>   |

## ssm-groups

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>ssm-groups [ <i>ip-addresses</i> ];</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</code><br><code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</code><br><code>[edit routing-options multicast]</code>  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>              | <p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the <b>ssm-groups</b> statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the <b>ssm-groups</b> statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p> |
| <b>Options</b>                  | <i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895</a></li></ul>  |

## ssm-map (Routing Options Multicast)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>ssm-map <i>ssm-map-name</i> {     policy [ <i>policy-names</i> ];     source [ <i>addresses</i> ]; }</pre>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | Configure SSM mapping.   |
| <b>Options</b>                  | <p><b><i>ssm-map-name</i></b>—Name of the SSM map.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SSM Mapping on page 4893</a></li> </ul>  |

## static (Routes)

---

|                     |  |
|---------------------|--|
| Syntax              | <pre>static {   defaults {     static-options;   }   rib-group group-name;   route destination-prefix {     next-hop address;     next-hop options;     qualified-next-hop address {       metric metric;       preference preference;     }     static-options;   } }</pre>   |
| Hierarchy Level     | [edit routing-options],<br>[edit routing-options rib <i>routing-table-name</i> ]   |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| Description         | Configure static routes to be installed in the routing table. You can specify any number of routes within a single <b>static</b> statement, and you can specify any number of <b>static</b> options in the configuration.  |
| Options             | <p><b>defaults</b>—Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the <b>static</b> statement. This part of the <b>static</b> statement is optional.</p> <p><b>route destination-prefix</b>—Destination of the static route.</p> <ul style="list-style-type: none"><li>• <b>defaults</b>—For the default route to the destination. This is equivalent to specifying an IP address of <b>0.0.0.0/0</b>.</li><li>• <b>destination-prefix/prefix-length</b>—<b>destination-prefix</b> is the network portion of the IP address, and <b>prefix-length</b> is the destination prefix length.</li><li>• <b>next-hop address</b>—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).</li><li>• <b>nsap-prefix</b>—<b>nsap-prefix</b> is the network service access point (NSAP) address for ISO.</li></ul> <p><b>next-hop options</b>—Additional information for how to manage forwarding of packets to the next hop.</p> <ul style="list-style-type: none"><li>• <b>discard</b>—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.</li></ul> |

- **iso-net**—Reach the next-hop routing device by specifying an ISO NSAP.
- **next-table *routing-table-name***—Name of the next routing table to the destination.
- **receive**—Install a receive route for this destination into the routing table.
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

**static-options**—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in **static-options**. Each of the options is explained separately.

- **(active | passive);**
- **(install | no-install);**
- **(metric | metric2 | metric3 | metric4) *value* <type type>;**
- **(preference | preference2 | color | color2) *preference* <type type>;**
- **(resolve | no-resolve);**
- **(no-retain | retain);**

The remaining statements are explained separately.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Static Routing on page 2904</a></li> </ul>         |

## subscriber-leave-timer

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>subscriber-leave-timer seconds;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <a href="#">interface interface-name</a> ],<br>[edit routing-options multicast <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.2.<br>Statement introduced in Junos OS Release 9.2 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.   |
| <b>Options</b>                  | <b>seconds</b> —Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.<br><b>Range:</b> 0 through 30<br><b>Default:</b> 0 seconds  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li></ul>   |



## tag (Routing Options)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>tag metric type number;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-options (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)],</p> <p>[edit routing-options rib <i>routing-table-name</i> (<a href="#">aggregate</a>   <a href="#">generate</a>   static) (defaults   route)]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>  |
| <b>Description</b>              | Associate a tag with a static, aggregate, or generated route.   |
| <b>Default</b>                  | No tag strings are associated with routes.  |
| <b>Options</b>                  | <p><i>metric</i>—Tag metric.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><i>type number</i>—Tag type.</p> <p><b>Range:</b> 1 through 16</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Examples: Configuring Static Routes</i></li> <li>• <i>Example: Summarizing Routes Through Route Aggregation</i></li> <li>• <i>Example: Conditionally Generating Static Routes</i></li> <li>• <a href="#">aggregate on page 2941</a></li> <li>• <a href="#">generate on page 2976</a></li> <li>• <i>static</i></li> </ul>  |

## threshold (Multicast Forwarding Cache)

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>threshold {<br/>    log-warning <i>value</i>;<br/>    suppress <i>value</i> &lt;reuse <i>value</i>&gt;;<br/>}</pre>   |
| Hierarchy Level          | <pre>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>  routing-options multicast forwarding-cache],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i><br/>  routing-options multicast forwarding-cache family (inet   inet6)],<br/>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],<br/>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache<br/>  family (inet   inet6)],<br/>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],<br/>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache<br/>  (inet   inet6)],<br/>[edit routing-options multicast forwarding-cache],<br/>[edit routing-options multicast forwarding-cache family (inet   inet6)]</pre> |
| Release Information      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| Description              | <p>Configure the global suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the <b>threshold</b> statement globally for the multicast forwarding cache or including the <b>family</b> statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.</p> <p>To confirm the configured threshold values, use the <b>show multicast forwarding-cache statistics</b> command.</p>  |
| Options                  | <p><b>reuse <i>value</i></b>—(Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the <b>suppress</b> value.</p> <p><b>Range:</b> 1 through 200,000</p> <p><b>suppress <i>value</i></b>—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the <b>reuse</b> value.</p> <p><b>Range:</b> 1 through 200,000</p> <p>The remaining statement is explained separately.</p>   |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Examples: Configuring the Multicast Forwarding Cache</i></li></ul>  |

## timeout (Flow Maps)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | timeout (never non-discard-entry-only   <i>minutes</i> );   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ],<br>[edit routing-options multicast <b>flow-map</b> <i>flow-map-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.   |
| <b>Description</b>              | Configure the timeout value for multicast forwarding cache entries associated with the flow map.  |
| <b>Options</b>                  | <b>minutes</b> —Length of time that the forwarding cache entry remains active.<br><b>Range:</b> 1 through 720<br><br><b>never non-discard-entry-only</b> —Specify that the forwarding cache entry always remain active. If you omit the <b>non-discard-entry-only</b> option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the <b>non-discard-entry-only</b> option, entries with forwarding states are kept forever, and entries with pruned states time out. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |

## timeout (Multicast)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | timeout <i>minutes</i> <family (inet   inet6)>;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],<br>[edit logical-systems <i>logical-system-name</i> routing-options multicast forwarding-cache],<br>[edit routing-instances <i>routing-instance-name</i> routing-options multicast forwarding-cache],<br>[edit routing-options multicast forwarding-cache]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.2.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Statement introduced in Junos OS Release 12.3 for ACX Series routers.  |
| <b>Description</b>              | Configure the timeout value for multicast forwarding cache entries.  |
| <b>Options</b>                  | <b>minutes</b> —Length of time that the forwarding cache limit remains active.<br><b>Range:</b> 1 through 720<br><br><b>family (inet   inet6)</b> —(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the <b>timeout</b> statement globally for the multicast forwarding cache or including the <b>family</b> statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.<br><br><b>Default:</b> By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Multicast Forwarding Cache</i></li></ul>   |

## traceoptions (Routing Options)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;disable&gt;; } </pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options],</p> <p>[edit routing-options flow],</p> <p>[edit routing-options multicast]</p>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p><b>nsr-synchronization</b> flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.</p> <p><b>nsr-synchronization</b> and <b>nsr-packet</b> flags for BFD sessions added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>nsr-synchronization</b> flag for RIP and RIPng added in Junos OS Release 9.0.</p> <p><b>nsr-synchronization</b> flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.</p> <p><b>nsr-synchronization</b> flag for PIM added in Junos OS Release 9.3.</p> <p><b>nsr-synchronization</b> flag for MPLS added in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>nsr-synchronization</b> flag for MSDP added in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> |
| <b>Description</b>         | <p>Define tracing operations that track all routing protocol functionality in the routing device.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>  |
| <b>Default</b>             | If you do not include this statement, no global tracing operations are performed.  |
| <b>Options</b>             | <p><b>Values:</b></p> <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place global routing protocol tracing output in the file <b>routing-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and</p>   |

so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization-detail**—(MPLS only) Tracing operations for nonstop active routing in detail
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Tracing Global Routing Protocol Operations*
- [Tracing Nonstop Active Routing Synchronization Events on page 2278](#)

## upstream-interface

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>upstream-interface [ <i>interface-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-igmp-proxy</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-mld-proxy</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>pim-to-igmp-proxy</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast <b>pim-to-mld-proxy</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-igmp-proxy</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>pim-to-mld-proxy</b>],</p> <p>[edit routing-options multicast <b>pim-to-igmp-proxy</b>],</p> <p>[edit routing-options multicast <b>pim-to-mld-proxy</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>   |
| <b>Description</b>              | <p>Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the <b>pim-to-igmp-proxy</b> statement), or into corresponding MLD report or leave messages (if you include the <b>pim-to-mld-proxy</b> statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.</p>  |
| <b>Options</b>                  | <p><b><i>interface-names</i></b>—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ( [ ] ).</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring PIM-to-IGMP Message Translation</i></li> <li>• <i>Configuring PIM-to-MLD Message Translation</i></li> </ul>  |



# Administration

- [Routine Monitoring on page 3055](#)
- [Operational Commands on page 3057](#)

## Routine Monitoring

- [Monitoring Routing Information on page 3055](#)
- [Verifying That Virtual Router Routing Instances Are Working on page 3056](#)

### Monitoring Routing Information

- Purpose** Use the monitoring functionality to view the `inet.0` routing table on the routing device.
- Action** To view the routing table, enter the following commands in the CLI interface:
- `show route terse`
  - `show route detail`
- Meaning** [Table 272 on page 3055](#) describes the different filters, their functions, and the associated actions.
- [Table 273 on page 3056](#) summarizes key output fields in the routing information display.

Table 272: Filtering Route Messages

| Field               | Function   | Your Action                                 |
|---------------------|--|---|
| Destination Address | Specifies the destination address of the route.  | Enter the destination address.              |
| Next hop address    | Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. | Enter the next hop address.                 |
| Best route          | Specifies only the best route available.   | Select the view details of the best route.  |
| Inactive routes     | Specifies the inactive routes.   | Select the view details of inactive routes. |
| Exact route         | Specifies the exact route.   | Select the view details of the exact route. |

Table 272: Filtering Route Messages (*continued*)

| Field         | Function   | Your Action   |
|---------------|--|---|
| Hidden routes | Specifies the hidden routes.                                     | Select the view details of hidden routes.                       |
| Search        | Applies the specified filter and displays the matching messages. | To apply the filter and display messages, click <b>Search</b> . |

Table 273: Summary of Key Routing Information Output Fields

| Field                  | Values   | Additional Information  |
|------------------------|--|---|
| Static Route Addresses | The list of static route addresses.  |   |
| Protocol               | Protocol from which the route was learned: <b>Static</b> , <b>Direct</b> , <b>Local</b> .                              |   |
| Preference             | The preference is the individual preference value for the route.   | The route preference is used as one of the route selection criteria.  |
| Next-Hop               | Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. | <p>If a next hop is listed as <b>Discard</b>, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the <b>discard</b> attribute has been set.</p> <p>If a next hop is listed as <b>Reject</b>, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected.</p> <p>If a next hop is listed as <b>Local</b>, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).</p> |
| Age                    | How long the route has been active.  |   |
| State                  | Flags for this route.  | There are many possible flags.  |

**Related Documentation** • [Configuring Static Routing on page 2904](#)

## Verifying That Virtual Router Routing Instances Are Working

**Purpose** After creating a virtual router routing instance, verify that it has been set up properly.

**Action** 1. Use the **show route instance** command to list all the routing instances and their properties:

```
user@switch> show route instance
```

| Instance                    | Type           | Active/holddown/hidden |
|-----------------------------|----------------|------------------------|
| Primary RIB                 |                |                        |
| master                      | forwarding     |                        |
| inet.0                      |                | 4/0/1                  |
| __juniper_private1__        | forwarding     |                        |
| __juniper_private1__.inet.0 |                | 1/0/3                  |
| __juniper_private2__        | forwarding     |                        |
| __juniper_private2__.inet.0 |                | 0/0/1                  |
| __juniper_private3__        | forwarding     |                        |
| __juniper_private3__.inet.0 |                | 1/0/2                  |
| __juniper_private4__        | forwarding     |                        |
| __juniper_private4__.inet.0 |                | 4/0/2                  |
| __master.anon__             | forwarding     |                        |
| r1                          | virtual-router |                        |
| r2                          | virtual-router |                        |

- Use the **show route forwarding-table** command to view the forwarding table information for each routing instance:

```

user@switch> show route forwarding-table
Routing table: r1---qfabric.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0      0              rjct 1628  1
0.0.0.0/32       perm  0      0              dscd 1626  1
224.0.0.0/4      perm  0      0              mdsc 1627  1
224.0.0.1/32     perm  0 224.0.0.1     mcst 1623  1
255.255.255.255/32 perm  0      0              bcst 1624  1

```

**Meaning** The output displays the routing table information and confirms that the virtual router routing instances have been created and the links are up.

**Related Documentation**

- [Configuring Virtual Router Routing Instances on page 2908](#)

## Operational Commands

- [clear ipv6 neighbors](#)
- [show as-path](#)
- [show as-path domain](#)
- [show as-path summary](#)
- [show ipv6 neighbors](#)
- [show ipv6 router-advertisement](#)
- [show route](#)
- [show route active-path](#)
- [show route all](#)

- `show route aspath-regex`
- `show route best`
- `show route brief`
- `show route community`
- `show route community-name`
- `show route damping`
- `show route detail`
- `show route exact`
- `show route export`
- `show route extensive`
- `show route flow validation`
- `show route forwarding-table`
- `show route inactive-path`
- `show route inactive-prefix`
- `show route instance`
- `show route label`
- `show route label-switched-path`
- `show route martians`
- `show route next-hop`
- `show route no-community`
- `show route protocol`
- `show route range`
- `show route receive-protocol`
- `show route resolution`
- `show route snooping`
- `show route source-gateway`
- `show route summary`
- `show route table`
- `show route terse`

## clear ipv6 neighbors

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | clear ipv6 neighbors<br><all   host <i>hostname</i> >   |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 12.2 for the QFX Series.  |
| <b>Description</b>              | Clear IPv6 neighbor cache information.  |
| <b>Options</b>                  | <p><b>none</b>—Clear all IPv6 neighbor cache information.</p> <p><b>all</b>—(Optional) Clear all IPv6 neighbor cache information.</p> <p><b>host <i>hostname</i></b>—(Optional) Clear the information for the specified IPv6 neighbors.</p> |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show ipv6 neighbors on page 3069</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">clear ipv6 neighbors on page 3059</a>   |
| <b>Output Fields</b>            | When you enter this command, you are provided feedback on the status of your request.   |

## Sample Output

### clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

## show as-path

---

|                             |   |
|-----------------------------|---|
| List of Syntax              | <a href="#">Syntax on page 3060</a><br><a href="#">Syntax (EX Series Switches) on page 3060</a>   |
| Syntax                      | <code>show as-path</code><br><code>&lt;brief   detail&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches) | <code>show as-path</code><br><code>&lt;brief   detail&gt;</code>  |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.  |
| Description                 | <p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the <b>as-path-prepend</b> action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p> |
| Options                     | <p><b>none</b>—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>  |
| Required Privilege Level    | view  |
| Related Documentation       | <ul style="list-style-type: none"><li>• <a href="#">show as-path summary on page 3067</a></li></ul>   |
| List of Sample Output       | <a href="#">show as-path on page 3061</a><br><a href="#">show as-path detail on page 3062</a>   |
| Output Fields               | <a href="#">Table 274 on page 3061</a> lists the output fields for the <b>show as-path</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 274: show as-path Output Fields

| Field Name            | Field Description   | Level of Output   |
|-----------------------|---|-------------------|
| <b>Total AS paths</b> | Total number of AS paths.   | <b>brief none</b> |
| <b>Bucket</b>         | Bucket number.  | All levels        |
| <b>Count</b>          | Number of AS path entries in this bucket.   | All levels        |
| <b>AS path</b>        | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> <li>• <b>Atomic</b>—Route is an aggregate of several route prefixes.</li> <li>• <b>Aggregat</b>—Routing device has summarized a range of prefixes.</li> </ul> | All levels        |
| <b>domain</b>         | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.  | <b>detail</b>     |
| <b>neighbor as</b>    | AS peer address.  | <b>detail</b>     |
| <b>length</b>         | Length of the AS path.  | <b>detail</b>     |
| <b>segments</b>       | Length of the AS segment descriptor.  | <b>detail</b>     |
| <b>references</b>     | Path reference count.   | <b>detail</b>     |

## Sample Output

### show as-path

```

user@host> show as-path
Total AS paths: 30382
Bucket 0      Count: 36
I
14203 2914 174 31752 I
14203 2914 701 21512 I
14203 2914 1239 26632 I
14203 2914 1239 29704 I
14203 2914 4323 10248 I
14203 2914 4766 23560 I
14203 2914 6395 32776 I
14203 2914 7911 11272 I
14203 2914 12180 18440 I
14203 2914 17408 17416 I
14203 2914 701 702 24586 I
14203 2914 1239 4657 9226 I
14203 2914 1239 7132 16394 I
14203 2914 1299 8308 34826 I
14203 2914 3320 5603 28682 I

```

```

14203 2914 3491 1680 33802 I
14203 2914 3549 7908 27658 I
14203 2914 3549 20804 30730 I
14203 2914 7018 2687 9226 I
14203 2914 174 9318 9318 23564 I
14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I

```

...

## show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I

```



```
    domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3320 5603 28682 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
    domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
    domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
    domain 1, neighbor as: 14203, length 7, segments 1, references 10
...

```

## show as-path domain

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3064</a><br><a href="#">Syntax (EX Series Switches) on page 3064</a>  |
| <b>Syntax</b>                      | show as-path domain<br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches)</b> | show as-path domain  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display autonomous system (AS) path domain information.  |
| <b>Options</b>                     | <b>none</b> —(Optional) Display AS path domain information for all routing instances.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show as-path domain on page 3066</a>   |
| <b>Output Fields</b>               | <a href="#">Table 275 on page 3064</a> lists the output fields for the <b>show as-path domain</b> command. Output fields are listed in the approximate order in which they appear  |

**Table 275: show as-path domain Output Fields**

| Field Name          | Field Description   |
|---------------------|---|
| <b>Domain</b>       | Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.  |
| <b>Primary</b>      | Primary AS number.  |
| <b>References</b>   | Path reference count.   |
| <b>Number Paths</b> | Number of known AS paths.   |
| <b>Flags</b>        | Information about the AS path: <ul style="list-style-type: none"> <li>• <b>ASLoop</b>—Path contains an AS loop.</li> <li>• <b>Atomic</b>—Path includes the ATOMIC_AGGREGATE path attribute.</li> <li>• <b>Local</b>—Path was created by local aggregation.</li> <li>• <b>Master</b>—Path was created by the master routing instance.</li> </ul> |
| <b>Local AS</b>     | AS number of the local routing device.  |

Table 275: show as-path domain Output Fields (*continued*)

| Field Name   | Field Description                                       |
|--------------|---|
| <b>Loops</b> | How many times this AS number can appear in an AS path. |

## Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

## show as-path summary

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3067</a><br><a href="#">Syntax (EX Series Switches) on page 3067</a>   |
| <b>Syntax</b>                      | <pre>show as-path summary &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>   |
| <b>Syntax (EX Series Switches)</b> | show as-path summary  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>                 | <p>Display autonomous system (AS) path summary information.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p> |
| <b>Options</b>                     | <p><b>none</b>—(Optional) Display AS path summary information for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>  |
| <b>Required Privilege Level</b>    | view  |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">show as-path on page 3060</a></li> </ul>   |
| <b>List of Sample Output</b>       | <a href="#">show as-path summary on page 3068</a>   |
| <b>Output Fields</b>               | <p><a href="#">Table 276 on page 3067</a> lists the output fields for the <b>show as-path summary</b> command. Output fields are listed in the approximate order in which they appear.</p>  |

**Table 276: show as-path summary Output Fields**

| Field Name      | Field Description                             |
|-----------------|---|
| <b>AS Paths</b> | Number of AS paths.                           |
| <b>Buckets</b>  | Number of hash buckets in use.                |
| <b>Max</b>      | Maximum number of AS path entries per bucket. |
| <b>Min</b>      | Minimum number of AS path entries per bucket. |
| <b>Avg</b>      | Average number of AS path entries per bucket. |

Table 276: show as-path summary Output Fields (*continued*)

| Field Name    | Field Description                                 |
|---------------|---|
| Std deviation | Standard deviation of AS path entries per bucket. |

## Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets  Max   Min   Avg   Std deviation
30425     1024     95    12    29    6.481419
```

## show ipv6 neighbors

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | show ipv6 neighbors  |
| <b>Release Information</b>      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.3 for EX Series switches.<br>Command introduced in Junos OS Release 12.2 for the QFX Series. |
| <b>Description</b>              | Display information about the IPv6 neighbor cache.   |
| <b>Options</b>                  | This command has no options.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ipv6 neighbors on page 3059</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show ipv6 neighbors on page 3069</a>   |
| <b>Output Fields</b>            | <a href="#">Table 277 on page 3069</a> describes the output fields for the <b>show ipv6 neighbors</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 277: show ipv6 neighbors Output Fields**

| Field Name        | Field Description  |
|-------------------|--|
| IPv6 Address      | Name of the IPv6 interface.  |
| Linklayer Address | Link-layer address.  |
| State             | State of the link: <b>up</b> , <b>down</b> , <b>incomplete</b> , <b>reachable</b> , <b>stale</b> , or <b>unreachable</b> . |
| Exp               | Number of seconds until the entry expires.   |
| Rtr               | Whether the neighbor is a routing device: <b>yes</b> or <b>no</b> .  |
| Secure            | Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: <b>yes</b> or <b>no</b> .              |
| Interface         | Name of the interface.   |

## Sample Output

### show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address          Linklayer Address  State      Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale      546 yes no
fe-1/2/0.1

```

|                                    |                   |       |     |     |    |
|------------------------------------|-------------------|-------|-----|-----|----|
| fe80::2a0:a514:0:24c<br>fe-1/2/0.1 | 00:05:85:8f:c8:bd | stale | 258 | yes | no |
| fe80::2a0:a514:0:64c<br>fe-1/2/1.5 | 00:05:85:8f:c8:bd | stale | 111 | yes | no |
| fe80::2a0:a514:0:a4c<br>fe-1/2/2.9 | 00:05:85:8f:c8:bd | stale | 327 | yes | no |



## show ipv6 router-advertisement

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>show ipv6 router-advertisement &lt;conflicts&gt; &lt;interface <i>interface</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;prefix <i>prefix/prefix length</i>&gt;</pre>  |
| <b>Release Information</b>      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.2 for the QFX Series.</p>  |
| <b>Description</b>              | Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.  |
| <b>Options</b>                  | <p><b>none</b>—Display all IPv6 router advertisement information for all interfaces.</p> <p><b>conflicts</b>—(Optional) Display only the IPv6 router advertisement information that is conflicting.</p> <p><b>interface <i>interface</i></b>—(Optional) Display IPv6 router advertisement information for the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>prefix <i>prefix/prefix length</i></b>—(Optional) Display IPv6 router advertisement information for the specified prefix.</p> |
| <b>Additional Information</b>   | The display identifies conflicting information by enclosing the value the router is advertising in brackets.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear ipv6 router-advertisement</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show ipv6 router-advertisement on page 3072</a><br><a href="#">show ipv6 router-advertisement conflicts on page 3073</a><br><a href="#">show ipv6 router-advertisement prefix on page 3073</a>   |
| <b>Output Fields</b>            | <p><a href="#">Table 278 on page 3071</a> describes the output fields for the <b>show ipv6 router-advertisement</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

Table 278: show ipv6 router-advertisement Output Fields

| Field Name          | Field Description   |
|---------------------|---|
| Interface           | Name of the interface.  |
| Advertisements sent | Number of router advertisements sent and the elapsed time since they were sent. |

Table 278: show ipv6 router-advertisement Output Fields (*continued*)

| Field Name              | Field Description   |
|-------------------------|---|
| Solicits received       | Number of solicitation messages received.   |
| Advertisements received | Number of router advertisements received.   |
| Advertisements from     | Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received. |
| Managed                 | Managed address configuration flag: 0 (stateless) or 1 (stateful).  |
| Other configuration     | Other stateful configuration flag: 0 (stateless) or 1 (stateful).   |
| Reachable time          | Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.             |
| Default lifetime        | Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router. |
| Retransmit timer        | Time between retransmitted Neighbor Solicitation messages, in milliseconds.   |
| Current hop limit       | Configured current hop limit.   |
| Prefix                  | Name and length of the prefix.  |
| Valid lifetime          | How long the prefix remains valid for onlink determination.   |
| Preferred lifetime      | How long the prefix generated by stateless autoconfiguration remains preferred.   |
| On link                 | Onlink flag: 0 (not onlink) or 1 (onlink).  |
| Autonomous              | Autonomous address configuration flag: 0 (not autonomous) or 1 (autonomous).  |

## Sample Output

### show ipv6 router-advertisement

```

user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
  Managed: 0
  Other configuration: 0 [1]
  Reachable time: 0 ms
  Default lifetime: 1800 sec

```

```
Retransmit timer: 0 ms
Current hop limit: 64
```

#### show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
  Other configuration: 0 [1]
```

#### show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
  Managed: 0
  Other configuration: 0
  Reachable time: 0 ms
  Default lifetime: 180 sec [1800 sec]
  Retransmit timer: 0 ms
  Current hop limit: 64
  Prefix: 8040:1::/64
    Valid lifetime: 2592000 sec
    Preferred lifetime: 604800 sec
    On link: 1
    Autonomous: 1
```

## show route

---

|                             |  |
|-----------------------------|--|
| List of Syntax              | <a href="#">Syntax on page 3074</a><br><a href="#">Syntax (EX Series Switches) on page 3074</a>  |
| Syntax                      | <pre>show route &lt;all&gt; &lt;destination-prefix&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;private&gt;</pre>  |
| Syntax (EX Series Switches) | <pre>show route &lt;all&gt; &lt;destination-prefix&gt; &lt;private&gt;</pre>   |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Option <b>private</b> introduced in Junos OS Release 9.5.<br>Option <b>private</b> introduced in Junos OS Release 9.5 for EX Series switches.   |
| Description                 | Display the active entries in the routing tables.  |
| Options                     | <p><b>none</b>—Display brief information about all active entries in the routing tables.</p> <p><b>all</b>—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p><b>destination-prefix</b>—(Optional) Display active entries for the specified address or range of addresses.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>private</b>—(Optional) Display information only about all private, or internal, routing tables.</p> |
| Required Privilege Level    | view   |
| Related Documentation       | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4311</a></li><li>• <a href="#">Example: Configuring RIPng</a></li><li>• <a href="#">Example: Configuring IS-IS</a></li><li>• <a href="#">Examples: Configuring Internal BGP Peering on page 3284</a></li><li>• <a href="#">Examples: Configuring External BGP Peering on page 3261</a></li><li>• <a href="#">Examples: Configuring OSPF Routing Policy on page 4164</a></li></ul>   |
| List of Sample Output       | <a href="#">show route on page 3077</a><br><a href="#">show route on page 3078</a>   |

[show route destination-prefix on page 3078](#)

[show route extensive on page 3078](#)

**Output Fields** [Table 279 on page 3075](#) describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

**Table 279: show route Output Fields**

| Field Name                 | Field Description  |
|----------------------------|--|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).   |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.  |
| <i>number routes</i>       | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly.</li> </ul> <p>However, if you have configured advertisement of multiple routes (with the <b>add-path</b> or <b>advertise-inactive</b> statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul> |
| <i>destination-prefix</i>  | <p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul>  |

Table 279: show route Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| [ <i>protocol, preference</i> ]                   | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• - —A hyphen indicates the last active route.</li> <li>• *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>   |
| <i>weeks:days</i><br><i>hours:minutes:seconds</i> | How long the route been known (for example, <b>2w4d 13:11:14</b> , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).  |
| metric  | Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.  |
| localpref   | Local preference value included in the route.  |
| from  | Interface from which the route was received.   |
| AS path   | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• I—IGP.</li> <li>• E—EGP.</li> <li>• ?—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• [ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• ( )—Parentheses enclose a confederation.</li> <li>• ( [ ] )—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |

Table 279: show route Output Fields (*continued*)

| Field Name              | Field Description  |
|-------------------------|--|
| <b>validation-state</b> | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>  |
| <b>to</b>               | <p>Next hop to the destination. An angle bracket (&gt;) indicates that the route is the selected route.</p> <p>If the destination is <b>Discard</b>, traffic is dropped.</p>   |
| <b>via</b>              | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> <li>• <b>lsp-path-name</b>—Name of the LSP used to reach the next hop.</li> <li>• <b>label-action</b>—MPLS label and operation occurring at the next hop. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label). For VPNs, expect to see multiple <b>push</b> operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).</li> </ul> |

## Sample Output

### show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    *[MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30

```

```

AS path: I
> to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
[BGP/170] 19:53:25, localpref 100, from 10.0.0.33
AS path: I
> to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

## show route

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```

user@host> show route 70.0.0.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

70.0.0.0/24      @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                  #[Multipath/255] 00:28:28, metric2 102
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)

```

## show route destination-prefix

```

user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12   *[Static/5] 2w4d 12:54:27
                  > to 192.168.167.254 via fxp0.0

```

## show route extensive

```

user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
    PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

    Next hop type: Indirect
    Address: 0x92455b8
    Next-hop reference count: 2
    Source: 10.0.0.30
    Protocol next hop: 10.0.0.40
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
      Local AS: 65500 Peer AS: 65500
    Age: 3 Metric2: 1
    Validation State: unverified
    Task: BGP_65500.10.0.0.30+179
    Announcement bits (2): 0-PIM.v1 1-mvpn global task
    AS path: I (Originator) Cluster list: 10.0.0.30
    AS path: Originator ID: 10.0.0.40
    Communities: target:65520:100

```



```
Import Accepted
Localpref: 100
Router ID: 10.0.0.30
Primary Routing Table bgp.mvpn.0
Indirect next hops: 1
  Protocol next hop: 10.0.0.40 Metric: 1
  Indirect next hop: 2 no-forward
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
10.0.0.40/32 Originating RIB: inet.3
  Metric: 1 Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 10.0.24.4 via lt-0/3/0.24
```

## show route active-path

---

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3080</a><br><a href="#">Syntax (EX Series Switches) on page 3080</a>   |
| <b>Syntax</b>                      | <code>show route active-path</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| <b>Syntax (EX Series Switches)</b> | <code>show route active-path</code><br><code>&lt;brief   detail   extensive   terse&gt;</code>  |
| <b>Release Information</b>         | Command introduced in Junos OS Release 8.0.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>                 | Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.  |
| <b>Options</b>                     | <b>none</b> —Display all active routes.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route active-path on page 3080</a><br><a href="#">show route active-path brief on page 3081</a><br><a href="#">show route active-path detail on page 3081</a><br><a href="#">show route active-path extensive on page 3082</a><br><a href="#">show route active-path terse on page 3084</a>  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32   *[IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      *[Direct/0] 00:18:36
                  > via so-2/1/3.0
```

```

100.1.2.2/32      *[Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21  *[Direct/0] 21:33:52
                  > via fxp0.0
192.168.70.19/32 *[Local/0] 21:33:52
                  Local via fxp0.0

```

### show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 3080](#).

### show route active-path detail

```

user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local

```

```
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

#### show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397
```

```

Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3

```

AS path: I

### show route active-path terse

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

| A | Destination      | P | Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---|------------------|---|-----|----------|----------|-------------|---------|
| * | 10.255.70.19/32  | D | 0   |          |          | >1o0.0      |         |
| * | 10.255.71.50/32  | I | 15  | 10       |          | >100.1.2.1  |         |
| * | 100.1.2.0/24     | D | 0   |          |          | >so-2/1/3.0 |         |
| * | 100.1.2.2/32     | L | 0   |          |          | Local       |         |
| * | 192.168.64.0/21  | D | 0   |          |          | >fxp0.0     |         |
| * | 192.168.70.19/32 | L | 0   |          |          | Local       |         |

## show route all

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3085</a><br><a href="#">Syntax (EX Series Switches) on page 3085</a>  |
| <b>Syntax</b>                      | show route all<br><logical-system (all   <i>logical-system-name</i> )>   |
| <b>Syntax (EX Series Switches)</b> | show route all   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display information about all routes in all routing tables, including private, or internal, tables.  |
| <b>Options</b>                     | <p><b>none</b>—Display information about all routes in all routing tables, including private, or internal, tables.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>   |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route all on page 3085</a>  |
| <b>Output Fields</b>               | In Junos OS Release 9.5 and later, only the output fields for the <b>show route all</b> command display all routing tables, including private, or hidden, routing tables. The output field table of the <b>show route</b> command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later. |

## Sample Output

### show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```

user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop

```

```
user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
            Receive
800017     *[VPLS/7] 1d 13:54:49
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
            > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
              Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
              Unusable
```



## show route aspath-regex

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3087</a><br><a href="#">Syntax (EX Series Switches) on page 3087</a>   |
| <b>Syntax</b>                      | show route aspath-regex <i>regular-expression</i><br><logical-system (all   <i>logical-system-name</i> )>   |
| <b>Syntax (EX Series Switches)</b> | show route aspath-regex <i>regular-expression</i>   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>                 | Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.   |
| <b>Options</b>                     | <p><i>regular-expression</i>—Regular expression that matches an entire AS path.</p> <p><i>logical-system</i> (all   <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>   |
| <b>Additional Information</b>      | <p>You can specify a regular expression as:</p> <ul style="list-style-type: none"> <li>• An individual AS number</li> <li>• A period wildcard used in place of an AS number</li> <li>• An AS path regular expression that is enclosed in parentheses</li> </ul> <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none"> <li>• <b>{<i>m,n</i>}</b>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term.</li> <li>• <b>{<i>m</i>}</b>—Exactly <i>m</i> repetitions of the AS path term.</li> <li>• <b>{<i>m</i>,}</b>—<i>m</i> or more repetitions of the AS path term.</li> <li>• <b>*</b>—Zero or more repetitions of an AS path term.</li> <li>• <b>+</b>—One or more repetitions of an AS path term.</li> <li>• <b>?</b>—Zero or one repetition of an AS path term.</li> <li>• <b><i>aspath_term</i>   <i>aspath_term</i></b>—Match one of the two AS path terms.</li> </ul> <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 ."</pre> |

|                          |  |
|--------------------------|--|
| Required Privilege Level | view   |
| Related Documentation    | <ul style="list-style-type: none"><li>• <i>Example: Using AS Path Regular Expressions</i></li></ul>  |
| List of Sample Output    | <a href="#">show route aspath-regex (Matching a Specific AS Number) on page 3088</a><br><a href="#">show route aspath-regex (Matching Any Path with Two AS Numbers) on page 3088</a> |
| Output Fields            | For information about output fields, see the output field table for the <a href="#">show route</a> command.  |

## Sample Output

### [show route aspath-regex \(Matching a Specific AS Number\)](#)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

### [show route aspath-regex \(Matching Any Path with Two AS Numbers\)](#)

```
user@host> show route aspath-regex ?.* 234 3561.*?

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

## show route best

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3089</a><br><a href="#">Syntax (EX Series Switches) on page 3089</a>   |
| <b>Syntax</b>                      | show route best <i>destination-prefix</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>   |
| <b>Syntax (EX Series Switches)</b> | show route best <i>destination-prefix</i><br><brief   detail   extensive   terse>   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>                 | Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>destination-prefix</b> —Address or range of addresses.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route best on page 3089</a><br><a href="#">show route best detail on page 3090</a><br><a href="#">show route best extensive on page 3091</a><br><a href="#">show route best terse on page 3091</a>   |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route best

```

user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSVP/7] 1d 13:20:13, metric 2

```

```

> via so-0/3/0.0, label-switched-path green-r1-r3

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8          *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0

```

### show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  *OSPF   Preference: 10
          Next-hop reference count: 9
          Next hop: 10.31.1.6 via ge-3/1/0.0, selected
          Next hop: via so-0/3/0.0
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:06      Metric: 2
          Area: 0.0.0.0
          Task: OSPF
          Announcement bits (2): 0-KRT 3-Resolve tree 2
          AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 5
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100016
          State: <Active Int>
          Local AS: 69
          Age: 1d 13:20:59      Metric: 2
          Task: RSVP
          Announcement bits (1): 1-Resolve tree 2
          AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp2.0, selected
          State: <Active Int>
          Age: 2d 1:44:20
          Task: IF
          AS path: I
  Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via fxp1.0, selected
          State: <NotBest Int>
          Inactive reason: No difference
          Age: 2d 1:44:20

```

Task: IF  
AS path: I

### show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 3090](#).

### show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  0 10           2           >10.31.1.6
                               so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.255.70.103/32  R  7           2           >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.0.0.0/8        D  0           0           >fxp2.0
                    D  0           0           >fxp1.0
```

## show route brief

---

|                             |   |
|-----------------------------|---|
| List of Syntax              | <a href="#">Syntax on page 3092</a><br><a href="#">Syntax (EX Series Switches) on page 3092</a>   |
| Syntax                      | <code>show route brief</code><br><code>&lt;destination-prefix&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches) | <code>show route brief</code><br><code>&lt;destination-prefix&gt;</code>  |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| Description                 | Display brief information about the active entries in the routing tables.   |
| Options                     | <b>none</b> —Display all active entries in the routing table.<br><br><b><i>destination-prefix</i></b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b><i>logical-system (all   <i>logical-system-name</i>)</i></b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level    | view  |
| List of Sample Output       | <a href="#">show route brief on page 3092</a>   |
| Output Fields               | For information about output fields, see the Output Field table of the <a href="#">show route</a> command.  |

## Sample Output

### show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                   Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                   > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                   > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                   > via fxp0.0
```

```
192.168.164.51/32 *[Local/0] 2w4d 13:11:14
                  Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14
                  > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16    *[Direct/0] 1w5d 20:30:28
                  > via fe-0/0/3.0
100.101.2.3/32   *[Local/0] 1w5d 20:30:28
                  Local via fe-0/0/3.0
224.0.0.5/32     *[OSPF/10] 1w5d 20:30:29, metric 1
                  MultiRecv
```

## show route community

---

|                             |  |
|-----------------------------|--|
| List of Syntax              | <a href="#">Syntax on page 3094</a><br><a href="#">Syntax (EX Series Switches) on page 3094</a>  |
| Syntax                      | <code>show route community <i>as-number:community-value</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches) | <code>show route community <i>as-number:community-value</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>  |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| Description                 | Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.   |
| Options                     | <p><b><i>as-number:community-value</i></b>—One or more community identifiers. <b><i>as-number</i></b> is the AS number, and <b><i>community-value</i></b> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| Additional Information      | Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.   |
| Required Privilege Level    | view   |
| Related Documentation       | <ul style="list-style-type: none"><li>• <a href="#">show route detail on page 3103</a></li></ul>   |
| List of Sample Output       | <a href="#">show route community on page 3094</a>  |
| Output Fields               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.   |

## Sample Output

### show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both
```



```
4.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8      *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 568 721 Incomplete
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16     *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                AS Path: {666} 234 2548 1673 1675 1747 IGP
                to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

## show route community-name

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3096</a><br><a href="#">Syntax (EX Series Switches) on page 3096</a>  |
| <b>Syntax</b>                      | <b>show route community-name</b> <i>community-name</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>   |
| <b>Syntax (EX Series Switches)</b> | <b>show route community-name</b> <i>community-name</i><br><brief   detail   extensive   terse>   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.  |
| <b>Options</b>                     | <i>community-name</i> —Name of the community.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route community-name on page 3096</a>   |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.   |

## Sample Output

### show route community-name

```

user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: 300 I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204

```

```

AS path: I
> to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
    *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
        AS path: 300 I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
        AS path: I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
        AS path: I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

## show route damping

---

|  |   |
|--|---|
| List of Syntax                           | <a href="#">Syntax on page 3098</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3098</a>  |
| Syntax                                   | <code>show route damping (decayed   history   suppressed)</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>   |
| Syntax (EX Series Switch and QFX Series) | <code>show route damping (decayed   history   suppressed)</code><br><code>&lt;brief   detail   extensive   terse&gt;</code>   |
| Release Information                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.  |
| Description                              | Display the BGP routes for which updates might have been reduced because of route flap damping.   |
| Options                                  | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.<br><br><b>decayed</b> —Display route damping entries that might no longer be valid, but are not suppressed.<br><br><b>history</b> —Display entries that have already been withdrawn, but have been logged.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>suppressed</b> —Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols. |
| Required Privilege Level                 | view  |
| Related Documentation                    | <ul style="list-style-type: none"><li>• <a href="#">clear bgp damping on page 3750</a></li><li>• <a href="#">show policy damping on page 3784</a></li></ul>   |
| List of Sample Output                    | <a href="#">show route damping decayed detail on page 3101</a><br><a href="#">show route damping history on page 3102</a><br><a href="#">show route damping history detail on page 3102</a>   |
| Output Fields                            | <a href="#">Table 280 on page 3099</a> lists the output fields for the <b>show route damping</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 280: show route damping Output Fields

| Field Name                                   | Field Description   | Level of Output         |
|--|---|-------------------------|
| <i>routing-table-name</i>                    | Name of the routing table—for example, <b>inet.0</b> .  | All levels              |
| <b>destinations</b>                          | Number of destinations for which there are routes in the routing table.   | All levels              |
| <b>number routes</b>                         | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b></li> <li>• <b>holdddown</b> (routes that are in a pending state before being declared inactive)</li> <li>• <b>hidden</b> (the routes are not used because of a routing policy)</li> </ul>  | All levels              |
| <b>destination-prefix (entry, announced)</b> | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.   | <b>detail extensive</b> |
| <b>[protocol, preference]</b>                | Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p> | All levels              |
| <b>Next-hop reference count</b>              | Number of references made to the next hop.  | <b>detail extensive</b> |
| <b>Source</b>                                | IP address of the route source.   | <b>detail extensive</b> |
| <b>Next hop</b>                              | Network layer address of the directly reachable neighboring system.   | <b>detail extensive</b> |
| <b>via</b>                                   | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b> .  | <b>detail extensive</b> |
| <b>Protocol next hop</b>                     | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.  | <b>detail extensive</b> |
| <b>Indirect next hop</b>                     | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.   | <b>detail extensive</b> |
| <b>State</b>                                 | Flags for this route. For a description of possible values for this field, see the output field table for the <a href="#">show route detail</a> command.  | <b>detail extensive</b> |

Table 280: show route damping Output Fields (*continued*)

| Field Name        | Field Description   | Level of Output  |
|-------------------|---|------------------|
| Local AS          | AS number of the local routing device.  | detail extensive |
| Peer AS           | AS number of the peer routing device.   | detail extensive |
| Age               | How long the route has been known.  | detail extensive |
| Metric            | Metric for the route.   | detail extensive |
| Task              | Name of the protocol that has added the route.  | detail extensive |
| Announcement bits | List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.  | detail extensive |
| AS path           | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• I—IGP.</li> <li>• E—EGP.</li> <li>• ?—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• [ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured.</li> <li>• { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• ( )—Parentheses enclose a confederation.</li> <li>• ( [ ] )—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels       |
| to                | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.   | brief none       |
| via               | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b> .  | brief none       |
| Communities       | Community path attribute for the route. See the output field table for the <a href="#">show route detail</a> command.   | detail extensive |
| Localpref         | Local preference value included in the route.   | All levels       |
| Router ID         | BGP router ID as advertised by the neighbor in the open message.  | detail extensive |

Table 280: show route damping Output Fields (*continued*)

| Field Name                     | Field Description  | Level of Output         |
|--------------------------------|--|-------------------------|
| <b>Merit (last update/now)</b> | Last updated and current figure-of-merit value.  | <b>detail extensive</b> |
| <b>damping-parameters</b>      | Name that identifies the damping parameters used, which is defined in the damping statement at the <b>[edit policy-options]</b> hierarchy level.                         | <b>detail extensive</b> |
| <b>Last update</b>             | Time of most recent change in path attributes.   | <b>detail extensive</b> |
| <b>First update</b>            | Time of first change in path attributes, which started the route damping process.  | <b>detail extensive</b> |
| <b>Flaps</b>                   | Number of times the route has gone up or down or its path attributes have changed.   | <b>detail extensive</b> |
| <b>Suppressed</b>              | ( <b>suppressed</b> keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it. | All levels              |
| <b>Reusable in</b>             | ( <b>suppressed</b> keyword only) Time when a suppressed route will again be available.  | All levels              |
| <b>Preference will be</b>      | ( <b>suppressed</b> keyword only) Preference value that will be applied to the route when it is again active.  | All levels              |

## Sample Output

### show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

        6-Resolve tree 2 7-Resolve tree 3
        AS path: 65490 65520 65525 65525 65525 65525 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:701
        Localpref: 100
        Router ID: 172.23.2.129
        Merit (last update/now): 1934/1790
        damping-parameters: damping-high

```

```
Last update:      00:03:28 First update:      00:06:40
Flaps: 2
```

### show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0
```

### show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:      00:01:05 First update:      00:01:05
        Flaps: 1
```



## show route detail

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3103</a><br><a href="#">Syntax (EX Series Switches) on page 3103</a>   |
| <b>Syntax</b>                      | show route detail<br><destination-prefix><br><logical-system (all   logical-system-name)>   |
| <b>Syntax (EX Series Switches)</b> | show route detail<br><destination-prefix>   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.   |
| <b>Description</b>                 | Display detailed information about the active entries in the routing tables.  |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table on all systems.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route detail on page 3112</a><br><a href="#">show route detail (with BGP Multipath) on page 3118</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3118</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3119</a>  |
| <b>Output Fields</b>               | Table 281 on page 3103 describes the output fields for the <b>show route detail</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 281: show route detail Output Fields

| Field Name                 | Field Description   |
|----------------------------|---|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).  |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.   |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active)</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |

Table 281: show route detail Output Fields (*continued*)

| Field Name                                     | Field Description   |
|--|---|
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example:10.0.0.1/24). The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| <b>label stacking</b>                          | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>  |
| <i>[protocol, preference]</i>                  | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+—</b>A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>- —</b>A hyphen indicates the last active route.</li> <li>• <b>*—</b>An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>   |
| <b>Level</b>                                   | <p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>  |
| <b>Route Distinguisher</b>                     | IP subnet augmented with a 64-bit prefix.   |
| <b>PMSI</b>                                    | Provider multicast service interface (MVPN routing table).  |
| <b>Next-hop type</b>                           | Type of next hop. For a description of possible values for this field, see <a href="#">Table 282 on page 3108</a> .   |

Table 281: show route detail Output Fields (*continued*)

| Field Name   | Field Description   |
|--|---|
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.  |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.   |
| <b>Source</b>  | IP address of the route source.   |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.   |
| <b>via</b>   | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path<br/>lsp-path-name</b>         | Name of the LSP used to reach the next hop.   |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).  |
| <b>Interface</b>                                     | (Local only) Local interface name.  |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.  |
| <b>Indirect next hop</b>                             | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.   |
| <b>State</b>   | State of the route (a route can be in more than one state). See <a href="#">Table 283 on page 3109</a> .  |
| <b>Local AS</b>                                      | AS number of the local routing device.  |
| <b>Age</b>   | How long the route has been known.  |
| <b>AIGP</b>  | Accumulated interior gateway protocol (AIGP) BGP attribute.   |
| <b>Metricn</b>                                       | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.   |

Table 281: show route detail Output Fields (*continued*)

| Field Name                 | Field Description   |
|----------------------------|---|
| <b>MED-plus-IGP</b>        | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.   |
| <b>TTL-Action</b>          | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>   |
| <b>Task</b>                | Name of the protocol that has added the route.  |
| <b>Announcement bits</b>   | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.  |
| <b>AS path</b>             | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893.</li> <li>• <b>[ ]</b>—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>validation-state</b>    | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>   |
| <b>FECs bound to route</b> | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.   |

Table 281: show route detail Output Fields (*continued*)

| Field Name                | Field Description   |
|---------------------------|---|
| Primary Upstream          | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. |
| RPF Nexthops              | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.   |
| Label                     | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                         |
| weight                    | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.  |
| VC Label                  | MPLS label assigned to the Layer 2 circuit virtual connection.  |
| MTU                       | Maximum transmission unit (MTU) of the Layer 2 circuit.   |
| VLAN ID                   | VLAN identifier of the Layer 2 circuit.   |
| Prefixes bound to route   | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.  |
| Communities               | Community path attribute for the route. See <a href="#">Table 284 on page 3111</a> for all possible values for this field.  |
| Layer2-info: encaps       | Layer 2 encapsulation (for example, VPLS).  |
| control flags             | Control flags: <b>none</b> or <b>Site Down</b> .  |
| mtu                       | Maximum transmission unit (MTU) information.  |
| Label-Base, range         | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.  |
| status vector             | Layer 2 VPN and VPLS network layer reachability information (NLRI).   |
| Accepted Multipath        | Current active path when BGP multipath is configured.   |
| Accepted MultipathContrib | Path currently contributing to BGP multipath.   |
| Localpref                 | Local preference value included in the route.   |
| Router ID                 | BGP router ID as advertised by the neighbor in the open message.  |
| Primary Routing Table     | In a routing table group, the name of the primary routing table in which the route resides.   |
| Secondary Tables          | In a routing table group, the name of one or more secondary tables in which the route resides.  |

Table 282 on page 3108 describes all possible values for the Next-hop Types output field.

**Table 282: Next-hop Types Output Field Values**

| Next-Hop Type                   | Description  |
|---------------------------------|--|
| <b>Broadcast (bcast)</b>        | Broadcast next hop.  |
| <b>Deny</b>                     | Deny next hop.   |
| <b>Discard</b>                  | Discard next hop.  |
| <b>Flood</b>                    | Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast. |
| <b>Hold</b>                     | Next hop is waiting to be resolved into a unicast or multicast type.   |
| <b>Indexed (idxd)</b>           | Indexed next hop.  |
| <b>Indirect (indr)</b>          | Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.   |
| <b>Interface</b>                | Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.   |
| <b>Local (locl)</b>             | Local address on an interface. This next-hop type causes packets with this destination address to be received locally.   |
| <b>Multicast (mcst)</b>         | Wire multicast next hop (limited to the LAN).  |
| <b>Multicast discard (mdsc)</b> | Multicast discard.   |
| <b>Multicast group (mgrp)</b>   | Multicast group member.  |
| <b>Receive (recv)</b>           | Receive.   |
| <b>Reject (rjct)</b>            | Discard. An ICMP unreachable message was sent.   |
| <b>Resolve (rslv)</b>           | Resolving next hop.  |
| <b>Routed multicast (mcrt)</b>  | Regular multicast next hop.  |

Table 282: Next-hop Types Output Field Values (*continued*)

| Next-Hop Type         | Description  |
|-----------------------|--|
| <b>Router</b>         | <p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> <li>• Must not be a direct or local subnet for the routing device.</li> <li>• Must have a next hop that is directly connected to the routing device.</li> </ul> |
| <b>Table</b>          | Routing table next hop.  |
| <b>Unicast (ucst)</b> | Unicast.   |
| <b>Unilist (ulst)</b> | List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.  |

Table 283 on page 3109 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 283: State Output Field Values

| Value  | Description  |
|--|--|
| <b>Accounting</b>                            | Route needs accounting.  |
| <b>Active</b>                                | Route is active.   |
| <b>Always Compare MED</b>                    | Path with a lower multiple exit discriminator (MED) is available.                |
| <b>AS path</b>                               | Shorter AS path is available.  |
| <b>Cisco Non-deterministic MED selection</b> | Cisco nondeterministic MED is enabled, and a path with a lower MED is available. |
| <b>Clone</b>                                 | Route is a clone.  |
| <b>Cluster list length</b>                   | Length of cluster list sent by the route reflector.                              |
| <b>Delete</b>                                | Route has been deleted.  |
| <b>Ex</b>                                    | Exterior route.  |
| <b>Ext</b>                                   | BGP route received from an external BGP neighbor.                                |

Table 283: State Output Field Values (*continued*)

| Value  | Description  |
|--|--|
| <b>FlashAll</b>  | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| <b>Hidden</b>  | Route not used because of routing policy.  |
| <b>IfCheck</b>   | Route needs forwarding RPF check.  |
| <b>IGP metric</b>  | Path through next hop with lower IGP metric is available.  |
| <b>Inactive reason</b>                                   | Flags for this route, which was not selected as best for a particular destination.   |
| <b>Initial</b>   | Route being added.   |
| <b>Int</b>   | Interior route.  |
| <b>Int Ext</b>   | BGP route received from an internal BGP peer or a BGP confederation peer.  |
| <b>Interior &gt; Exterior &gt; Exterior via Interior</b> | Direct, static, IGP, or EBGP path is available.  |
| <b>Local Preference</b>                                  | Path with a higher local preference value is available.  |
| <b>Martian</b>   | Route is a martian (ignored because it is obviously invalid).  |
| <b>MartianOK</b>   | Route exempt from martian filtering.   |
| <b>Next hop address</b>                                  | Path with lower metric next hop is available.  |
| <b>No difference</b>                                     | Path from neighbor with lower IP address is available.   |
| <b>NoReadvrt</b>   | Route not to be advertised.  |
| <b>NotBest</b>   | Route not chosen because it does not have the lowest MED.  |
| <b>Not Best in its group</b>                             | Incoming BGP AS is not the best of a group (only one AS can be the best).  |
| <b>NotInstall</b>  | Route not to be installed in the forwarding table.   |
| <b>Number of gateways</b>                                | Path with a greater number of next hops is available.  |
| <b>Origin</b>  | Path with a lower origin code is available.  |
| <b>Pending</b>   | Route pending because of a hold-down configured on another route.  |



Table 283: State Output Field Values (*continued*)

| Value                                 | Description   |
|---------------------------------------|---|
| <b>Release</b>                        | Route scheduled for release.  |
| <b>RIB preference</b>                 | Route from a higher-numbered routing table is available.  |
| <b>Route Distinguisher</b>            | 64-bit prefix added to IP subnets to make them unique.  |
| <b>Route Metric or MED comparison</b> | Route with a lower metric or MED is available.  |
| <b>Route Preference</b>               | Route with lower preference value is available  |
| <b>Router ID</b>                      | Path through a neighbor with lower ID is available.   |
| <b>Secondary</b>                      | Route not a primary route.  |
| <b>Unusable path</b>                  | Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> <li>• The route is damped.</li> <li>• The route is rejected by an import policy.</li> <li>• The route is unresolved.</li> </ul> |
| <b>Update source</b>                  | Last tiebreaker is the lowest IP address value.   |

Table 284 on page 3111 describes the possible values for the Communities output field.

Table 284: Communities Output Field Values

| Value   | Description   |
|---|---|
| <i>area-number</i>                                      | 4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.  |
| <b>bandwidth: local AS number:link-bandwidth-number</b> | Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute. |
| <b>domain-id</b>  | Unique configurable number that identifies the OSPF domain.   |
| <b>domain-id-vendor</b>                                 | Unique configurable number that further identifies the OSPF domain.   |
| <i>link-bandwidth-number</i>                            | Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).   |
| <i>local AS number</i>                                  | Local AS number: from 1 through 65,535.   |
| <i>options</i>  | 1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.   |

Table 284: Communities Output Field Values (*continued*)

| Value                                | Description   |
|--------------------------------------|---|
| <b>origin</b>                        | (Used with VPNs) Identifies where the route came from.  |
| <b>ospf-route-type</b>               | 1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses. |
| <b>route-type-vendor</b>             | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x8000</b> . The format is <b>area-number:ospf-route-type:options</b> .  |
| <b>rte-type</b>                      | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x0306</b> . The format is <b>area-number:ospf-route-type:options</b> .  |
| <b>target</b>                        | Defines which VPN the route participates in; <b>target</b> has the format <b>32-bit IP address:16-bit number</b> . For example, 10.19.0.0:100.  |
| <b>unknown IANA</b>                  | Incoming IANA codes with a value between <b>0x1</b> and <b>0x7fff</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.  |
| <b>unknown OSPF vendor community</b> | Incoming IANA codes with a value above <b>0x8000</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.   |

## Sample Output

### show route detail

```

user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10

```

```

Next-hop reference count: 1
Next hop: via so-0/3/0.0, selected
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:30:17 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: IGMP
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

```

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS: 69
          Age: 1:25:49    Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS: 69
          Age: 1:25:49    Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
          State: <Active Int>
          Local AS: 69
          Age: 1:31:44
          Task: IF
          AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
  *MPLS   Preference: 0
          Next hop type: Receive
          Next-hop reference count: 6
          State: <Active Int>
          Local AS: 69
          Age: 1:31:45    Metric: 1
          Task: MPLS
          Announcement bits (1): 0-KRT
          AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

299840 (1 entry, 1 announced)
```

```

TSI:
KRT in-kerne 299840 /52 -> {indirect(1048575)}
    *RSVP Preference: 7/2
        Next hop type: Flood
        Address: 0x9174a30
        Next-hop reference count: 4
        Next hop type: Router, Next hop index: 798
        Address: 0x9174c28
        Next-hop reference count: 2
        Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
        Label-switched-path R2-to-R4-2p2mp
        Label operation: Pop
        Next hop type: Router, Next hop index: 1048574
        Address: 0x92544f0
        Next-hop reference count: 2
        Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
        Label-switched-path R2-to-R200-p2mp
        Label operation: Pop
        Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
        Label operation: Pop
        State: <Active Int>
        Age: 1:29 Metric: 1
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I...

800010 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:29:30
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:29:30 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected

```

```
State: <Active Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

ff02::2/128 (1 entry, 1 announced)
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::d/128 (1 entry, 1 announced)
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::16/128 (1 entry, 1 announced)
*MLD Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:43
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:31:44
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS: 69 Peer AS: 69
        Age: 1:25:49 Metric2: 1
        AIGP 210
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-green-l2vpn
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Label-base: 800008, range: 8
        Localpref: 100
        Router ID: 10.255.70.103
        Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
        mtu: 0
        Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
        Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp
        Label operation: Push 100000[0]
        Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
        State: <Active Int>
        Local AS: 99
        Age: 10:21
        Task: l2 circuit

```

```
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512
```

### show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 262142
        Address: 0x901a010
        Next-hop reference count: 2
        Source: 10.1.1.2
        Next hop: 10.1.1.2 via ge-0/3/0.1, selected
        Next hop: 10.1.1.6 via ge-0/3/0.5
        State: <Active Ext>
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.2+59955
        Announcement bits (1): 0-KRT
        AS path: 2 I
        Accepted Multipath
        Localpref: 100
        Router ID: 1.1.1.2
  BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 678
        Address: 0x8f97520
        Next-hop reference count: 9
        Source: 10.1.1.6
        Next hop: 10.1.1.6 via ge-0/3/0.5, selected
        State: <NotBest Ext>
        Inactive reason: Not Best in its group - Active preferred
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.6+58198
        AS path: 2 I
        Accepted MultipathContrib
        Localpref: 100
        Router ID: 1.1.1.3
```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Next-hop reference count: 3
        Address: 0x9097d90
        Next hop: via vt-0/1/0.1
        Next-hop index: 661
        Label operation: Pop
        Address: 0x9172130
        Next hop: via so-0/0/3.0
        Next-hop index: 654
        Label operation: Swap 299872
        State: **Active Int>
        Local AS: 1001
```



```

Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
    Primary Upstream : 1.1.1.3:0--1.1.1.2:0
      RPF Nexthops :
        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
    Backup Upstream : 1.1.1.3:0--1.1.1.6:0
      RPF Nexthops :
        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

## show route exact

---

|                             |  |
|-----------------------------|--|
| List of Syntax              | <a href="#">Syntax on page 3120</a><br><a href="#">Syntax (EX Series Switches) on page 3120</a>  |
| Syntax                      | <code>show route exact <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>   |
| Syntax (EX Series Switches) | <code>show route exact <i>destination-prefix</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>   |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| Description                 | Display only the routes that exactly match the specified address or range of addresses.  |
| Options                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b><i>destination-prefix</i></b> —Address or range of addresses.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level    | view   |
| List of Sample Output       | <a href="#">show route exact on page 3120</a><br><a href="#">show route exact detail on page 3120</a><br><a href="#">show route exact extensive on page 3121</a><br><a href="#">show route exact terse on page 3121</a>  |
| Output Fields               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.   |

## Sample Output

### show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0
```

### show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
```

```
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2d 3:30:26
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

#### show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I
```

#### show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 207.17.136.0/24  S  5                >192.168.71.254
```

## show route export

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3122</a><br><a href="#">Syntax (EX Series Switches) on page 3122</a>   |
| <b>Syntax</b>                      | <pre>show route export &lt;brief   detail&gt; &lt;instance &lt;instance-name&gt;   routing-table-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show route export &lt;brief   detail&gt; &lt;instance &lt;instance-name&gt;   routing-table-name&gt;</pre>   |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>                 | Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.  |
| <b>Options</b>                     | <p><b>none</b>—(Same as <b>brief</b>.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance &lt;instance-name&gt;</b>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>routing-table-name</b>—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route export inet</b> command).</p> |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route export on page 3123</a><br><a href="#">show route export detail on page 3123</a><br><a href="#">show route export instance detail on page 3123</a>   |
| <b>Output Fields</b>               | <a href="#">Table 285 on page 3122</a> lists the output fields for the <b>show route export</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 285: show route export Output Fields

| Field Name                 | Field Description   | Level of Output   |
|----------------------------|---|-------------------|
| Table or <i>table-name</i> | Name of the routing tables that either import or export routes.   | All levels        |
| Routes                     | Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one. | <b>brief</b> none |

Table 285: show route export Output Fields (*continued*)

| Field Name    | Field Description   | Level of Output   |
|---------------|---|-------------------|
| Export        | Whether the table is currently exporting routes to other tables: <b>Y</b> or <b>N</b> (Yes or No).  | <b>brief</b> none |
| Import        | Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)  | <b>detail</b>     |
| Flags         | ( <b>instance</b> keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> <li><b>config auto-policy</b>—The policy was deduced from the configured IGP export policies.</li> <li><b>cleanup</b>—Configuration information for this instance is no longer valid.</li> <li><b>config</b>—The instance was explicitly configured.</li> </ul> | <b>detail</b>     |
| Options       | ( <b>instance</b> keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> <li><b>unicast</b>—Indicates <i>instance.inet.0</i>.</li> <li><b>multicast</b>—Indicates <i>instance.inet.2</i>.</li> <li><b>unicast multicast</b>—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>.</li> </ul>    | <b>detail</b>     |
| Import policy | ( <b>instance</b> keyword only) Policy that <b>route export</b> uses to construct the import-export matrix. Not displayed if the instance type is <b>vrf</b> .  | <b>detail</b>     |
| Instance      | ( <b>instance</b> keyword only) Name of the routing instance.   | <b>detail</b>     |
| Type          | ( <b>instance</b> keyword only) Type of routing instance: <b>forwarding</b> , <b>non-forwarding</b> , or <b>vrf</b> .   | <b>detail</b>     |

## Sample Output

### show route export

```

user@host> show route export
Table      Export      Routes
inet.0     N            0
black.inet.0 Y           3
red.inet.0 Y            4

```

### show route export detail

```

user@host> show route export detail
inet.0                                Routes:      0
black.inet.0                          Routes:      3
  Import: [ inet.0 ]
red.inet.0                            Routes:      4
  Import: [ inet.0 ]

```

### show route export instance detail

```

user@host> show route export instance detail
Instance: master                      Type: forwarding
Flags: <config auto-policy> Options: <unicast multicast>
Import policy: [ (ospf-master-from-red || isis-master-from-black) ]

```

Instance: black  
Instance: red

Type: non-forwarding  
Type: non-forwarding

## show route extensive

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3125</a><br><a href="#">Syntax (EX Series Switches) on page 3125</a>  |
| <b>Syntax</b>                      | show route extensive<br><destination-prefix><br><logical-system (all   logical-system-name)>   |
| <b>Syntax (EX Series Switches)</b> | show route extensive<br><destination-prefix>   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display extensive information about the active entries in the routing tables.  |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.   |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route extensive on page 3132</a><br><a href="#">show route extensive (Access Route) on page 3138</a><br><a href="#">show route extensive (BGP PIC Edge) on page 3139</a><br><a href="#">show route extensive (FRR and LFA) on page 3139</a><br><a href="#">show route extensive (Route Reflector) on page 3140</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3140</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3141</a> |
| <b>Output Fields</b>               | Table 286 on page 3125 describes the output fields for the <b>show route extensive</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 286: show route extensive Output Fields**

| Field Name                 | Field Description   |
|----------------------------|---|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).                        |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table. |

Table 286: show route extensive Output Fields (*continued*)

| Field Name                                     | Field Description   |
|--|---|
| <i>number routes</i>                           | <p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive).</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul>   |
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example: 10.0.0.1/24). The <b>entry</b> value is the number of route for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| <b>TSI</b>                                     | Protocol header information.  |
| <b>label stacking</b>                          | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>  |
| <b>[protocol, preference]</b>                  | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>  |



Table 286: show route extensive Output Fields (*continued*)

| Field Name   | Field Description   |
|--|---|
| <b>Level</b>   | (IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.   |
| <b>Route Distinguisher</b>                           | IP subnet augmented with a 64-bit prefix.   |
| <b>PMSI</b>  | Provider multicast service interface (MVPN routing table).  |
| <b>Next-hop type</b>                                 | Type of next hop. For a description of possible values for this field, see the Output Field table in the <a href="#">show route detail</a> command.   |
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.  |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.   |
| <b>Source</b>  | IP address of the route source.   |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.   |
| <b>via</b>   | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path lsp-path-name</b>             | Name of the LSP used to reach the next hop.   |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).  |
| <b>Offset</b>  | Whether the metric has been increased or decreased by an offset value.  |
| <b>Interface</b>                                     | (Local only) Local interface name.  |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.  |

Table 286: show route extensive Output Fields (*continued*)

| Field Name                    | Field Description  |
|-------------------------------|--|
| <b><i>label-operation</i></b> | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).   |
| <b>Indirect next hops</b>     | <p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain <b>Indirect next hop: weight</b> follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> <li>• 0x1 indicates active next hops.</li> <li>• 0x4000 indicates passive next hops.</li> </ul> |
| <b>State</b>                  | State of the route (a route can be in more than one state). See the Output Field table in the <a href="#">show route detail</a> command.   |
| <b>Session ID</b>             | The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).  |
| <b>Weight</b>                 | <p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see <a href="#">show route table</a>.</p>   |

Table 286: show route extensive Output Fields (*continued*)

| Field Name      | Field Description   |
|-----------------|---|
| Inactive reason | <p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> <li>• <b>Active preferred</b>—Currently active route was selected over this route.</li> <li>• <b>Always compare MED</b>—Path with a lower multiple exit discriminator (MED) is available.</li> <li>• <b>AS path</b>—Shorter AS path is available.</li> <li>• <b>Cisco Non-deterministic MED selection</b>—Cisco nondeterministic MED is enabled and a path with a lower MED is available.</li> <li>• <b>Cluster list length</b>—Path with a shorter cluster list length is available.</li> <li>• <b>Forwarding use only</b>—Path is only available for forwarding purposes.</li> <li>• <b>IGP metric</b>—Path through the next hop with a lower IGP metric is available.</li> <li>• <b>IGP metric type</b>—Path with a lower OSPF link-state advertisement type is available.</li> <li>• <b>Interior &gt; Exterior &gt; Exterior via Interior</b>—Direct, static, IGP, or EBGP path is available.</li> <li>• <b>Local preference</b>—Path with a higher local preference value is available.</li> <li>• <b>Next hop address</b>—Path with a lower metric next hop is available.</li> <li>• <b>No difference</b>—Path from a neighbor with a lower IP address is available.</li> <li>• <b>Not Best in its group</b>—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).</li> <li>• <b>Number of gateways</b>—Path with a higher number of next hops is available.</li> <li>• <b>Origin</b>—Path with a lower origin code is available.</li> <li>• <b>OSPF version</b>—Path does not support the indicated OSPF version.</li> <li>• <b>RIB preference</b>—Route from a higher-numbered routing table is available.</li> <li>• <b>Route distinguisher</b>—64-bit prefix added to IP subnets to make them unique.</li> <li>• <b>Route metric or MED comparison</b>—Route with a lower metric or MED is available.</li> <li>• <b>Route preference</b>—Route with a lower preference value is available.</li> <li>• <b>Router ID</b>—Path through a neighbor with a lower ID is available.</li> <li>• <b>Unusable path</b>—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved.</li> <li>• <b>Update source</b>—Last tiebreaker is the lowest IP address value.</li> </ul> |
| Local AS        | Autonomous system (AS) number of the local routing device.  |
| Age             | How long the route has been known.  |
| AIGP            | Accumulated interior gateway protocol (AIGP) BGP attribute.   |
| Metric          | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.   |
| MED-plus-IGP    | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.   |
| TTL-Action      | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>   |

Table 286: show route extensive Output Fields (*continued*)

| Field Name                           | Field Description   |
|--------------------------------------|---|
| <b>Task</b>                          | Name of the protocol that has added the route.  |
| <b>Announcement bits</b>             | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.  |
| <b>AS path</b>                       | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>validation-state</b>              | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>  |
| <b>FECs bound to route</b>           | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.   |
| <b>AS path: I &lt;Originator&gt;</b> | (For route reflected output only) Originator ID attribute set by the route reflector.   |
| <b>Primary Upstream</b>              | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.   |
| <b>RPF Nexthops</b>                  | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.   |

Table 286: show route extensive Output Fields (*continued*)

| Field Name              | Field Description  |
|-------------------------|--|
| Label                   | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                      |
| weight                  | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.   |
| VC Label                | MPLS label assigned to the Layer 2 circuit virtual connection.   |
| MTU                     | Maximum transmission unit (MTU) of the Layer 2 circuit.  |
| VLAN ID                 | VLAN identifier of the Layer 2 circuit.  |
| Cluster list            | (For route reflected output only) Cluster ID sent by the route reflector.  |
| Originator ID           | (For route reflected output only) Address of router that originally sent the route to the route reflector.   |
| Prefixes bound to route | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.   |
| Communities             | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.  |
| Layer2-info: encaps     | Layer 2 encapsulation (for example, VPLS).   |
| control flags           | Control flags: <b>none</b> or Site Down.   |
| mtu                     | Maximum transmission unit (MTU) information.   |
| Label-Base, range       | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.   |
| status vector           | Layer 2 VPN and VPLS network layer reachability information (NLRI).  |
| Localpref               | Local preference value included in the route.  |
| Router ID               | BGP router ID as advertised by the neighbor in the open message.   |
| Primary Routing Table   | In a routing table group, the name of the primary routing table in which the route resides.  |
| Secondary Tables        | In a routing table group, the name of one or more secondary tables in which the route resides.   |
| Originating RIB         | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix. |
| Node path count         | Number of nodes in the path.   |

Table 286: show route extensive Output Fields (*continued*)

| Field Name                 | Field Description   |
|----------------------------|---|
| <b>Forwarding nexthops</b> | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it. |

## Sample Output

### show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 29
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 1:34:06
        Task: RT
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

10.31.1.0/30 (2 entries, 1 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 2
        Next hop: via so-0/3/0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:32:40
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I
    OSPF Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/3/0.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Local AS: 69
        Age: 1:32:40 Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

10.31.1.1/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 7
        Interface: so-0/3/0.0
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:32:43
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I

```

```

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.31.2.0/30 -> {10.31.1.6}
    *OSPF   Preference: 10
            Next-hop reference count: 9
            Next hop: via so-0/3/0.0
            Next hop: 10.31.1.6 via ge-3/1/0.0, selected
            State: <Active Int>
            Local AS:    69
            Age: 1:32:19   Metric: 2
            Area: 0.0.0.0
            Task: OSPF
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.2/32 -> {}
    *PIM    Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:08
            Task: PIM Recv
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.22/32 -> {}
    *IGMP   Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:06
            Task: IGMP
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP   Preference: 7
            Next-hop reference count: 6
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 100096
            State: <Active Int>
            Local AS:    69
            Age: 1:28:12   Metric: 2
            Task: RSVP
            Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
            AS path: I

```

```
10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS: 69
          Age: 1:28:12    Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
          State: <Active Int>
          Local AS: 69
          Age: 1:34:07
          Task: IF
          AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
  *MPLS   Preference: 0
          Next hop type: Receive
          Next-hop reference count: 6
          State: <Active Int>
          Local AS: 69
          Age: 1:34:08    Metric: 1
          Task: MPLS
          Announcement bits (1): 0-KRT
          AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299840 (1 entry, 1 announced)
TSI:
KRT in-kernel 299840 /52 -> {indirect(1048575)}
  *RSVP   Preference: 7/2
          Next hop type: Flood
          Address: 0x9174a30
          Next-hop reference count: 4
          Next hop type: Router, Next hop index: 798
          Address: 0x9174c28
          Next-hop reference count: 2
          Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
          Label-switched-path R2-to-R4-2p2mp
```



```

Label operation: Pop
Next hop type: Router, Next hop index: 1048574
Address: 0x92544f0
Next-hop reference count: 2
Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
Label-switched-path R2-to-R200-p2mp
Label operation: Pop
Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
Label operation: Pop
State: <Active Int>
Age: 1:29      Metric: 1
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I...

```

800010 (1 entry, 1 announced)

TSI:

```

KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: via vt-3/2/0.32769, selected
    Label operation: Pop
    State: <Active Int>
    Age: 1:31:53
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

```

vt-3/2/0.32769 (1 entry, 1 announced)

TSI:

```

KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:31:53      Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Push 800012
      Indirect next hop: 87272e4 1048574
      Indirect path forwarding next hops: 1
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
        Metric: 2                      Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.31.1.6 via ge-3/1/0.0

```

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)

```
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Local AS: 69
  Age: 1:34:07
  Task: IF
  AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
```

```

Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:34:07
Task: IF
AS path: I

```

```
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
```

```

*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 1:28:12 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-l2vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

```

```
10.255.71.52:1:1:1/96 (1 entry, 1 announced)
```

```
TSI:
```

```
Page 0 idx 0 Type 1 val 8699540
```

```

*L2VPN Preference: 170/-1
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8, status-vector: 0x9F

```

```
10.255.71.52:1:5:1/96 (1 entry, 1 announced)
```

```
TSI:
```

```
Page 0 idx 0 Type 1 val 8699528
```

```

*L2VPN Preference: 170/-101
Next-hop reference count: 5
Protocol next hop: 10.255.71.52
Indirect next hop: 0 -
State: <Active Int Ext>
Age: 1:34:03 Metric2: 1
Task: green-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
Label-base: 800008, range: 8, status-vector: 0x9F

```

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)

```
*L2CKT Preference: 7
  Next hop: via so-1/1/2.0 weight 1, selected
  Label-switched-path my-lsp
  Label operation: Push 100000[0]
  Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
  State: <Active Int>
  Local AS: 99
  Age: 10:21
  Task: l2 circuit
  Announcement bits (1): 0-LDP
  AS path: I
  VC Label 100000, MTU 1500, VLAN ID 512
```

55.0.0.0/24 (1 entry, 1 announced)

TSI:

KRT queued (pending) add

55.0.0.0/24 -> {Push 300112}

```
*BGP Preference: 170/-101
  Next hop type: Router
  Address: 0x925c208
  Next-hop reference count: 2
  Source: 10.0.0.9
  Next hop: 10.0.0.9 via ge-1/2/0.15, selected
  Label operation: Push 300112
  Label TTL action: prop-ttl
  State: <Active Ext>
  Local AS: 7019 Peer AS: 13979
  Age: 1w0d 23:06:56
  AIGP: 25
  Task: BGP_13979.10.0.0.9+56732
  Announcement bits (1): 0-KRT
  AS path: 13979 7018 I
  Accepted
  Route Label: 300112
  Localpref: 100
  Router ID: 10.9.9.1
```

#### show route extensive (Access Route)

user@host> show route 13.160.0.102 extensive

inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)

13.160.0.102/32 (1 entry, 1 announced)

TSI:

KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}

OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern

```
*Access Preference: 13
  Next-hop reference count: 78472
  Next hop: 13.160.0.2 via fe-0/0/0.0, selected
  State: <Active Int>
```

Age: 12

```
Task: RPD Unix Domain Server./var/run/rpd_serv.local
Announcement bits (2): 0-KRT 1-OSPFv2
AS path: I
```

## show route extensive (BGP PIC Edge)

```

user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
TSI:
KRT in-kerne1 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
Page 0 idx 0 Type 1 val 9219e30
  Nexthop: Self
  AS path: [2] 3 I
  Communities: target:2:1
  Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
..
    #Multipath Preference: 255
      Next hop type: Indirect
      Address: 0x93f4010
      Next-hop reference count: 2
..
      Protocol next hop: 1.1.1.4
      Push 299824
      Indirect next hop: 944c000 1048574 INH Session ID: 0x3
      Indirect next hop: weight 0x1
      Protocol next hop: 1.1.1.5
      Push 299824
      Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
      Indirect next hop: weight 0x4000
      State: <ForwardingOnly Int Ext>
      Inactive reason: Forwarding use only
      Age: 25      Metric2: 15
      Validation State: unverified
      Task: RT
      Announcement bits (1): 0-KRT
      AS path: 3 I
      Communities: target:2:1

```

## show route extensive (FRR and LFA)

```

user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
  20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll
TSI:
KRT in-kerne1 20.31.2.0/24 -> {Push 299776, Push 299792}
  *RSVP Preference: 7/1
    Next hop type: Router, Next hop index: 1048574
    Address: 0xbbbc010
    Next-hop reference count: 5
    Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299776
    Label TTL action: prop-ttl
    Session Id: 0x201
    Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299792
    Label TTL action: prop-ttl
    Session Id: 0x202
    State: Active Int
    Local AS: 100
    Age: 5:31 Metric: 2

```

```

Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

### show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
    Protocol next hop: 207.17.136.192 Metric: 0
    Indirect next hop: 84ac908 40
    Indirect path forwarding next hops: 0
    Next hop type: Discard

```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0
Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>

```

```

Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

## show route flow validation

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3142</a><br><a href="#">Syntax (EX Series Switches) on page 3142</a>  |
| <b>Syntax</b>                      | show route flow validation<br><brief   detail><br><ip-prefix><br><table table-name><br><logical-system (all   logical-system-name)>  |
| <b>Syntax (EX Series Switches)</b> | show route flow validation<br><brief   detail><br><ip-prefix><br><table table-name>  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display flow route information.  |
| <b>Options</b>                     | <p><b>none</b>—Display flow route information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p><b>ip-prefix</b>—(Optional) IP address for the flow route.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>table table-name</b>—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route flow validation inet</b> command).</p> |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route flow validation on page 3143</a>  |
| <b>Output Fields</b>               | <a href="#">Table 287 on page 3142</a> lists the output fields for the <b>show route flow validation</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 287: show route flow validation Output Fields

| Field Name                | Field Description                                | Level of Output |
|---------------------------|--|-----------------|
| <i>routing-table-name</i> | Name of the routing table (for example, inet.0). | All levels      |
| <i>prefix</i>             | Route address.                                   | All levels      |
| Active unicast route      | Active route in the routing table.               | All levels      |



Table 287: show route flow validation Output Fields (*continued*)

| Field Name                  | Field Description   | Level of Output |
|-----------------------------|---|-----------------|
| Dependent flow destinations | Number of flows for which there are routes in the routing table.        | All levels      |
| Origin                      | Source of the route flow.   | All levels      |
| Neighbor AS                 | Autonomous system identifier of the neighbor.                           | All levels      |
| Flow destination            | Number of entries and number of destinations that match the route flow. | All levels      |
| Unicast best match          | Destination that is the best match for the route flow.                  | All levels      |
| Flags                       | Information about the route flow.                                       | All levels      |

## Sample Output

### show route flow validation

```

user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent

```

## show route forwarding-table

---

|  |   |
|--|---|
| <b>List of Syntax</b>                                | <a href="#">Syntax on page 3144</a><br><a href="#">Syntax (MX Series Routers) on page 3144</a><br><a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 3144</a>  |
| <b>Syntax</b>  | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre>   |
| <b>Syntax (MX Series Routers)</b>                    | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;bridge-domain (all   domain-name)&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;learning-vlan-id learning-vlan-id&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre> |
| <b>Syntax (TX Matrix and TX Matrix Plus Routers)</b> | <pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;matching matching&gt; &lt;label name&gt; &lt;lcc number&gt; &lt;multicast&gt; &lt;table routing-instance-name&gt; &lt;vpn vpn&gt;</pre>   |
| <b>Release Information</b>                           | Command introduced before Junos OS Release 7.4.<br>Option <b>bridge-domain</b> introduced in Junos OS Release 7.5<br>Option <b>learning-vlan-id</b> introduced in Junos OS Release 8.4  |

Options **all** and **vlan** introduced in Junos OS Release 9.6.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



**NOTE:** The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

**Options** **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

**detail | extensive | summary**—(Optional) Display the specified level of output.

**all**—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

**bridge-domain (all | bridge-domain-name)**—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

**ccc interface-name**—(Optional) Display route entries for the specified circuit cross-connect interface.

**destination destination-prefix**—(Optional) Destination prefix.

**family family**—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

**interface-name interface-name**—(Optional) Display routing table entries for the specified interface.

**label name**—(Optional) Display route entries for the specified label.

**lcc number**—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**learning-vlan-id** *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

**matching** *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

**multicast**—(Optional) Display routing table entries for multicast routes.

**table** (**default** | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

**vlan** (**all** | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

**vpn** *vpn*—(Optional) Display routing table entries for a specified VPN.

**Required Privilege Level**

view

**List of Sample Output**

[show route forwarding-table on page 3149](#)  
[show route forwarding-table detail on page 3150](#)  
[show route forwarding-table destination extensive \(Weights and Balances\) on page 3150](#)  
[show route forwarding-table extensive on page 3151](#)  
[show route forwarding-table extensive \(RPF\) on page 3152](#)  
[show route forwarding-table family mpls on page 3153](#)  
[show route forwarding-table family vpls on page 3153](#)  
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 3153](#)  
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 3154](#)  
[show route forwarding-table family vpls extensive on page 3154](#)  
[show route forwarding-table table default on page 3155](#)  
[show route forwarding-table table logical-system-name/routing-instance-name on page 3156](#)

[show route forwarding-table vpn on page 3157](#)

**Output Fields** [Table 288 on page 3147](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 288: show route forwarding-table Output Fields**

| Field Name              | Field Description  | Level of Output         |
|-------------------------|--|-------------------------|
| Logical system          | Name of the logical system. This field is displayed if you specify the <b>table logical-system-name/routing-instance-name</b> option on a device that is configured for and supports logical systems.  | All levels              |
| Routing table           | Name of the routing table (for example, inet, inet6, mpls).  | All levels              |
| Address family          | Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).   | All levels              |
| Destination             | Destination of the route.  | <b>detail extensive</b> |
| Route Type (Type)       | How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li><b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li><b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li><b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li><b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li><b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li><b>ignore (ignr)</b>—Ignore this route.</li> <li><b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li><b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li><b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul> | All levels              |
| Route Reference (RtRef) | Number of routes to reference.   | <b>detail extensive</b> |
| Flags                   | Route type flags: <ul style="list-style-type: none"> <li><b>none</b>—No flags are enabled.</li> <li><b>accounting</b>—Route has accounting enabled.</li> <li><b>cached</b>—Cache route.</li> <li><b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li><b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li><b>rt nh decoupled</b>—Route has been decoupled from the next hop to the destination.</li> <li><b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li><b>static</b>—Static route.</li> </ul>   | <b>extensive</b>        |
| Next hop                | IP address of the next hop to the destination.   | <b>detail extensive</b> |

Table 288: show route forwarding-table Output Fields (*continued*)

| Field Name                 | Field Description  | Level of Output              |
|----------------------------|--|------------------------------|
| Next hop Type (Type)       | <p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>discard (dscd)</b>—Discard.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcrst)</b>—Regular multicast next hop.</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b>—Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul> | <b>detail extensive</b>      |
| Index                      | Software index of the next hop that is used to route the traffic for a given prefix.   | <b>detail extensive none</b> |
| Route interface-index      | Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.   | <b>extensive</b>             |
| Reference (NhRef)          | Number of routes that refer to this next hop.  | <b>detail extensive none</b> |
| Next-hop interface (Netif) | Interface used to reach the next hop.  | <b>detail extensive none</b> |
| Weight                     | Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the <b>Balance</b> field description).  | <b>extensive</b>             |
| Balance                    | Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.   | <b>extensive</b>             |
| RPF interface              | List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when <b>rpf-check</b> is configured on the interface.  | <b>extensive</b>             |

## Sample Output

### show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  46   4
0.0.0.0/32       perm  0                               dscd  44   1
1.1.1.0/24       ifdn  0                               rslv  608  1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0           recv  606  1 ge-2/0/1.0
1.1.1.1/32       user  0                               rjct  46   4
1.1.1.1/32       intf  0 1.1.1.1           locl  607  2
1.1.1.1/32       iddn  0 1.1.1.1           locl  607  2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff bcst  605  1 ge-2/0/1.0
10.0.0.0/24      intf  0                               rslv  616  1 ge-2/0/0.0
10.0.0.0/32      dest  0 10.0.0.0          recv  614  1 ge-2/0/0.0
10.0.0.1/32      intf  0 10.0.0.1          locl  615  2
10.0.0.1/32      dest  0 10.0.0.1          locl  615  2
10.0.0.255/32    dest  0 10.0.0.255        bcst  613  1 ge-2/0/0.0
10.1.1.0/24      ifdn  0                               rslv  612  1 ge-2/0/1.0
10.1.1.0/32      iddn  0 10.1.1.0          recv  610  1 ge-2/0/1.0
10.1.1.1/32      user  0                               rjct  46   4
10.1.1.1/32      intf  0 10.1.1.1          locl  611  2
10.1.1.1/32      iddn  0 10.1.1.1          locl  611  2
10.1.1.255/32    iddn  0 ff:ff:ff:ff:ff:ff bcst  609  1 ge-2/0/1.0
10.206.0.0/16    user  0 10.209.63.254      ucst  419  20 fxp0.0
10.209.0.0/16    user  1 0:12:1e:ca:98:0    ucst  419  20 fxp0.0
10.209.0.0/18    intf  0                               rslv  418  1 fxp0.0
10.209.0.0/32    dest  0 10.209.0.0        recv  416  1 fxp0.0
10.209.2.131/32  intf  0 10.209.2.131      locl  417  2
10.209.2.131/32  dest  0 10.209.2.131      locl  417  2
10.209.17.55/32  dest  0 0:30:48:5b:78:d2   ucst  435  1 fxp0.0
10.209.63.42/32  dest  0 0:23:7d:58:92:ca   ucst  434  1 fxp0.0
10.209.63.254/32 dest  0 0:12:1e:ca:98:0    ucst  419  20 fxp0.0
10.209.63.255/32 dest  0 10.209.63.255     bcst  415  1 fxp0.0
10.227.0.0/16    user  0 10.209.63.254      ucst  419  20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  27   1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                               locl  28   1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct   6   1
ff00::/8         perm  0                               mdsc   4   1
ff02::1/128      perm  0 ff02::1           mcst   3   1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  16   1
100004(top)fe-0/0/1.0

```

## show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:90:69:8e:b1:1b ucst  132   4 fxp0.0
default          perm   0                               rjct   14    1
10.1.1.0/24      intf   0 ff.3.0.21         ucst  322   1 so-5/3/0.0
10.1.1.0/32      dest   0 10.1.1.0          recv  324   1 so-5/3/0.0
10.1.1.1/32      intf   0 10.1.1.1          locl  321   1
10.1.1.255/32    dest   0 10.1.1.255        bcst  323   1 so-5/3/0.0
10.21.21.0/24    intf   0 ff.3.0.21         ucst  326   1 so-5/3/0.0
10.21.21.0/32    dest   0 10.21.21.0        recv  328   1 so-5/3/0.0
10.21.21.1/32    intf   0 10.21.21.1        locl  325   1
10.21.21.255/32  dest   0 10.21.21.255      bcst  327   1 so-5/3/0.0
127.0.0.1/32     intf   0 127.0.0.1         locl  320   1
172.17.28.19/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0
172.17.28.44/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   46    1
10.0.0.0/8       intf   0                               rslv  136   1 fxp1.0
10.0.0.0/32      dest   0 10.0.0.0          recv  134   1 fxp1.0
10.0.0.4/32      intf   0 10.0.0.4          locl  135   2
10.0.0.4/32      dest   0 10.0.0.4          locl  135   2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   22    1
ff00::/8         perm   0                               mdsc   21    1
ff02::1/128      perm   0 ff02::1          mcst   17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  28    1

```

## show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```



```

Flags: sent to PFE
Next-hop type: unicast          Index: 262143  Reference: 1
Nexthop: 4.4.4.4
Next-hop type: unicast          Index: 335      Reference: 2
Next-hop interface: so-1/1/0.0  Weight: 22     Balance: 3
Nexthop: 145.12.1.2
Next-hop type: unicast          Index: 337     Reference: 2
Next-hop interface: so-0/1/2.0  Weight: 33     Balance: 33

```

### show route forwarding-table extensive

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2                      Route interface-index: 0
Flags: sent to PFE
Nexthop: 0:90:69:8e:b1:1b
Next-hop type: unicast                  Index: 132      Reference: 4
Next-hop interface: fxp0.0

Destination: default
Route type: permanent
Route reference: 0                      Route interface-index: 0
Flags: none
Next-hop type: reject                   Index: 14       Reference: 1

Destination: 127.0.0.1/32
Route type: interface
Route reference: 0                      Route interface-index: 0
Flags: sent to PFE
Nexthop: 127.0.0.1
Next-hop type: local                     Index: 320      Reference: 1

...

Routing table: private1__inet [Index 1]
Internet:

Destination: default
Route type: permanent
Route reference: 0                      Route interface-index: 0
Flags: sent to PFE
Next-hop type: reject                   Index: 46       Reference: 1

Destination: 10.0.0.0/8
Route type: interface
Route reference: 0                      Route interface-index: 3
Flags: sent to PFE
Next-hop type: resolve                  Index: 136      Reference: 1
Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination: default
Route type: permanent

```

```
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE
Next-hop type: reject                           Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE
Next-hop type: reject                           Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE
Next-hop type: multicast discard                Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE
Next-hop type: reject                           Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local                            Index: 75      Reference: 1

...
```

#### show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}
```

```
user@host> show route forwarding-table extensive
```

```
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0                               Route interface-index: 67
```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast          Index: 328      Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

### show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6          swap 100001 fe-1/1/0.0
800002           user  0                  Pop          vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                  user  0                  indr  351    4
                  Push 800000, Push 100002(top)

so-0/0/0.0

```

### show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48 <<<<<Remote CE
                  dymn  0                  indr  351    4
                  Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48 <<<<<Local CE
                  dymn  0                  ucst  354    2 fe-0/1/0.0

```

### show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop          Type Index   NhRef Netif
default          perm  0
lsi.1048832      intf  0
                  4.4.3.2          indr 1048574  4
                  Push 262145        621    2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0                  ucst  590    5 ge-2/3/9.0
0x30003/51       user  0                  comp  627    2
ge-2/3/9.0       intf  0                  ucst  590    5 ge-2/3/9.0
ge-3/1/3.0       intf  0                  ucst  619    4 ge-3/1/3.0
0x30002/51       user  0                  comp  600    2
0x30001/51       user  0                  comp  597    2

```

**show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)**

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats

```

| Destination          | Type | RtRef | Next hop | Type | Index   | NhRef | Netif      |
|----------------------|------|-------|----------|------|---------|-------|------------|
| default              | perm | 0     |          | dscd | 519     | 1     |            |
| 1si.1048834          | intf | 0     |          | indr | 1048574 | 4     |            |
|                      |      |       | 4.4.3.2  | Push | 262145  | 592   | 2          |
| ge-3/0/0.0           |      |       |          |      |         |       |            |
| 00:19:e2:25:d0:01/48 | user | 0     |          | ucst | 590     | 5     | ge-2/3/9.0 |
| 0x30003/51           | user | 0     |          | comp | 630     | 2     |            |
| ge-2/3/9.0           | intf | 0     |          | ucst | 590     | 5     | ge-2/3/9.0 |
| ge-3/1/3.0           | intf | 0     |          | ucst | 591     | 4     | ge-3/1/3.0 |
| 0x30002/51           | user | 0     |          | comp | 627     | 2     |            |
| 0x30001/51           | user | 0     |          | comp | 624     | 2     |            |

**show route forwarding-table family vpls extensive**

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

```

Destination: default

|                                |                              |
|--------------------------------|------------------------------|
| Route type: dynamic            | Route interface-index: 72    |
| Route reference: 0             |                              |
| Flags: sent to PFE             |                              |
| Next-hop type: flood           | Index: 289      Reference: 1 |
| Next-hop type: unicast         | Index: 291      Reference: 3 |
| Next-hop interface: fe-0/1/3.0 |                              |
| Next-hop type: unicast         | Index: 290      Reference: 3 |
| Next-hop interface: fe-0/1/2.0 |                              |

Destination: default

|                        |                              |
|------------------------|------------------------------|
| Route type: permanent  | Route interface-index: 0     |
| Route reference: 0     |                              |
| Flags: none            |                              |
| Next-hop type: discard | Index: 341      Reference: 1 |

Destination: fe-0/1/2.0

|                                |                              |
|--------------------------------|------------------------------|
| Route type: dynamic            | Route interface-index: 69    |
| Route reference: 0             |                              |
| Flags: sent to PFE             |                              |
| Next-hop type: flood           | Index: 293      Reference: 1 |
| Next-hop type: indirect        | Index: 363      Reference: 4 |
| Next-hop type: Push 800016     |                              |
| Next-hop interface: at-1/0/1.0 |                              |
| Next-hop type: indirect        | Index: 301      Reference: 5 |
| Next hop: 10.31.3.2            |                              |
| Next-hop type: Push 800000     |                              |
| Next-hop interface: fe-0/1/1.0 |                              |
| Next-hop type: unicast         | Index: 291      Reference: 3 |
| Next-hop interface: fe-0/1/3.0 |                              |

Destination: fe-0/1/3.0

|                      |                              |
|----------------------|------------------------------|
| Route type: dynamic  | Route interface-index: 70    |
| Route reference: 0   |                              |
| Flags: sent to PFE   |                              |
| Next-hop type: flood | Index: 292      Reference: 1 |

```

Next-hop type: indirect          Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0                Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640    Byte count:      675786
Route used as source
  Packet count:      6894    Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0                Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96      Byte count:      8079
Route used as source:
  Packet count:      296      Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0                Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

### show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13      ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0               rslv  688  1 fe-0/1/3.0
10.0.60.12/32    dest  0 10.0.60.12      recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22 ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14      locl  687  2
10.0.60.14/32    dest  0 10.0.60.14      locl  687  2
10.0.60.15/32    dest  0 10.0.60.15      bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13      ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21      ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0       recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0               rjct  36   2
10.0.80.2/32     intf  0 10.0.80.2       locl  675  1

```

```

10.0.80.3/32      dest    0 10.0.80.3      bcst  677    1 so-0/0/1.0
10.0.90.12/30     intf    0                rslv  684    1 fe-0/1/0.0
10.0.90.12/32     dest    0 10.0.90.12    recv  682    1 fe-0/1/0.0
10.0.90.14/32     intf    0 10.0.90.14     locl  683    2
10.0.90.14/32     dest    0 10.0.90.14     locl  683    2
10.0.90.15/32     dest    0 10.0.90.15     bcst  681    1 fe-0/1/0.0
10.5.0.0/16       user    0 192.168.187.126 ucst  324    15 fxp0.0
10.10.0.0/16      user    0 192.168.187.126 ucst  324    15 fxp0.0
10.13.10.0/23     user    0 192.168.187.126 ucst  324    15 fxp0.0
10.84.0.0/16      user    0 192.168.187.126 ucst  324    15 fxp0.0
10.150.0.0/16     user    0 192.168.187.126 ucst  324    15 fxp0.0
10.157.64.0/19    user    0 192.168.187.126 ucst  324    15 fxp0.0
10.209.0.0/16     user    0 192.168.187.126 ucst  324    15 fxp0.0

```

...

Routing table: default.iso

ISO:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 60    | 1     |       |

Routing table: default.inet6

Internet6:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 44    | 1     |       |
| ::/128      | perm | 0     |          | dscd | 42    | 1     |       |
| ff00::/8    | perm | 0     |          | mdsc | 43    | 1     |       |
| ff02::1/128 | perm | 0     | ff02::1  | mcst | 39    | 1     |       |

Routing table: default.mpls

MPLS:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | dscd | 50    | 1     |       |

### show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

| Destination        | Type | RtRef | Next hop                                       | Type | Index | NhRef | Netif      |
|--------------------|------|-------|--|------|-------|-------|------------|
| default            | perm | 0     |  | rjct | 563   | 1     |            |
| 0.0.0.0/32         | perm | 0     |  | dscd | 561   | 2     |            |
| 1.0.0.1/32         | user | 0     |  | dscd | 561   | 2     |            |
| 2.0.2.0/24         | intf | 0     |  | rslv | 771   | 1     | ge-1/2/0.3 |
| 2.0.2.0/32         | dest | 0     | 2.0.2.0  | recv | 769   | 1     | ge-1/2/0.3 |
| 2.0.2.1/32         | intf | 0     | 2.0.2.1  | locl | 770   | 2     |            |
| 2.0.2.1/32         | dest | 0     | 2.0.2.1  | locl | 770   | 2     |            |
| 2.0.2.2/32         | dest | 0     | 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0 | ucst | 789   | 1     | ge-1/2/0.3 |
| 2.0.2.255/32       | dest | 0     | 2.0.2.255                                      | bcst | 768   | 1     | ge-1/2/0.3 |
| 224.0.0.0/4        | perm | 1     |  | mdsc | 562   | 1     |            |
| 224.0.0.1/32       | perm | 0     | 224.0.0.1                                      | mcst | 558   | 1     |            |
| 255.255.255.255/32 | perm | 0     |  | bcst | 559   | 1     |            |

Logical system: R4

Routing table: vpn-red.iso

ISO:

| Destination | Type | RtRef | Next hop | Type | Index | NhRef | Netif |
|-------------|------|-------|----------|------|-------|-------|-------|
| default     | perm | 0     |          | rjct | 608   | 1     |       |

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0              rjct  708   1
::/128           perm  0              dscd  706   1
ff00::/8         perm  0              mdsc  707   1
ff02::1/128      perm  0 ff02::1        mcst  704   1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0              dscd  638

```

### show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination      Type RtRef Nexthop      Type Index NhRef Netif
default          perm  0              rjct   4    4
10.39.10.20/30   intf  0 ff.3.0.21         ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21       locl   36    1
10.255.14.172/32 user  0              ucst   69    2
so-0/0/0.0
10.255.14.175/32 user  0              indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2              mdsc   5    3
224.0.0.1/32     perm  0 224.0.0.1         mcst   1    8
224.0.0.5/32     user  1 224.0.0.5         mcst   1    8
255.255.255.255/32 perm  0              bcst   2    3

```

## show route inactive-path

---

|                             |   |
|-----------------------------|---|
| List of Syntax              | <a href="#">Syntax on page 3158</a><br><a href="#">Syntax (EX Series Switches) on page 3158</a>   |
| Syntax                      | <code>show route inactive-path</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches) | <code>show route inactive-path</code><br><code>&lt;brief   detail   extensive   terse&gt;</code>  |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| Description                 | Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.   |
| Options                     | <b>none</b> —Display all inactive routes.<br><br><b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level    | view  |
| List of Sample Output       | <a href="#">show route inactive-path on page 3158</a><br><a href="#">show route inactive-path detail on page 3159</a><br><a href="#">show route inactive-path extensive on page 3160</a><br><a href="#">show route inactive-path terse on page 3160</a>   |
| Output Fields               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```



```

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                    AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

#### show route inactive-path detail

```

user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
         Next-hop reference count: 1
         Next hop: via so-0/3/0.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Local AS: 1
         Age: 3:58:24   Metric: 1
         Area: 0.0.0.0
         Task: OSPF
         AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
  Direct Preference: 0
         Next hop type: Interface
         Next-hop reference count: 1
         Next hop: via fxp1.0, selected
         State: <NotBest Int>
         Inactive reason: No difference
         Age: 4:40:52
         Task: IF
         AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)

```

```
BGP      Preference: 170/-101
        Next-hop reference count: 6
        Source: 10.12.80.1
        Next hop: 10.12.80.1 via ge-6/3/2.0, selected
        State: <Ext>
        Inactive reason: Route Preference
        Peer AS: 100
        Age: 4:39:13
        Task: BGP_100.10.12.80.1+179
        AS path: 100 I
        Localpref: 100
        Router ID: 10.0.0.0
```

### show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 3159](#).

### show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
10.12.100.12/30    0 10      1          >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
10.0.0.0/8         D  0          >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
10.12.80.0/30      B 170     100        >10.12.80.1    100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

## show route inactive-prefix

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3161</a><br><a href="#">Syntax (EX Series Switches) on page 3161</a>  |
| <b>Syntax</b>                      | <pre>show route inactive-prefix &lt;brief   detail   extensive   terse&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>   |
| <b>Syntax (EX Series Switches)</b> | <pre>show route inactive-prefix &lt;brief   detail   extensive   terse&gt;</pre>   |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>   |
| <b>Description</b>                 | Display inactive route destinations in each routing table.   |
| <b>Options</b>                     | <p><b>none</b>—Display all inactive route destination.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route inactive-prefix on page 3161</a><br><a href="#">show route inactive-prefix detail on page 3161</a><br><a href="#">show route inactive-prefix extensive on page 3162</a><br><a href="#">show route inactive-prefix terse on page 3162</a>  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.   |

## Sample Output

### show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0
```

### show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
```

```
127.0.0.1/32 (1 entry, 0 announced)
  Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Hidden Martian Int>
    Age: 4:51
    Task: IF
    AS path: I00:04:54
      > via lo0.0
```

### `show route inactive-prefix extensive`

The output for the `show route inactive-prefix extensive` command is identical to that of the `show route inactive-path detail` command. For sample output, see [show route inactive-prefix detail on page 3161](#).

### `show route inactive-prefix terse`

```
user@host> show route inactive-prefix terse
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---------------|-------|----------|----------|----------|---------|
| 127.0.0.1/32  | D 0   |          |          | >lo0.0   |         |

## show route instance

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3163</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3163</a>   |
| <b>Syntax</b>                                     | <pre>show route instance &lt;brief   detail   summary&gt; &lt;instance-name&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;operational&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show route instance &lt;brief   detail   summary&gt; &lt;instance-name&gt; &lt;operational&gt;</pre>  |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>                                | Display routing instance information.  |
| <b>Options</b>                                    | <p><b>none</b>—(Same as <b>brief</b>) Display standard information about all routing instances.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>. (These options are not available with the <b>operational</b> keyword.)</p> <p><b>instance-name</b>—(Optional) Display information for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show route instance cust1</b> command).</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>operational</b>—(Optional) Display operational routing instances.</p> |
| <b>Required Privilege Level</b>                   | view   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling</a></li> <li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4144</a></li> </ul>  |
| <b>List of Sample Output</b>                      | <a href="#">show route instance on page 3165</a><br><a href="#">show route instance detail (Graceful Restart Complete) on page 3165</a><br><a href="#">show route instance detail (Graceful Restart Incomplete) on page 3167</a><br><a href="#">show route instance detail (VPLS Routing Instance) on page 3168</a><br><a href="#">show route instance operational on page 3169</a><br><a href="#">show route instance summary on page 3169</a>  |

**Output Fields** Table 90 on page 1499 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

**Table 289: show route instance Output Fields**

| Field Name                       | Field Description  | Level of Output                  |
|----------------------------------|--|----------------------------------|
| Instance or <i>instance-name</i> | Name of the routing instance.  | All levels                       |
| Operational Routing Instances    | ( <b>operational</b> keyword only) Names of all operational routing instances.   | —                                |
| Type                             | Type of routing instance: <b>forwarding</b> , <b>l2vpn</b> , <b>no-forwarding</b> , <b>vpls</b> , <b>virtual-router</b> , or <b>vrf</b> .  | All levels                       |
| State                            | State of the routing instance: <b>active</b> or <b>inactive</b> .  | <b>brief detail</b> none         |
| Interfaces                       | Name of interfaces belonging to this routing instance.   | <b>brief detail</b> none         |
| Restart State                    | Status of graceful restart for this instance: <b>Pending</b> or <b>Complete</b> .  | <b>detail</b>                    |
| Path selection timeout           | Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is <b>300</b> .   | <b>detail</b>                    |
| Tables                           | Tables (and number of routes) associated with this routing instance.   | <b>brief detail</b> none         |
| Route-distinguisher              | Unique route distinguisher associated with this routing instance.  | <b>detail</b>                    |
| Vrf-import                       | VPN routing and forwarding instance import policy name.  | <b>detail</b>                    |
| Vrf-export                       | VPN routing and forwarding instance export policy name.  | <b>detail</b>                    |
| Vrf-import-target                | VPN routing and forwarding instance import target community name.  | <b>detail</b>                    |
| Vrf-export-target                | VPN routing and forwarding instance export target community name.  | <b>detail</b>                    |
| Fast-reroute-priority            | Fast reroute priority setting for a VPLS routing instance: <b>high</b> , <b>medium</b> , or <b>low</b> . The default is <b>low</b> .   | <b>detail</b>                    |
| Restart State                    | Restart state: <ul style="list-style-type: none"> <li><b>Pending;protocol-name</b>—List of protocols that have not yet completed graceful restart for this routing table.</li> <li><b>Complete</b>—All protocols have restarted for this routing table.</li> </ul> | <b>detail</b>                    |
| Primary rib                      | Primary table for this routing instance.   | <b>brief</b> none <b>summary</b> |
| Active/holddown/hidden           | Number of active, hold-down, and hidden routes.  | All levels                       |

## Sample Output

### show route instance

```

user@host> show route instance
Instance              Type
Primary RIB
master                forwarding
inet.0                16/0/1
iso.0                 1/0/0
mpls.0                0/0/0
inet6.0               2/0/0
l2circuit.0          0/0/0
__juniper_private1__ forwarding
__juniper_private1__.inet.0 12/0/0
__juniper_private1__.inet6.0 1/0/0

```

### show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0              : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3              : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0               : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0              : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0         : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0             : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0         : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf           State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0    : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf           State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:

```

```
BGP-L.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
BGP-L.mpls.0          : 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
L2VPN:
Router ID: 0.0.0.0
Type: l2vpn           State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
LDP:
Router ID: 10.69.105.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
LDP.inet.0            : 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
OSPF:
Router ID: 10.69.101.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
OSPF.inet.0           : 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP:
Router ID: 10.69.102.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
RIP.inet.0            : 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC:
Router ID: 10.69.100.1
Type: vrf             State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
```



```
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
  Restart Complete
```

### show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding          State: Active
  Restart State: Pending    Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf                State: Active
    Restart State: Pending    Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0       : 6 routes (5 active, 0 holddown, 0 hidden)
      Restart Pending: VPN
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf                State: Active
    Restart State: Pending    Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0          : 6 routes (5 active, 0 holddown, 0 hidden)
      Restart Pending: VPN
      BGP-L.mpls.0          : 2 routes (2 active, 0 holddown, 0 hidden)
      Restart Pending: VPN
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn              State: Active
    Restart State: Pending    Path selection timeout: 300
    Interfaces:
      t3-0/0/0.512
    Route-distinguisher: 10.255.14.176:512
    Vrf-import: [ L2VPN-import ]
```

```
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0      : 2 routes (2 active, 0 holddown, 0 hidden)
  Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf          State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0       : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf          State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0      : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf          State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0       : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf          State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0    : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

#### show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls          State: Active
  Interfaces:
    lsi.1048833
```

```

1si.1048832
fe-0/1/0.513
Route-distinguisher: 10.255.37.65:1
Vrf-import: [ __vrf-import-test-vpls-internal__ ]
Vrf-export: [ __vrf-export-test-vpls-internal__ ]
Vrf-import-target: [ target:300:1 ]
Vrf-export-target: [ target:300:1 ]
Fast-reroute-priority: high
Tables:
  test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)

```

### show route instance operational

```

user@host> show route instance operational
Operational Routing Instances:

master
default

```

### show route instance summary

```

user@host> show route instance summary

```

| Instance | Type       | Primary rib      | Active/holddown/hidden |
|----------|------------|------------------|------------------------|
| master   | forwarding | inet.0           | 15/0/1                 |
|          |            | iso.0            | 1/0/0                  |
|          |            | mpls.0           | 35/0/0                 |
|          |            | l3vpn.0          | 0/0/0                  |
|          |            | inet6.0          | 2/0/0                  |
|          |            | l2vpn.0          | 0/0/0                  |
|          |            | l2circuit.0      | 0/0/0                  |
| BGP-INET | vrf        | BGP-INET.inet.0  | 5/0/0                  |
|          |            | BGP-INET.iso.0   | 0/0/0                  |
|          |            | BGP-INET.inet6.0 | 0/0/0                  |
| BGP-L    | vrf        | BGP-L.inet.0     | 5/0/0                  |
|          |            | BGP-L.iso.0      | 0/0/0                  |
|          |            | BGP-L.mpls.0     | 4/0/0                  |
|          |            | BGP-L.inet6.0    | 0/0/0                  |
| L2VPN    | l2vpn      | L2VPN.inet.0     | 0/0/0                  |
|          |            | L2VPN.iso.0      | 0/0/0                  |
|          |            | L2VPN.inet6.0    | 0/0/0                  |
|          |            | L2VPN.l2vpn.0    | 2/0/0                  |
| LDP      | vrf        | LDP.inet.0       | 4/0/0                  |
|          |            | LDP.iso.0        | 0/0/0                  |
|          |            | LDP.mpls.0       | 0/0/0                  |
|          |            | LDP.inet6.0      | 0/0/0                  |
|          |            | LDP.l2circuit.0  | 0/0/0                  |
| OSPF     | vrf        | OSPF.inet.0      | 7/0/0                  |
|          |            | OSPF.iso.0       | 0/0/0                  |
|          |            | OSPF.inet6.0     | 0/0/0                  |
| RIP      | vrf        | RIP.inet.0       | 6/0/0                  |
|          |            | RIP.iso.0        | 0/0/0                  |
|          |            | RIP.inet6.0      | 0/0/0                  |
| STATIC   | vrf        | STATIC.inet.0    | 4/0/0                  |

|                |       |
|----------------|-------|
| STATIC.iso.0   | 0/0/0 |
| STATIC.inet6.0 | 0/0/0 |

## show route label

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3171</a><br><a href="#">Syntax (EX Series Switches) on page 3171</a>  |
| <b>Syntax</b>                      | show route label <i>label</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches)</b> | show route label <i>label</i><br><brief   detail   extensive   terse>  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.  |
| <b>Description</b>                 | Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.  |
| <b>Options</b>                     | <p><i>label</i>—Value of the MPLS label.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>   |
| <b>Required Privilege Level</b>    | view   |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li><a href="#">Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs</a></li> </ul>   |
| <b>List of Sample Output</b>       | <a href="#">show route label terse on page 3171</a><br><a href="#">show route label on page 3172</a><br><a href="#">show route label detail on page 3172</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3172</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3173</a><br><a href="#">show route label extensive on page 3173</a> |
| <b>Output Fields</b>               | For information about output fields, see the output field table for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route label terse

```

user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

```

| A Destination | P Prf | Metric 1 | Metric 2 | Next hop    | AS path |
|---------------|-------|----------|----------|-------------|---------|
| * 100016      | V 170 |          |          | >10.12.80.1 |         |

### show route label

user@host> show route label 100016

```

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016                *[VPN/170] 03:25:41
                    > to 10.12.80.1 via ge-6/3/2.0, Pop

```

### show route label detail

user@host> show route label 100016 detail

```

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
    *VPN      Preference: 170
              Next-hop reference count: 2
              Source: 10.12.80.1
              Next hop: 10.12.80.1 via ge-6/3/2.0, selected
              Label operation: Pop
              State: <Active Int Ext>
              Local AS: 1
              Age: 3:23:31
              Task: BGP.0.0.0.0+179
              Announcement bits (1): 0-KRT
              AS path: 100 I
              Ref Cnt: 2

```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

user@host> show route label 299872 detail

```

mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Next-hop reference count: 3
              Address: 0x9097d90
              Next hop: via vt-0/1/0.1
              Next-hop index: 661
              Label operation: Pop
              Address: 0x9172130
              Next hop: via so-0/0/3.0
              Next-hop index: 654
              Label operation: Swap 299872
              State: **Active Int>
              Local AS: 1001
              Age: 8:20      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

**show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)**

```

user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
          Next hop type: Flood
          Address: 0x2735208
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 1397
          Address: 0x2735d2c
          Next-hop reference count: 3
          Next hop: 1.3.8.2 via ge-1/2/22.0
          Label operation: Pop
          Load balance label: None;
          Next hop type: Router, Next hop index: 1395
          Address: 0x2736290
          Next-hop reference count: 3
          Next hop: 1.3.4.2 via ge-1/2/18.0
          Label operation: Pop
          Load balance label: None;
          State: <Active Int AckRequest MulticastRPF>
          Local AS: 10
          Age: 54:05      Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
          Primary Upstream : 1.1.1.3:0--1.1.1.2:0
          RPF Nexthops :
              ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
              ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
          Backup Upstream : 1.1.1.3:0--1.1.1.6:0
          RPF Nexthops :
              ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
              ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

**show route label extensive**

The output for the **show route label extensive** command is identical to that of the **show route label detail** command. For sample output, see [show route label detail on page 3172](#).

## show route label-switched-path

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3174</a><br><a href="#">Syntax (EX Series Switches) on page 3174</a>   |
| <b>Syntax</b>                      | <b>show route label-switched-path</b> <i>path-name</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches)</b> | <b>show route label-switched-path</b> <i>path-name</i><br><brief   detail   extensive   terse>  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.   |
| <b>Description</b>                 | Display the routes used in an MPLS label-switched path (LSP).   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><i>path-name</i> —LSP tunnel name.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route label-switched-path on page 3174</a>   |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                      |

## Sample Output

### show route label-switched-path

```

user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32          * [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          * [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32          * [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path abc
> to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

```



```
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                   > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

## show route martians

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3176</a><br><a href="#">Syntax (EX Series Switches) on page 3176</a>   |
| <b>Syntax</b>                      | <pre>show route martians &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;table <i>routing-table-name</i>&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show route martians &lt;table <i>routing-table-name</i>&gt;</pre>  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>                 | Display the martian (invalid and ignored) entries associated with each routing table.   |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about route martians for all routing tables.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>table <i>routing-table-name</i></b>—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, <b>inet.0</b> and <b>inet6.0</b> are both displayed when you run the <b>show route martians table inet</b> command).</p> |
| <b>Required Privilege Level</b>    | view  |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li><a href="#">Example: Configuring Martian Addresses</a></li> </ul>  |
| <b>List of Sample Output</b>       | <a href="#">show route martians on page 3177</a>  |
| <b>Output Fields</b>               | <p><a href="#">Table 290 on page 3176</a> lists the output fields for the <b>show route martians</b> command. Output fields are listed in the approximate order in which they appear</p>  |

**Table 290: show route martians Output Fields**

| Field Name                | Field Description   |
|---------------------------|---|
| <i>table-name</i>         | Name of the route table in which the route martians reside. |
| <i>destination-prefix</i> | Route destination.  |
| <i>match value</i>        | Route match parameter.                                      |
| <i>status</i>             | Status of the route: <b>allowed</b> or <b>disallowed</b> .  |

## Sample Output

### show route martians

```

user@host> show route martians

inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

...

```

## show route next-hop

---

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3178</a><br><a href="#">Syntax (EX Series Switches) on page 3178</a>  |
| <b>Syntax</b>                      | <code>show route next-hop <i>next-hop</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| <b>Syntax (EX Series Switches)</b> | <code>show route next-hop <i>next-hop</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display the entries in the routing table that are being sent to the specified next-hop address.  |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><b><i>next-hop</i></b> —Next-hop address. |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route next-hop on page 3178</a><br><a href="#">show route next-hop detail on page 3179</a><br><a href="#">show route next-hop extensive on page 3181</a><br><a href="#">show route next-hop terse on page 3182</a>  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.                                     |

## Sample Output

### show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16     *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12     *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16    *[Static/5] 06:26:25
```

```

> to 192.168.71.254 via fxp0.0
192.168.102.0/23  *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.0/24  *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
> to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

#### show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2

```

```
AS path: I

192.168.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

## show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

### show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.10.0.0/16     S 5                >192.168.71.254
* 10.209.0.0/16    S 5                >192.168.71.254
* 172.16.0.0/12    S 5                >192.168.71.254

```



```
* 192.168.0.0/16      S   5                >192.168.71.254
* 192.168.102.0/23   S   5                >192.168.71.254
* 207.17.136.0/24    S   5                >192.168.71.254
* 207.17.136.192/32  S   5                >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

## show route no-community

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3184</a><br><a href="#">Syntax (EX Series Switches) on page 3184</a>  |
| <b>Syntax</b>                      | show route no-community<br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches)</b> | show route no-community<br><brief   detail   extensive   terse>  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>                 | Display the route entries in each routing table that are not associated with any community.  |
| <b>Options</b>                     | <p><b>none</b>—(Same as <b>brief</b>) Display the route entries in each routing table that are not associated with any community.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route no-community on page 3184</a><br><a href="#">show route no-community detail on page 3185</a><br><a href="#">show route no-community extensive on page 3185</a><br><a href="#">show route no-community terse on page 3186</a>  |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.   |

## Sample Output

### show route no-community

```

user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
> to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
> via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
> to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2

```

```

> to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32 * [OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0
> via so-0/3/2.0
10.255.71.241/32 * [OSPF/10] 00:05:14, metric 1
> via so-0/1/2.0
10.255.71.242/32 * [OSPF/10] 00:05:19, metric 1
> via so-0/3/2.0
12.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/3/2.0
14.1.1.0/24      * [OSPF/10] 00:00:08, metric 3
> to 35.1.1.2 via ge-3/1/0.0
                  via so-0/1/2.0
                  via so-0/3/2.0
16.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

### show route no-community detail

```

user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

### show route no-community extensive

```

user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

```

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

### show route no-community terse

```
user@host> show route no-community terse
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A   | Destination      | P | Prf | Metric 1 | Metric 2 | Next hop        | AS path |
|-----|------------------|---|-----|----------|----------|-----------------|---------|
| *   | 10.10.0.0/16     | S | 5   |          |          | >192.168.71.254 |         |
| *   | 10.209.0.0/16    | S | 5   |          |          | >192.168.71.254 |         |
| *   | 10.255.71.52/32  | D | 0   |          |          | >100.0          |         |
| *   | 10.255.71.63/32  | O | 10  | 1        |          | >35.1.1.2       |         |
| *   | 10.255.71.64/32  | O | 10  | 2        |          | >35.1.1.2       |         |
| *   | 10.255.71.240/32 | O | 10  | 2        |          | so-0/1/2.0      |         |
|     |                  |   |     |          |          | >so-0/3/2.0     |         |
| *   | 10.255.71.241/32 | O | 10  | 1        |          | >so-0/1/2.0     |         |
| *   | 10.255.71.242/32 | O | 10  | 1        |          | >so-0/3/2.0     |         |
| *   | 12.1.1.0/24      | O | 10  | 2        |          | >so-0/3/2.0     |         |
| *   | 14.1.1.0/24      | O | 10  | 3        |          | >35.1.1.2       |         |
|     |                  |   |     |          |          | so-0/1/2.0      |         |
|     |                  |   |     |          |          | so-0/3/2.0      |         |
| *   | 16.1.1.0/24      | O | 10  | 2        |          | >so-0/1/2.0     |         |
| ... |                  |   |     |          |          |                 |         |

## show route protocol

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3187</a><br><a href="#">Syntax (EX Series Switches) on page 3187</a>  |
| <b>Syntax</b>                      | <pre>show route protocol <i>protocol</i> &lt;brief   detail   extensive   terse&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show route protocol <i>protocol</i> &lt;brief   detail   extensive   terse&gt;</pre>  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2.</p> <p><b>ospf2</b> and <b>ospf3</b> options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p><b>flow</b> option introduced in Junos OS Release 10.0.</p> <p><b>flow</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p>   |
| <b>Description</b>                 | Display the route entries in the routing table that were learned from a particular protocol.   |
| <b>Options</b>                     | <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>protocol</i></b>—Protocol from which the route was learned:</p> <ul style="list-style-type: none"> <li>• <b>access</b>—Access route for use by DHCP application</li> <li>• <b>access-internal</b>—Access-internal route for use by DHCP application</li> <li>• <b>aggregate</b>—Locally generated aggregate route</li> <li>• <b>arp</b>—Route learned through the Address Resolution Protocol</li> <li>• <b>atmvpn</b>—Asynchronous Transfer Mode virtual private network</li> <li>• <b>bgp</b>—Border Gateway Protocol</li> <li>• <b>ccc</b>—Circuit cross-connect</li> <li>• <b>direct</b>—Directly connected route</li> <li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li> <li>• <b>esis</b>—End System-to-Intermediate System</li> <li>• <b>flow</b>—Locally defined flow-specification route</li> <li>• <b>frr</b>—Precomputed protection route or backup route used when a link goes down</li> <li>• <b>isis</b>—Intermediate System-to-Intermediate System</li> <li>• <b>ldp</b>—Label Distribution Protocol</li> <li>• <b>l2circuit</b>—Layer 2 circuit</li> </ul> |

- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



**NOTE:** EX Series switches run a subset of these protocols. See the switch CLI for details.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423</a></li></ul>  |
| <b>List of Sample Output</b>    | <ul style="list-style-type: none"><li>• <a href="#">show route protocol access on page 3189</a></li><li>• <a href="#">show route protocol access-internal extensive on page 3189</a></li><li>• <a href="#">show route protocol arp on page 3189</a></li><li>• <a href="#">show route protocol bgp on page 3190</a></li><li>• <a href="#">show route protocol bgp detail on page 3190</a></li><li>• <a href="#">show route protocol bgp extensive on page 3190</a></li><li>• <a href="#">show route protocol bgp terse on page 3191</a></li><li>• <a href="#">show route protocol direct on page 3191</a></li><li>• <a href="#">show route protocol frr on page 3192</a></li><li>• <a href="#">show route protocol l2circuit detail on page 3192</a></li><li>• <a href="#">show route protocol l2vpn extensive on page 3193</a></li><li>• <a href="#">show route protocol ldp on page 3194</a></li><li>• <a href="#">show route protocol ldp extensive on page 3194</a></li><li>• <a href="#">show route protocol ospf (Layer 3 VPN) on page 3195</a></li><li>• <a href="#">show route protocol ospf detail on page 3196</a></li></ul> |

[show route protocol rip on page 3196](#)  
[show route protocol rip detail on page 3196](#)  
[show route protocol ripng table inet6 on page 3197](#)  
[show route protocol static detail on page 3197](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### show route protocol access

```

user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0

```

### show route protocol access-internal extensive

```

user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I

```

### show route protocol arp

```

user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable

```

```

20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.11/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.12/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
20.20.1.13/32     [ARP/4294967293] 00:04:33, from 20.20.1.1
                  Unusable
...

```

### show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21    *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                  AS path: 10458 14203 2914 4788 4788 I
                  > to 192.168.167.254 via fxp0.0

```

### show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24    (1 entry, 1 announced)
   *BGP           Preference: 170/-101
                   Next hop type: Indirect
                   Next-hop reference count: 1006436
                   Source: 192.168.69.71
                   Next hop type: Router, Next hop index: 324
                   Next hop: 192.168.167.254 via fxp0.0, selected
                   Protocol next hop: 192.168.69.71
                   Indirect next hop: 8e166c0 342
                   State: <Active Ext>
                   Local AS: 69 Peer AS: 10458
                   Age: 6d 10:42:42 Metric2: 0
                   Task: BGP_10458.192.168.69.71+179
                   Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1

   AS path: 10458 14203 2914 4788 4788 I
   Communities: 2914:410 2914:2403 2914:3400
   Accepted
   Localpref: 100
   Router ID: 207.17.136.192

```

### show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1

```



```

*BGP      Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 1006502
          Source: 192.168.69.71
          Next hop type: Router, Next hop index: 324
          Next hop: 192.168.167.254 via fxp0.0, selected
          Protocol next hop: 192.168.69.71
          Indirect next hop: 8e166c0 342
          State: <Active Ext>
          Local AS: 69 Peer AS: 10458
          Age: 6d 10:44:45 Metric2: 0
          Task: BGP_10458.192.168.69.71+179
          Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
          AS path: 10458 14203 2914 4788 4788 I
          Communities: 2914:410 2914:2403 2914:3400
          Accepted
          Localpref: 100
          Router ID: 207.17.136.192
          Indirect next hops: 1
            Protocol next hop: 192.168.69.71
            Indirect next hop: 8e166c0 342
            Indirect path forwarding next hops: 1
              Next hop type: Router
              Next hop: 192.168.167.254 via fxp0.0
            192.168.0.0/16 Originating RIB: inet.0
              Node path count: 1
              Forwarding nexthops: 1
                Nexthop: 192.168.167.254 via fxp0.0

```

### show route protocol bgp terse

```

user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
192.168.64.0/21    B 170      100          >100.1.3.2    10023 21 I

```

### show route protocol direct

```

user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24          *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24       *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
> via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21

```

```

> via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0

```

### show route protocol frr

```

user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

### show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop          Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

```

```

ge-2/0/0.0 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000, Push 100000(top)[0] Offset: -4
    Protocol next hop: 10.245.255.63
    Push 100000 Offset: -4
    Indirect next hop: 86af0c0 298
    State: <Active Int>
    Local AS: 99
    Age: 9:52
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

### show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected

```

```
Label operation: Push 800000 Offset: -4
Protocol next hop: 10.255.14.220
Push 800000 Offset: -4
  Indirect next hop: 85142a0 288
State: <Active Int>
Local AS: 69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0
```

### show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000
```

### show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
```

```

Age: 1d 23:03:58      Metric: 1
Task: LDP
Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Pop
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
    *LDP      Preference: 9
              Next-hop reference count: 2
              Next hop: via t1-4/0/0.0, selected
              Label operation: Swap 100000
              State: <Active Int>
              Local AS: 65500
              Age: 1d 23:03:58      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 192.168.16.1/32

```

### show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      * [OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2

```

```

> via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
> via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
> via t3-3/2/0.0
224.0.0.5/32     *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
10.255.14.173/32  *[OSPF/10] 00:05:43, metric 1
> via so-0/2/2.0
224.0.0.5/32     *[OSPF/10] 20:26:20, metric 1

```

### show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
  OSPF   Preference: 10
        Nexthop: via so-0/2/2.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Age: 6:25      Metric: 1
        Area: 0.0.0.0
        Task: VPN-AB-OSPF
        AS path: I
        Communities: Route-Type:0.0.0.0:1:0

...

```

### show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  *[RIP/100] 20:24:34, metric 2
> to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32     *[RIP/100] 00:03:59, metric 1

```

### show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
  *RIP   Preference: 100
        Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
        State: <Active Int>
        Age: 20:25:02  Metric: 2
        Task: VPN-AB-RIPv2
        Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179

```

```
AS path: I
Route learned from 10.39.1.22 expires in 96 seconds
```

### show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

### show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
```

State: <Active NoReadvrt Int Ext>  
Age: 7w3d 21:24:25  
Validation State: unverified  
Task: RT  
Announcement bits (1): 0-KRT  
AS path: I



## show route range

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3199</a><br><a href="#">Syntax (EX Series Switches) on page 3199</a>   |
| <b>Syntax</b>                      | <pre>show route range &lt;brief   detail   extensive   terse&gt; &lt;destination-prefix&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show route range &lt;brief   detail   extensive   terse&gt; &lt;destination-prefix&gt;</pre>   |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>                 | Display routing table entries using a prefix range.   |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about all routing table entries using a prefix range.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>destination-prefix</b>—Destination and prefix mask for the range.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route range on page 3199</a><br><a href="#">show route range destination-prefix on page 3200</a><br><a href="#">show route range detail on page 3200</a><br><a href="#">show route range extensive on page 3201</a><br><a href="#">show route range terse on page 3202</a>   |
| <b>Output Fields</b>               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route range

```
user@host> show route range
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
```

```
10.255.71.14/32      > to 192.168.71.254 via fxp0.0
                    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:30:01
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

### show route range destination-prefix

```
user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16      *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21     *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32    *[Local/0] 00:31:14
                    Local via fxp0.0
192.168.102.0/23    *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

### show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
```

```

State: <Active Int>
Age: 30:05
Task: IF
AS path: I

172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

...

```

### show route range extensive

```

user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kerne1 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22

```

```

Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Age: 30:17
Task: RT
Announcement bits (1): 0-KRT
AS path: I

```

```
...
```

### show route range terse

```
user@host> show route range terse
```

```
inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination       | P | Prf | Metric 1 | Metric 2 | Next hop        | AS path |
|---|-------------------|---|-----|----------|----------|-----------------|---------|
| * | 10.10.0.0/16      | S | 5   |          |          | >192.168.71.254 |         |
| * | 10.209.0.0/16     | S | 5   |          |          | >192.168.71.254 |         |
| * | 10.255.71.14/32   | D | 0   |          |          | >lo0.0          |         |
| * | 172.16.0.0/12     | S | 5   |          |          | >192.168.71.254 |         |
| * | 192.168.0.0/16    | S | 5   |          |          | >192.168.71.254 |         |
| * | 192.168.64.0/21   | D | 0   |          |          | >fxp0.0         |         |
| * | 192.168.71.14/32  | L | 0   |          |          | Local           |         |
| * | 192.168.102.0/23  | S | 5   |          |          | >192.168.71.254 |         |
| * | 207.17.136.0/24   | S | 5   |          |          | >192.168.71.254 |         |
| * | 207.17.136.192/32 | S | 5   |          |          | >192.168.71.254 |         |

```
__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|-------------|---|-----|----------|----------|----------|---------|
| * | 10.0.0.0/8  | D | 0   |          |          | >fxp2.0  |         |
|   |             | D | 0   |          |          | >fxp1.0  |         |
| * | 10.0.0.4/32 | L | 0   |          |          | Local    |         |

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination   | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|---|---|-----|----------|----------|----------|---------|
|   | 47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152 |   |     |          |          |          |         |
| * |   | D | 0   |          |          | >lo0.0   |         |

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination                  | P | Prf | Metric 1 | Metric 2 | Next hop | AS path |
|---|------------------------------|---|-----|----------|----------|----------|---------|
|   | abcd::10:255:71:14/128       |   |     |          |          |          |         |
| * |                              | D | 0   |          |          | >lo0.0   |         |
|   | fe80::280:42ff:fe11:226f/128 |   |     |          |          |          |         |
| * |                              | D | 0   |          |          | >lo0.0   |         |

```
__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

| A | Destination                  | P | Prf | Metric 1 | Metric 2 | Next hop   | AS path |
|---|------------------------------|---|-----|----------|----------|------------|---------|
|   | fe80::280:42ff:fe11:226f/128 |   |     |          |          |            |         |
| * |                              | D | 0   |          |          | >lo0.16385 |         |

## show route receive-protocol

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3203</a><br><a href="#">Syntax (EX Series Switches) on page 3203</a>   |
| <b>Syntax</b>                      | show route receive-protocol <i>protocol neighbor-address</i><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches)</b> | show route receive-protocol <i>protocol neighbor-address</i><br><brief   detail   extensive   terse>  |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>                 | Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.   |
| <b>Options</b>                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>protocol neighbor-address</i></b> —Protocol transmitting the route ( <b>bgp</b> , <b>dvmrp</b> , <b>msdp</b> , <b>pim</b> , <b>rip</b> , or <b>ripng</b> ) and address of the neighboring router from which the route entry was received.  |
| <b>Additional Information</b>      | The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.   |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route receive-protocol bgp on page 3206</a><br><a href="#">show route receive-protocol bgp extensive on page 3206</a><br><a href="#">show route receive-protocol bgp table extensive on page 3206</a><br><a href="#">show route receive-protocol bgp logical-system extensive on page 3207</a><br><a href="#">show route receive-protocol bgp detail (Layer 2 VPN) on page 3208</a><br><a href="#">show route receive-protocol bgp extensive (Layer 2 VPN) on page 3208</a><br><a href="#">show route receive-protocol bgp (Layer 3 VPN) on page 3209</a><br><a href="#">show route receive-protocol bgp detail (Layer 3 VPN) on page 3209</a><br><a href="#">show route receive-protocol bgp extensive (Layer 3 VPN) on page 3210</a> |
| <b>Output Fields</b>               | <a href="#">Table 291 on page 3203</a> describes the output fields for the <b>show route receive-protocol</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 291: show route receive-protocol Output Fields

| Field Name                | Field Description                              | Level of Output |
|---------------------------|--|-----------------|
| <i>routing-table-name</i> | Name of the routing table—for example, inet.0. | All levels      |

Table 291: show route receive-protocol Output Fields (*continued*)

| Field Name                                      | Field Description  | Level of Output         |
|---|--|-------------------------|
| <i>number destinations</i>                      | Number of destinations for which there are routes in the routing table.  | All levels              |
| <i>number routes</i>                            | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b></li> <li>• <b>holddown</b> (routes that are in pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> | All levels              |
| Prefix  | Destination prefix.  | none <b>brief</b>       |
| MED   | Multiple exit discriminator value included in the route.   | none <b>brief</b>       |
| <i>destination-prefix</i><br>(entry, announced) | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.  | <b>detail extensive</b> |
| Route Distinguisher                             | 64-bit prefix added to IP subnets to make them unique.   | <b>detail extensive</b> |
| Label-Base, range                               | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.   | <b>detail extensive</b> |
| VPN Label                                       | Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.  | <b>detail extensive</b> |
| Next hop  | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.  | All levels              |
| Localpref or Lclpref                            | Local preference value included in the route.  | All levels              |

Table 291: show route receive-protocol Output Fields (*continued*)

| Field Name          | Field Description  | Level of Output  |
|---------------------|--|------------------|
| AS path             | <p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893.</li> <li>• <b>[ ]</b>—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels       |
| Cluster list        | (For route reflected output only) Cluster ID sent by the route reflector.  | detail extensive |
| Originator ID       | (For route reflected output only) Address of routing device that originally sent the route to the route reflector.   | detail extensive |
| Communities         | Community path attribute for the route. See the Output Field table in the <a href="#">show route detail</a> command for all possible values for this field.  | detail extensive |
| AIGP                | Accumulated interior gateway protocol (AIGP) BGP attribute.  | detail extensive |
| Attrset AS          | Number, local preference, and path of the AS that originated the route. These values are stored in the <b>Attrset</b> attribute at the originating routing device.   | detail extensive |
| Layer2-info: encaps | Layer 2 encapsulation (for example, VPLS).   | detail extensive |
| control flags       | Control flags: <b>none</b> or <b>Site Down</b> .   | detail extensive |
| mtu                 | Maximum transmission unit (MTU) of the Layer 2 circuit.  | detail extensive |

## Sample Output

### show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
10.22.1.0/24     10.255.245.215    0        100      I
10.22.2.0/24     10.255.245.215    0        100      I
```

### show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

### show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29
  Localpref: 100
  AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
  AS path: AS4 PA[2]: 33437 393219
  AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
  Communities: 2914:420
```



**show route receive-protocol bgp logical-system extensive**

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300096
  Nexthop: 10.0.0.9
  AS path: 13979 I
  AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300144
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300160
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
```

**show route receive-protocol bgp detail (Layer 2 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

**show route receive-protocol bgp extensive (Layer 2 VPN)**

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lc1pref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100

```

```

AS path: I
Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

### show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

### show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

```

* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

#### show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  195.1.2.0/24 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

## show route resolution

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3211</a><br><a href="#">Syntax (EX Series Switches) on page 3211</a>  |
| <b>Syntax</b>                      | <pre>show route resolution &lt;brief   detail   extensive   summary&gt; &lt;index <i>index</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>prefix</i>&gt; &lt;table <i>routing-table-name</i>&gt; &lt;unresolved&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show route resolution &lt;brief   detail   extensive   summary&gt; &lt;index <i>index</i>&gt; &lt;<i>prefix</i>&gt; &lt;table <i>routing-table-name</i>&gt; &lt;unresolved&gt;</pre>  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>   |
| <b>Description</b>                 | Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.   |
| <b>Options</b>                     | <p><b>none</b>—Display standard information about all entries in the next-hop resolution database.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>index <i>index</i></b>—(Optional) Show the index of the resolution tree.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>prefix network/destination-prefix</i></b>—(Optional) Display database entries for the specified address.</p> <p><b>table <i>routing-table-name</i></b>—(Optional) Display information about a particular routing table (for example, <b>inet.0</b>) where policy-based export is currently enabled.</p> <p><b>unresolved</b>—(Optional) Display routes that could not be resolved.</p> |
| <b>Required Privilege Level</b>    | view   |
| <b>Related Documentation</b>       | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Route Resolution on PE Routers</a></li> </ul>  |
| <b>List of Sample Output</b>       | <a href="#">show route resolution detail on page 3212</a><br><a href="#">show route resolution summary on page 3213</a><br><a href="#">show route resolution unresolved on page 3213</a>   |

**Output Fields** Table 292 on page 3212 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

**Table 292: show route resolution Output Fields**

| Field Name                         | Field Description   |
|------------------------------------|---|
| <i>routing-table-name</i>          | Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.   |
| <b>Tree index</b>                  | Tree index identifier.  |
| <b>Nodes</b>                       | Number of nodes in the tree.  |
| <b>Reference count</b>             | Number of references made to the next hop.  |
| <b>Contributing routing tables</b> | Routing tables used for next-hop resolution.  |
| <b>Originating RIB</b>             | Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of <b>inet.0</b> resolving through <b>inet.0</b> and <b>inet.3</b> , this field indicates which routing table, <b>inet.0</b> or <b>inet.3</b> , provided the best path for a particular prefix. |
| <b>Metric</b>                      | Metric associated with the forwarding next hop.   |
| <b>Node path count</b>             | Number of nodes in the path.  |
| <b>Forwarding next hops</b>        | Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.   |

## Sample Output

### show route resolution detail

```

user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0

```

```

Node path count: 1
Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
Metric: 2 Node path count: 1
Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1

```

#### show route resolution summary

```

user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3

```

#### show route resolution unresolved

```

user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
  Protocol Nexthop: 10.255.71.238 Push 800000
  Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
  Protocol Nexthop: 10.255.70.103 Push 800008
  Indirect nexthop: 0 -
Tree Index 2

```

## show route snooping

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show route snooping</code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;all&gt;</code><br><code>&lt;best address/prefix&gt;</code><br><code>&lt;exact address&gt;</code><br><code>&lt;range prefix-range&gt;</code><br><code>&lt;summary&gt;</code><br><code>&lt;table table-name&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 8.5.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Display the entries in the routing table that were learned from snooping.   |
| <b>Options</b>                  | <p><b>none</b>—Display the entries in the routing table that were learned from snooping.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b>.</p> <p><b>all</b>—(Optional) Display all entries, including hidden entries.</p> <p><b>best address/prefix</b>—(Optional) Display the longest match for the provided address and optional prefix.</p> <p><b>exact address/prefix</b>—(Optional) Display exact matches for the provided address and optional prefix.</p> <p><b>range prefix-range</b>—(Optional) Display information for the provided address range.</p> <p><b>summary</b>—(Optional) Display route snooping summary statistics.</p> <p><b>table table-name</b>—(Optional) Display information for the named table.</p> |
| <b>Required Privilege Level</b> | view  |
| <b>List of Sample Output</b>    | <a href="#">show route snooping detail on page 3214</a>   |
| <b>Output Fields</b>            | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
  *IGMP    Preference: 0
           Next hop type: MultiRecv
           Next-hop reference count: 4
           State: <Active NoReadvrt Int>
```



```

Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113

```

```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```

```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
```

```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

```
AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 9
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

## show route source-gateway

---

|                             |   |
|-----------------------------|---|
| List of Syntax              | <a href="#">Syntax on page 3222</a><br><a href="#">Syntax (EX Series Switches) on page 3222</a>   |
| Syntax                      | <code>show route source-gateway address</code><br><brief   detail   extensive   terse><br><logical-system (all   <i>logical-system-name</i> )>  |
| Syntax (EX Series Switches) | <code>show route source-gateway address</code><br><brief   detail   extensive   terse>  |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.   |
| Description                 | Display the entries in the routing table that were learned from a particular address. The <b>Source</b> field in the <code>show route detail</code> command output lists the source for each route, if known.   |
| Options                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to <b>brief</b> .<br><br><b>address</b> —IP address of the system.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level    | view  |
| List of Sample Output       | <a href="#">show route source-gateway on page 3222</a><br><a href="#">show route source-gateway detail on page 3223</a><br><a href="#">show route source-gateway extensive on page 3225</a>   |
| Output Fields               | For information about output fields, see the output field tables for the <a href="#">show route</a> command, the <a href="#">show route detail</a> command, the <a href="#">show route extensive</a> command, or the <a href="#">show route terse</a> command.  |

## Sample Output

### show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
```



```

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

```

#### show route source-gateway detail

```

user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

```

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-green-12vpn
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS,
          control flags:, mtu: 0
          Label-base: 800008, range: 8
          Localpref: 100
          Router ID: 10.255.70.103
          Primary Routing Table bgp.12vpn.0
```

red.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-red-12vpn
          AS path: I
          Communities: target:11111:2 Layer2-info: encaps:VPLS,
          control flags:Site-Down, mtu: 0
          Label-base: 800016, range: 8
          Localpref: 0
          Router ID: 10.255.70.103
          Primary Routing Table bgp.12vpn.0
```

bgp.12vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:14:00 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS, control
          flags:, mtu: 0
```

```

Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:14:00 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Secondary Tables: red.l2vpn.0

```

#### show route source-gateway extensive

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0

```

```
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-1
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-l2vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Secondary Tables: green.l2vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-1
```

```
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.12vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop: via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0
```

## show route summary

|                                    |  |
|------------------------------------|--|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3228</a><br><a href="#">Syntax (EX Series Switches) on page 3228</a>  |
| <b>Syntax</b>                      | <pre>show route summary &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;table <i>routing-table-name</i>&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | show route summary   |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>   |
| <b>Description</b>                 | <p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the <b>show route summary</b> command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the <b>show route summary</b> command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p> |
| <b>Options</b>                     | <p><b>none</b>—Display summary statistics about the entries in the routing table.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>table <i>routing-table-name</i></b>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, <b>inet.0</b> and <b>inet6.0</b> are both displayed when you run the <b>show route summary table inet</b> command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p>      |
| <b>Required Privilege Level</b>    | view   |
| <b>List of Sample Output</b>       | <a href="#">show route summary on page 3229</a><br><a href="#">show route summary table on page 3230</a><br><a href="#">show route summary table (with Route Limits Configured for the Routing Table) on page 3230</a>   |
| <b>Output Fields</b>               | <p><a href="#">Table 293 on page 3228</a> lists the output fields for the <b>show route summary</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

**Table 293: show route summary Output Fields**

| Field Name | Field Description                    |
|------------|--------------------------------------|
| Router ID  | Address of the local routing device. |

Table 293: show route summary Output Fields (*continued*)

| Field Name                | Field Description   |
|---------------------------|---|
| <i>routing-table-name</i> | Name of the routing table (for example, <b>inet.0</b> ).  |
| <b>destinations</b>       | Number of destinations for which there are routes in the routing table.   |
| <b>routes</b>             | Number of routes in the routing table: <ul style="list-style-type: none"> <li><b>active</b>—Number of routes that are active.</li> <li><b>holddown</b>—Number of routes that are in the hold-down state before being declared inactive.</li> <li><b>hidden</b>—Number of routes that are not used because of routing policy.</li> </ul>   |
| <b>Limit/Threshold</b>    | Displays the configured route limits for the routing table set with the <b>maximum-prefixes</b> and the <b>maximum-paths</b> statements. If you do not configure route limits for the routing table, the show output does not display this information. <ul style="list-style-type: none"> <li><b>destinations</b>—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message.</li> <li><b>routes</b>—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.</li> </ul> |
| <b>Direct</b>             | Routes on the directly connected network.   |
| <b>Local</b>              | Local routes.   |
| <i>protocol-name</i>      | Name of the protocol from which the route was learned. For example, <b>OSPF</b> , <b>RSVP</b> , and <b>Static</b> .   |

## Sample Output

### show route summary

```

user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
      Direct:      6 routes,      5 active
      Local:       4 routes,      4 active
      OSPF:        5 routes,      4 active
      Static:      7 routes,      7 active
      IGMP:        1 routes,      1 active
      PIM:         2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
      RSVP:        2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

```
Restart Complete
  Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
  MPLS:       3 routes,      3 active
  VPLS:       4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
  Direct:      2 routes,      2 active
  PIM:         2 routes,      2 active
  MLD:         1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         2 routes,      2 active
  L2VPN:       2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         2 routes,      2 active
  L2VPN:       1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
  BGP:         4 routes,      4 active
```

#### show route summary table

```
user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
  Direct:      6 routes,      5 active
  Local:       9 routes,      9 active
  OSPF:        3 routes,      1 active
  Static:     13 routes,     13 active
  ICMP:        1 routes,      1 active
  PIM:         2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Multicast:    1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
  Local:        1 routes,      1 active
  PIM:          2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Multicast:    1 routes,      1 active
```

#### show route summary table (with Route Limits Configured for the Routing Table)

```
user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
  Direct:      2 routes,      2 active
  Local:       1 routes,      1 active
```



|       |           |          |
|-------|-----------|----------|
| OSPF: | 4 routes, | 3 active |
| BGP:  | 4 routes, | 4 active |
| IGMP: | 1 routes, | 1 active |
| PIM:  | 2 routes, | 2 active |

## show route table

---

|                             |  |
|-----------------------------|--|
| List of Syntax              | <a href="#">Syntax on page 3232</a><br><a href="#">Syntax (EX Series Switches) on page 3232</a>  |
| Syntax                      | <code>show route table <i>routing-table-name</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>   |
| Syntax (EX Series Switches) | <code>show route table <i>routing-table-name</i></code><br><code>&lt;brief   detail   extensive   terse&gt;</code>   |
| Release Information         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.  |
| Description                 | Display the route entries in a particular routing table.   |
| Options                     | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><i>routing-table-name</i></b> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route table inet</b> command).   |
| Required Privilege Level    | view   |
| Related Documentation       | <ul style="list-style-type: none"><li>• <a href="#">show route summary on page 3228</a></li></ul>  |
| List of Sample Output       | <a href="#">show route table bgp.l2.vpn on page 3233</a><br><a href="#">show route table bgp.l3vpn.0 on page 3233</a><br><a href="#">show route table bgp.l3vpn.0 detail on page 3233</a><br><a href="#">show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 3235</a><br><a href="#">show route table inet.0 on page 3235</a><br><a href="#">show route table inet6.0 on page 3235</a><br><a href="#">show route table inet6.3 on page 3236</a><br><a href="#">show route table inetflow detail on page 3236</a><br><a href="#">show route table l2circuit.0 on page 3236</a><br><a href="#">show route table mpls on page 3237</a><br><a href="#">show route table mpls extensive on page 3237</a><br><a href="#">show route table mpls.0 on page 3237</a><br><a href="#">show route table mpls.0 detail (PTX Series) on page 3238</a><br><a href="#">show route table mpls.0 extensive (PTX Series) on page 3238</a><br><a href="#">show route table mpls.0 (RSVP Route—Transit LSP) on page 3239</a><br><a href="#">show route table vpls_1 detail on page 3239</a><br><a href="#">show route table vpn-a on page 3240</a> |

[show route table vpn-a.mdt.0 on page 3240](#)  
[show route table VPN-A detail on page 3240](#)  
[show route table VPN-AB.inet.0 on page 3241](#)  
[show route table VPN\\_blue.mvpn-inet6.0 on page 3241](#)  
[show route table vrf1.mvpn.0 extensive on page 3242](#)  
[show route table inetflow detail on page 3242](#)

**Output Fields** For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

## Sample Output

### show route table bgp.l2.vpn

```

user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

### show route table bgp.l3vpn.0

```

user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)

```

### show route table bgp.l3vpn.0 detail

```

user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Route Distinguisher: 10.255.245.12:1
        Source: 10.255.245.12
        Next hop: 192.168.208.66 via fe-0/0/0.0, selected
        Label operation: Push 182449
        Protocol next hop: 10.255.245.12
        Push 182449
        Indirect next hop: 863a630 297
        State: <Active Int Ext>
        Local AS: 35 Peer AS: 35
        Age: 12:19 Metric2: 1
        Task: BGP_35.10.255.245.12+179

```

```
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

Communities: 2914:420 target:11111:1 origin:56:78
VPN Label: 182449
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
```

```

Age: 12:19      Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496

6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

### show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```

user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                *[RTarget/5] 00:03:14
                  Type Proxy
                    for 10.255.165.103
                    for 10.255.166.124
                  Local

```

### show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0        *[Static/5] 00:51:57
                  > to 111.222.5.254 via fxp0.0
1.0.0.1/32       *[Direct/0] 00:51:58
                  > via at-5/3/0.0
1.0.0.2/32       *[Local/0] 00:51:58
                  Local
12.12.12.21/32   *[Local/0] 00:51:57
                  Reject
13.13.13.13/32   *[Direct/0] 00:51:58
                  > via t3-5/2/1.0
13.13.13.14/32   *[Local/0] 00:51:58
                  Local
13.13.13.21/32   *[Local/0] 00:51:58
                  Local
13.13.13.22/32   *[Direct/0] 00:33:59
                  > via t3-5/2/0.0
127.0.0.1/32     [Direct/0] 00:51:58
                  > via lo0.0
111.222.5.0/24   *[Direct/0] 00:51:58
                  > via fxp0.0
111.222.5.81/32  *[Local/0] 00:51:58
                  Local

```

### show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34

```

```
>Local
```

```
fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

### show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
    *[LDP/9] 00:00:22, metric 1
    > via so-1/0/0.0
::10.255.245.196/128
    *[LDP/9] 00:00:08, metric 1
    > via so-1/0/0.0, Push 100008
```

### show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP Preference: 170/-101
        Next-hop reference count: 2
        State: <Active Ext>
        Local AS: 65002 Peer AS: 65000
        Age: 4
        Task: BGP_65000.10.12.99.5+3792
        Announcement bits (1): 0-Flow
        AS path: 65000 I
        Communities: traffic-rate:0:0
        Validation state: Accept, Originator: 10.12.99.5
        Via: 10.12.44.0/24, Active
        Localpref: 100
        Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow Preference: 5
        Next-hop reference count: 2
        State: <Active>
        Local AS: 65002
        Age: 6:30
        Task: RT Flow
        Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
        AS path: I
        Communities: 1:1
```

### show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    > via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
```

```

* [L2CKT/7] 00:50:47
> via so-0/1/2.0, Push 100049
  via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
* [LDP/9] 00:50:14
  Discard

```

### show route table mpls

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          * [MPLS/0] 00:13:55, metric 1
           Receive
1          * [MPLS/0] 00:13:55, metric 1
           Receive
2          * [MPLS/0] 00:13:55, metric 1
           Receive
1024       * [VPN/0] 00:04:18
           to table red.inet.0, Pop

```

### show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP      Preference: 9
             Next hop: via so-1/0/0.0, selected
             Pop
             State: <Active Int>
             Age: 29:50      Metric: 1
             Task: LDP
             Announcement bits (1): 0-KRT
             AS path: I
             Prefixes bound to route: 10.0.0.194/32

```

### show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          * [MPLS/0] 00:45:09, metric 1
           Receive
1          * [MPLS/0] 00:45:09, metric 1
           Receive
2          * [MPLS/0] 00:45:09, metric 1
           Receive
100000     * [L2VPN/7] 00:43:04
           > via so-0/1/0.1, Pop
100001     * [L2VPN/7] 00:43:03
           > via so-0/1/0.2, Pop      Offset: 4
100002     * [LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100002(S=0) * [LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100003     * [LDP/9] 00:43:22, metric 1
           > via so-0/1/2.0, Swap 100002

```

```

100004          via so-0/1/3.0, Swap 100002
                *[LDP/9] 00:43:16, metric 1
                via so-0/1/2.0, Swap 100049
                > via so-0/1/3.0, Swap 100049
so-0/1/0.1      *[L2VPN/7] 00:43:04
                > via so-0/1/2.0, Push 100001, Push 100049(top)
                via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2      *[L2VPN/7] 00:43:03
                via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

### show route table mpls.0 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 21 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC
    AS path: I

```

### show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0 /32 -> {composite(570)}
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>

```



```

Age: 47          Metric2: 1
Validation State: unverified
Task: Common L2 VC
Announcement bits (2): 0-KRT 2-Common L2 VC
AS path: I
Composite next hops: 1
  Protocol next hop: 10.255.255.1 Metric: 1
  Label operation: Push 299872 Offset: 252
  Label TTL action: no-prop-ttl
  Load balance label: Label 299872:Flow label PUSH;
  Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
  Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 3.0.0.1 via ge-0/0/1.0
    Session Id: 0x1
  10.255.255.1/32 Originating RIB: inet.3
    Metric: 1                      Node path count: 1
    Forwarding nexthops: 1
      Nexthop: 3.0.0.1 via ge-0/0/1.0

```

### show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 00:37:31, metric 1
           Receive
1          *[MPLS/0] 00:37:31, metric 1
           Receive
2          *[MPLS/0] 00:37:31, metric 1
           Receive
13         *[MPLS/0] 00:37:31, metric 1
           Receive
300352     *[RSVP/7/1] 00:08:00, metric 1
           > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
           > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384     *[RSVP/7/2] 00:05:20, metric 1
           > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
           > to 8.64.1.106 via ge-1/0/0.0, Pop

```

### show route table vpls\_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I

```

```
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

### show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

### show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2
```

### show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP Preference: 170/-101
    Route Distinguisher: 10.255.179.13:200
    Next hop type: Indirect
    Next-hop reference count: 5
    Source: 10.255.179.13
    Next hop type: Router, Next hop index: 732
    Next hop: 10.39.1.14 via fe-0/3/0.0, selected
    Label operation: Push 299824, Push 299824(top)
    Protocol next hop: 10.255.179.13
    Push 299824
    Indirect next hop: 8f275a0 1048574
    State: (Secondary Active Int Ext)
    Local AS: 1 Peer AS: 1
    Age: 3:41:06 Metric: 1 Metric2: 1
    Task: BGP_1.10.255.179.13+64309
    Announcement bits (2): 0-KRT 1-BGP RT Background
    AS path: I
    Communities: target:1:200 rte-type:0.0.0.0:1:0
```

```

Import Accepted
VPN Label: 299824 TTL Action: vrf-ttl-propagate
Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.13vpn.0

```

### show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

### show route table VPN\_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                  AS path: I
                  > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                  *[MVPN/70] 00:57:23, metric2 1
                  Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
                  *[PIM/105] 00:02:37
                  Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
                  *[MVPN/70] 00:02:37, metric2 1
                  Indirect

```

**show route table vrf1.mvpn.0 extensive**

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
    *MVPN    Preference: 70
              PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
    Next hop type: Indirect
    Address: 0xbb2c944
    Next-hop reference count: 360
    Protocol next hop: 10.255.50.77
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Active Int Ext>
    Age: 53:03      Metric2: 1
    Validation State: unverified
    Task: mvpn global task
    Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

    AS path: I

```

**show route table inetflow detail**

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Next-hop reference count: 2
              State: <Active Ext>
              Local AS: 65002 Peer AS: 65000
              Age: 4
              Task: BGP_65000.10.12.99.5+3792
              Announcement bits (1): 0-Flow
              AS path: 65000 I
              Communities: traffic-rate:0:0
              Validation state: Accept, Originator: 10.12.99.5
              Via: 10.12.44.0/24, Active
              Localpref: 100
              Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow     Preference: 5
              Next-hop reference count: 2
              State: <Active>
              Local AS: 65002
              Age: 6:30
              Task: RT Flow
              Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
              AS path: I
              Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    >    via ge-1/2/1.5

```

```

1.1.1.2:100:1:0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1:0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
    *[VPLS/7] 1d 03:11:02, metric2 1
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
    *[LDP/9] 1d 03:11:02
    Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
    State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
    @BGP Preference: 170/-1
    Route Distinguisher: 2:1
    Next hop type: Indirect
    Address: 0x258059e4
    Next-hop reference count: 2
    Source: 2.2.0.0
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
    Local AS: 2 Peer AS: 2
    Age: 23 Metric2: 35
    Validation State: unverified
    Task: BGP_2.2.2.0.0+34549
    AS path: I
    Communities: target:2:1
    Import Accepted
    VPN Label: 16
    Localpref: 0
    Router ID: 2.2.0.0
    Primary Routing Table bgp.13vpn.0
    Composite next hops: 1
        Protocol next hop: 2.2.0.0 Metric: 35
        Push 16
        Composite next hop: 0x25805988 - INH Session ID: 0x193c
        Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 10.1.1.1 via ge-1/1/9.0

```

```

                Session Id: 0x17d8
                2.2.0.0/32 Originating RIB: inet.3
                Metric: 35                      Node path count: 1
                Forwarding nexthops: 1
                Nexthop: 10.1.1.1 via ge-1/1/9.0
BGP Preference: 170/-1
Route Distinguisher: 2:1
Next hop type: Indirect
Address: 0x9347028
Next-hop reference count: 3
Source: 2.3.0.0
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0, selected
Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

Inactive reason: Not Best in its group - IGP metric
Local AS:      2 Peer AS:      2
Age: 3:34      Metric2: 70
Validation State: unverified
Task: BGP_2.2.3.0.0+32805
Announcement bits (2): 0-KRT 1-BGP_RT_Background
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.3.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
    Protocol next hop: 2.3.0.0 Metric: 70
    Push 16
    Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
    Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.1.4.2 via ge-1/0/0.0
        Session Id: 0x17d9
        2.3.0.0/32 Originating RIB: inet.3
        Metric: 70                      Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
Next hop type: Indirect
Address: 0x24afca30
Next-hop reference count: 1
Next hop type: Router
Next hop: 10.1.1.1 via ge-1/1/9.0, selected
Label operation: Push 707633
Label TTL action: prop-ttl
Session Id: 0x17d8
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0

```

Label operation: Push 634278  
Label TTL action: prop-ttl  
Session Id: 0x17d9  
Protocol next hop: 2.2.0.0  
Push 16  
Composite next hop: 0x25805988 - INH Session ID: 0x193c  
Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

Protocol next hop: 2.3.0.0  
Push 16  
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da  
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000

State: <ForwardingOnly Int Ext>  
Inactive reason: Forwarding use only  
Age: 23 Metric2: 35  
Validation State: unverified  
Task: RT  
AS path: I  
Communities: target:2:1

## show route terse

**List of Syntax** [Syntax on page 3246](#)  
[Syntax \(EX Series Switches\) on page 3246](#)

**Syntax** show route terse  
 <logical-system (all | *logical-system-name*)>

**Syntax (EX Series Switches)** show route terse

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Display a high-level summary of the routes in the routing table.



**NOTE:** For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the [show route extensive](#) command.

**Options** **none**—Display a high-level summary of the routes in the routing table.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**Required Privilege Level** view

**List of Sample Output** [show route terse on page 3248](#)

**Output Fields** [Table 294 on page 3246](#) describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

**Table 294: show route terse Output Fields**

| Field Name                 | Field Description   |
|----------------------------|---|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).  |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.   |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active)</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li><b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |



Table 294: show route terse Output Fields (*continued*)

| Field Name       | Field Description  |
|------------------|--|
| <i>route key</i> | Key for the state of the route: <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul>  |
| <b>A</b>         | Active route. An asterisk (*) indicates this is the active route.  |
| <b>V</b>         | Validation status of the route: <ul style="list-style-type: none"> <li>• <b>?</b>—Not evaluated. Indicates that the route was not learned through BGP.</li> <li>• <b>I</b>—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>N</b>—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>V</b>—Valid. Indicates that the prefix and autonomous system pair are found in the database.</li> </ul> |
| Destination      | Destination of the route.  |
| <b>P</b>         | Protocol through which the route was learned: <ul style="list-style-type: none"> <li>• <b>A</b>—Aggregate</li> <li>• <b>B</b>—BGP</li> <li>• <b>C</b>—CCC</li> <li>• <b>D</b>—Direct</li> <li>• <b>G</b>—GMPLS</li> <li>• <b>I</b>—IS-IS</li> <li>• <b>L</b>—L2CKT, L2VPN, LDP, Local</li> <li>• <b>K</b>—Kernel</li> <li>• <b>M</b>—MPLS, MSDP</li> <li>• <b>O</b>—OSPF</li> <li>• <b>P</b>—PIM</li> <li>• <b>R</b>—RIP, RIPng</li> <li>• <b>S</b>—Static</li> <li>• <b>T</b>—Tunnel</li> </ul>   |
| <b>Prf</b>       | Preference value of the route. In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.  |
| Metric 1         | First metric value in the route. For routes learned from BGP, this is the MED metric.  |
| Metric 2         | Second metric value in the route. For routes learned from BGP, this is the IGP metric.   |

Table 294: show route terse Output Fields (*continued*)

| Field Name | Field Description  |
|------------|--|
| Next hop   | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.  |
| AS path    | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>I—IGP.</li> <li>E—EGP.</li> <li>?—Incomplete; typically, the AS path was aggregated.</li> </ul> |

## Sample Output

### show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* ? 1.0.1.1/32        0 10      1           >10.0.0.2      I
?                               B 170      100           >10.0.0.2      I
unverified
* ? 1.1.1.1/32        D 0           >10.0.0.2      200 I
* V 2.2.0.2/32        B 170     110           >10.0.0.2
valid
* ? 10.0.0.0/30       D 0           >1t-1/2/0.1    I
?                               B 170     100           >10.0.0.2
unverified
* ? 10.0.0.1/32       L 0           Local          I
* ? 10.0.0.4/30       B 170     100           >10.0.0.2
unverified
* ? 10.0.0.8/30       B 170     100           >10.0.0.2
unverified
* I 172.16.1.1/32     B 170      90           >10.0.0.2      200 I
invalid
* N 192.168.2.3/32    B 170     100           >10.0.0.2      200 I
unknown
* ? 224.0.0.5/32      O 10      1           MultiRecv

```

## CHAPTER 40

# Troubleshooting

- [Troubleshooting Procedures on page 3249](#)

## Troubleshooting Procedures

---

- [Troubleshooting Virtual Routing Instances on page 3249](#)

## Troubleshooting Virtual Routing Instances

- [Direct Routes Not Leaked Between Routing Instances on page 3249](#)

### Direct Routes Not Leaked Between Routing Instances

---

**Problem** **Description:** Direct routes are not exported (leaked) between virtual routing instances. For example, consider the following scenario:

- Switch with two virtual routing instances:
  - Routing instance 1 connects to downstream device through interface xe-0/0/1.
  - Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the switch connects to the upstream device over a direct route and these routes are not leaked between instances.



**NOTE:** You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the switch over indirect routes.

**Solution** This is expected behavior.

**Related Documentation**

- [Understanding Virtual Router Routing Instances on page 2898](#)
- [Configuring Virtual Router Routing Instances on page 2908](#)

- [rib-group on page 3031](#)

## PART 11

# Border Gateway Protocol

- [Overview on page 3253](#)
- [Configuration on page 3261](#)
- [Administration on page 3749](#)



## CHAPTER 41

# Overview

- [BGP Overview on page 3253](#)

### **BGP Overview**

---

- [Understanding BGP on page 3254](#)
- [BGP Routes Overview on page 3256](#)
- [BGP Messages Overview on page 3257](#)
- [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP on page 3258](#)

## Understanding BGP

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes MP\_REACH\_NLRI and MP\_UNREACH\_NLRI, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- [Autonomous Systems on page 3254](#)
- [AS Paths and Attributes on page 3254](#)
- [External and Internal BGP on page 3255](#)
- [Multiple Instances of BGP on page 3255](#)

---

### Autonomous Systems

An *autonomous system* (AS) is a set of routers that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routers. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

---

### AS Paths and Attributes

The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate



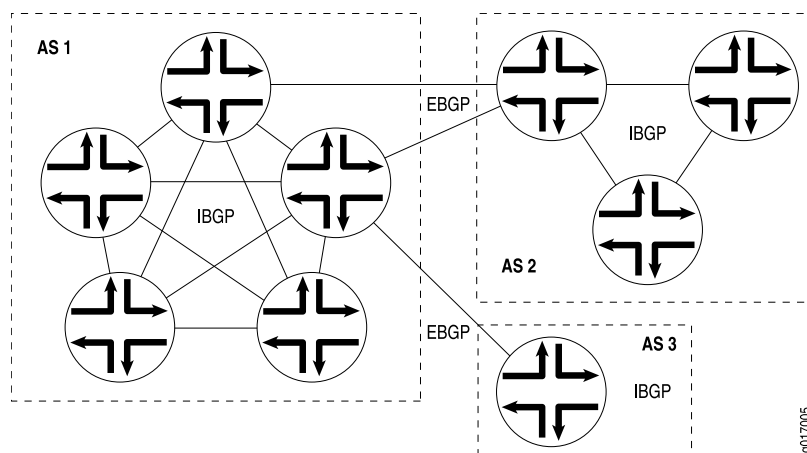
routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

### External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*.

Figure 59 on page 3255 illustrates ASs, IBGP, and EBGP.

Figure 59: ASs, EBGP, and IBGP



A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an IGP to resolve forwarding addresses. They also propagate external routes among all other internal routers running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local router.

### Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.



**NOTE:** When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

- Related Documentation**
- [BGP Routes Overview on page 3256](#)
  - [BGP Messages Overview on page 3257](#)

## BGP Routes Overview

A BGP route is a destination, described as an IP address prefix, and information that describes the path to the destination.

The following information describes the path:

- AS path, which is a list of numbers of the ASs that a route passes through to reach the local router. The first number in the path is that of the last AS in the path—the AS closest to the local router. The last number in the path is the AS farthest from the local router, which is generally the origin of the path.
- Path attributes, which contain additional information about the AS path that is used in routing policy.

BGP peers advertise routes to each other in update messages.

BGP stores its routes in the Junos OS routing table (**inet.0**). The routing table stores the following information about BGP routes:

- Routing information learned from update messages received from peers

- Local routing information that BGP applies to routes because of local policies
- Information that BGP advertises to BGP peers in update messages

For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

The BGP router that first advertises a route assigns it one of the following values to identify its origin. During route selection, the lowest origin value is preferred.

- **0**—The router originally learned the route through an IGP (OSPF, IS-IS, or a static route).
- **1**—The router originally learned the route through an EGP (most likely BGP).
- **2**—The route's origin is unknown.

**Related  
Documentation**

- [Understanding BGP Path Selection](#)
- [Example: Advertising Multiple Paths in BGP on page 3495](#)

## BGP Messages Overview

All BGP messages have the same fixed-size header, which contains a marker field that is used for both synchronization and authentication, a length field that indicates the length of the packet, and a type field that indicates the message type (for example, open, update, notification, keepalive, and so on).

This section discusses the following topics:

- [Open Messages on page 3257](#)
- [Update Messages on page 3258](#)
- [Keepalive Messages on page 3258](#)
- [Notification Messages on page 3258](#)

### Open Messages

After a TCP connection is established between two BGP systems, they exchange BGP open messages to create a BGP connection between them. Once the connection is established, the two systems can exchange BGP messages and data traffic.

Open messages consist of the BGP header plus the following fields:

- **Version**—The current BGP version number is 4.
- **Local AS number**—You configure this by including the **autonomous-system** statement at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level.
- **Hold time**—Proposed hold-time value. You configure the local hold time with the BGP **hold-time** statement.
- **BGP identifier**—IP address of the BGP system. This address is determined when the system starts and is the same for every local interface and every BGP peer. You can

configure the BGP identifier by including the **router-id** statement at the **[edit routing-options]** or **[edit logical-systems *logical-system-name* routing-options]** hierarchy level. By default, BGP uses the IP address of the first interface it finds in the router.

- Parameter field length and the parameter itself—These are optional fields.

---

### Update Messages

BGP systems send update messages to exchange network reachability information. BGP systems use this information to construct a graph that describes the relationships among all known ASs.

Update messages consist of the BGP header plus the following optional fields:

- Unfeasible routes length—Length of the withdrawn routes field
- Withdrawn routes—IP address prefixes for the routes being withdrawn from service because they are no longer deemed reachable
- Total path attribute length—Length of the path attributes field; it lists the path attributes for a feasible route to a destination
- Path attributes—Properties of the routes, including the path origin, the multiple exit discriminator (MED), the originating system's preference for the route, and information about aggregation, communities, confederations, and route reflection
- Network layer reachability information (NLRI)—IP address prefixes of feasible routes being advertised in the update message

---

### Keepalive Messages

BGP systems exchange keepalive messages to determine whether a link or host has failed or is no longer available. Keepalive messages are exchanged often enough so that the hold timer does not expire. These messages consist only of the BGP header.

---

### Notification Messages

BGP systems send notification messages when an error condition is detected. After the message is sent, the BGP session and the TCP connection between the BGP systems are closed. Notification messages consist of the BGP header plus the error code and subcode, and data that describes the error.

#### Related Documentation

- [Understanding BGP on page 3254](#)
- [BGP Routes Overview on page 3256](#)

## Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

BGP peers advertise routes to each other in update messages. BGP stores its routes in the Junos OS routing table (**inet.0**). For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

Instead of advertising only the active path to a destination, you can configure BGP to advertise multiple paths to the destination. Within an autonomous system (AS), the availability of multiple exit points to reach a destination provides the following benefits:

- **Fault tolerance**—Path diversity leads to reduction in restoration time after failure. For instance, a border after receiving multiple paths to the same destination can precompute a backup path and have it ready so that when the primary path becomes invalid, the border routing device can use the backup to quickly restore connectivity. Without a backup path, the restoration time depends on BGP reconvergence, which includes withdraw and advertisement messages in the network before a new best path can be learned.
- **Load balancing**—The availability of multiple paths to reach the same destination enables load balancing of traffic, if the routing within the AS meets certain constraints.
- **Maintenance**—The availability of alternate exit points allows for graceful maintenance operation of routers.

The following limitations apply to advertising multiple routes in BGP:

- Address families supported:
  - IPv4 unicast (**family inet unicast**)
  - IPv6 unicast (**family inet6 unicast**)
  - IPv4 labeled unicast (**family inet labeled-unicast**)
  - IPv6 labeled unicast (**family inet6 labeled-unicast**)
- Internal BGP (IBGP) peers only. No support on external BGP (EBGP) peers.
- Master instance only. No support for routing instances.
- Graceful restart and nonstop active routing (NSR) are supported.
- No BGP Monitoring Protocol (BMP) support.
- No support for EBGP sessions between confederations.
- Prefix policies enable you to filter routes on a router that is configured to advertise multiple paths to a destination. Prefix policies can only match prefixes. They cannot match route attributes, and they cannot change the attributes of routes.

**Related  
Documentation**

- *Understanding BGP Path Selection*
- [Example: Advertising Multiple Paths in BGP on page 3495](#)



## CHAPTER 42

# Configuration

- [Basic BGP Configuration on page 3261](#)
- [BGP Path Attribute Configuration on page 3310](#)
- [BGP Policy Configuration on page 3421](#)
- [BGP BFD Configuration on page 3462](#)
- [BGP Load Balancing Configuration on page 3476](#)
- [IBGP Scaling Configuration on page 3547](#)
- [BGP Security Configuration on page 3570](#)
- [BGP Flap Configuration on page 3591](#)
- [BGP Monitoring Configuration on page 3619](#)
- [Configuration Statements on page 3627](#)

### Basic BGP Configuration

---

- [Examples: Configuring External BGP Peering on page 3261](#)
- [Examples: Configuring Internal BGP Peering on page 3284](#)
- [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

### Examples: Configuring External BGP Peering

- [Understanding External BGP Peering Sessions on page 3261](#)
- [Example: Configuring External BGP Point-to-Point Peer Sessions on page 3262](#)
- [Example: Configuring External BGP on Logical Systems with IPv6 Interfaces on page 3269](#)


#### Understanding External BGP Peering Sessions

---

To establish point-to-point connections between peer autonomous systems (ASs), you configure a BGP session on each interface of a point-to-point link. Generally, such sessions are made at network exit points with neighboring hosts outside the AS.

[Figure 60 on page 3262](#) shows an example of a BGP peering session.

Figure 60: BGP Peering Session



In [Figure 60 on page 3262](#), Router A is a gateway router for AS 3, and Router B is a gateway router for AS 10. For traffic internal to either AS, an interior gateway protocol (IGP) is used (OSPF, for instance). To route traffic between peer ASs, a BGP session is used.

You arrange BGP routing devices into groups of peers. Different peer groups can have different group types, AS numbers, and route reflector cluster identifiers.

To define a BGP group that recognizes only the specified BGP systems as peers, statically configure all the system's peers by including one or more **neighbor** statements. The peer neighbor's address can be either an IPv6 or IPv4 address.

As the number of external BGP (EBGP) groups increases, the ability to support a large number of BGP sessions might become a scaling issue. The preferred way to configure a large number of BGP neighbors is to configure a few groups consisting of multiple neighbors per group. Supporting fewer EBGP groups generally scales better than supporting a large number of EBGP groups. This becomes more evident in the case of hundreds of EBGP groups when compared with a few EBGP groups with multiple peers in each group.

After the BGP peers are established, BGP routes are not automatically advertised by the BGP peers. At each BGP-enabled device, policy configuration is required to export the local, static, or IGP-learned routes into the BGP RIB and then advertise them as BGP routes to the other peers. BGP's advertisement policy, by default, does not advertise any non-BGP routes (such as local routes) to peers.

### Example: Configuring External BGP Point-to-Point Peer Sessions

This example shows how to configure BGP point-to-point peer sessions.

- [Requirements on page 3262](#)
- [Overview on page 3263](#)
- [Configuration on page 3263](#)
- [Verification on page 3265](#)

#### Requirements


Before you begin, if the default BGP policy is not adequate for your network, configure routing policies to filter incoming BGP routes and to advertise BGP routes.



### Overview

Figure 61 on page 3263 shows a network with BGP peer sessions. In the sample network, Device E in AS 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22 and have IP addresses 10.10.10.2, 10.10.10.6, and 10.10.10.10. Peer D resides in AS 79, at IP address 10.21.7.2. This example shows the configuration on Device E.

Figure 61: Typical Network with BGP Peer Sessions



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces ge-1/2/0 unit 0 description to-A
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-0/0/1 unit 5 description to-B
set interfaces ge-0/0/1 unit 5 family inet address 10.10.10.5/30
set interfaces ge-0/1/0 unit 9 description to-C
set interfaces ge-0/1/0 unit 9 family inet address 10.10.10.9/30
set interfaces ge-1/2/1 unit 21 description to-D
set interfaces ge-1/2/1 unit 21 family inet address 10.21.7.1/30
set protocols bgp group external-peers type external
set protocols bgp group external-peers peer-as 22
set protocols bgp group external-peers neighbor 10.10.10.2
set protocols bgp group external-peers neighbor 10.10.10.6
set protocols bgp group external-peers neighbor 10.10.10.10
set protocols bgp group external-peers neighbor 10.21.7.2 peer-as 79
set routing-options autonomous-system 17
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Configure the interfaces to Peers A, B, C, and D.

```
[edit interfaces]
user@E# set ge-1/2/0 unit 0 description to-A
user@E# set ge-1/2/0 unit 0 family inet address 10.10.10.1/30
user@E# set ge-0/0/1 unit 5 description to-B
user@E# set ge-0/0/1 unit 5 family inet address 10.10.10.5/30
user@E# set ge-0/1/0 unit 9 description to-C
user@E# set ge-0/1/0 unit 9 family inet address 10.10.10.9/30
user@E# set ge-1/2/1 unit 21 description to-D
user@E# set ge-1/2/1 unit 21 family inet address 10.21.7.1/30
```

2. Set the autonomous system (AS) number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

3. Create the BGP group, and add the external neighbor addresses.

```
[edit protocols bgp group external-peers]
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10
```

4. Specify the autonomous system (AS) number of the external AS.

```
[edit protocols bgp group external-peers]
user@E# set peer-as 22
```

5. Add Peer D, and set the AS number at the individual neighbor level.

The neighbor configuration overrides the group configuration. So, while **peer-as 22** is set for all the other neighbors in the group, **peer-as 79** is set for neighbor 10.21.7.2.

```
[edit protocols bgp group external-peers]
user@E# set neighbor 10.21.7.2 peer-as 79
```

6. Set the peer type to external BGP (EBGP).

```
[edit protocols bgp group external-peers]
user@E# set type external
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@E# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-A;
    family inet {
      address 10.10.10.1/30;
```

```

    }
  }
}
ge-0/0/1 {
  unit 5 {
    description to-B;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
ge-0/1/0 {
  unit 9 {
    description to-C;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
ge-1/2/1 {
  unit 21 {
    description to-D;
    family inet {
      address 10.21.7.1/30;
    }
  }
}

[edit]
user@E# show protocols
bgp {
  group external-peers {
    type external;
    peer-as 22;
    neighbor 10.10.10.2;
    neighbor 10.10.10.6;
    neighbor 10.10.10.10;
    neighbor 10.21.7.2 {
      peer-as 79;
    }
  }
}

[edit]
user@E# show routing-options
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3266](#)
- [Verifying BGP Groups on page 3268](#)
- [Verifying BGP Summary Information on page 3268](#)

### *Verifying BGP Neighbors*

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From operational mode, run the **show bgp neighbor** command.

```
user@E> show bgp neighbor
Peer: 10.10.10.2+179 AS 22      Local: 10.10.10.1+65406 AS 17
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.10.10.2      Local ID: 10.10.10.1      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 0
  BFD: disabled, down
  Local Interface: ge-1/2/0.0
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        0
    Accepted prefixes:        0
    Suppressed due to damping: 0
    Advertised prefixes:      0
  Last traffic (seconds): Received 10  Sent 6  Checked 1
  Input messages: Total 8522  Updates 1  Refreshes 0  Octets 161922
  Output messages: Total 8433  Updates 0  Refreshes 0  Octets 160290
  Output Queue[0]: 0

Peer: 10.10.10.6+54781 AS 22  Local: 10.10.10.5+179 AS 17
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.10.10.6      Local ID: 10.10.10.1      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 1
  BFD: disabled, down
  Local Interface: ge-0/0/1.5
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
```

```

Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 12   Sent 6   Checked 33
Input messages: Total 8527   Updates 1   Refreshes 0   Octets 162057
Output messages: Total 8430   Updates 0   Refreshes 0   Octets 160233
Output Queue[0]: 0

Peer: 10.10.10.10+55012 AS 22 Local: 10.10.10.9+179 AS 17
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.10.10.10 Local ID: 10.10.10.1 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 2
BFD: disabled, down
Local Interface: fe-0/1/0.9
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 15   Sent 6   Checked 37
Input messages: Total 8527   Updates 1   Refreshes 0   Octets 162057
Output messages: Total 8429   Updates 0   Refreshes 0   Octets 160214
Output Queue[0]: 0

Peer: 10.21.7.2+61867 AS 79 Local: 10.21.7.1+179 AS 17
Type: External State: Established Flags: <ImportEval Sync>

```

```

Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.21.7.2          Local ID: 10.10.10.1      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 3
BFD: disabled, down
Local Interface: ge-1/2/1.21
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 79)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 28   Sent 24   Checked 47
Input messages: Total 8521   Updates 1   Refreshes 0   Octets 161943
Output messages: Total 8427   Updates 0   Refreshes 0   Octets 160176
Output Queue[0]: 0

```

### Verifying BGP Groups

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From operational mode, run the **show bgp group** command.

```

user@E> show bgp group
Group Type: External                      Local AS: 17
Name: external-peers   Index: 0           Flags: <>
Holdtime: 0
Total peers: 4          Established: 4
10.10.10.2+179
10.10.10.6+54781
10.10.10.10+55012
10.21.7.2+61867
inet.0: 0/0/0/0

Groups: 1   Peers: 4   External: 4   Internal: 0   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0      0           0           0           0        0    0       0

```

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From operational mode, run the **show bgp summary** command.

```
user@E> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed  History  Damp State   Pending
inet.0          0          0          0          0          0          0          0
Peer           AS           InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.10.2      22          8559      8470      0        0 2d 16:12:56
0/0/0/0         0/0/0/0
10.10.10.6      22          8566      8468      0        0 2d 16:12:12
0/0/0/0         0/0/0/0
10.10.10.10     22          8565      8466      0        0 2d 16:11:31
0/0/0/0         0/0/0/0
10.21.7.2       79          8560      8465      0        0 2d 16:10:58
0/0/0/0         0/0/0/0
```

### Example: Configuring External BGP on Logical Systems with IPv6 Interfaces

This example shows how to configure external BGP (EBGP) point-to-point peer sessions on logical systems with IPv6 interfaces.

- [Requirements on page 3269](#)
- [Overview on page 3269](#)
- [Configuration on page 3270](#)
- [Verification on page 3279](#)

#### Requirements

In this example, no special configuration beyond device initialization is required.

#### Overview

Junos OS supports EBGP peer sessions by means of IPv6 addresses. An IPv6 peer session can be configured when an IPv6 address is specified in the **neighbor** statement. This example uses EUI-64 to generate IPv6 addresses that are automatically applied to the interfaces. An EUI-64 address is an IPv6 address that uses the IEEE EUI-64 format for the interface identifier portion of the address (the last 64 bits).



**NOTE:** Alternatively, you can configure EBGP sessions using manually assigned 128-bit IPv6 addresses.

If you use 128-bit link-local addresses for the interfaces, you must include the **local-interface** statement. This statement is valid only for 128-bit IPv6 link-local addresses and is mandatory for configuring an IPv6 EBGP link-local peer session.

Configuring EBGP peering using link-local addresses is only applicable for directly connected interfaces. There is no support for multihop peering.


After your interfaces are up, you can use the **show interfaces terse** command to view the EUI-64-generated IPv6 addresses on the interfaces. You must use these generated

addresses in the BGP **neighbor** statements. This example demonstrates the full end-to-end procedure.

In this example, Frame Relay interface encapsulation is applied to the logical tunnel (lt) interfaces. This is a requirement because only Frame Relay encapsulation is supported when IPv6 addresses are configured on the lt interfaces.

Figure 62 on page 3270 shows a network with BGP peer sessions. In the sample network, Router R1 has five logical systems configured. Device E in autonomous system (AS) 17 has BGP peer sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22. This example shows the step-by-step configuration on Logical System A and Logical System E.

Figure 62: Typical Network with BGP Peer Sessions



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device A

```
set logical-systems A interfaces lt-0/1/0 unit 1 description to-E
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation frame-relay
set logical-systems A interfaces lt-0/1/0 unit 1 dlci 1
```



```

set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 25
set logical-systems A interfaces lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64
  eui-64
set logical-systems A interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set logical-systems A protocols bgp group external-peers type external
set logical-systems A protocols bgp group external-peers peer-as 17
set logical-systems A protocols bgp group external-peers neighbor
  2001:db8:0:1:2a0:a502:0:19da
set logical-systems A routing-options router-id 1.1.1.1
set logical-systems A routing-options autonomous-system 22

```

**Device B**

```

set logical-systems B interfaces lt-0/1/0 unit 6 description to-E
set logical-systems B interfaces lt-0/1/0 unit 6 encapsulation frame-relay
set logical-systems B interfaces lt-0/1/0 unit 6 dlci 6
set logical-systems B interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems B interfaces lt-0/1/0 unit 6 family inet6 address 2001:db8:0:2::/64
  eui-64
set logical-systems B interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set logical-systems B protocols bgp group external-peers type external
set logical-systems B protocols bgp group external-peers peer-as 17
set logical-systems B protocols bgp group external-peers neighbor
  2001:db8:0:2:2a0:a502:0:5da
set logical-systems B routing-options router-id 2.2.2.2
set logical-systems B routing-options autonomous-system 22

```

**Device C**

```

set logical-systems C interfaces lt-0/1/0 unit 10 description to-E
set logical-systems C interfaces lt-0/1/0 unit 10 encapsulation frame-relay
set logical-systems C interfaces lt-0/1/0 unit 10 dlci 10
set logical-systems C interfaces lt-0/1/0 unit 10 peer-unit 9
set logical-systems C interfaces lt-0/1/0 unit 10 family inet6 address 2001:db8:0:3::/64
  eui-64
set logical-systems C interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set logical-systems C protocols bgp group external-peers type external
set logical-systems C protocols bgp group external-peers peer-as 17
set logical-systems C protocols bgp group external-peers neighbor
  2001:db8:0:3:2a0:a502:0:9da
set logical-systems C routing-options router-id 3.3.3.3
set logical-systems C routing-options autonomous-system 22

```

**Device D**

```

set logical-systems D interfaces lt-0/1/0 unit 7 description to-E
set logical-systems D interfaces lt-0/1/0 unit 7 encapsulation frame-relay
set logical-systems D interfaces lt-0/1/0 unit 7 dlci 7
set logical-systems D interfaces lt-0/1/0 unit 7 peer-unit 21
set logical-systems D interfaces lt-0/1/0 unit 7 family inet6 address 2001:db8:0:4::/64
  eui-64
set logical-systems D interfaces lo0 unit 4 family inet6 address 2001:db8::4/128
set logical-systems D protocols bgp group external-peers type external
set logical-systems D protocols bgp group external-peers peer-as 17
set logical-systems D protocols bgp group external-peers neighbor
  2001:db8:0:4:2a0:a502:0:15da
set logical-systems D routing-options router-id 4.4.4.4
set logical-systems D routing-options autonomous-system 79

```

**Device E**

```

set logical-systems E interfaces lt-0/1/0 unit 5 description to-B
set logical-systems E interfaces lt-0/1/0 unit 5 encapsulation frame-relay

```

```

set logical-systems E interfaces lt-0/1/0 unit 5 dlci 6
set logical-systems E interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems E interfaces lt-0/1/0 unit 5 family inet6 address 2001:db8:0:2::/64
    eui-64
set logical-systems E interfaces lt-0/1/0 unit 9 description to-C
set logical-systems E interfaces lt-0/1/0 unit 9 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 9 dlci 10
set logical-systems E interfaces lt-0/1/0 unit 9 peer-unit 10
set logical-systems E interfaces lt-0/1/0 unit 9 family inet6 address 2001:db8:0:3::/64
    eui-64
set logical-systems E interfaces lt-0/1/0 unit 21 description to-D
set logical-systems E interfaces lt-0/1/0 unit 21 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 21 dlci 7
set logical-systems E interfaces lt-0/1/0 unit 21 peer-unit 7
set logical-systems E interfaces lt-0/1/0 unit 21 family inet6 address 2001:db8:0:4::/64
    eui-64
set logical-systems E interfaces lt-0/1/0 unit 25 description to-A
set logical-systems E interfaces lt-0/1/0 unit 25 encapsulation frame-relay
set logical-systems E interfaces lt-0/1/0 unit 25 dlci 1
set logical-systems E interfaces lt-0/1/0 unit 25 peer-unit 1
set logical-systems E interfaces lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64
    eui-64
set logical-systems E interfaces lo0 unit 5 family inet6 address 2001:db8::5/128
set logical-systems E protocols bgp group external-peers type external
set logical-systems E protocols bgp group external-peers peer-as 22
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:1:2a0:a502:0:1da
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:2:2a0:a502:0:6da
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:3:2a0:a502:0:ada
set logical-systems E protocols bgp group external-peers neighbor
    2001:db8:0:4:2a0:a502:0:7da peer-as 79
set logical-systems E routing-options router-id 5.5.5.5
set logical-systems E routing-options autonomous-system 17

```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Run the **show interfaces terse** command to verify that the physical router has a logical tunnel (lt) interface.

```

user@R1> show interfaces terse
Interface           Admin Link Proto  Local          Remote
...
lt-0/1/0             up    up
...

```

2. On Logical System A, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System E.

```

user@R1> set cli logical-system A
Logical system: A
[edit]

```

```

user@R1:A> edit
Entering configuration mode
[edit]
user@R1:A# edit interfaces
[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 encapsulation frame-relay
user@R1:A# set lt-0/1/0 unit 1 dlci 1
user@R1:A# set lt-0/1/0 unit 1 peer-unit 25

```

3. On Logical System A, configure the network address for the link to Peer E, and configure a loopback interface.

```

[edit interfaces]
user@R1:A# set lt-0/1/0 unit 1 description to-E
user@R1:A# set lt-0/1/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:A# set lo0 unit 1 family inet6 address 2001:db8::1/128

```

4. On Logical System E, configure the interface encapsulation, peer-unit number, and DLCI to reach Logical System A.

```

user@R1> set cli logical-system E
Logical system: E
[edit]
user@R1:E> edit
Entering configuration mode
[edit]
user@R1:E# edit interfaces
[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 encapsulation frame-relay
user@R1:E# set lt-0/1/0 unit 25 dlci 1
user@R1:E# set lt-0/1/0 unit 25 peer-unit 1

```

5. On Logical System E, configure the network address for the link to Peer A, and configure a loopback interface.

```

[edit interfaces]
user@R1:E# set lt-0/1/0 unit 25 description to-A
user@R1:E# set lt-0/1/0 unit 25 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1:E# set lo0 unit 5 family inet6 address 2001:db8::5/128

```

6. Run the **show interfaces terse** command to see the IPv6 addresses that are generated by EUI-64.

The 2001 addresses are used in this example in the BGP **neighbor** statements.



**NOTE:** The fe80 addresses are link-local addresses and are not used in this example.

```

user@R1:A> show interfaces terse
Interface          Admin Link Proto  Local              Remote
Logical system: A

betsy@tp8:A> show interfaces terse
Interface          Admin Link Proto  Local              Remote
lt-0/1/0
lt-0/1/0.1         up    up    inet6    2001:db8:0:1:2a0:a502:0:1da/64

```

```

                                fe80::2a0:a502:0:1da/64
1o0
1o0.1                        up    up    inet6  2001:db8::1
                                fe80::2a0:a50f:fc56:1da

user@R1:E> show interfaces terse
Interface                Admin Link Proto  Local                Remote
1t-0/1/0
1t-0/1/0.25              up    up    inet6  2001:db8:0:1:2a0:a502:0:19da/64
                                fe80::2a0:a502:0:19da/64
1o0
1o0.5                      up    up    inet6  2001:db8::5
                                fe80::2a0:a50f:fc56:1da

```

7. Repeat the interface configuration on the other logical systems.

### Configuring the External BGP Sessions

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. On Logical System A, create the BGP group, and add the external neighbor address.
 

```
[edit protocols bgp group external-peers]
user@R1:A# set neighbor 2001:db8:0:1:2a0:a502:0:19da
```
2. On Logical System E, create the BGP group, and add the external neighbor address.
 

```
[edit protocols bgp group external-peers]
user@R1:E# set neighbor 2001:db8:0:1:2a0:a502:0:1da
```
3. On Logical System A, specify the autonomous system (AS) number of the external AS.
 

```
[edit protocols bgp group external-peers]
user@R1:A# set peer-as 17
```
4. On Logical System E, specify the autonomous system (AS) number of the external AS.
 

```
[edit protocols bgp group external-peers]
user@R1:E# set peer-as 22
```
5. On Logical System A, set the peer type to EBGP.
 

```
[edit protocols bgp group external-peers]
user@R1:A# set type external
```
6. On Logical System E, set the peer type to EBGP.
 

```
[edit protocols bgp group external-peers]
user@R1:E# set type external
```
7. On Logical System A, set the autonomous system (AS) number and router ID.
 

```
[edit routing-options]
user@R1:A# set router-id 1.1.1.1
user@R1:A# set autonomous-system 22
```

8. On Logical System E, set the AS number and router ID.

```
[edit routing-options]
user@R1:E# set router-id 5.5.5.5
user@R1:E# set autonomous-system 17
```

9. Repeat these steps for Peers A, B, C, and D.

**Results** From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show logical-systems
A {
  interfaces {
    lt-0/1/0 {
      unit 1 {
        description to-E;
        encapsulation frame-relay;
        dlci 1;
        peer-unit 25;
        family inet6 {
          address 2001:db8:0:1::/64 {
            eui-64;
          }
        }
      }
    }
  }
  lo0 {
    unit 1 {
      family inet6 {
        address 2001:db8::1/128;
      }
    }
  }
  protocols {
    bgp {
      group external-peers {
        type external;
        peer-as 17;
        neighbor 2001:db8:0:1:2a0:a502:0:19da;
      }
    }
    routing-options {
      router-id 1.1.1.1;
      autonomous-system 22;
    }
  }
}
B {
  interfaces {
    lt-0/1/0 {
      unit 6 {
        description to-E;
        encapsulation frame-relay;
```

```
        dlci 6;
        peer-unit 5;
        family inet6 {
            address 2001:db8:0:2::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 2 {
        family inet6 {
            address 2001:db8::2/128;
        }
    }
}
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 17;
            neighbor 2001:db8:0:2:2a0:a502:0:5da;
        }
    }
    routing-options {
        router-id 2.2.2.2;
        autonomous-system 22;
    }
}
C {
    interfaces {
        lt-0/1/0 {
            unit 10 {
                description to-E;
                encapsulation frame-relay;
                dlci 10;
                peer-unit 9;
                family inet6 {
                    address 2001:db8:0:3::/64 {
                        eui-64;
                    }
                }
            }
        }
    }
    lo0 {
        unit 3 {
            family inet6 {
                address 2001:db8::3/128;
            }
        }
    }
}
}
protocols {
    bgp {
        group external-peers {
```

```

        type external;
        peer-as 17;
        neighbor 2001:db8:0:3:2a0:a502:0:9da;
    }
}
routing-options {
    router-id 3.3.3.3;
    autonomous-system 22;
}
D {
    interfaces {
        lt-0/1/0 {
            unit 7 {
                description to-E;
                encapsulation frame-relay;
                dlci 7;
                peer-unit 21;
                family inet6 {
                    address 2001:db8:0:4::/64 {
                        eui-64;
                    }
                }
            }
        }
    }
    lo0 {
        unit 4 {
            family inet6 {
                address 2001:db8::4/128;
            }
        }
    }
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 17;
            neighbor 2001:db8:0:4:2a0:a502:0:15da;
        }
    }
    routing-options {
        router-id 4.4.4.4;
        autonomous-system 79;
    }
}
E {
    interfaces {
        lt-0/1/0 {
            unit 5 {
                description to-B;
                encapsulation frame-relay;
                dlci 6;
                peer-unit 6;
                family inet6 {

```

```
        address 2001:db8:0:2::/64 {
            eui-64;
        }
    }
}
unit 9 {
    description to-C;
    encapsulation frame-relay;
    dlci 10;
    peer-unit 10;
    family inet6 {
        address 2001:db8:0:3::/64 {
            eui-64;
        }
    }
}
unit 21 {
    description to-D;
    encapsulation frame-relay;
    dlci 7;
    peer-unit 7;
    family inet6 {
        address 2001:db8:0:4::/64 {
            eui-64;
        }
    }
}
unit 25 {
    description to-A;
    encapsulation frame-relay;
    dlci 1;
    peer-unit 1;
    family inet6 {
        address 2001:db8:0:1::/64 {
            eui-64;
        }
    }
}
lo0 {
    unit 5 {
        family inet6 {
            address 2001:db8::5/128;
        }
    }
}
}
protocols {
    bgp {
        group external-peers {
            type external;
            peer-as 22;
            neighbor 2001:db8:0:1:2a0:a502:0:1da;
            neighbor 2001:db8:0:2:2a0:a502:0:6da;
            neighbor 2001:db8:0:3:2a0:a502:0:ada;
            neighbor 2001:db8:0:4:2a0:a502:0:7da {
```



```

        peer-as 79;
    }
}
}
routing-options {
    router-id 5.5.5.5;
    autonomous-system 17;
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3279](#)
- [Verifying BGP Groups on page 3282](#)
- [Verifying BGP Summary Information on page 3282](#)
- [Checking the Routing Table on page 3282](#)

### Verifying BGP Neighbors

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From operational mode, run the **show bgp neighbor** command.

```

user@R1:E> show bgp neighbor
Peer: 2001:db8:0:1:2a0:a502:0:1da+54987 AS 22 Local:
2001:db8:0:1:2a0:a502:0:19da+179 AS 17
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: Open Message Error
  Options: <Preference PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Error: 'Open Message Error' Sent: 20 Recv: 0
  Peer ID: 1.1.1.1      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 0
  BFD: disabled, down
  Local Interface: lt-0/1/0.25
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet6-unicast
  NLRI of received end-of-rib markers: inet6-unicast
  NLRI of all end-of-rib markers sent: inet6-unicast
  Peer supports 4 byte AS extension (peer-as 22)
  Peer does not support Addpath
  Table inet6.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync

```

```

Active prefixes:          0
Received prefixes:       0
Accepted prefixes:       0
Suppressed due to damping: 0
Advertised prefixes:     0
Last traffic (seconds): Received 7   Sent 18   Checked 81
Input messages:  Total 1611  Updates 1       Refreshes 0       Octets 30660
Output messages: Total 1594  Updates 0       Refreshes 0       Octets 30356
Output Queue[0]: 0

```

```

Peer: 2001:db8:0:2:2a0:a502:0:6da+179 AS 22 Local:
2001:db8:0:2:2a0:a502:0:5da+55502 AS 17
Type: External   State: Established   Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: Open Message Error
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Error: 'Open Message Error' Sent: 26 Recv: 0
Peer ID: 2.2.2.2          Local ID: 5.5.5.5          Active Holdtime: 90
Keepalive Interval: 30    Peer index: 2
BFD: disabled, down
Local Interface: lt-0/1/0.5
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast
NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet6.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:          0
Received prefixes:       0
Accepted prefixes:       0
Suppressed due to damping: 0
Advertised prefixes:     0
Last traffic (seconds): Received 15   Sent 8    Checked 8
Input messages:  Total 1610  Updates 1       Refreshes 0       Octets 30601
Output messages: Total 1645  Updates 0       Refreshes 0       Octets 32417
Output Queue[0]: 0

```

```

Peer: 2001:db8:0:3:2a0:a502:0:ada+55983 AS 22 Local:
2001:db8:0:3:2a0:a502:0:9da+179 AS 17
Type: External   State: Established   Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 3.3.3.3          Local ID: 5.5.5.5          Active Holdtime: 90
Keepalive Interval: 30    Peer index: 3
BFD: disabled, down
Local Interface: lt-0/1/0.9
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast

```

```

NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 22)
Peer does not support Addpath
Table inet6.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 21   Sent 21   Checked 67
Input messages: Total 1610 Updates 1     Refreshes 0     Octets 30641
Output messages: Total 1587 Updates 0     Refreshes 0     Octets 30223
Output Queue[0]: 0

Peer: 2001:db8:0:4:2a0:a502:0:7da+49255 AS 79 Local:
2001:db8:0:4:2a0:a502:0:15da+179 AS 17
Type: External   State: Established   Flags: <Sync>
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Options: <Preference PeerAS Refresh>
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 4.4.4.4           Local ID: 5.5.5.5           Active Holdtime: 90
Keepalive Interval: 30     Peer index: 1
BFD: disabled, down
Local Interface: lt-0/1/0.21
NLRI for restart configured on peer: inet6-unicast
NLRI advertised by peer: inet6-unicast
NLRI for this session: inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet6-unicast
NLRI of received end-of-rib markers: inet6-unicast
NLRI of all end-of-rib markers sent: inet6-unicast
Peer supports 4 byte AS extension (peer-as 79)
Peer does not support Addpath
Table inet6.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        0
  Accepted prefixes:        0
  Suppressed due to damping: 0
  Advertised prefixes:      0
Last traffic (seconds): Received 6     Sent 17   Checked 25
Input messages: Total 1615 Updates 1     Refreshes 0     Octets 30736
Output messages: Total 1593 Updates 0     Refreshes 0     Octets 30337
Output Queue[0]: 0

```

**Meaning** IPv6 unicast network layer reachability information (NLRI) is being exchanged between the neighbors.

**Verifying BGP Groups**

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From operational mode, run the **show bgp group** command.

```
user@R1:~> show bgp group
Group Type: External                               Local AS: 17
  Name: external-peers  Index: 0                   Flags: <>
  Holdtime: 0
  Total peers: 4      Established: 4
  2001:db8:0:1:2a0:a502:0:1da+54987
  2001:db8:0:2:2a0:a502:0:6da+179
  2001:db8:0:3:2a0:a502:0:ada+55983
  2001:db8:0:4:2a0:a502:0:7da+49255
  inet6.0: 0/0/0/0

Groups: 1 Peers: 4 External: 4 Internal: 0 Down peers: 0 Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet6.0           0          0          0          0          0          0
inet6.2           0          0          0          0          0          0
```

**Meaning** The group type is external, and the group has four peers.

**Verifying BGP Summary Information**

**Purpose** Verify that the BGP that the peer relationships are established.

**Action** From operational mode, run the **show bgp summary** command.

```
user@R1:~> show bgp summary
Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet6.0           0          0          0          0          0          0
inet6.2           0          0          0          0          0          0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
2001:db8:0:1:2a0:a502:0:1da      22    1617    1600      0      0
  12:07:00 Establ
    inet6.0: 0/0/0/0
2001:db8:0:2:2a0:a502:0:6da      22    1616    1651      0      0
  12:06:56 Establ
    inet6.0: 0/0/0/0
2001:db8:0:3:2a0:a502:0:ada      22    1617    1594      0      0
  12:04:32 Establ
    inet6.0: 0/0/0/0
2001:db8:0:4:2a0:a502:0:7da      79    1621    1599      0      0
  12:07:00 Establ
    inet6.0: 0/0/0/0
```

**Meaning** The Down peers: 0 output shows that the BGP peers are in the established state.

**Checking the Routing Table**

**Purpose** Verify that the inet6.0 routing table is populated with local and direct routes.

**Action** From operational mode, run the **show route** command.

```

user@R1:E> show route
inet6.0: 15 destinations, 18 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::5/128    *[Direct/0] 12:41:18
                  > via lo0.5
2001:db8:0:1::/64  *[Direct/0] 14:40:01
                  > via lt-0/1/0.25
2001:db8:0:1:2a0:a502:0:19da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.25
2001:db8:0:2::/64  *[Direct/0] 14:40:02
                  > via lt-0/1/0.5
2001:db8:0:2:2a0:a502:0:5da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.5
2001:db8:0:3::/64  *[Direct/0] 14:40:02
                  > via lt-0/1/0.9
2001:db8:0:3:2a0:a502:0:9da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.9
2001:db8:0:4::/64  *[Direct/0] 14:40:01
                  > via lt-0/1/0.21
2001:db8:0:4:2a0:a502:0:15da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.21
fe80::/64          *[Direct/0] 14:40:02
                  > via lt-0/1/0.5
                  [Direct/0] 14:40:02
                  > via lt-0/1/0.9
                  [Direct/0] 14:40:01
                  > via lt-0/1/0.21
                  [Direct/0] 14:40:01
                  > via lt-0/1/0.25
fe80::2a0:a502:0:5da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.5
fe80::2a0:a502:0:9da/128
                  *[Local/0] 14:40:02
                  Local via lt-0/1/0.9
fe80::2a0:a502:0:15da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.21
fe80::2a0:a502:0:19da/128
                  *[Local/0] 14:40:01
                  Local via lt-0/1/0.25
fe80::2a0:a50f:fc56:1da/128
                  *[Direct/0] 12:41:18
                  > via lo0.5

```

**Meaning** The inet6.0 routing table contains local and direct routes. To populate the routing table with other types of routes, you must configure routing policies.

**Related Documentation**

- [Examples: Configuring Internal BGP Peering on page 3284](#)
- [BGP Configuration Overview](#)


## Examples: Configuring Internal BGP Peering

- [Understanding Internal BGP Peering Sessions on page 3284](#)
- [Example: Configuring Internal BGP Peer Sessions on page 3285](#)
- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3296](#)

### Understanding Internal BGP Peering Sessions

When two BGP-enabled devices are in the same autonomous system (AS), the BGP session is called an *internal* BGP session, or IBGP session. BGP uses the same message types on IBGP and external BGP (EBGP) sessions, but the rules for when to send each message and how to interpret each message differ slightly. For this reason, some people refer to IBGP and EBGP as two separate protocols.

Figure 63: Internal and External BGP



In [Figure 63 on page 3284](#), Device Jackson, Device Memphis, and Device Biloxi have IBGP peer sessions with each other. Likewise, Device Miami and Device Atlanta have IBGP peer sessions between each other.

The purpose of IBGP is to provide a means by which EBGP route advertisements can be forwarded throughout the network. In theory, to accomplish this task you could redistribute all of your EBGP routes into an interior gateway protocol (IGP), such as OSPF or IS-IS. This, however, is not recommended in a production environment because of the large number of EBGP routes in the Internet and because of the way that IGPs operate. In short, with that many routes the IGP churns or crashes.

Generally, the loopback interface (lo0) is used to establish connections between IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.

While IBGP neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every

other device through neighbor peer relationships. The **neighbor** statement creates the mesh. Because of the full mesh requirement of IBGP, you must configure individual peering sessions between all IBGP devices in the AS. The full mesh need not be physical links. Rather, the configuration on each routing device must create a full mesh of peer sessions (using multiple **neighbor** statements).



**NOTE:** The requirement for a full mesh is waived if you configure a confederation or route reflection.

To understand the full-mesh requirement, consider that an IBGP-learned route cannot be readvertised to another IBGP peer. The reason for preventing the readvertisement of IBGP routes and requiring the full mesh is to avoid routing loops within an AS. The AS path attribute is the means by which BGP routing devices avoid loops. The path information is examined for the local AS number only when the route is received from an EBGP peer. Because the attribute is only modified across AS boundaries, this system works well. However, the fact that the attribute is only modified across AS boundaries presents an issue inside the AS. For example, suppose that routing devices A, B, and C are all in the same AS. Device A receives a route from an EBGP peer and sends the route to Device B, which installs it as the active route. The route is then sent to Device C, which installs it locally and sends it back to Device A. If Device A installs the route, a loop is formed within the AS. The routing devices are not able to detect the loop because the AS path attribute is not modified during these advertisements. Therefore, the BGP protocol designers decided that the only assurance of never forming a routing loop was to prevent an IBGP peer from advertising an IBGP-learned route within the AS. For route reachability, the IBGP peers are fully meshed.

IBGP supports multihop connections, so IBGP neighbors can be located anywhere within the AS and often do not share a link. A recursive route lookup resolves the loopback peering address to an IP forwarding next hop. The lookup service is provided by static routes or an IGP such as OSPF, or BGP routes.

### Example: Configuring Internal BGP Peer Sessions

This example shows how to configure internal BGP peer sessions.

- [Requirements on page 3285](#)
- [Overview on page 3285](#)
- [Configuration on page 3287](#)
- [Verification on page 3294](#)

#### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

#### **Overview**

In this example, you configure internal BGP (IBGP) peer sessions. The loopback interface (lo0) is used to establish connections between IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address,

the IBGP peer session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peer session also goes up and down. Thus, if the device has link redundancy, the loopback interface provides fault tolerance in case the physical interface or one of the links goes down.

When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The **local-address** statement enables you to specify the source information in BGP update messages. If you omit the **local-address** statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally results in the egress interface address being the expected source of update messages. When this happens, the peer session is not established because a mismatch exists between the expected source address (the egress interface of the peer) and the actual source (the loopback interface of the peer). To make sure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.

Because IBGP supports multihop connections, IBGP neighbors can be located anywhere within the autonomous system (AS) and often do not share a link. A recursive route lookup resolves the loopback peer address to an IP forwarding next hop. In this example, this service is provided by OSPF. Although interior gateway protocol (IGP) neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every other device through neighbor peer relationships. The **neighbor** statement creates the mesh.



**NOTE:** The requirement for a full mesh is waived if you configure a confederation or route reflection.

---


After the BGP peers are established, BGP routes are not automatically advertised by the BGP peers. At each BGP-enabled device, policy configuration is required to export the local, static, or IGP-learned routes into the BGP routing information base (RIB) and then advertise them as BGP routes to the other peers. BGP's advertisement policy, by default, does not advertise any non-BGP routes (such as local routes) to peers.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

Figure 64 on page 3287 shows a typical network with internal peer sessions.



Figure 64: Typical Network with IBGP Sessions

**Configuration**

- [Configuring Device A on page 3288](#)
- [Configuring Device B on page 3290](#)
- [Configuring Device C on page 3292](#)

**CLI Quick Configuration**

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device A**

```

set interfaces ge-0/1/0 unit 1 description to-B
set interfaces ge-0/1/0 unit 1 family inet address 10.10.10.1/30
set interfaces lo0 unit 1 family inet address 192.168.6.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to B and C"
set protocols bgp group internal-peers local-address 192.168.6.5
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.1
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.168.6.5
set routing-options autonomous-system 17

```

**Device B**

```

set interfaces ge-0/1/0 unit 2 description to-A
set interfaces ge-0/1/0 unit 2 family inet address 10.10.10.2/30
set interfaces ge-0/1/1 unit 5 description to-C
set interfaces ge-0/1/1 unit 5 family inet address 10.10.10.5/30
set interfaces lo0 unit 2 family inet address 192.163.6.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to A and C"
set protocols bgp group internal-peers local-address 192.163.6.4
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols bgp group internal-peers neighbor 192.168.6.5

```

```
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.2
set protocols ospf area 0.0.0.0 interface ge-0/1/1.5
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.163.6.4
set routing-options autonomous-system 17
```

Device C

```
set interfaces ge-0/1/0 unit 6 description to-B
set interfaces ge-0/1/0 unit 6 family inet address 10.10.10.6/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers description "connections to A and B"
set protocols bgp group internal-peers local-address 192.168.40.4
set protocols bgp group internal-peers export send-direct
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-0/1/0.6
set policy-options policy-statement send-direct term 2 from protocol direct
set policy-options policy-statement send-direct term 2 then accept
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

### Configuring Device A

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 1]
user@A# set description to-B
user@A# set family inet address 10.10.10.1/30
```

```
[edit interfaces]
user@A# set lo0 unit 1 family inet address 192.168.6.5/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@A# set type internal
user@A# set description "connections to B and C"
user@A# set local-address 192.168.6.5
user@A# set export send-direct
user@A# set neighbor 192.163.6.4
user@A# set neighbor 192.168.40.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@A# set interface lo0.1 passive
user@A# set interface ge-0/1/0.1
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@A# set from protocol direct
user@A# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@A# set router-id 192.168.6.5
user@A# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@A# show interfaces
ge-0/1/0 {
  unit 1 {
    description to-B;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@A# show protocols
bgp {
  group internal-peers {
    type internal;
    description "connections to B and C";
    local-address 192.168.6.5;
    export send-direct;
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
```

```
}
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface ge-0/1/0.1;
  }
}

user@A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Configuring Device B**

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure internal BGP peer sessions on Device B:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 2]
user@B# set description to-A
user@B# set family inet address 10.10.10.2/30
```

```
[edit interfaces ge-0/1/1]
user@B# set unit 5 description to-C
user@B# set unit 5 family inet address 10.10.10.5/30
```

```
[edit interfaces]
user@B# set lo0 unit 2 family inet address 192.163.6.4/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@B# set type internal
user@B# set description "connections to A and C"
user@B# set local-address 192.163.6.4
user@B# set export send-direct
user@B# set neighbor 192.168.40.4
user@B# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@B# set interface lo0.2 passive
user@B# set interface ge-0/1/0.2
user@B# set interface ge-0/1/1.5
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@B# set from protocol direct
user@B# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@B# set router-id 192.163.6.4
user@B# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
ge-0/1/0 {
  unit 2 {
    description to-A;
    family inet {
      address 10.10.10.2/30;
    }
  }
}
ge-0/1/1 {
  unit 5 {
    description to-C;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.163.6.4/32;
    }
  }
}

user@B# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@B# show protocols
bgp {
  group internal-peers {
    type internal;
```

```
        description "connections to A and C";
        local-address 192.163.6.4;
        export send-direct;
        neighbor 192.168.40.4;
        neighbor 192.168.6.5;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.2 {
            passive;
        }
        interface ge-0/1/0.2;
        interface ge-0/1/1.5;
    }
}
```

```
user@B# show routing-options
router-id 192.163.6.4;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

### *Configuring Device C*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device C:

1. Configure the interfaces.

```
[edit interfaces ge-0/1/0 unit 6]
user@C# set description to-B
user@C# set family inet address 10.10.10.6/30

[edit interfaces]
user@C# set lo0 unit 3 family inet address 192.168.40.4/32
```

2. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@C# set type internal
user@C# set description "connections to A and B"
user@C# set local-address 192.168.40.4
user@C# set export send-direct
user@C# set neighbor 192.163.6.4
user@C# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@C# set interface lo0.3 passive
user@C# set interface ge-0/1/0.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@C# set from protocol direct
user@C# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@C# set router-id 192.168.40.4
user@C# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@C# show interfaces
ge-0/1/0 {
  unit 6 {
    description to-B;
    family inet {
      address 10.10.10.6/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.40.4/32;
    }
  }
}

user@C# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@C# show protocols
bgp {
  group internal-peers {
    type internal;
    description "connections to A and B";
    local-address 192.168.40.4;
    export send-direct;
    neighbor 192.163.6.4;
    neighbor 192.168.6.5;
  }
}
ospf {
```

```
area 0.0.0.0 {  
    interface lo0.3 {  
        passive;  
    }  
    interface ge-0/1/0.6;  
}  
}
```

```
user@C# show routing-options  
router-id 192.168.40.4;  
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3294](#)
- [Verifying BGP Groups on page 3295](#)
- [Verifying BGP Summary Information on page 3296](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3296](#)

### **Verifying BGP Neighbors**

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From operational mode, enter the **show bgp neighbor** command.

```
user@A> show bgp neighbor  
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+58852 AS 17  
Type: Internal    State: Established    Flags: Sync  
Last State: OpenConfirm    Last Event: RecvKeepAlive  
Last Error: None  
Export: [ send-direct ]  
Options: Preference LocalAddress Refresh  
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170  
Number of flaps: 0  
Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90  
Keepalive Interval: 30    Peer index: 0  
BFD: disabled, down  
NLRI for restart configured on peer: inet-unicast  
NLRI advertised by peer: inet-unicast  
NLRI for this session: inet-unicast  
Peer supports Refresh capability (2)  
Restart time configured on the peer: 120  
Stale routes from peer are kept for: 300  
Restart time requested by this peer: 120  
NLRI that peer supports restart for: inet-unicast  
NLRI that restart is negotiated for: inet-unicast  
NLRI of received end-of-rib markers: inet-unicast  
NLRI of all end-of-rib markers sent: inet-unicast  
Peer supports 4 byte AS extension (peer-as 17)  
Peer does not support Addpath  
Table inet.0 Bit: 10000  
RIB State: BGP restart is complete  
Send state: in sync
```



```

Active prefixes:          0
Received prefixes:       3
Accepted prefixes:       3
Suppressed due to damping: 0
Advertised prefixes:     2
Last traffic (seconds): Received 25   Sent 19   Checked 67
Input messages:  Total 2420   Updates 4       Refreshes 0       Octets 46055
Output messages: Total 2411   Updates 2       Refreshes 0       Octets 45921
Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17   Local: 192.168.6.5+56466 AS 17
Type: Internal   State: Established   Flags: Sync
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: None
Export: [ send-direct ]
Options: Preference LocalAddress Refresh
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4   Local ID: 192.168.6.5       Active Holdtime: 90
Keepalive Interval: 30   Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:          0
Received prefixes:       2
Accepted prefixes:       2
Suppressed due to damping: 0
Advertised prefixes:     2
Last traffic (seconds): Received 7   Sent 21   Checked 24
Input messages:  Total 2412   Updates 2       Refreshes 0       Octets 45867
Output messages: Total 2409   Updates 2       Refreshes 0       Octets 45883
Output Queue[0]: 0

```

### Verifying BGP Groups

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From operational mode, enter the **show bgp group** command.

```

user@A> show bgp group
Group Type: Internal   AS: 17                               Local AS: 17
Name: internal-peers   Index: 0                               Flags: <Export Eval>
Export: [ send-direct ]
Holdtime: 0
Total peers: 2         Established: 2
192.163.6.4+179
192.168.40.4+179

```

```
inet.0: 0/5/5/0
```

| Groups: | 1 | Peers:    | 2 | External: | 0          | Internal: | 2       | Down peers: | 0 | Flaps:  | 0 |
|---------|---|-----------|---|-----------|------------|-----------|---------|-------------|---|---------|---|
| Table   |   | Tot Paths |   | Act Paths | Suppressed |           | History | Damp State  |   | Pending |   |
| inet.0  |   | 5         |   | 0         | 0          |           | 0       | 0           |   | 0       |   |

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From operational mode, enter the **show bgp summary** command.

```
user@A> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 5 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4 17 2441 2432 0 0 18:18:52
0/3/3/0 0/0/0/0
192.168.40.4 17 2432 2430 0 0 18:18:48
0/2/2/0 0/0/0/0
```

### Verifying That BGP Routes Are Installed in the Routing Table

**Purpose** Verify that the export policy configuration is causing the BGP routes to be installed in the routing tables of the peers.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@A> show route protocol bgp
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
10.10.10.4/30 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
[BGP/170] 07:07:12, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
192.163.6.4/32 [BGP/170] 07:09:57, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
192.168.40.4/32 [BGP/170] 07:07:12, localpref 100, from 192.168.40.4
AS path: I
> to 10.10.10.2 via ge-0/1/0.1
```

### Example: Configuring Internal BGP Peering Sessions on Logical Systems

This example shows how to configure internal BGP peer sessions on logical systems.

- [Requirements on page 3297](#)
- [Overview on page 3297](#)

- [Configuration on page 3297](#)
- [Verification on page 3304](#)

### Requirements

In this example, no special configuration beyond device initialization is required.


### Overview

In this example, you configure internal BGP (IBGP) peering sessions.

In the sample network, the devices in AS 17 are fully meshed in the group **internal-peers**. The devices have loopback addresses 192.168.6.5, 192.163.6.4, and 192.168.40.4.

[Figure 65 on page 3297](#) shows a typical network with internal peer sessions.

**Figure 65: Typical Network with IBGP Sessions**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems A interfaces lt-0/1/0 unit 1 description to-B
set logical-systems A interfaces lt-0/1/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-0/1/0 unit 1 peer-unit 2
set logical-systems A interfaces lt-0/1/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-0/1/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
```

```
set logical-systems A routing-options autonomous-system 17
set logical-systems B interfaces lt-0/1/0 unit 2 description to-A
set logical-systems B interfaces lt-0/1/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-0/1/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-0/1/0 unit 5 description to-C
set logical-systems B interfaces lt-0/1/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-0/1/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-0/1/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-0/1/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17
set logical-systems C interfaces lt-0/1/0 unit 6 description to-B
set logical-systems C interfaces lt-0/1/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-0/1/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-0/1/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-0/1/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17
```

### **Device A**

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure internal BGP peer sessions on Device A:

1. Configure the interfaces.  

```
[edit logical-systems A interfaces lt-0/1/0 unit 1]
user@R1# set description to-B
user@R1# set encapsulation ethernet
user@R1# set peer-unit 2
user@R1# set family inet address 10.10.10.1/30
```

```

user@R1# set family inet address 192.168.6.5/32
user@R1# up
user@R1# up
[edit logical-systems A interfaces]
user@R1# set lo0 unit 1 family inet address 192.168.6.5/32
user@R1# exit
[edit]
user@R1# edit logical-systems B interfaces lt-0/1/0
[edit logical-systems B interfaces lt-0/1/0]
user@R1# set unit 2 description to-A
user@R1# set unit 2 encapsulation ethernet
user@R1# set unit 2 peer-unit 1
user@R1# set unit 2 family inet address 10.10.10.2/30
user@R1# set unit 5 description to-C
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 6
user@R1# set family inet address 10.10.10.5/30
user@R1# up
[edit logical-systems B interfaces]
user@R1# set lo0 unit 2 family inet address 192.163.6.4/32
user@R1# exit
[edit]
user@R1# edit logical-systems C interfaces lt-0/1/0 unit 6
[edit logical-systems C interfaces lt-0/1/0 unit 6]
set description to-B
set encapsulation ethernet
set peer-unit 5
set family inet address 10.10.10.6/30
user@R1# up
user@R1# up
[edit logical-systems C interfaces]
set lo0 unit 3 family inet address 192.168.40.4/32

```

## 2. Configure BGP.

On Logical System A, the **neighbor** statements are included for both Device B and Device C, even though Logical System A is not directly connected to Device C.

```

[edit logical-systems A protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.6.5
user@R1# set export send-direct
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.40.4

```

```

[edit logical-systems B protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.163.6.4
user@R1# set export send-direct
user@R1# set neighbor 192.168.40.4
user@R1# set neighbor 192.168.6.5

```

```

[edit logical-systems C protocols bgp group internal-peers]
user@R1# set type internal
user@R1# set local-address 192.168.40.4
user@R1# set export send-direct

```

```
user@R1# set neighbor 192.163.6.4
user@R1# set neighbor 192.168.6.5
```

3. Configure OSPF.

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface lt-0/1/0.1
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.2 passive
user@R1# set interface lt-0/1/0.2
user@R1# set interface lt-0/1/0.5
```

```
[edit logical-systems A protocols ospf area 0.0.0.0]
user@R1# set interface lo0.3 passive
user@R1# set interface lt-0/1/0.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit logical-systems A policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems B policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

```
[edit logical-systems C policy-options policy-statement send-direct term 2]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit logical-systems A routing-options]
user@R1# set router-id 192.168.6.5
user@R1# set autonomous-system 17
```

```
[edit logical-systems B routing-options]
user@R1# set router-id 192.163.6.4
user@R1# set autonomous-system 17
```

```
[edit logical-systems C routing-options]
user@R1# set router-id 192.168.40.4
user@R1# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show logical-systems** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show logical-systems
A {
```

```

interfaces {
  lt-0/1/0 {
    unit 1 {
      description to-B;
      encapsulation ethernet;
      peer-unit 2;
      family inet {
        address 10.10.10.1/30;
      }
    }
  }
  lo0 {
    unit 1 {
      family inet {
        address 192.168.6.5/32;
      }
    }
  }
}
protocols {
  bgp {
    group internal-peers {
      type internal;
      local-address 192.168.6.5;
      export send-direct;
      neighbor 192.163.6.4;
      neighbor 192.168.40.4;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface lo0.1 {
        passive;
      }
      interface lt-0/1/0.1;
    }
  }
}
policy-options {
  policy-statement send-direct {
    term 2 {
      from protocol direct;
      then accept;
    }
  }
}
routing-options {
  router-id 192.168.6.5;
  autonomous-system 17;
}
}
B {
  interfaces {
    lt-0/1/0 {
      unit 2 {
        description to-A;

```

```
        encapsulation ethernet;
        peer-unit 1;
        family inet {
            address 10.10.10.2/30;
        }
    }
    unit 5 {
        description to-C;
        encapsulation ethernet;
        peer-unit 6;
        family inet {
            address 10.10.10.5/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.163.6.4/32;
        }
    }
}
}
protocols {
    bgp {
        group internal-peers {
            type internal;
            local-address 192.163.6.4;
            export send-direct;
            neighbor 192.168.40.4;
            neighbor 192.168.6.5;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface lo0.2 {
                passive;
            }
            interface lt-0/1/0.2;
            interface lt-0/1/0.5;
        }
    }
}
policy-options {
    policy-statement send-direct {
        term 2 {
            from protocol direct;
            then accept;
        }
    }
}
routing-options {
    router-id 192.163.6.4;
    autonomous-system 17;
}
}
```



```

C {
  interfaces {
    lt-0/1/0 {
      unit 6 {
        description to-B;
        encapsulation ethernet;
        peer-unit 5;
        family inet {
          address 10.10.10.6/30;
        }
      }
    }
    lo0 {
      unit 3 {
        family inet {
          address 192.168.40.4/32;
        }
      }
    }
  }
  protocols {
    bgp {
      group internal-peers {
        type internal;
        local-address 192.168.40.4;
        export send-direct;
        neighbor 192.163.6.4;
        neighbor 192.168.6.5;
      }
    }
    ospf {
      area 0.0.0.0 {
        interface lo0.3 {
          passive;
        }
        interface lt-0/1/0.6;
      }
    }
  }
  policy-options {
    policy-statement send-direct {
      term 2 {
        from protocol direct;
        then accept;
      }
    }
  }
  routing-options {
    router-id 192.168.40.4;
    autonomous-system 17;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3304](#)
- [Verifying BGP Groups on page 3305](#)
- [Verifying BGP Summary Information on page 3305](#)
- [Verifying That BGP Routes Are Installed in the Routing Table on page 3306](#)

### Verifying BGP Neighbors

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From the operational mode, enter the **show bgp neighbor** command.

```
user@R1> show bgp neighbor logical-system A
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+58852 AS 17
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-direct ]
  Options: <Preference LocalAddress Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 17)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        3
    Accepted prefixes:        3
    Suppressed due to damping: 0
    Advertised prefixes:      2
  Last traffic (seconds): Received 16    Sent 1    Checked 63
  Input messages: Total 15713 Updates 4    Refreshes 0    Octets 298622
  Output messages: Total 15690 Updates 2    Refreshes 0    Octets 298222
  Output Queue[0]: 0

Peer: 192.168.40.4+179 AS 17    Local: 192.168.6.5+56466 AS 17
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
```

```

Export: [ send-direct ]
Options: <Preference LocalAddress Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.40.4    Local ID: 192.168.6.5    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 1
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        2
  Accepted prefixes:        2
  Suppressed due to damping: 0
  Advertised prefixes:      2
Last traffic (seconds): Received 15    Sent 22    Checked 68
Input messages: Total 15688 Updates 2    Refreshes 0    Octets 298111
Output messages: Total 15688 Updates 2    Refreshes 0    Octets 298184
Output Queue[0]: 0

```

### Verifying BGP Groups

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From the operational mode, enter the **show bgp group** command.

```

user@A> show bgp group logical-system A
Group Type: Internal    AS: 17                      Local AS: 17
Name: internal-peers   Index: 0                    Flags: <Export Eval>
Export: [ send-direct ]
Holdtime: 0
Total peers: 2          Established: 2
192.163.6.4+179
192.168.40.4+179
inet.0: 0/5/5/0

Groups: 1  Peers: 2  External: 0  Internal: 2  Down peers: 0  Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet.0          5          0          0          0          0          0

```

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From the operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary logical-system A

```

```

Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      5          0          0          0        0      0      0
Peer        AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4  17      15723    15700     0        0 4d 22:13:15
0/3/3/0      0/0/0/0
192.168.40.4 17      15698    15699     0        0 4d 22:13:11
0/2/2/0      0/0/0/0

```

### Verifying That BGP Routes Are Installed in the Routing Table

**Purpose** Verify that the export policy configuration is working.

**Action** From the operational mode, enter the **show route protocol bgp** command.

```

user@A> show route protocol bgp logical-system A
inet.0: 7 destinations, 12 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30      [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
10.10.10.4/30      [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
                  [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
192.163.6.4/32     [BGP/170] 4d 11:05:55, localpref 100, from 192.163.6.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1
192.168.40.4/32    [BGP/170] 4d 11:03:10, localpref 100, from 192.168.40.4
                  AS path: I
                  > to 10.10.10.2 via lt-0/1/0.1

```

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3261](#)

## Configuring BGP Monitoring Protocol Version 3

BGP Monitoring Protocol (BMP) allows the Junos OS to send the BGP route information from the router to a monitoring application on a separate device. The monitoring application is called the BMP monitoring station or BMP station. To deploy BMP in your network, you need to configure BMP on each router and you also need to configure at least one BMP station. This procedure describes how to configure BMP on a router.

You can specify these settings for all BMP stations by configuring the statements described here at the **[edit routing-options bmp]** hierarchy level. You can also configure settings for specific BMP stations by configuring these statements at the **[edit routing-options bmp station station-name]** hierarchy level.

The following procedure describes how to configure BMP version 3 on the router:

1. Specify the name or address for the BMP monitoring station by configuring the **station-address** statement. You can specify one or the other but not both. The address must be a valid IPv4 or IPv6 address.

```
station-address (station-address | station-name);
```

2. Specify the authentication algorithm used to encrypt authentication between the BMP-enabled router and the BMP station using the **authentication-algorithm** statement.

```
authentication-algorithm algorithm;
```

You can specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
  - **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
  - **md5**—Message digest 5.
3. Specify an MD5 authentication key (password) using the **authentication-key** statement.

```
authentication-key key;
```

4. Specify the authentication key chain using the **authentication-key-chain** statement.

```
authentication-key-chain key-chain;
```

The authentication key chain itself needs to be configured at the **[edit security authentication-key-chains key-chain]** hierarchy level. For a detailed example, see [“Example: Configuring Route Authentication for BGP” on page 3572](#).

5. Specify how to handle a BMP station flap by configuring the **hold-down** statement. A flap is when the TCP session unexpectedly switches from established to non-established. The BMP station can be prevented from attempting to reconnect to the device for a specified period of time.

```
hold-down {
  seconds;
  flaps number;
  period seconds;
}
```

You can specify the following options for the **hold-down** statement:

- **seconds**—Specify the time in seconds to wait before allowing the BMP station to reconnect to the device.
  - **flaps number**—Specify the number of BMP station flaps allowed before terminating the connection to the BMP station and triggering the hold down timer.
  - **period seconds**—Specify the time in seconds between BMP station flaps before terminating the connection to the BMP station and triggering the hold down timer.
6. (Optional) Specify an initiation message to be sent to the BMP station using the **initiation-message** statement. This statement allows you to provide some information to the BMP station system administrator (for example, a contact phone number).

**initiation-message** *text*;

7. Specify the connection mode for the connection between the BMP-enabled router and the BMP station using the **connection-mode** statement. The connection mode can be **active** or **passive**:
- **active**—BMP initiates the connection to the BMP station. If you configure active mode, you must also configure a station port using the **station-port** statement. However, you must not configure a local port (active mode).
  - **passive**—BMP does not initiate a connection the BMP station. However, it does listen for a connection request from active BMP stations and will connect if a station is available. If you configure passive mode, you must not configure a station port. However, you must configure a local port using the **local-port** statement (passive mode).

**connection-mode** (active | passive);

8. Specify the port number for the BMP monitoring station by configuring the **station-port** statement. See also **connection-mode**.

**station-port** *port*;

9. Specify the listening port for the BMP station connection using the **local-port** statement. See also **connection-mode**.

If you change the local port, the BMP station connection flaps when you commit the configuration.

**local-port**

10. (Optional) Specify the IPv4 or IPv6 address for the BMP connection on the device using the **local-address** statement. For both active and passive connections, configure a loopback local address. This provides a consistent local endpoint, is useful for debugging, and assures greater reliability for the BMP connection since it is not tied to a single router interface.

For passive mode, specifying a local address is required. It also provides some security against a malicious BMP connection. For active mode, we also recommend configuring a local address to help ensure reliability.

If you change the local address, the BMP station connection flaps when you commit the configuration.

**local-address** *address*;

11. BMP monitoring is enabled by default. You can explicitly enable BMP monitoring or disable it. You can also selectively enable or disable BMP monitoring at various hierarchy levels (for example, `[edit protocols bgp group group-name]` or `[edit protocols bgp group group-name neighbor address]`). If you disable BMP monitoring, withdrawal messages are sent for any previously advertised routes. These are followed by a down message. If you enable BMP monitoring, an up message is sent first and then the route advertisements follow.

`monitor` (enable | disable);

12. Specify the dispatch priority for BMP by configuring the `priority` statement. The dispatch priority controls the frequency with which the device is able to forward BMP messages to BMP stations. You can configure the dispatch priority as either **high**, **medium**, or **low**.

`priority` (high | medium | low);

13. Specify whether BMP should send pre-policy route monitoring messages, post-policy route monitoring messages, both types of messages, or none at all. The pre-policy can be configured to exclude routes that are non-feasible for the decision process (for example, a route loop) by including the **non-feasible** option for the **pre-policy** statement. This represents the view of the BGP routes before running the import policy.

The post-policy can be configured to exclude routes that are not eligible for the decision process (for example, protocol nexthop not resolved) by including the **exclude-non-eligible** option for the **post-policy** statement. This represents the view of the BGP routes after running the import policy. If the import policy has rejected the BGP route, the route does not exist in the post policy view.

You can explicitly disable route monitoring by specifying the **none** option for the **route-monitoring** statement. This is the default behavior.

```
route-monitoring {
  none;
  post-policy {
    exclude-non-eligible;
  }
  pre-policy {
    exclude-non-feasible;
  }
}
```

14. Configure how often statistics messages are sent to the BMP monitoring station by specifying the number of seconds between message transmissions using **statistics-timeout** statement. If you configure a value of 0, no statistics messages are sent.

`statistics-timeout` *seconds*;

#### Related Documentation

- [Example: Configuring Route Authentication for BGP on page 3572](#)

## BGP Path Attribute Configuration

---

- [Example: Configuring BGP Local Preference on page 3310](#)
- [Examples: Configuring BGP MED on page 3323](#)
- [Examples: Configuring BGP Local AS on page 3362](#)
- [Example: Configuring the Accumulated IGP Attribute for BGP on page 3382](#)

### Example: Configuring BGP Local Preference

- [Understanding the BGP Local Preference on page 3310](#)
- [Example: Configuring the Local Preference Value for BGP Routes on page 3310](#)

#### Understanding the BGP Local Preference

---

Internal BGP (IBGP) sessions use a metric called the *local preference*, which is carried in IBGP update packets in the path attribute LOCAL\_PREF. When an autonomous system (AS) has multiple routes to another AS, the local preference indicates the degree of preference for one route over the other routes. The route with the highest local preference value is preferred.

The LOCAL\_PREF path attribute is always advertised to IBGP peers and to neighboring confederations. It is never advertised to external BGP (EBGP) peers. The default behavior is to not modify the LOCAL\_PREF path attribute if it is present.

The LOCAL\_PREF path attribute applies at export time only, when the routes are exported from the routing table into BGP.

If a BGP route is received without a LOCAL\_PREF attribute, the route is stored in the routing table and advertised by BGP as if it were received with a LOCAL\_PREF value of 100. A non-BGP route that is advertised by BGP is advertised with a LOCAL\_PREF value of 100 by default.

#### Example: Configuring the Local Preference Value for BGP Routes

---

This example shows how to configure local preference in internal BGP (IBGP) peer sessions.

- [Requirements on page 3310](#)
- [Overview on page 3310](#)
- [Configuration on page 3311](#)
- [Verification on page 3321](#)

##### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

##### **Overview**


To change the local preference metric advertised in the path attribute, you must include the **local-preference** statement, specifying a value from 0 through 4,294,967,295 ( $2^{32} - 1$ ).



There are several reasons you might want to prefer one path over another. For example, compared to other paths, one path might be less expensive to use, might have higher bandwidth, or might be more stable.

Figure 66 on page 3311 shows a typical network with internal peer sessions and multiple exit points to a neighboring AS.

Figure 66: Typical Network with IBGP Sessions and Multiple Exit Points



To reach Device R4, Device R1 can take a path through either Device R2 or Device R3. By default, the local preference is 100 for either route. When the local preferences are equal, Junos OS has rules for breaking the tie and choosing a path. (See *Understanding BGP Path Selection*.) In this example, the active route is through Device R2 because the router ID of Device R2 is lower than the router ID of Device R3. The following example shows how to override the default behavior with an explicit setting for the local preference. The example configures a local preference of 300 on Device R3, thereby making Device R3 the preferred path to reach Device R4.

### Configuration

- [Configuring Device R1 on page 3313](#)
- [Configuring Device R2 on page 3315](#)
- [Configuring Device R3 on page 3317](#)
- [Configuring Device R4 on page 3319](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
```

```
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

Device R2

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```

Device R3

```
set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1
```

Device R4

```
set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
```

```

set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3
set protocols bgp group external neighbor 24.24.24.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.
 

```

[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24

[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24

[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32

```
2. Configure BGP.
 

```

[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1

```
3. Configure OSPF.
 

```

[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2

```
4. Configure a policy that accepts direct routes.



**NOTE:** Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct

```

```
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
```

```
user@R1# set autonomous-system 123
```

```
user@R1# set router-id 192.168.1.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
```

```
fe-1/2/0 {
```

```
  unit 1 {
```

```
    family inet {
```

```
      address 12.12.12.1/24;
```

```
    }
```

```
  }
```

```
fe-1/2/1 {
```

```
  unit 2 {
```

```
    family inet {
```

```
      address 13.13.13.1/24;
```

```
    }
```

```
  }
```

```
}
```

```
lo0 {
```

```
  unit 1 {
```

```
    family inet {
```

```
      address 192.168.1.1/32;
```

```
    }
```

```
  }
```

```
}
```

```
user@R1# show policy-options
```

```
policy-statement send-direct {
```

```
  term 1 {
```

```
    from protocol direct;
```

```
    then accept;
```

```
  }
```

```
}
```

```
user@R1# show protocols
```

```
bgp {
```

```
  group internal {
```

```
    type internal;
```

```
    local-address 192.168.1.1;
```

```
    export send-direct;
```

```
    neighbor 192.168.2.1;
```

```
    neighbor 192.168.3.1;
```

```
  }
```

```
}
```

```
ospf {
```

```
  area 0.0.0.0 {
```

```

interface lo0.1 {
    passive;
}
interface fe-1/2/0.1;
interface fe-1/2/1.2;
}
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.
 

```

[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24

[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24

[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32

```
2. Configure BGP.
 

```

[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1

[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4

```
3. Configure OSPF.
 

```

[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4

```
4. Configure a policy that accepts direct routes.



**NOTE:** Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.1/32;
    }
  }
}

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R2# show protocols
bgp {
  group internal {
    type internal;
  }
}
```

```

    local-address 192.168.2.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.3.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 24.24.24.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
    interface fe-1/2/0.3;
    interface fe-1/2/1.4;
  }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R3

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.
 

```

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24

[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32

```
2. Configure BGP.
 

```

[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1

```

```
[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.



**NOTE:** Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R3# set from protocol direct
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 123
user@R3# set router-id 192.168.3.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
  unit 5 {
    family inet {
      address 13.13.13.3/24;
    }
  }
}
fe-1/2/1 {
  unit 6 {
    family inet {
      address 34.34.34.3/24;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.3.1/32;
    }
  }
}
```



```

}

user@R3# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R3# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.3.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.2.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 34.34.34.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.3 {
      passive;
    }
    interface fe-1/2/0.5;
    interface fe-1/2/1.6;
  }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.
 

```

[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

```

```
[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
```

2. Configure BGP.

```
[edit protocols bgp group external]
user@R4# set type external
user@R4# set export send-direct
user@R4# set peer-as 123
user@R4# set neighbor 34.34.34.3
user@R4# set neighbor 24.24.24.2
```

3. Configure a policy that accepts direct routes.



**NOTE:** Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept
```

4. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 7 {
    family inet {
      address 24.24.24.4/24;
    }
  }
}
fe-1/2/1 {
  unit 8 {
    family inet {
      address 34.34.34.4/24;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.4.1/32;
    }
  }
}
```

```

user@R4# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R4# show protocols
bgp {
  group external {
    type external;
    export send-direct;
    peer-as 123;
    neighbor 34.34.34.3;
    neighbor 24.24.24.2;
  }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3321](#)
- [Altering the Local Preference to Change the Path Selection on page 3322](#)
- [Rechecking the Active Path From Device R1 to Device R4 on page 3322](#)

### Checking the Active Path From Device R1 to Device R4

**Purpose** Verify that the active path from Device R1 to Device R4 goes through Device R2.

**Action** From operational mode, enter the **show route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 11 destinations, 18 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32     [BGP/170] 00:11:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1

```

```
192.168.3.1/32      [BGP/170] 00:11:48, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32      *[BGP/170] 00:05:14, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
                  [BGP/170] 00:05:14, localpref 100, from 192.168.3.1
                  AS path: 4 I
                  > to 13.13.13.3 via fe-1/2/1.2
```

**Meaning** The asterisk (\*) shows that the preferred path is through Device R2. In the default configuration, Device R2 has a lower router ID than Device R3. The router ID is controlling the path selection.

#### *Altering the Local Preference to Change the Path Selection*

**Purpose** Change the path so that it goes through Device R3.

**Action** From configuration mode, enter the **set local-preference 300** command.

```
[edit protocols bgp group internal]
user@R3# set local-preference 300
user@R3# commit
```

#### *Rechecking the Active Path From Device R1 to Device R4*

**Purpose** Verify that the active path from Device R1 to Device R4 goes through Device R3.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 11 destinations, 17 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32      [BGP/170] 00:16:48, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32      [BGP/170] 00:00:22, localpref 300, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32      *[BGP/170] 00:00:21, localpref 300, from 192.168.3.1
                  AS path: 4 I
                  > to 13.13.13.3 via fe-1/2/1.2
```

**Meaning** The asterisk (\*) shows that the preferred path is through Device R3. In the altered configuration, Device R3 has a higher local preference than Device R2. The local preference is controlling the path selection.

**Related Documentation**

- [Examples: Configuring Internal BGP Peering on page 3284](#)
- [BGP Configuration Overview](#)

## Examples: Configuring BGP MED

- [Understanding the MED Attribute on page 3323](#)
- [Example: Configuring the MED Attribute Directly on page 3325](#)
- [Example: Configuring the MED Using Route Filters on page 3338](#)
- [Example: Configuring the MED Using Communities on page 3351](#)
- [Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3352](#)

### Understanding the MED Attribute

The BGP multiple exit discriminator (MED, or MULTI\_EXIT\_DISC) is a non-transitive attribute, meaning that it is not propagated throughout the Internet, but only to adjacent autonomous systems (ASs). The MED attribute is optional, meaning that it is not always sent with the BGP updates. The purpose of MED is to influence how other ASs enter your AS to reach a certain prefix.

The MED attribute has a value that is referred to as a *metric*. If all other factors in determining an exit point are equal, the exit point with the lowest metric is preferred.

If a MED is received over an external BGP link, it is propagated over internal links to other BGP-enabled devices within the AS.

BGP update messages include a MED metric if the route was learned from BGP and already had a MED metric associated with it, or if you configure the MED metric in the configuration file.


A MED metric is advertised with a route according to the following general rules:

- A more specific metric overrides a less specific metric. That is, a group-specific metric overrides a global BGP metric, and a peer-specific metric overrides a global BGP or group-specific metric.
- A metric defined with a routing policy overrides a metric defined with the **metric-out** statement.
- If any metric is defined, it overrides a metric received in a route.
- If the received route does not have an associated MED metric, and if you do not explicitly configure a metric value, no metric is advertised. When you do not explicitly configure a metric value, the MED value is equivalent to zero (0) when advertising an active route.

Because the AS path rather than the number of hops between hosts is the primary criterion for BGP route selection, an AS with multiple connections to a peer AS can have multiple equivalent AS paths. When the routing table contains two routes to the same host in a neighboring AS, a MED metric assigned to each route can determine which to include in the forwarding table. The MED metric you assign can force traffic through a particular exit point in an AS.

[Figure 67 on page 3324](#) illustrates how MED metrics are used to determine route selection.

**Figure 67: Default MED Example**



[Figure 67 on page 3324](#) shows AS 1 and AS 2 connected by two separate BGP links to Routers C and D. Host E in AS 1 is located nearer to Router C. Host F, also in AS 1, is located nearer to Router D. Because the AS paths are equivalent, two routes exist for each host, one through Router C and one through Router D. To force all traffic destined for Host E through Router C, the network administrator for AS 1 assigns a MED metric for each router to Host E at its exit point. A MED metric of 10 is assigned to the route to Host E through Router C, and a MED metric of 20 is assigned to the route to Host E through Router D. BGP routers in AS 2 select the route with the lower MED metric for the forwarding table.

By default, only the MEDs of routes that have the same peer ASs are compared. However, you can configure the routing table path selection options listed in [Table 295 on page 3325](#) to compare MEDs in different ways. The MED options are not mutually exclusive and can be configured in combination or independently. For the MED options to take effect, you must configure them uniformly all through your network. The MED option or options you configure determine the route selected. Thus we recommend that you carefully evaluate your network for preferred routes before configuring the MED options.

Table 295: MED Options for Routing Table Path Selection

| Option (Name)   | Function   | Use  |
|---|--|--|
| Always comparing MEDs<br>( <b>always-compare-med</b> )                          | Ensures that the MEDs for paths from peers in different ASs are always compared in the route selection process.  | Useful when all enterprises participating in a network agree on a uniform policy for setting MEDs. For example, in a network shared by two ISPs, both must agree that a certain path is the better path to configure the MED values correctly. |
| Adding IGP cost to MED ( <b>med-plus-igp</b> )                                  | <p>Before comparing MED values for path selection, adds to the MED the cost of the IGP route to the BGP next-hop destination.</p> <p>This option replaces the MED value for the router, but does not affect the IGP metric comparison. As a result, when multiple routes have the same value after the MED-plus-IGP comparison, and route selection continues, the IGP route metric is also compared, even though it was added to the MED value and compared earlier in the selection process.</p>                                   | Useful when the downstream AS requires the complete cost of a certain route that is received across multiple ASs.  |
| Applying Cisco IOS nondeterministic behavior ( <b>cisco-non-deterministic</b> ) | <p>Specifies the nondeterministic behavior of the Cisco IOS software:</p> <ul style="list-style-type: none"> <li>The active path is always first. All nonactive but eligible paths follow the active path and are maintained in the order in which they were received. Ineligible paths remain at the end of the list.</li> <li>When a new path is added to the routing table, path comparisons are made among all routes, including those paths that must never be selected because they lose the MED tie-breaking rule.</li> </ul> | We recommend that you do not configure this option, because the nondeterministic behavior sometimes prevents the system from properly comparing the MEDs between paths.  |

### Example: Configuring the MED Attribute Directly

This example shows how to configure a multiple exit discriminator (MED) metric to advertise in BGP update messages.

- [Requirements on page 3325](#)
- [Overview on page 3326](#)
- [Configuration on page 3327](#)
- [Verification on page 3337](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.

### Overview

To directly configure a MED metric to advertise in BGP update messages, include the **metric-out** statement:

**metric-out** (*metric* | **minimum-igp** *offset* | **igp** **delay-med-update** | *offset*);

**metric** is the primary metric on all routes sent to peers. It can be a value in the range from 0 through 4,294,967,295 ( $2^{32} - 1$ ).

The following optional settings are also supported:


- **minimum-igp**—Sets the metric to the minimum metric value calculated in the interior gateway protocol (IGP) to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value.
- **igp**—Sets the metric to the most recent metric value calculated in the IGP to get to the BGP next hop.
- **delay-med-update**—Delays sending MED updates when the MED value increases. Include the **delay-med-update** statement when you configure the **igp** statement. The default interval to delay sending updates, unless the MED is lower or another attribute associated with the route has changed is 10 minutes. Include the **med-igp-update-interval** *minutes* statement at the **[edit routing-options]** hierarchy level to modify the default interval.
- **offset**—Specifies a value for **offset** to increase or decrease the metric that is used from the metric value calculated in the IGP. The metric value is offset by the value specified. The metric calculated in the IGP (by specifying either **igp** or **igp-minimum**) is increased if the **offset** value is positive. The metric calculated in the IGP (by specifying either **igp** or **igp-minimum**) is decreased if the **offset** value is negative.

**offset** can be a value in the range from  $-2^{31}$  through  $2^{31} - 1$ . Note that the adjusted metric can never go below 0 or above  $2^{32} - 1$ .

Figure 68 on page 3327 shows a typical network with internal peer sessions and multiple exit points to a neighboring autonomous system (AS).



Figure 68: Typical Network with IBGP Sessions and Multiple Exit Points



Device R4 has multiple loopback interfaces configured to simulate advertised prefixes. The extra loopback interface addresses are 44.44.44.44/32 and 144.144.144.144/32. This example shows how to configure Device R4 to advertise a MED value of 30 to Device R3 and a MED value of 20 to Device R2. This causes all of the devices in AS 123 to prefer the path through Device R2 to reach AS 4.

### Configuration

- [Configuring Device R1 on page 3329](#)
- [Configuring Device R2 on page 3331](#)
- [Configuring Device R3 on page 3333](#)
- [Configuring Device R4 on page 3335](#)

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept

```

```
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

**Device R2**

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```

**Device R3**

```
set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1
```

**Device R4**

```
set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set interfaces lo0 unit 4 family inet address 44.44.44.44/32
set interfaces lo0 unit 4 family inet address 144.144.144.144/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3 metric-out 30
set protocols bgp group external neighbor 24.24.24.2 metric-out 20
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
```

```
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1
```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24
```

```
[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 12.12.12.1/24;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.1.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.1.1;
    export send-direct;
    neighbor 192.168.2.1;
    neighbor 192.168.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/1.2;
  }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;
```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24
```

```
[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24
```

```
[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32
```

2. Configure BGP.

```
[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1
```

```
[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4
```

3. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.1/32;
    }
  }
}

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R2# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.2.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.3.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 24.24.24.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
  }
}
```

```

    }
    interface fe-1/2/0.3;
    interface fe-1/2/1.4;
  }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R3

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.

```

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

```

```

[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24

```

```

[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1

```

```

[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5
user@R3# set interface fe-1/2/1.6

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R3# set from protocol direct
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R3# set autonomous-system 123
user@R3# set router-id 192.168.3.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
  unit 5 {
    family inet {
      address 13.13.13.3/24;
    }
  }
}
fe-1/2/1 {
  unit 6 {
    family inet {
      address 34.34.34.3/24;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.3.1/32;
    }
  }
}

user@R3# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R3# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.3.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.2.1;
  }
  group external {
```



```

        type external;
        export send-direct;
        peer-as 4;
        neighbor 34.34.34.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.3 {
            passive;
        }
        interface fe-1/2/0.5;
        interface fe-1/2/1.6;
    }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```

[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

```

```

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

```

```

[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
user@R4# set family inet address 44.44.44.44/32
user@R4# set family inet address 144.144.144.144/32

```

Device R4 has multiple loopback interface addresses to simulate advertised prefixes.

2. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R4# set from protocol direct
user@R4# set then accept

```

3. Configure BGP.

```

[edit protocols bgp group external]
user@R4# set type external

```

```
user@R4# set export send-direct
user@R4# set peer-as 123
```

4. Configure a MED value of 30 for neighbor Device R3, and a MED value of 20 for neighbor Device R2.

```
[edit protocols bgp group external]
user@R4# set neighbor 34.34.34.3 metric-out 30
user@R4# set neighbor 24.24.24.2 metric-out 20
```

This configuration causes autonomous system (AS) 123 (of which Device R1, Device R2, and Device R3 are members) to prefer the path through Device R2 to reach AS 4.

5. Configure the router ID and AS number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 7 {
    family inet {
      address 24.24.24.4/24;
    }
  }
}
fe-1/2/1 {
  unit 8 {
    family inet {
      address 34.34.34.4/24;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.4.1/32;
      address 44.44.44.44/32;
      address 144.144.144.144/32;
    }
  }
}

user@R4# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
```

```

user@R4# show protocols
bgp {
  group external {
    type external;
    export send-direct;
    peer-as 123;
    neighbor 34.34.34.3 {
      metric-out 30;
    }
    neighbor 24.24.24.2 {
      metric-out 20;
    }
  }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the Active Path From Device R1 to Device R4 on page 3337](#)
- [Verifying That Device R4 Is Sending Its Routes Correctly on page 3338](#)

### Checking the Active Path From Device R1 to Device R4

**Purpose** Verify that the active path goes through Device R2.

**Action** From operational mode, enter the **show route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
44.44.44.44/32     *[BGP/170] 01:41:11, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
144.144.144.144/32 *[BGP/170] 00:08:13, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.2.1/32     [BGP/170] 3d 22:52:38, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1

```

```

192.168.3.1/32      [BGP/170] 3d 03:15:16, localpref 100, from 192.168.3.1
                   AS path: I
                   > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32      *[BGP/170] 01:41:11, MED 20, localpref 100, from 192.168.2.1
                   AS path: 4 I
                   > to 12.12.12.2 via fe-1/2/0.1

```

**Meaning** The asterisk (\*) shows that the preferred path is through Device R2. The reason for the path selection is listed as MED 20.

### *Verifying That Device R4 Is Sending Its Routes Correctly*

**Purpose** Make sure that Device R4 is sending update messages with a value of 20 to Device R2 and a value of 30 to Device R3.

**Action** From operational mode, enter the **show route advertising-protocol bgp 24.24.24.2** command.

```

user@R4> show route advertising-protocol bgp 24.24.24.2
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop      MED      Lclpref  AS path
* 24.24.24.0/24         Self                20              I
* 34.34.34.0/24         Self                20              I
* 44.44.44.44/32        Self                20              I
* 144.144.144.144/32    Self                20              I
* 192.168.4.1/32        Self                20              I

```

```

user@R4> show route advertising-protocol bgp 34.34.34.3
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop      MED      Lclpref  AS path
* 24.24.24.0/24         Self                30              I
* 34.34.34.0/24         Self                30              I
* 44.44.44.44/32        Self                30              I
* 144.144.144.144/32    Self                30              I
* 192.168.4.1/32        Self                30              I

```

**Meaning** The MED column shows that Device R4 is sending the correct MED values to its two external BGP (EBGP) neighbors.

### Example: Configuring the MED Using Route Filters

This example shows how to configure a policy that uses route filters to modify the multiple exit discriminator (MED) metric to advertise in BGP update messages.

- [Requirements on page 3338](#)
- [Overview on page 3339](#)
- [Configuration on page 3339](#)
- [Verification on page 3350](#)

#### **Requirements**


No special configuration beyond device initialization is required before you configure this example.

### Overview

To configure a route-filter policy that modifies the advertised MED metric in BGP update messages, include the **metric** statement in the policy action.

Figure 69 on page 3339 shows a typical network with internal peer sessions and multiple exit points to a neighboring autonomous system (AS).

Figure 69: Typical Network with IBGP Sessions and Multiple Exit Points



Device R4 has multiple loopback interfaces configured to simulate advertised prefixes. The extra loopback interface addresses are 44.44.44.44/32 and 144.144.144.144/32. This example shows how to configure Device R4 to advertise a MED value of 30 to Device R3 for all routes except 144.144.144.144. For 144.144.144.144, a MED value of 10 is advertised to Device 3. A MED value of 20 is advertised to Device R2, regardless of the route prefix.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 12.12.12.1/24
set interfaces fe-1/2/1 unit 2 family inet address 13.13.13.1/24
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.1.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.2

```

```
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.1.1
```

**Device R2**

```
set interfaces fe-1/2/0 unit 3 family inet address 12.12.12.2/24
set interfaces fe-1/2/1 unit 4 family inet address 24.24.24.2/24
set interfaces lo0 unit 2 family inet address 192.168.2.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.2.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.3.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 24.24.24.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.2.1
```

**Device R3**

```
set interfaces fe-1/2/0 unit 5 family inet address 13.13.13.3/24
set interfaces fe-1/2/1 unit 6 family inet address 34.34.34.3/24
set interfaces lo0 unit 3 family inet address 192.168.3.1/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.3.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.1.1
set protocols bgp group internal neighbor 192.168.2.1
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 4
set protocols bgp group external neighbor 34.34.34.4
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.5
set protocols ospf area 0.0.0.0 interface fe-1/2/1.6
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 123
set routing-options router-id 192.168.3.1
```

**Device R4**

```
set interfaces fe-1/2/0 unit 7 family inet address 24.24.24.4/24
set interfaces fe-1/2/1 unit 8 family inet address 34.34.34.4/24
set interfaces lo0 unit 4 family inet address 192.168.4.1/32
set interfaces lo0 unit 4 family inet address 44.44.44.44/32
set interfaces lo0 unit 4 family inet address 144.144.144.144/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 123
set protocols bgp group external neighbor 34.34.34.3 export med-10
set protocols bgp group external neighbor 34.34.34.3 export med-30
```

```

set protocols bgp group external neighbor 24.24.24.2 metric-out 20
set policy-options policy-statement med-10 from route-filter 144.144.144.144/32 exact
set policy-options policy-statement med-10 then metric 10
set policy-options policy-statement med-10 then accept
set policy-options policy-statement med-30 from route-filter 0.0.0.0/0 longer
set policy-options policy-statement med-30 then metric 30
set policy-options policy-statement med-30 then accept
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 4
set routing-options router-id 192.168.4.1

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 1]
user@R1# set family inet address 12.12.12.1/24

```

```

[edit interfaces fe-1/2/1 unit 2]
user@R1# set family inet address 13.13.13.1/24

```

```

[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.1.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.1.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.2.1
user@R1# set neighbor 192.168.3.1

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R1# set interface lo0.1 passive
user@R1# set interface fe-1/2/0.1
user@R1# set interface fe-1/2/1.2

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept

```

5. Configure the router ID and autonomous system (AS) number.

```

[edit routing-options]

```

```
user@R1# set autonomous-system 123
user@R1# set router-id 192.168.1.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 12.12.12.1/24;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 13.13.13.1/24;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.1.1;
    export send-direct;
    neighbor 192.168.2.1;
    neighbor 192.168.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/1.2;
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
```



```

    }
}

user@R1# show routing-options
autonomous-system 123;
router-id 192.168.1.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 3]
user@R2# set family inet address 12.12.12.21/24

```

```

[edit interfaces fe-1/2/1 unit 4]
user@R2# set family inet address 24.24.24.2/24

```

```

[edit interfaces lo0 unit 2]
user@R2# set family inet address 192.168.2.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R2# set type internal
user@R2# set local-address 192.168.2.1
user@R2# set export send-direct
user@R2# set neighbor 192.168.1.1
user@R2# set neighbor 192.168.3.1

```

```

[edit protocols bgp group external]
user@R2# set type external
user@R2# set export send-direct
user@R2# set peer-as 4
user@R2# set neighbor 24.24.24.4

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R2# set interface lo0.2 passive
user@R2# set interface fe-1/2/0.3
user@R2# set interface fe-1/2/1.4

```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept

```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 123
user@R2# set router-id 192.168.2.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 12.12.12.2/24;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 24.24.24.2/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.1/32;
    }
  }
}

user@R2# show protocols
bgp {
  group internal {
    type internal;
    local-address 192.168.2.1;
    export send-direct;
    neighbor 192.168.1.1;
    neighbor 192.168.3.1;
  }
  group external {
    type external;
    export send-direct;
    peer-as 4;
    neighbor 24.24.24.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
    interface fe-1/2/0.3;
```

```

        interface fe-1/2/1.4;
    }
}

user@R2# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}

user@R2# show routing-options
autonomous-system 123;
router-id 192.168.2.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R3

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 5]
user@R3# set family inet address 13.13.13.3/24

```

```

[edit interfaces fe-1/2/1 unit 6]
user@R3# set family inet address 34.34.34.3/24

```

```

[edit interfaces lo0 unit 3]
user@R3# set family inet address 192.168.3.1/32

```

2. Configure BGP.

```

[edit protocols bgp group internal]
user@R3# set type internal
user@R3# set local-address 192.168.3.1
user@R3# set export send-direct
user@R3# set neighbor 192.168.1.1
user@R3# set neighbor 192.168.2.1

```

```

[edit protocols bgp group external]
user@R3# set type external
user@R3# set export send-direct
user@R3# set peer-as 4
user@R3# set neighbor 34.34.34.4

```

3. Configure OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0.3 passive
user@R3# set interface fe-1/2/0.5

```

```
user@R3# set interface fe-1/2/1.6
```

4. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
```

```
user@R3# set from protocol direct
```

```
user@R3# set then accept
```

5. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
```

```
user@R3# set autonomous-system 123
```

```
user@R3# set router-id 192.168.3.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
```

```
fe-1/2/0 {  
  unit 5 {  
    family inet {  
      address 13.13.13.3/24;  
    }  
  }  
}  
fe-1/2/1 {  
  unit 6 {  
    family inet {  
      address 34.34.34.3/24;  
    }  
  }  
}  
lo0 {  
  unit 3 {  
    family inet {  
      address 192.168.3.1/32;  
    }  
  }  
}
```

```
user@R3# show protocols
```

```
bgp {  
  group internal {  
    type internal;  
    local-address 192.168.3.1;  
    export send-direct;  
    neighbor 192.168.1.1;  
    neighbor 192.168.2.1;  
  }  
  group external {  
    type external;  
    export send-direct;  
  }  
}
```

```

        peer-as 4;
        neighbor 34.34.34.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.3 {
            passive;
        }
        interface fe-1/2/0.5;
        interface fe-1/2/1.6;
    }
}

user@R3# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}

user@R3# show routing-options
autonomous-system 123;
router-id 192.168.3.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the device interfaces.

```

[edit interfaces fe-1/2/0 unit 7]
user@R4# set family inet address 24.24.24.4/24

```

```

[edit interfaces fe-1/2/1 unit 8]
user@R4# set family inet address 34.34.34.4/24

```

```

[edit interfaces lo0 unit 4]
user@R4# set family inet address 192.168.4.1/32
user@R4# set family inet address 44.44.44.44/32
user@R4# set family inet address 144.144.144.144/32

```

Device R4 has multiple loopback interface addresses to simulate advertised prefixes.

2. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-direct term 1]

```

```
user@R4# set from protocol direct
user@R4# set then accept
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@R4# set type external
user@R4# set export send-direct
user@R4# set peer-as 123
```

4. Configure the two MED policies.

```
[edit policy-options]
set policy-statement med-10 from route-filter 144.144.144.144/32 exact
set policy-statement med-10 then metric 10
set policy-statement med-10 then accept
```

```
set policy-statement med-30 from route-filter 0.0.0.0/0 longer
set policy-statement med-30 then metric 30
set policy-statement med-30 then accept
```

5. Configure the two EBGP neighbors, applying the two MED policies to Device R3, and a MED value of 20 to Device R2.

```
[edit protocols bgp group external]
user@R4# set neighbor 34.34.34.3 export med-10
user@R4# set neighbor 34.34.34.3 export med-30
user@R4# set neighbor 24.24.24.2 metric-out 20
```

6. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R4# set autonomous-system 4
user@R4# set router-id 192.168.4.1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 7 {
    family inet {
      address 24.24.24.4/24;
    }
  }
}
fe-1/2/1 {
  unit 8 {
    family inet {
      address 34.34.34.4/24;
    }
  }
}
lo0 {
  unit 4 {
```

```

        family inet {
            address 192.168.4.1/32;
            address 44.44.44.44/32;
            address 144.144.144.144/32;
        }
    }
}

user@R4# show protocols
bgp {
    group external {
        type external;
        export send-direct;
        peer-as 123;
        neighbor 24.24.24.2 {
            metric-out 20;
        }
        neighbor 34.34.34.3 {
            export [ med-10 med-30 ];
        }
    }
}

user@R4# show policy-options
policy-statement med-10 {
    from {
        route-filter 144.144.144.144/32 exact;
    }
    then {
        metric 10;
        accept;
    }
}
policy-statement med-30 {
    from {
        route-filter 0.0.0.0/0 longer;
    }
    then {
        metric 30;
        accept;
    }
}
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}

user@R4# show routing-options
autonomous-system 4;
router-id 192.168.4.1;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the Active Path from Device R1 to Device R4 on page 3350](#)
- [Verifying That Device R4 Is Sending Its Routes Correctly on page 3350](#)

### Checking the Active Path from Device R1 to Device R4

**Purpose** Verify that the active path goes through Device R2.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R1> show route protocol bgp
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

12.12.12.0/24      [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
13.13.13.0/24      [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
24.24.24.0/24      [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
34.34.34.0/24      [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
44.44.44.44/32     *[BGP/170] 00:06:03, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
144.144.144.144/32 *[BGP/170] 00:06:03, MED 10, localpref 100, from 192.168.3.1
                  AS path: 4 I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.2.1/32      [BGP/170] 4d 01:13:32, localpref 100, from 192.168.2.1
                  AS path: I
                  > to 12.12.12.2 via fe-1/2/0.1
192.168.3.1/32      [BGP/170] 3d 05:36:10, localpref 100, from 192.168.3.1
                  AS path: I
                  > to 13.13.13.3 via fe-1/2/1.2
192.168.4.1/32     *[BGP/170] 00:06:03, MED 20, localpref 100, from 192.168.2.1
                  AS path: 4 I
                  > to 12.12.12.2 via fe-1/2/0.1
```

**Meaning** The output shows that the preferred path to the routes advertised by Device R4 is through Device R2 for all routes except 144.144.144.144/32. For 144.144.144.144/32, the preferred path is through Device R3.

### Verifying That Device R4 Is Sending Its Routes Correctly

**Purpose** Make sure that Device R4 is sending update messages with a value of 20 to Device R2 and a value of 30 to Device R3.

**Action** From operational mode, enter the **show route advertising-protocol bgp** command.

```
user@R4> show route advertising-protocol bgp 24.24.24.2
```



```
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref   AS path
* 24.24.24.0/24         Self              20                I
* 34.34.34.0/24         Self              20                I
* 44.44.44.44/32        Self              20                I
* 144.144.144.144/32    Self              20                I
* 192.168.4.1/32        Self              20                I
```

```
user@R4> show route advertising-protocol bgp 34.34.34.3
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref   AS path
* 24.24.24.0/24         Self              30                I
* 34.34.34.0/24         Self              30                I
* 44.44.44.44/32        Self              30                I
* 144.144.144.144/32    Self              10                I
* 192.168.4.1/32        Self              30                I
```

**Meaning** The MED column shows that Device R4 is sending the correct MED values to its two EBGp neighbors.

### Example: Configuring the MED Using Communities

Set the multiple exit discriminator (MED) metric to 20 for all routes from a particular community.

```
[edit]
routing-options {
  router-id 10.0.0.1;
  autonomous-system 23;
}
policy-options {
  policy-statement from-otago {
    from community otago;
    then metric 20;
  }
  community otago members [56:2379 23:46944];
}
protocols {
  bgp {
    import from-otago;
    group 23 {
      type external;
      peer-as 56;
      neighbor 192.168.0.1 {
        traceoptions {
          file bgp-log-peer;
          flag packets;
        }
        log-updown;
      }
    }
  }
}
```

### Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates

---

This example shows how to associate the multiple exit discriminator (MED) path attribute with the interior gateway protocol (IGP) metric, and configure a timer to delay update of the MED attribute.

- [Requirements on page 3352](#)
- [Overview on page 3352](#)
- [Configuration on page 3354](#)
- [Verification on page 3360](#)

#### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

#### **Overview**

BGP can be configured to advertise the MED attribute for a route based on the IGP distance of its internal BGP (IBGP) route next-hop. The IGP metric enables internal routing to follow the shortest path according to the administrative setup. In some deployments, it might be ideal to communicate IGP shortest-path knowledge to external BGP (EBGP) peers in a neighboring autonomous system (AS). This allows those EBGP peers to forward traffic into your AS using the shortest paths possible.

Routes learned from an EBGP peer usually have a next hop on a directly connected interface, and thus the IGP value is equal to zero. Zero is the value advertised. The IGP metric is a nonzero value when a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the **multihop** command. In these scenarios, it might make sense to associate the MED value with the IGP metric by including the **metric-out minimum-igp** or **metric-out igp** option.

The drawback of associating the MED with the IGP metric is the risk of excessive route advertisements when there are IGP instabilities in the network. Configuring a delay for the MED update provides a mechanism to reduce route advertisements in such scenarios. The delay works by slowing down MED updates when the IGP metric for the next hop changes. The approach uses a timer to periodically advertise MED updates. When the timer expires, the MED attribute for routes with **metric-out igp delay-updates** configured is updated to the current IGP metric of the next hop. The BGP-enabled device sends out advertisements for routes for which the MED attribute has changed.

The **delay-updates** option identifies the BGP groups (or peers) for which the MED updates must be suppressed. The time for advertising MED updates is set to 10 minutes by default. You can increase the interval up to 600 minutes by including the **med-igp-update-interval** statement in the **routing-options** configuration.



**NOTE:** If you have nonstop active routing (NSR) enabled and a switchover occurs, the delayed MED updates might be advertised as soon as the switchover occurs.


When you configure the **metric-out igp** option, the IGP metric directly tracks the IGP cost to the IBGP peer. When the IGP cost goes down, so does the advertised MED value. Conversely, when the IGP cost goes up, the MED value goes up as well.

When you configure the **metric-out minimum-igp** option, the advertised MED value changes only when the IGP cost to the IBGP peer goes down. An increase in the IGP cost does not affect the MED value. The router monitors and remembers the lowest IGP cost until the routing process (rpd) is restarted. The BGP peer sends an update only if the MED is lower than the previously advertised value or another attribute associated with the route has changed, or if the BGP peer is responding to a refresh route request.

This example uses the **metric** statement in the OSPF configuration to demonstrate that when the IGP metric changes, the MED also changes after the configured delay interval. The OSPF metric can range from 1 through 65,535.

[Figure 70 on page 3354](#) shows the sample topology.

Figure 70: Topology for Delaying the MED Update



In this example, the MED value advertised by Device R1 is associated with the IGP running in AS 1. The MED value advertised by Device R1 impacts the decisions of the neighboring AS (AS 2) when AS 2 is forwarding traffic into AS 1.

#### Configuration

- [Configuring Device R1 on page 3358](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1 set interfaces fe-1/2/0 unit 2 description R1->R2
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.1/30
set interfaces fe-1/2/1 unit 7 description R1->R4
set interfaces fe-1/2/1 unit 7 family inet address 172.16.0.1/30
set interfaces lo0 unit 1 family inet address 192.168.0.1/32
set protocols bgp group internal type internal
```

```

set protocols bgp group internal local-address 192.168.0.1
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.2
set protocols bgp group internal neighbor 192.168.0.3
set protocols bgp group external type external
set protocols bgp group external metric-out igp delay-med-update
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.2
set protocols ospf area 0.0.0.0 interface fe-1/2/0.2 metric 600
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options med-igp-update-interval 12
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 1

```

**Device R2**

```

set interfaces fe-1/2/0 unit 1 description R2->R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 4 description R2->R3
set interfaces fe-1/2/1 unit 4 family inet address 10.0.2.2/30
set interfaces lo0 unit 2 family inet address 192.168.0.2/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.2
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.1
set protocols bgp group internal neighbor 192.168.0.3
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 1

```

**Device R3**

```

set interfaces fe-1/2/0 unit 3 description R3->R2
set interfaces fe-1/2/0 unit 3 family inet address 10.0.2.1/30
set interfaces fe-1/2/1 unit 5 description R3->R5
set interfaces fe-1/2/1 unit 5 family inet address 172.16.0.5/30
set interfaces lo0 unit 3 family inet address 192.168.0.3/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.3
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.1
set protocols bgp group internal neighbor 192.168.0.2
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.6
set protocols ospf area 0.0.0.0 interface fe-1/2/0.3
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.3
set routing-options autonomous-system 1

```

**Device R4**

```
set interfaces fe-1/2/0 unit 8 description R4->R1
set interfaces fe-1/2/0 unit 8 family inet address 172.16.0.2/30
set interfaces fe-1/2/1 unit 9 description R4->R5
set interfaces fe-1/2/1 unit 9 family inet address 10.0.4.1/30
set interfaces fe-1/2/2 unit 13 description R4->R6
set interfaces fe-1/2/2 unit 13 family inet address 172.16.0.9/30
set interfaces lo0 unit 4 family inet address 192.168.0.4/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.4
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.5
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external neighbor 172.16.0.10 peer-as 3
set protocols bgp group external neighbor 172.16.0.1 peer-as 1
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.4
set routing-options autonomous-system 2
```

**Device R5**

```
set interfaces fe-1/2/0 unit 6 description R5->R3
set interfaces fe-1/2/0 unit 6 family inet address 172.16.0.6/30
set interfaces fe-1/2/1 unit 10 description R5->R4
set interfaces fe-1/2/1 unit 10 family inet address 10.0.4.2/30
set interfaces fe-1/2/2 unit 11 description R5->R8
set interfaces fe-1/2/2 unit 11 family inet address 172.16.0.13/30
set interfaces lo0 unit 5 family inet address 192.168.0.5/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.5
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.4
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external neighbor 172.16.0.5 peer-as 1
set protocols bgp group external neighbor 172.16.0.14 peer-as 3
set protocols ospf area 0.0.0.0 interface fe-1/2/1.10
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.5
set routing-options autonomous-system 2
```

**Device R6**

```
set interfaces fe-1/2/0 unit 14 description R6->R4
set interfaces fe-1/2/0 unit 14 family inet address 172.16.0.10/30
set interfaces fe-1/2/1 unit 15 description R6->R7
set interfaces fe-1/2/1 unit 15 family inet address 10.0.6.1/30
set interfaces lo0 unit 6 family inet address 192.168.0.6/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.6
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.7
set protocols bgp group internal neighbor 192.168.0.8
set protocols bgp group external type external
```

```

set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.9 peer-as 2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.15
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.6
set routing-options autonomous-system 3

```

**Device R7**

```

set interfaces fe-1/2/0 unit 16 description R7->R6
set interfaces fe-1/2/0 unit 16 family inet address 10.0.6.2/30
set interfaces fe-1/2/1 unit 17 description R7->R8
set interfaces fe-1/2/1 unit 17 family inet address 10.0.7.2/30
set interfaces lo0 unit 7 family inet address 192.168.0.7/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.7
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.6
set protocols bgp group internal neighbor 192.168.0.8
set protocols ospf area 0.0.0.0 interface fe-1/2/0.16
set protocols ospf area 0.0.0.0 interface fe-1/2/1.17
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.7
set routing-options autonomous-system 3

```

**Device R8**

```

set interfaces fe-1/2/0 unit 12 description R8->R5
set interfaces fe-1/2/0 unit 12 family inet address 172.16.0.14/30
set interfaces fe-1/2/1 unit 18 description R8->R7
set interfaces fe-1/2/1 unit 18 family inet address 10.0.7.1/30
set interfaces lo0 unit 8 family inet address 192.168.0.8/32
set protocols bgp group internal type internal
set protocols bgp group internal local-address 192.168.0.8
set protocols bgp group internal export send-direct
set protocols bgp group internal neighbor 192.168.0.6
set protocols bgp group internal neighbor 192.168.0.7
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 2
set protocols bgp group external neighbor 172.16.0.13 peer-as 2
set protocols ospf area 0.0.0.0 interface fe-1/2/1.18
set protocols ospf area 0.0.0.0 interface lo0.8 passive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options router-id 192.168.0.8
set routing-options autonomous-system 3

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```
[edit interfaces fe-1/2/0 unit 2]
user@R1# set description R1->R2
user@R1# set family inet address 10.0.0.1/30
```

```
[edit interfaces fe-1/2/1 unit 7]
user@R1# set description R1->R4
user@R1# set family inet address 172.16.0.1/30
```

```
[edit interfaces lo0 unit 1]
user@R1# set family inet address 192.168.0.1/32
```

2. Configure IBGP.

```
[edit protocols bgp group internal]
user@R1# set type internal
user@R1# set local-address 192.168.0.1
user@R1# set export send-direct
user@R1# set neighbor 192.168.0.2
user@R1# set neighbor 192.168.0.3
```

3. Configure EBGP.

```
[edit protocols bgp group external]
user@R1# set type external
user@R1# set export send-direct
user@R1# set peer-as 2
user@R1# set neighbor 172.16.0.2
```

4. Associate the MED value with the IGP metric.

```
[edit protocols bgp group external]
user@R1# set metric-out igp delay-med-update
```

The default for the MED update is 10 minutes when you include the **delay-med-update** option. When you exclude the **delay-med-update** option, the MED update occurs immediately after the IGP metric changes.

5. (Optional) Configure the update interval for the MED update.

```
[edit routing-options]
user@R1# set med-igp-update-interval 12
```

You can configure the interval from 10 minutes through 600 minutes.

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/0.2 metric 600
user@R1# set interface lo0.1 passive
```



The **metric** statement is used here to demonstrate what happens when the IGP metric changes.

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

8. Configure the router ID and autonomous system (AS) number.

```
[edit routing-options]
user@R1# set router-id 192.168.0.1
user@R1# set autonomous-system 1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 2 {
    description R1->R2;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
fe-1/2/1 {
  unit 7 {
    description R1->R4;
    family inet {
      address 172.16.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
```

```
bgp {
  group internal {
    type internal;
    local-address 192.168.0.1;
    export send-direct;
    neighbor 192.168.0.2;
    neighbor 192.168.0.3;
  }
  group external {
    type external;
    metric-out igp delay-med-update;
    export send-direct;
    peer-as 2;
    neighbor 172.16.0.2;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.2 {
      metric 600;
    }
    interface lo0.1 {
      passive;
    }
  }
}
```

```
user@R1# show routing-options
med-igp-update-interval 12;
router-id 192.168.0.1;
autonomous-system 1;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration steps on the other devices in the topology, as needed for your network.

### **Verification**

Confirm that the configuration is working properly.

- [Checking the BGP Advertisements on page 3360](#)
- [Verifying That the MED Value Changes When the OSPF Metric Changes on page 3361](#)
- [Testing the minimum-igp Setting on page 3361](#)

### **Checking the BGP Advertisements**

**Purpose** Verify that Device R1 is advertising to Device R4 a BGP MED value that reflects the IGP metric.

**Action** From operational mode, enter the **show route advertising-protocol bgp** command.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix                Nexthop        MED      Lc1pref  AS path
* 10.0.0.0/30           Self           0         I         I
* 172.16.0.0/30         Self           0         I         I
```

|                  |      |     |   |
|------------------|------|-----|---|
| * 172.16.0.4/30  | Self | 601 | I |
| * 192.168.0.1/32 | Self | 0   | I |

**Meaning** The 601 value in the MED column shows that the MED value has been updated to reflect the configured OSPF metric.

#### *Verifying That the MED Value Changes When the OSPF Metric Changes*

**Purpose** Make sure that when you raise the OSPF metric to 700, the MED value is updated to reflect this change.

**Action** From configuration mode, enter the **set protocols ospf area 0 interface fe-1/2/0.2 metric 700** command.

```
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 700
user@R1# commit
```

After waiting 12 minutes (the configured delay period), enter the **show route advertising-protocol bgp** command from operational mode.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix            Nexthop      MED      Lclpref  AS path
* 10.0.0.0/30       Self         0         I
* 172.16.0.0/30     Self         0         I
* 172.16.0.4/30     Self         701        I
* 192.168.0.1/32    Self         0         I
```

**Meaning** The 701 value in the MED column shows that the MED value has been updated to reflect the configured OSPF metric.

#### *Testing the minimum-igp Setting*

**Purpose** Change the configuration to use the **minimum-igp** statement instead of the **igp** statement. When you increase the OSPF metric, the MED value remains unchanged, but when you decrease the OSPF metric, the MED value reflects the new OSPF metric.

**Action** From configuration mode, delete the **igp** statement, add the **minimum-igp** statement, and increase the OSPF metric.

```
user@R1# delete protocols bgp group external metric-out igp
user@R1# set protocols bgp group external metric-out minimum-igp
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 800
user@R1# commit
```

From operational mode, enter the **show route advertising-protocol bgp** command to make sure that the MED value does not change.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix            Nexthop      MED      Lclpref  AS path
* 10.0.0.0/30       Self         0         I
* 172.16.0.0/30     Self         0         I
* 172.16.0.4/30     Self         701        I
* 192.168.0.1/32    Self         0         I
```

From configuration mode, decrease the OSPF metric.

```
user@R1# set protocols ospf area 0 interface fe-1/2/0.2 metric 20
user@R1# commit
```

From operational mode, enter the **show route advertising-protocol bgp** command to make sure that the MED value does change.

```
user@R1> show route advertising-protocol bgp 172.16.0.2
inet.0: 19 destinations, 33 routes (19 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.0.0.0/30           Self             0                I
* 172.16.0.0/30          Self             0                I
* 172.16.0.4/30          Self             21               I
* 192.168.0.1/32         Self             0                I
```

**Meaning** When the **minimum-igp** statement is configured, the MED value changes only when a shorter path is available.

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3261](#)
- [BGP Configuration Overview](#)

## Examples: Configuring BGP Local AS

- [Understanding the BGP Local AS Attribute on page 3362](#)
- [Example: Configuring a Local AS for EBGp Sessions on page 3367](#)
- [Example: Configuring a Private Local AS for EBGp Sessions on page 3377](#)

### Understanding the BGP Local AS Attribute

When an Internet service provider (ISP) acquires a network that belongs to a different autonomous system (AS), there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. Sometimes customers do not want to or are not immediately able to modify their peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a *local AS*.

Using a local AS number permits the routing devices in an acquired network to appear to belong to the former AS.

For example, ISP A, with an AS of 200, acquires ISP B, with an AS of 250. ISP B has a customer, ISP C, that does not want to change its configuration. After ISP B becomes part of ISP A, a local AS number of 250 is configured for use in EBGp peer sessions with ISP C. Consequently, the local AS number of 250 is either prepended before or used instead of the global AS number of 200 in the AS path used to export routes to direct external peers in ISP C.

If the route is received from an internal BGP (IBGP) peer, the AS path includes the local AS number prepended before the global AS number.

The local AS number is used instead of the global AS number if the route is an external route, such as a static route or an interior gateway protocol (IGP) route that is imported into BGP. If the route is external and you want the global AS number to be included in the AS path, you can apply a routing policy that uses **as-path-expand** or **as-path-prepend**. Use the **as-path-expand** policy action to place the global AS number behind the local AS number. Use the **as-path-prepend** policy action to place the global AS number in front of the local AS number.

For example:

```

user@R2# show policy-options
policy-statement prepend-global {
  term 1 {
    from protocol static;
    then {
      as-path-prepend 200; # or use as-path-expand
      accept;
    }
  }
}

user@R2# show protocols bgp
group ext {
  export prepend-global;
  type external;
  local-as 250;
  neighbor 10.0.0.1 {
    peer-as 100;
  }
  neighbor 10.1.0.2 {
    peer-as 300;
  }
}

user@R2# show routing-options
static {
  route 1.1.1.1/32 next-hop 10.0.0.1;
}
autonomous-system 200;

user@R3# run show route 1.1.1.1 protocol bgp
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[BGP/170] 00:05:11, localpref 100
                   AS path: 200 250 I, validation-state: unverified
                   > to 10.1.0.1 via 1t-1/2/0.4

```

In a Layer 3 VPN scenario, in which a provider edge (PE) device uses external BGP (EBGP) to peer with a customer edge (CE) device, the **local-as** statement behaves differently than in the non-VPN scenario. In the VPN scenario, the global AS number defined in the master instance is prepended to the AS path by default. To override this behavior, you can configure the **no-prepend-global-as** in the routing-instance BGP configuration on the PE device, as shown here:

```

user@R2# show routing-instances

```

```
red {  
  instance-type vrf;  
  interface fe-1/2/0.2;  
  route-distinguisher 2:1;  
  vrf-target target:2:1;  
  protocols {  
    bgp {  
      group toR1 {  
        type external;  
        peer-as 1;  
        local-as 200 no-prepend-global-as;  
        neighbor 10.1.1.1;  
      }  
    }  
  }  
}
```

The Junos operating system (Junos OS) implementation of the local AS attribute supports the following options:


- **Local AS with private option**—When you use the **private** option, the local AS is used during the establishment of the BGP session with an EBGP neighbor but is hidden in the AS path sent to other EBGP peers. Only the global AS is included in the AS path sent to external peers.

The **private** option is useful for establishing local peering with routing devices that remain configured with their former AS or with a specific customer that has not yet modified its peer arrangements. The local AS is used to establish the BGP session with the EBGP neighbor but is hidden in the AS path sent to external peers in another AS.

Include the **private** option so that the local AS is not prepended before the global AS in the AS path sent to external peers. When you specify the **private** option, the local AS is prepended only in the AS path sent to the EBGP neighbor.

For example, in [Figure 71 on page 3365](#), Router 1 and Router 2 are in AS 64496, Router 4 is in AS 64511, and Router 3 is in AS 64510. Router 2 formerly belonged to AS 64497, which has merged with another network and now belongs to AS 64496. Because Router 3 still peers with Router 2 using its former AS (64497), Router 2 needs to be configured with a local AS of 64497 in order to maintain peering with Router 3. Configuring a local AS of 64497 permits Router 2 to add AS 64497 when advertising routes to Router 3. Router 3 sees an AS path of 64497 64496 for the prefix 10/8.

**Figure 71: Local AS Configuration**



To prevent Router 2 from adding the local AS number in its announcements to other peers, use the **local-as 64497 private** statement. This statement configures Router 2 to not include local AS 64497 when announcing routes to Router 1 and to Router 4. In this case, Router 4 sees an AS path of 64496 64510 for the prefix 10.222/16.

- **Local AS with alias option**—In Junos OS Release 9.5 and later, you can configure a local AS as an alias. During the establishment of the BGP open session, the AS used in the open message alternates between the local AS and the global AS. If the local AS is used to connect with the EBGP neighbor, then only the local AS is prepended to the AS path when the BGP peer session is established. If the global AS is used to connect with the EBGP neighbor, then only the global AS is prepended to the AS path when the BGP peer session is established. The use of the **alias** option also means that

the local AS is not prepended to the AS path for any routes learned from that EBGp neighbor. Therefore, the local AS remains hidden from other external peers.

Configuring a local AS with the **alias** option is especially useful when you are migrating the routing devices in an acquired network to the new AS. During the migration process, some routing devices might be configured with the new AS while others remain configured with the former AS. For example, it is good practice to start by first migrating to the new AS any routing devices that function as route reflectors. However, as you migrate the route reflector clients incrementally, each route reflector has to peer with routing devices configured with the former AS, as well as peer with routing devices configured with the new AS. To establish local peer sessions, it can be useful for the BGP peers in the network to use both the local AS and the global AS. At the same time, you want to hide this local AS from external peers and use only the global AS in the AS path when exporting routes to another AS. In this kind of situation, configure the **alias** option.

Include the **alias** option to configure the local AS as an alias to the global AS configured at the **[edit routing-options]** hierarchy level. When you configure a local AS as an alias, during the establishment of the BGP open session, the AS used in the open message alternates between the local AS and the global AS. The local AS is prepended to the AS path only when the peer session with an EBGp neighbor is established using that local AS. The local AS is hidden in the AS path sent to any other external peers. Only the global AS is prepended to the AS path when the BGP session is established using the global AS.



**NOTE:** The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

---

- **Local AS with option not to prepend the global AS**—In Junos OS Release 9.6 and later, you can configure a local AS with the option not to prepend the global AS. Only the local AS is included in the AS path sent to external peers.

Use the **no-prepend-global-as** option when you want to strip the global AS number from outbound BGP updates in a virtual private network (VPN) scenario. This option is useful in a VPN scenario in which you want to hide the global AS from the VPN.

Include the **no-prepend-global-as** option to have the global AS configured at the **[edit routing-options]** hierarchy level removed from the AS path sent to external peers. When you use this option, only the local AS is included in the AS path for the routes sent to a customer edge (CE) device.

- **Number of loops option**—The local AS feature also supports specifying the number of times that detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

For the **loops number** statement, you can configure 1 through 10.





**NOTE:** If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the `local-as` statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the `local-as` statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

### Example: Configuring a Local AS for EBGP Sessions

This example shows how to configure a local autonomous system (AS) for a BGP peer so that both the global AS and the local AS are used in BGP inbound and outbound updates.

- [Requirements on page 3367](#)
- [Overview on page 3367](#)
- [Configuration on page 3368](#)
- [Verification on page 3374](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.

#### Overview

Use the `local-as` statement when ISPs merge and want to preserve a customer's configuration, particularly the AS with which the customer is configured to establish a peer relationship. The `local-as` statement simulates the AS number already in place in customer routers, even if the ISP's router has moved to a different AS.


This example shows how to use the `local-as` statement to configure a local AS. The `local-as` statement is supported for BGP at the global, group, and neighbor hierarchy levels.

When you configure the `local-as` statement, you must specify an AS number. You can specify a number from 1 through 4,294,967,295 in plain-number format. In Junos OS Release 9.1 and later, the range for AS numbers is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value from 0.0 through 65535.65535 in AS-dot notation format. Junos

OS continues to support 2-byte AS numbers. The 2-byte AS number range is 1 through 65,535 (this is a subset of the 4-byte range).

Figure 72 on page 3368 shows the sample topology.

**Figure 72: Topology for Configuring the Local AS**



In this example, Device R2 formerly belonged to AS 250 and now is in AS 200. Device R1 and Device R3 are configured to peer with AS 250 instead of with the new AS number (AS 200). Device R2 has the new AS number configured with the **autonomous-system 200** statement. To enable the peering sessions to work, the **local-as 250** statement is added in the BGP configuration. Because **local-as 250** is configured, Device R2 includes both the global AS (200) and the local AS (250) in its BGP inbound and outbound updates.

#### Configuration

- [Configuring Device R1 on page 3369](#)
- [Configuring Device R2 on page 3371](#)
- [Configuring Device R3 on page 3373](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1  set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
           set interfaces lo0 unit 1 family inet address 192.168.0.1/32
           set protocols bgp group ext type external
           set protocols bgp group ext export send-direct
           set protocols bgp group ext export send-static
           set protocols bgp group ext peer-as 250
           set protocols bgp group ext neighbor 10.0.0.2
           set policy-options policy-statement send-direct term 1 from protocol direct
           set policy-options policy-statement send-direct term 1 then accept
           set policy-options policy-statement send-static term 1 from protocol static
           set policy-options policy-statement send-static term 1 then accept
           set routing-options static route 10.1.0.0/30 next-hop 10.0.0.2
           set routing-options autonomous-system 100
```

```
Device R2  set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
           set interfaces fe-1/2/1 unit 3 family inet address 10.1.0.1/30
           set interfaces lo0 unit 2 family inet address 192.168.0.2/32
           set protocols bgp group ext type external
```

```

set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext local-as 250
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols bgp group ext neighbor 10.1.0.2 peer-as 300
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options autonomous-system 200

```

**Device R3**

```

set interfaces fe-1/2/0 unit 4 family inet address 10.1.0.2/30
set interfaces lo0 unit 3 family inet address 192.168.0.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 250
set protocols bgp group ext neighbor 10.1.0.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.0.0.0/30 next-hop 10.1.0.1
set routing-options autonomous-system 300

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.  

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 192.168.0.1/32

```
2. Configure external BGP (EBGP).  

```

[edit protocols bgp group ext]
user@R1# set type external
user@R1# set export send-direct
user@R1# set export send-static
user@R1# set peer-as 250
user@R1# set neighbor 10.0.0.2

```
3. Configure the routing policy.  

```

[edit policy-options]
user@R1# set policy-statement send-direct term 1 from protocol direct
user@R1# set policy-statement send-direct term 1 then accept
user@R1# set policy-statement send-static term 1 from protocol static
user@R1# set policy-statement send-static term 1 then accept

```

4. Configure a static route to the remote network between Device R2 and Device R3.

```
[edit routing-options]
user@R1# set static route 10.1.0.0/30 next-hop 10.0.0.2
```

5. Configure the global AS number.

```
[edit routing-options]
user@R1# set autonomous-system 100
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group ext {
    type external;
    export [ send-direct send-static ];
    peer-as 250;
    neighbor 10.0.0.2;
  }
}

user@R1# show routing-options
static {
```

```

    route 10.1.0.0/30 next-hop 10.0.0.2;
}
autonomous-system 100;

```

When you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.
 

```

[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 3 family inet address 10.1.0.1/30

user@R2# set lo0 unit 2 family inet address 192.168.0.2/32

```
2. Configure EBGP.
 

```

[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set export send-static
user@R2# set neighbor 10.0.0.1 peer-as 100
user@R2# set neighbor 10.1.0.2 peer-as 300

```
3. Configure the local autonomous system (AS) number.
 

```

[edit protocols bgp group ext]
user@R2# set local-as 250

```
4. Configure the global AS number.
 

```

[edit routing-options]
user@R2# set autonomous-system 200

```
5. Configure the routing policy.
 

```

[edit policy-options]
user@R2# set policy-statement send-direct term 1 from protocol direct
user@R2# set policy-statement send-direct term 1 then accept
user@R2# set policy-statement send-static term 1 from protocol static
user@R2# set policy-statement send-static term 1 then accept

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R2# show interfaces
fe-1/2/0 {
  unit 2 {

```

```
        family inet {
            address 10.0.0.2/30;
        }
    }
}
fe-1/2/1 {
    unit 3 {
        family inet {
            address 10.1.0.1/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.0.2/32;
        }
    }
}

user@R2# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}
policy-statement send-static {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R2# show protocols
bgp {
    group ext {
        type external;
        export [ send-direct send-static ];
        local-as 250;
        neighbor 10.0.0.1 {
            peer-as 100;
        }
        neighbor 10.1.0.2 {
            peer-as 300;
        }
    }
}

user@R2# show routing-options
autonomous-system 200;
```

When you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R3

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Configure the interfaces.  

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 4 family inet address 10.1.0.2/30

user@R3# set lo0 unit 3 family inet address 192.168.0.3/32
```
2. Configure EBGP.  

```
[edit protocols bgp group ext]
user@R3# set type external
user@R3# set export send-direct
user@R3# set export send-static
user@R3# set peer-as 250
user@R3# set neighbor 10.1.0.1
```
3. Configure the global autonomous system (AS) number.  

```
[edit routing-options]
user@R3# set autonomous-system 300
```
4. Configure a static route to the remote network between Device R1 and Device R2.  

```
[edit routing-options]
user@R3# set static route 10.0.0.0/30 next-hop 10.1.0.1
```
5. Configure the routing policy.  

```
[edit policy-options]
user@R3# set policy-statement send-direct term 1 from protocol direct
user@R3# set policy-statement send-direct term 1 then accept
user@R3# set policy-statement send-static term 1 from protocol static
user@R3# set policy-statement send-static term 1 then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/2/0 {
  unit 4 {
    family inet {
      address 10.1.0.2/30;
    }
  }
}
lo0 {
  unit 3 {
```

```
family inet {
    address 192.168.0.3/32;
}
}

user@R3# show policy-options
policy-statement send-direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}
policy-statement send-static {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R3# show protocols
bgp {
    group ext {
        type external;
        export [ send-direct send-static ];
        peer-as 250;
        neighbor 10.1.0.1;
    }
}

user@R3# show routing-options
static {
    route 10.0.0.0/30 next-hop 10.1.0.1;
}
autonomous-system 300;
```

When you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Checking the Local and Global AS Settings on page 3374](#)
- [Checking the BGP Peering Sessions on page 3376](#)
- [Verifying the BGP AS Paths on page 3376](#)

### **Checking the Local and Global AS Settings**

**Purpose** Make sure that Device R2 has the local and global AS settings configured.

**Action** From operational mode, enter the **show bgp neighbors** command.

```
user@R2> show bgp neighbors
Peer: 10.0.0.1+179 AS 100      Local: 10.0.0.2+61036 AS 250
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
```



```

Export: [ send-direct send-static ]
Options: <Preference PeerAS LocalAS Refresh>
Holdtime: 90 Preference: 170 Local AS: 250 Local System AS: 200
Number of flaps: 0
Peer ID: 192.168.0.1      Local ID: 192.168.0.2      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 0
BFD: disabled, down
Local Interface: fe-1/2/0.2
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 100)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        3
  Accepted prefixes:        2
  Suppressed due to damping: 0
  Advertised prefixes:      4
Last traffic (seconds): Received 6    Sent 14    Checked 47
Input messages: Total 258    Updates 3    Refreshes 0    Octets 4969
Output messages: Total 258    Updates 2    Refreshes 0    Octets 5037
Output Queue[0]: 0

Peer: 10.1.0.2+179 AS 300      Local: 10.1.0.1+52296 AS 250
Type: External    State: Established    Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ send-direct send-static ]
Options: <Preference PeerAS LocalAS Refresh>
Holdtime: 90 Preference: 170 Local AS: 250 Local System AS: 200
Number of flaps: 0
Peer ID: 192.168.0.3      Local ID: 192.168.0.2      Active Holdtime: 90
Keepalive Interval: 30      Peer index: 1
BFD: disabled, down
Local Interface: fe-1/2/1.3
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 300)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        3
  Accepted prefixes:        2

```

```

    Suppressed due to damping:    0
    Advertised prefixes:          4
    Last traffic (seconds): Received 19   Sent 26   Checked 9
    Input messages:  Total 256   Updates 3     Refreshes 0   Octets 4931
    Output messages: Total 256   Updates 2     Refreshes 0   Octets 4999
    Output Queue[0]: 0

```

**Meaning** The Local AS: 250 and Local System AS: 200 output shows that Device R2 has the expected settings. Additionally, the output shows that the options list includes LocalAS.

### *Checking the BGP Peering Sessions*

**Purpose** Ensure that the sessions are established and that the local AS number 250 is displayed.

**Action** From operational mode, enter the **show bgp summary** command.

```

user@R1> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      4          2          0          0        0      0        0
Peer        AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.2      250      232      233       0        4    1:42:37
2/4/4/0      0/0/0/0

```

```

user@R3> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0      4          2          0          0        0      0        0
Peer        AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.0.1      250      235      236       0        4    1:44:25
2/4/4/0      0/0/0/0

```

**Meaning** Device R1 and Device R3 appear to be peering with a device in AS 250, even though Device R2 is actually in AS 200.

### *Verifying the BGP AS Paths*

**Purpose** Make sure that the routes are in the routing tables and that the AS paths show the local AS number 250.

**Action** From configuration mode, enter the **set route protocol bgp** command.

```

user@R1> show route protocol bgp
inet.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30      [BGP/170] 01:46:44, localpref 100
                 AS path: 250 I
                 > to 10.0.0.2 via fe-1/2/0.1
10.1.0.0/30      [BGP/170] 01:46:44, localpref 100
                 AS path: 250 I
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.0.2/32   *[BGP/170] 01:46:44, localpref 100
                 AS path: 250 I
                 > to 10.0.0.2 via fe-1/2/0.1

```

```

192.168.0.3/32      *[BGP/170] 01:46:40, localpref 100
                   AS path: 250 300 I
                   > to 10.0.0.2 via fe-1/2/0.1

user@R3> show route protocol bgp

inet.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/30        [BGP/170] 01:47:10, localpref 100
                   AS path: 250 I
                   > to 10.1.0.1 via fe-1/2/0.4
10.1.0.0/30        [BGP/170] 01:47:10, localpref 100
                   AS path: 250 I
                   > to 10.1.0.1 via fe-1/2/0.4
192.168.0.1/32     *[BGP/170] 01:47:10, localpref 100
                   AS path: 250 100 I
                   > to 10.1.0.1 via fe-1/2/0.4
192.168.0.2/32     *[BGP/170] 01:47:10, localpref 100
                   AS path: 250 I
                   > to 10.1.0.1 via fe-1/2/0.4

```

**Meaning** The output shows that Device R1 and Device R3 appear to have routes with AS paths that include AS 250, even though Device R2 is actually in AS 200.

### Example: Configuring a Private Local AS for EBGp Sessions

This example shows how to configure a private local autonomous system (AS) number. The local AS is considered to be private because it is advertised to peers that use the local AS number for peering, but is hidden in the announcements to peers that can use the global AS number for peering.

- [Requirements on page 3377](#)
- [Overview on page 3377](#)
- [Configuration on page 3378](#)
- [Verification on page 3381](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.

#### Overview

Use the **local-as** statement when ISPs merge and want to preserve a customer's configuration, particularly the AS with which the customer is configured to establish a peer relationship. The **local-as** statement simulates the AS number already in place in customer routers, even if the ISP's router has moved to a different AS.

When you use the **private** option, the local AS is used during the establishment of the BGP session with an external BGP (EBGP) neighbor, but is hidden in the AS path sent to other EBGp peers. Only the global AS is included in the AS path sent to external peers.


The **private** option is useful for establishing local peering with routing devices that remain configured with their former AS or with a specific customer that has not yet modified its

peer arrangements. The local AS is used to establish the BGP session with the EBGP neighbor, but is hidden in the AS path sent to external peers in another AS.

Include the **private** option so that the local AS is not prepended before the global AS in the AS path sent to external peers. When you specify the **private** option, the local AS is prepended only in the AS path sent to the EBGP neighbor.

Figure 73 on page 3378 shows the sample topology.

**Figure 73: Topology for Configuring a Private Local AS**



Device R1 is in AS 64496. Device R2 is in AS 64510. Device R3 is in AS 64511. Device R4 is in AS 64512. Device R1 formerly belonged to AS 64497, which has merged with another network and now belongs to AS 64496. Because Device R3 still peers with Device R1, using its former AS, 64497, Device R1 needs to be configured with a local AS of 64497 in order to maintain peering with Device R3. Configuring a local AS of 64497 permits Device R1 to add AS 64497 when advertising routes to Device R3. Device R3 sees an AS path of 64497 64496 for the prefix 10.1.1.2/32, which is Device R2's loopback interface. Device R4, which is behind Device R3, sees an AS path of 64511 64497 64496 64510 to Device R2's loopback interface. To prevent Device R1 from adding the local AS number in its announcements to other peers, this example includes the **local-as 64497 private** statement. The **private** option configures Device R1 to not include the local AS 64497 when announcing routes to Device R2. Device R2 sees an AS path of 64496 64511 to Device R3 and an AS path of 64496 64511 64512 to Device R4. The **private** option in Device R1's configuration causes the AS number 64497 to be missing from the AS paths that Device R1 readvertises to Device R2.

Device R2 is hiding the private local AS from all the routers, except Device R3. The **private** option applies to the routes that Device R1 receives (learns) from Device R3 and that Device R1, in turn, readvertises to other routers. When these routes learned from Device R3 are readvertised by Device R1 to Device R2, the private local AS is missing from the AS path advertised to Device R2.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 3 family inet address 192.168.1.1/24
set interfaces fe-1/2/1 unit 5 family inet address 192.168.10.1/24
set interfaces lo0 unit 2 family inet address 10.1.1.1/32
set protocols bgp group external-AS64511 type external
set protocols bgp group external-AS64511 peer-as 64511
set protocols bgp group external-AS64511 local-as 64497
set protocols bgp group external-AS64511 local-as private
set protocols bgp group external-AS64511 neighbor 192.168.1.2
set protocols bgp group external-AS64510 type external
set protocols bgp group external-AS64510 peer-as 64510
set protocols bgp group external-AS64510 neighbor 192.168.10.2
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64496

```

**Device R2**

```

set interfaces fe-1/2/0 unit 6 family inet address 192.168.10.2/24
set interfaces lo0 unit 3 family inet address 10.1.1.2/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 64496
set protocols bgp group external neighbor 192.168.10.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64510

```

**Device R3**

```

set interfaces fe-1/2/0 unit 4 family inet address 192.168.1.2/24
set interfaces fe-1/2/1 unit 7 family inet address 192.168.5.1/24
set interfaces lo0 unit 4 family inet address 10.1.1.3/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external neighbor 192.168.1.1 peer-as 64497
set protocols bgp group external neighbor 192.168.5.2 peer-as 64512
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64511

```

**Device R4**

```

set interfaces fe-1/2/0 unit 8 family inet address 192.168.5.2/24
set interfaces lo0 unit 5 family inet address 10.1.1.4/32
set protocols bgp group external type external
set protocols bgp group external export send-direct
set protocols bgp group external peer-as 64511
set protocols bgp group external neighbor 192.168.5.1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 64512

```

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.  

```
[edit interfaces fe-1/2/0 unit 3]
user@R1# set family inet address 192.168.1.1/24

[edit interfaces fe-1/2/1 unit 5]
user@R1# set family inet address 192.168.10.1/24

[edit interfaces lo0 unit 2]
user@R1# set family inet address 10.1.1.1/32
```
2. Configure the EBGP peering session with Device R2.  

```
[edit protocols bgp group external-AS64510]
user@R1# set type external
user@R1# set peer-as 64510
user@R1# set neighbor 192.168.10.2
```
3. Configure the EBGP peering session with Device R3.  

```
[edit protocols bgp group external-AS64511]
user@R1# set type external
user@R1# set peer-as 64511
user@R1# set local-as 64497
user@R1# set local-as private
user@R1# set neighbor 192.168.1.2
```
4. Configure the routing policy.  

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```
5. Configure the global autonomous system (AS) number.  

```
[edit routing-options]
user@R1# set autonomous-system 64496
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 192.168.1.1/24;
    }
  }
}
```

```

    }
  }
  fe-1/2/1 {
    unit 5 {
      family inet {
        address 192.168.10.1/24;
      }
    }
  }
  lo0 {
    unit 2 {
      family inet {
        address 10.1.1.1/32;
      }
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
bgp {
  group external-AS64511 {
    type external;
    peer-as 64511;
    local-as 64497 private;
    neighbor 192.168.1.2;
  }
  group external-AS64510 {
    type external;
    peer-as 64510;
    neighbor 192.168.10.2;
  }
}

user@R1# show routing-options
autonomous-system 64496;

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the configuration as needed for the other devices in the topology.

### Verification

Confirm that the configuration is working properly.

- [Checking Device R2's AS Paths on page 3382](#)
- [Checking Device R3's AS Paths on page 3382](#)

### ***Checking Device R2's AS Paths***

**Purpose** Make sure that Device R2 does not have AS 64497 in its AS paths to Device R3 and Device R4.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R2> show route protocol bgp
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.3/32      *[BGP/170] 01:33:11, localpref 100
                  AS path: 64496 64511 I
                  > to 192.168.10.1 via fe-1/2/0.6
10.1.1.4/32      *[BGP/170] 01:33:11, localpref 100
                  AS path: 64496 64511 64512 I
                  > to 192.168.10.1 via fe-1/2/0.6
192.168.5.0/24   *[BGP/170] 01:49:15, localpref 100
                  AS path: 64496 64511 I
                  > to 192.168.10.1 via fe-1/2/0.6
```

**Meaning** Device R2's AS paths do not include AS 64497.

### ***Checking Device R3's AS Paths***

**Purpose** Make sure that Device R3 does not have AS 64497 in its AS path to Device R4.

**Action** From operational mode, enter the **show route protocol bgp** command.

```
user@R3> show route protocol bgp
inet.0: 7 destinations, 8 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2/32      *[BGP/170] 01:35:11, localpref 100
                  AS path: 64497 64496 64510 I
                  > to 192.168.1.1 via fe-1/2/0.4
10.1.1.4/32      *[BGP/170] 01:35:11, localpref 100
                  AS path: 64512 I
                  > to 192.168.5.2 via fe-1/2/1.7
192.168.5.0/24   [BGP/170] 01:51:15, localpref 100
                  AS path: 64512 I
                  > to 192.168.5.2 via fe-1/2/1.7
```

**Meaning** Device R3's route to Device R2 (prefix 10.1.1.2) includes both the local and the global AS configured on Device R1 (64497 and 64496, respectively).

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3261](#)
- [BGP Configuration Overview](#)

## **Example: Configuring the Accumulated IGP Attribute for BGP**

- [Understanding the Accumulated IGP Attribute for BGP on page 3383](#)
- [Example: Configuring the Accumulated IGP Attribute for BGP on page 3383](#)



### Understanding the Accumulated IGP Attribute for BGP

The interior gateway protocols (IGPs) are designed to handle routing within a single domain or an autonomous system (AS). Each link is assigned a particular value called a metric. The distance between the two nodes is calculated as a sum of all the metric values of links along the path. The IGP selects the shortest path between two nodes based on distance.

BGP is designed to provide routing over a large number of independent ASs with limited or no coordination among respective administrations. BGP does not use metrics in the path selection decisions.

The accumulated IGP (AIGP) metric attribute for BGP enables deployment in which a single administration can run several contiguous BGP ASs. Such deployments allow BGP to make routing decisions based on the IGP metric. In such networks, it is possible for BGP to select paths based on metrics as is done by IGPs. In this case, BGP chooses the shortest path between two nodes, even though the nodes might be in two different ASs.

The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. The Juniper Networks® Junos® operating system (Junos OS) currently supports the AIGP attribute for two BGP address families, **family inet labeled-unicast** and **family inet6 labeled-unicast**.

AIGP impacts the BGP best-route decision process. The AIGP attribute preference rule is applied after the local-preference rule. The AIGP distance is compared to break a tie. The BGP best-route decision process also impacts the way the interior cost rule is applied if the resolving next hop has an AIGP attribute. Without AIGP enabled, the interior cost of a route is based on the calculation of the metric to the next hop for the route. With AIGP enabled, the resolving AIGP distance is added to the interior cost.

The AIGP attribute is an optional non-transitive BGP path attribute and is specified in Internet draft draft-ietf-idr-aigp-06, *The Accumulated IGP Metric Attribute for BGP*.

### Example: Configuring the Accumulated IGP Attribute for BGP

This example shows how to configure the accumulated IGP (AIGP) metric attribute for BGP.

- [Requirements on page 3383](#)
- [Overview on page 3384](#)
- [Configuration on page 3385](#)
- [Verification on page 3415](#)

#### Requirements

This example uses the following hardware and software components:

- Seven BGP-speaking devices.
- Junos OS Release 12.1 or later.

### Overview

The AIGP attribute enables deployments in which a single administration can run several contiguous BGP autonomous systems (ASs). Such deployments allow BGP to make routing decisions based on the IGP metric. With AIGP enabled, BGP can select paths based on IGP metrics. This enables BGP to choose the shortest path between two nodes, even though the nodes might be in different ASs. The AIGP attribute is particularly useful in networks that use tunneling to deliver a packet to its BGP next hop. This example shows AIGP configured with MPLS label-switched paths.

To enable AIGP, you include the **aigp** statement in the BGP configuration on a protocol family basis. Configuring AIGP on a particular family enables sending and receiving of the AIGP attribute on that family. By default, AIGP is disabled. An AIGP-disabled neighbor does not send an AIGP attribute and silently discards a received AIGP attribute.

Junos OS supports AIGP for **family inet labeled-unicast** and **family inet6 labeled-unicast**. The **aigp** statement can be configured for a given family at the global BGP, group, or neighbor level.

By default, the value of the AIGP attribute for a local prefix is zero. An AIGP-enabled neighbor can originate an AIGP attribute for a given prefix by export policy, using the **aigp-originate** policy action. The value of the AIGP attribute reflects the IGP distance to the prefix. Alternatively, you can specify a value, by using the **aigp-originate distance distance** policy action. The configurable range is 0 through 4,294,967,295. Only one node needs to originate an AIGP attribute. The AIGP attribute is retained and readvertised if the neighbors are AIGP enabled with the **aigp** statement in the BGP configuration.

The policy action to originate the AIGP attribute has the following requirements:

- Neighbor must be AIGP enabled.
- Policy must be applied as an export policy.
- Prefix must have no current AIGP attribute.
- Prefix must export with next-hop self.
- Prefix must reside within the AIGP domain. Typically, a loopback IP address is the prefix to originate.


The policy is ignored if these requirements are not met.

### Topology Diagram

[Figure 74 on page 3385](#) shows the topology used in this example. OSPF is used as the interior gateway protocol (IGP). Internal BGP (IBGP) is configured between Device PE1 and Device PE4. External BGP (EBGP) is configured between Device PE7 and Device PE1, between Device PE4 and Device PE3, and between Device PE4 and Device PE2. Devices PE4, PE2, and PE3 are configured for multihop. Device PE4 selects a path based on the AIGP value and then readvertises the AIGP value based on the AIGP and policy configuration. Device PE1 readvertises the AIGP value to Device PE7, which is in another administrative domain. Every device has two loopback interface addresses: 10.9.9.x is used for BGP peering and the router ID, and 10.100.1.x is used for the BGP next hop.

The network between Device PE1 and PE3 has IBGP peering and multiple OSPF areas. The external link to Device PE7 is configured to show that the AIGP attribute is readvertised to a neighbor outside of the administrative domain, if that neighbor is AIGP enabled.

**Figure 74: Advertisement of Multiple Paths in BGP**



For origination of an AIGP attribute, the BGP next hop is required to be itself. If the BGP next hop remains unchanged, the received AIGP attribute is readvertised, as is, to another AIGP neighbor. If the next hop changes, the received AIGP attribute is readvertised with an increased value to another AIGP neighbor. The increase in value reflects the IGP distance to the previous BGP next hop. To demonstrate, this example uses loopback interface addresses for Device PE4's EBGP peering sessions with Device PE2 and Device PE3. Multihop is enabled on these sessions so that a recursive lookup is performed to determine the point-to-point interface. Because the next hop changes, the IGP distance is added to the AIGP distance.

#### Configuration

- [Configuring Device P1 on page 3391](#)
- [Configuring Device P2 on page 3394](#)
- [Configuring Device PE4 on page 3397](#)
- [Configuring Device PE1 on page 3402](#)
- [Configuring Device PE2 on page 3406](#)
- [Configuring Device PE3 on page 3410](#)
- [Configuring Device PE7 on page 3413](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device P1 set interfaces fe-1/2/0 unit 1 description P1-to-PE1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.2/30
set interfaces fe-1/2/0 unit 1 family mpls
set interfaces fe-1/2/1 unit 4 description P1-to-P2
```

```
set interfaces fe-1/2/1 unit 4 family inet address 10.0.0.29/30
set interfaces fe-1/2/1 unit 4 family mpls
set interfaces fe-1/2/2 unit 8 description P1-to-PE4
set interfaces fe-1/2/2 unit 8 family inet address 10.0.0.17/30
set interfaces fe-1/2/2 unit 8 family mpls
set interfaces lo0 unit 3 family inet address 10.9.9.2/32
set interfaces lo0 unit 3 family inet address 10.100.1.2/32
set protocols rsvp interface fe-1/2/0.1
set protocols rsvp interface fe-1/2/2.8
set protocols rsvp interface fe-1/2/1.4
set protocols mpls label-switched-path P1-to-P2 to 10.9.9.3
set protocols mpls label-switched-path P1-to-PE1 to 10.9.9.1
set protocols mpls label-switched-path P1-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.1
set protocols mpls interface fe-1/2/2.8
set protocols mpls interface fe-1/2/1.4
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.2
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group internal neighbor 10.9.9.4
set protocols ospf area 0.0.0.1 interface fe-1/2/0.1 metric 1
set protocols ospf area 0.0.0.1 interface fe-1/2/1.4 metric 1
set protocols ospf area 0.0.0.0 interface fe-1/2/2.8 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.2 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.2 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.2 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.2 metric 1
set routing-options router-id 10.9.9.2
set routing-options autonomous-system 13979
```

**Device P2**

```
set interfaces fe-1/2/0 unit 3 description P2-to-PE1
set interfaces fe-1/2/0 unit 3 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 3 family mpls
set interfaces fe-1/2/1 unit 5 description P2-to-P1
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.30/30
set interfaces fe-1/2/1 unit 5 family mpls
set interfaces fe-1/2/2 unit 6 description P2-to-PE4
set interfaces fe-1/2/2 unit 6 family inet address 10.0.0.13/30
set interfaces fe-1/2/2 unit 6 family mpls
set interfaces lo0 unit 5 family inet address 10.9.9.3/32
set interfaces lo0 unit 5 family inet address 10.100.1.3/32
set protocols rsvp interface fe-1/2/1.5
set protocols rsvp interface fe-1/2/2.6
set protocols rsvp interface fe-1/2/0.3
set protocols mpls label-switched-path P2-to-PE1 to 10.9.9.1
set protocols mpls label-switched-path P2-to-P1 to 10.9.9.2
set protocols mpls label-switched-path P2-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/1.5
set protocols mpls interface fe-1/2/2.6
set protocols mpls interface fe-1/2/0.3
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.3
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
```

```

set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group internal neighbor 10.9.9.4
set protocols ospf area 0.0.0.0 interface fe-1/2/2.6 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.3 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.3 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.3 passive
set protocols ospf area 0.0.0.0 interface 10.100.1.3 metric 1
set routing-options router-id 10.9.9.3
set routing-options autonomous-system 13979

```

**Device PE4**

```

set interfaces fe-1/2/0 unit 7 description PE4-to-P2
set interfaces fe-1/2/0 unit 7 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 7 family mpls
set interfaces fe-1/2/1 unit 9 description PE4-to-P1
set interfaces fe-1/2/1 unit 9 family inet address 10.0.0.18/30
set interfaces fe-1/2/1 unit 9 family mpls
set interfaces fe-1/2/2 unit 10 description PE4-to-PE2
set interfaces fe-1/2/2 unit 10 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 10 family mpls
set interfaces fe-1/0/2 unit 12 description PE4-to-PE3
set interfaces fe-1/0/2 unit 12 family inet address 10.0.0.25/30
set interfaces fe-1/0/2 unit 12 family mpls
set interfaces lo0 unit 7 family inet address 10.9.9.4/32
set interfaces lo0 unit 7 family inet address 10.100.1.4/32
set protocols rsvp interface fe-1/2/0.7
set protocols rsvp interface fe-1/2/1.9
set protocols rsvp interface fe-1/2/2.10
set protocols rsvp interface fe-1/0/2.12
set protocols mpls label-switched-path PE4-to-PE2 to 10.9.9.5
set protocols mpls label-switched-path PE4-to-PE3 to 10.9.9.6
set protocols mpls label-switched-path PE4-to-P1 to 10.9.9.2
set protocols mpls label-switched-path PE4-to-P2 to 10.9.9.3
set protocols mpls interface fe-1/2/0.7
set protocols mpls interface fe-1/2/1.9
set protocols mpls interface fe-1/2/2.10
set protocols mpls interface fe-1/0/2.12
set protocols bgp export next-hop
set protocols bgp export aigp
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.4
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal neighbor 10.9.9.1
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.4
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external peer-as 7018
set protocols bgp group external neighbor 10.9.9.5
set protocols bgp group external neighbor 10.9.9.6
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9 metric 1
set protocols ospf area 0.0.0.0 interface fe-1/2/0.7 metric 1
set protocols ospf area 0.0.0.0 interface 10.9.9.4 passive
set protocols ospf area 0.0.0.0 interface 10.9.9.4 metric 1
set protocols ospf area 0.0.0.0 interface 10.100.1.4 passive

```

```

set protocols ospf area 0.0.0.0 interface 10.100.1.4 metric 1
set protocols ospf area 0.0.0.2 interface fe-1/2/2.10 metric 1
set protocols ospf area 0.0.0.3 interface fe-1/0/2.12 metric 1
set policy-options policy-statement aigp term 10 from protocol static
set policy-options policy-statement aigp term 10 from route-filter 44.0.0.0/24 exact
set policy-options policy-statement aigp term 10 then aigp-originate distance 200
set policy-options policy-statement aigp term 10 then next-hop 10.100.1.4
set policy-options policy-statement aigp term 10 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.4
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.4/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.4/32
    exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.4
set policy-options policy-statement next-hop term 20 then accept
set routing-options static route 44.0.0.0/24 discard
set routing-options router-id 10.9.9.4
set routing-options autonomous-system 13979

```

**Device PE1**

```

set interfaces fe-1/2/0 unit 0 description PE1-to-P1
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 2 description PE1-to-P2
set interfaces fe-1/2/1 unit 2 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 2 family mpls
set interfaces fe-1/2/2 unit 14 description PE1-to-PE7
set interfaces fe-1/2/2 unit 14 family inet address 10.0.0.9/30
set interfaces lo0 unit 1 family inet address 10.9.9.1/32
set interfaces lo0 unit 1 family inet address 10.100.1.1/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.2
set protocols rsvp interface fe-1/2/2.14
set protocols mpls label-switched-path PE1-to-P1 to 10.9.9.2
set protocols mpls label-switched-path PE1-to-P2 to 10.9.9.3
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.2
set protocols mpls interface fe-1/2/2.14
set protocols bgp group internal type internal
set protocols bgp group internal local-address 10.9.9.1
set protocols bgp group internal family inet labeled-unicast aigp
set protocols bgp group internal export SET_EXPORT_ROUTES
set protocols bgp group internal vpn-apply-export
set protocols bgp group internal neighbor 10.9.9.4
set protocols bgp group internal neighbor 10.9.9.2
set protocols bgp group internal neighbor 10.9.9.3
set protocols bgp group external type external
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external peer-as 7019
set protocols bgp group external neighbor 10.0.0.10
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0 metric 1
set protocols ospf area 0.0.0.1 interface fe-1/2/1.2 metric 1
set protocols ospf area 0.0.0.1 interface 10.9.9.1 passive
set protocols ospf area 0.0.0.1 interface 10.9.9.1 metric 1

```

```

set protocols ospf area 0.0.0.1 interface 10.100.1.1 passive
set protocols ospf area 0.0.0.1 interface 10.100.1.1 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set routing-options router-id 10.9.9.1
set routing-options autonomous-system 13979

```

**Device PE2**

```

set interfaces fe-1/2/0 unit 11 description PE2-to-PE4
set interfaces fe-1/2/0 unit 11 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 11 family mpls
set interfaces lo0 unit 9 family inet address 10.9.9.5/32 primary
set interfaces lo0 unit 9 family inet address 10.100.1.5/32
set protocols rsvp interface fe-1/2/0.11
set protocols mpls label-switched-path PE2-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.11
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.5
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export next-hop
set protocols bgp group external export aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external vpn-apply-export
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.9.9.4
set protocols ospf area 0.0.0.2 interface 10.9.9.5 passive
set protocols ospf area 0.0.0.2 interface 10.9.9.5 metric 1
set protocols ospf area 0.0.0.2 interface 10.100.1.5 passive
set protocols ospf area 0.0.0.2 interface 10.100.1.5 metric 1
set protocols ospf area 0.0.0.2 interface fe-1/2/0.11 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol static
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.5
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set policy-options policy-statement aigp term 10 from route-filter 55.0.0.0/24 exact
set policy-options policy-statement aigp term 10 then aigp-originate distance 20
set policy-options policy-statement aigp term 10 then next-hop 10.100.1.5
set policy-options policy-statement aigp term 10 then accept
set policy-options policy-statement aigp term 20 from route-filter 99.0.0.0/24 exact
set policy-options policy-statement aigp term 20 then aigp-originate distance 30
set policy-options policy-statement aigp term 20 then next-hop 10.100.1.5
set policy-options policy-statement aigp term 20 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.5
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.5/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.5/32
  exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.5
set policy-options policy-statement next-hop term 20 then accept

```

```
set routing-options static route 99.0.0.0/24 discard
set routing-options static route 55.0.0.0/24 discard
set routing-options router-id 10.9.9.5
set routing-options autonomous-system 7018
```

**Device PE3**

```
set interfaces fe-1/2/0 unit 13 description PE3-to-PE4
set interfaces fe-1/2/0 unit 13 family inet address 10.0.0.26/30
set interfaces fe-1/2/0 unit 13 family mpls
set interfaces lo0 unit 11 family inet address 10.9.9.6/32
set interfaces lo0 unit 11 family inet address 10.100.1.6/32
set protocols rsvp interface fe-1/2/0.13
set protocols mpls label-switched-path PE3-to-PE4 to 10.9.9.4
set protocols mpls interface fe-1/2/0.13
set protocols bgp group external type external
set protocols bgp group external multihop ttl 2
set protocols bgp group external local-address 10.9.9.6
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export next-hop
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external vpn-apply-export
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.9.9.4
set protocols ospf area 0.0.0.3 interface 10.9.9.6 passive
set protocols ospf area 0.0.0.3 interface 10.9.9.6 metric 1
set protocols ospf area 0.0.0.3 interface 10.100.1.6 passive
set protocols ospf area 0.0.0.3 interface 10.100.1.6 metric 1
set protocols ospf area 0.0.0.3 interface fe-1/2/0.13 metric 1
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol static
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.6
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set policy-options policy-statement next-hop term 10 from protocol bgp
set policy-options policy-statement next-hop term 10 then next-hop 10.100.1.6
set policy-options policy-statement next-hop term 10 then accept
set policy-options policy-statement next-hop term 20 from protocol direct
set policy-options policy-statement next-hop term 20 from route-filter 10.9.9.6/32 exact
set policy-options policy-statement next-hop term 20 from route-filter 10.100.1.6/32
  exact
set policy-options policy-statement next-hop term 20 then next-hop 10.100.1.6
set policy-options policy-statement next-hop term 20 then accept
set routing-options router-id 10.9.9.6
set routing-options autonomous-system 7018
```

**Device PE7**

```
set interfaces fe-1/2/0 unit 15 description PE7-to-PE1
set interfaces fe-1/2/0 unit 15 family inet address 10.0.0.10/30
set interfaces lo0 unit 13 family inet address 10.9.9.7/32
set interfaces lo0 unit 13 family inet address 10.100.1.7/32
set protocols bgp group external type external
set protocols bgp group external family inet labeled-unicast aigp
set protocols bgp group external export SET_EXPORT_ROUTES
set protocols bgp group external peer-as 13979
set protocols bgp group external neighbor 10.0.0.9
set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol direct
```



```

set policy-options policy-statement SET_EXPORT_ROUTES term 10 from protocol bgp
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then next-hop
  10.100.1.7
set policy-options policy-statement SET_EXPORT_ROUTES term 10 then accept
set routing-options router-id 10.9.9.7
set routing-options autonomous-system 7019

```

### Configuring Device P1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device P1:

1. Configure the interfaces.

```

[edit interfaces]
user@P1# set fe-1/2/0 unit 1 description P1-to-PE1
user@P1# set fe-1/2/0 unit 1 family inet address 10.0.0.2/30
user@P1# set fe-1/2/0 unit 1 family mpls
user@P1# set fe-1/2/1 unit 4 description P1-to-P2
user@P1# set fe-1/2/1 unit 4 family inet address 10.0.0.29/30
user@P1# set fe-1/2/1 unit 4 family mpls
user@P1# set fe-1/2/2 unit 8 description P1-to-PE4
user@P1# set fe-1/2/2 unit 8 family inet address 10.0.0.17/30
user@P1# set fe-1/2/2 unit 8 family mpls
user@P1# set lo0 unit 3 family inet address 10.9.9.2/32
user@P1# set lo0 unit 3 family inet address 10.100.1.2/32

```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```

[edit protocols]
user@P1# set rsvp interface fe-1/2/0.1
user@P1# set rsvp interface fe-1/2/2.8
user@P1# set rsvp interface fe-1/2/1.4
user@P1# set mpls label-switched-path P1-to-P2 to 10.9.9.3
user@P1# set mpls label-switched-path P1-to-PE1 to 10.9.9.1
user@P1# set mpls label-switched-path P1-to-PE4 to 10.9.9.4
user@P1# set mpls interface fe-1/2/0.1
user@P1# set mpls interface fe-1/2/2.8
user@P1# set mpls interface fe-1/2/1.4

```

3. Configure BGP.

```

[edit protocols bgp group internal]
user@P1# set type internal
user@P1# set local-address 10.9.9.2
user@P1# set neighbor 10.9.9.1
user@P1# set neighbor 10.9.9.3
user@P1# set neighbor 10.9.9.4

```

4. Enable AIGP.

```

[edit protocols bgp group internal]
user@P1# set family inet labeled-unicast aigp

```

5. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@P1# set area 0.0.0.1 interface fe-1/2/0.1 metric 1
user@P1# set area 0.0.0.1 interface fe-1/2/1.4 metric 1
user@P1# set area 0.0.0.0 interface fe-1/2/2.8 metric 1
user@P1# set area 0.0.0.0 interface 10.9.9.2 passive
user@P1# set area 0.0.0.0 interface 10.9.9.2 metric 1
user@P1# set area 0.0.0.0 interface 10.100.1.2 passive
user@P1# set area 0.0.0.0 interface 10.100.1.2 metric 1
```

6. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@P1# set router-id 10.9.9.2
user@P1# set autonomous-system 13979
```

7. If you are done configuring the device, commit the configuration.

```
user@P1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P1# show interfaces
fe-1/2/0 {
  unit 1 {
    description P1-to-PE1;
    family inet {
      address 10.0.0.2/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 4 {
    description P1-to-P2;
    family inet {
      address 10.0.0.29/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 8 {
    description P1-to-PE4;
    family inet {
      address 10.0.0.17/30;
    }
    family mpls;
  }
}
lo0 {
  unit 3 {
    family inet {
```

```

        address 10.9.9.2/32;
        address 10.100.1.2/32;
    }
}
}

user@P1# show protocols
rsvp {
    interface fe-1/2/0.1;
    interface fe-1/2/2.8;
    interface fe-1/2/1.4;
}
mpls {
    label-switched-path P1-to-P2 {
        to 10.9.9.3;
    }
    label-switched-path P1-to-PE1 {
        to 10.9.9.1;
    }
    label-switched-path P1-to-PE4 {
        to 10.9.9.4;
    }
    interface fe-1/2/0.1;
    interface fe-1/2/2.8;
    interface fe-1/2/1.4;
}
bgp {
    group internal {
        type internal;
        local-address 10.9.9.2;
        family inet {
            labeled-unicast {
                aigp;
            }
        }
        neighbor 10.9.9.1;
        neighbor 10.9.9.3;
        neighbor 10.9.9.4;
    }
}
ospf {
    area 0.0.0.1 {
        interface fe-1/2/0.1 {
            metric 1;
        }
        interface fe-1/2/1.4 {
            metric 1;
        }
    }
    area 0.0.0.0 {
        interface fe-1/2/2.8 {
            metric 1;
        }
        interface 10.9.9.2 {
            passive;
            metric 1;
        }
    }
}

```

```
    }  
    interface 10.100.1.2 {  
        passive;  
        metric 1;  
    }  
}  
}
```

```
user@P1# show routing-options  
router-id 10.9.9.2;  
autonomous-system 13979;
```

### Configuring Device P2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device P2:

1. Configure the interfaces.

```
[edit interfaces]  
user@P2# set fe-1/2/0 unit 3 description P2-to-PE1  
user@P2# set fe-1/2/0 unit 3 family inet address 10.0.0.6/30  
user@P2# set fe-1/2/0 unit 3 family mpls  
user@P2# set fe-1/2/1 unit 5 description P2-to-P1  
user@P2# set fe-1/2/1 unit 5 family inet address 10.0.0.30/30  
user@P2# set fe-1/2/1 unit 5 family mpls  
user@P2# set fe-1/2/2 unit 6 description P2-to-PE4  
user@P2# set fe-1/2/2 unit 6 family inet address 10.0.0.13/30  
user@P2# set fe-1/2/2 unit 6 family mpls  
user@P2# set lo0 unit 5 family inet address 10.9.9.3/32  
user@P2# set lo0 unit 5 family inet address 10.100.1.3/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]  
user@P2# set rsvp interface fe-1/2/1.5  
user@P2# set rsvp interface fe-1/2/2.6  
user@P2# set rsvp interface fe-1/2/0.3  
user@P2# set mpls label-switched-path P2-to-PE1 to 10.9.9.1  
user@P2# set mpls label-switched-path P2-to-P1 to 10.9.9.2  
user@P2# set mpls label-switched-path P2-to-PE4 to 10.9.9.4  
user@P2# set mpls interface fe-1/2/1.5  
user@P2# set mpls interface fe-1/2/2.6  
user@P2# set mpls interface fe-1/2/0.3
```

3. Configure BGP.

```
[edit protocols bgp group internal]  
user@P2# set type internal  
user@P2# set local-address 10.9.9.3  
user@P2# set neighbor 10.9.9.1  
user@P2# set neighbor 10.9.9.2  
user@P2# set neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp group internal]
user@P2# set family inet labeled-unicast aigp
```

5. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@P2# set area 0.0.0.0 interface fe-1/2/2.6 metric 1
user@P2# set area 0.0.0.0 interface 10.9.9.3 passive
user@P2# set area 0.0.0.0 interface 10.9.9.3 metric 1
user@P2# set area 0.0.0.0 interface 10.100.1.3 passive
user@P2# set area 0.0.0.0 interface 10.100.1.3 metric 1
```

6. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@P2# set router-id 10.9.9.3
user@P2# set autonomous-system 13979
```

7. If you are done configuring the device, commit the configuration.

```
user@P2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P2# show interfaces
fe-1/2/0 {
  unit 3 {
    description P2-to-PE1;
    family inet {
      address 10.0.0.6/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 5 {
    description P2-to-P1;
    family inet {
      address 10.0.0.30/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 6 {
    description P2-to-PE4;
    family inet {
      address 10.0.0.13/30;
    }
    family mpls;
  }
}
lo0 {
```

```
    unit 5 {
      family inet {
        address 10.9.9.3/32;
        address 10.100.1.3/32;
      }
    }
  }

user@P2# show protocols
rsvp {
  interface fe-1/2/1.5;
  interface fe-1/2/2.6;
  interface fe-1/2/0.3;
}
mpls {
  label-switched-path P2-to-PE1 {
    to 10.9.9.1;
  }
  label-switched-path P2-to-P1 {
    to 10.9.9.2;
  }
  label-switched-path P2-to-PE4 {
    to 10.9.9.4;
  }
  interface fe-1/2/1.5;
  interface fe-1/2/2.6;
  interface fe-1/2/0.3;
}
bgp {
  group internal {
    type internal;
    local-address 10.9.9.3;
    family inet {
      labeled-unicast {
        aigp;
      }
    }
    neighbor 10.9.9.1;
    neighbor 10.9.9.2;
    neighbor 10.9.9.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/2.6 {
      metric 1;
    }
    interface 10.9.9.3 {
      passive;
      metric 1;
    }
    interface 10.100.1.3 {
      passive;
      metric 1;
    }
  }
}
```

```

}
user@P2# show routing-options
router-id 10.9.9.3;
autonomous-system 13979;

```

### Configuring Device PE4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE4:

1. Configure the interfaces.

```

[edit interfaces]
user@PE4# set fe-1/2/0 unit 7 description PE4-to-P2
user@PE4# set fe-1/2/0 unit 7 family inet address 10.0.0.14/30
user@PE4# set fe-1/2/0 unit 7 family mpls
user@PE4# set fe-1/2/1 unit 9 description PE4-to-P1
user@PE4# set fe-1/2/1 unit 9 family inet address 10.0.0.18/30
user@PE4# set fe-1/2/1 unit 9 family mpls
user@PE4# set fe-1/2/2 unit 10 description PE4-to-PE2
user@PE4# set fe-1/2/2 unit 10 family inet address 10.0.0.21/30
user@PE4# set fe-1/2/2 unit 10 family mpls
user@PE4# set fe-1/0/2 unit 12 description PE4-to-PE3
user@PE4# set fe-1/0/2 unit 12 family inet address 10.0.0.25/30
user@PE4# set fe-1/0/2 unit 12 family mpls
user@PE4# set lo0 unit 7 family inet address 10.9.9.4/32
user@PE4# set lo0 unit 7 family inet address 10.100.1.4/32

```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```

[edit protocols]
user@PE4# set rsvp interface fe-1/2/0.7
user@PE4# set rsvp interface fe-1/2/1.9
user@PE4# set rsvp interface fe-1/2/2.10
user@PE4# set rsvp interface fe-1/0/2.12
user@PE4# set mpls label-switched-path PE4-to-PE2 to 10.9.9.5
user@PE4# set mpls label-switched-path PE4-to-PE3 to 10.9.9.6
user@PE4# set mpls label-switched-path PE4-to-P1 to 10.9.9.2
user@PE4# set mpls label-switched-path PE4-to-P2 to 10.9.9.3
user@PE4# set mpls interface fe-1/2/0.7
user@PE4# set mpls interface fe-1/2/1.9
user@PE4# set mpls interface fe-1/2/2.10
user@PE4# set mpls interface fe-1/0/2.12

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE4# set export next-hop
user@PE4# set export aigp
user@PE4# set group internal type internal
user@PE4# set group internal local-address 10.9.9.4
user@PE4# set group internal neighbor 10.9.9.1
user@PE4# set group internal neighbor 10.9.9.3

```

```
user@PE4# set group internal neighbor 10.9.9.2
user@PE4# set group external type external
user@PE4# set group external multihop ttl 2
user@PE4# set group external local-address 10.9.9.4
user@PE4# set group external peer-as 7018
user@PE4# set group external neighbor 10.9.9.5
user@PE4# set group external neighbor 10.9.9.6
```

4. Enable AIGP.

```
[edit protocols bgp]
user@PE4# set group external family inet labeled-unicast aigp
user@PE4# set group internal family inet labeled-unicast aigp
```

5. Originate a prefix, and configure an AIGP distance.

By default, a prefix is originated using the current IGP distance. Optionally, you can configure a distance for the AIGP attribute, using the **distance** option, as shown here.

```
[edit policy-options policy-statement aigp term 10]
user@PE4# set from protocol static
user@PE4# set from route-filter 44.0.0.0/24 exact
user@PE4# set then aigp-originate distance 200
user@PE4# set then next-hop 10.100.1.4
user@PE4# set then accept
```

6. Enable the policies.

```
[edit policy-options policy-statement next-hop]
user@PE4# set term 10 from protocol bgp
user@PE4# set term 10 then next-hop 10.100.1.4
user@PE4# set term 10 then accept
user@PE4# set term 20 from protocol direct
user@PE4# set term 20 from route-filter 10.9.9.4/32 exact
user@PE4# set term 20 from route-filter 10.100.1.4/32 exact
user@PE4# set term 20 then next-hop 10.100.1.4
user@PE4# set term 20 then accept
```

7. Configure a static route.

```
[edit routing-options]
user@PE4# set static route 44.0.0.0/24 discard
```

8. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf]
user@PE4# set area 0.0.0.0 interface fe-1/2/1.9 metric 1
user@PE4# set area 0.0.0.0 interface fe-1/2/0.7 metric 1
user@PE4# set area 0.0.0.0 interface 10.9.9.4 passive
user@PE4# set area 0.0.0.0 interface 10.9.9.4 metric 1
user@PE4# set area 0.0.0.0 interface 10.100.1.4 passive
user@PE4# set area 0.0.0.0 interface 10.100.1.4 metric 1
user@PE4# set area 0.0.0.2 interface fe-1/2/2.10 metric 1
user@PE4# set area 0.0.0.3 interface fe-1/0/2.12 metric 1
```

9. Configure the router ID and the autonomous system number.

```
[edit routing-options]
```



```

user@PE4# set router-id 10.9.9.4
user@PE4# set autonomous-system 13979

```

10. If you are done configuring the device, commit the configuration.

```

user@PE4# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE4# show interfaces
fe-1/0/2 {
  unit 12 {
    description PE4-to-PE3;
    family inet {
      address 10.0.0.25/30;
    }
    family mpls;
  }
}
fe-1/2/0 {
  unit 7 {
    description PE4-to-P2;
    family inet {
      address 10.0.0.14/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 9 {
    description PE4-to-P1;
    family inet {
      address 10.0.0.18/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 10 {
    description PE4-to-PE2;
    family inet {
      address 10.0.0.21/30;
    }
    family mpls;
  }
}
lo0 {
  unit 7 {
    family inet {
      address 10.9.9.4/32;
      address 10.100.1.4/32;
    }
  }
}

```

```
    }  
  }  
user@PE4# show policy-options  
policy-statement aigp {  
  term 10 {  
    from {  
      protocol static;  
      route-filter 44.0.0.0/24 exact;  
    }  
    then {  
      aigp-originate distance 200;  
      next-hop 10.100.1.4;  
      accept;  
    }  
  }  
}  
policy-statement next-hop {  
  term 10 {  
    from protocol bgp;  
    then {  
      next-hop 10.100.1.4;  
      accept;  
    }  
  }  
  term 20 {  
    from {  
      protocol direct;  
      route-filter 10.9.9.4/32 exact;  
      route-filter 10.100.1.4/32 exact;  
    }  
    then {  
      next-hop 10.100.1.4;  
      accept;  
    }  
  }  
}  
user@PE4# show protocols  
rsvp {  
  interface fe-1/2/0.7;  
  interface fe-1/2/1.9;  
  interface fe-1/2/2.10;  
  interface fe-1/0/2.12;  
}  
mpls {  
  label-switched-path PE4-to-PE2 {  
    to 10.9.9.5;  
  }  
  label-switched-path PE4-to-PE3 {  
    to 10.9.9.6;  
  }  
  label-switched-path PE4-to-P1 {  
    to 10.9.9.2;  
  }  
  label-switched-path PE4-to-P2 {  
    to 10.9.9.3;  
  }  
}
```

```

    }
    interface fe-1/2/0.7;
    interface fe-1/2/1.9;
    interface fe-1/2/2.10;
    interface fe-1/0/2.12;
  }
  bgp {
    export [ next-hop aigp ];
    group internal {
      type internal;
      local-address 10.9.9.4;
      family inet {
        labeled-unicast {
          aigp;
        }
      }
      neighbor 10.9.9.1;
      neighbor 10.9.9.3;
      neighbor 10.9.9.2;
    }
    group external {
      type external;
      multihop {
        ttl 2;
      }
      local-address 10.9.9.4;
      family inet {
        labeled-unicast {
          aigp;
        }
      }
      peer-as 7018;
      neighbor 10.9.9.5;
      neighbor 10.9.9.6;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface fe-1/2/1.9 {
        metric 1;
      }
      interface fe-1/2/0.7 {
        metric 1;
      }
      interface 10.9.9.4 {
        passive;
        metric 1;
      }
      interface 10.100.1.4 {
        passive;
        metric 1;
      }
    }
    area 0.0.0.2 {
      interface fe-1/2/2.10 {
        metric 1;
      }
    }
  }

```

```
    }  
  }  
  area 0.0.0.3 {  
    interface fe-1/0/2.12 {  
      metric 1;  
    }  
  }  
}  
  
user@PE4# show routing-options  
static {  
  route 44.0.0.0/24 discard;  
}  
router-id 10.9.9.4;  
autonomous-system 13979;
```

### *Configuring Device PE1*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.

```
[edit interfaces]  
user@PE1# set fe-1/2/0 unit 0 description PE1-to-P1  
user@PE1# set fe-1/2/0 unit 0 family inet address 10.0.0.1/30  
user@PE1# set fe-1/2/0 unit 0 family mpls  
user@PE1# set fe-1/2/1 unit 2 description PE1-to-P2  
user@PE1# set fe-1/2/1 unit 2 family inet address 10.0.0.5/30  
user@PE1# set fe-1/2/1 unit 2 family mpls  
user@PE1# set fe-1/2/2 unit 14 description PE1-to-PE7  
user@PE1# set fe-1/2/2 unit 14 family inet address 10.0.0.9/30  
user@PE1# set lo0 unit 1 family inet address 10.9.9.1/32  
user@PE1# set lo0 unit 1 family inet address 10.100.1.1/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]  
user@PE1# set rsvp interface fe-1/2/0.0  
user@PE1# set rsvp interface fe-1/2/1.2  
user@PE1# set rsvp interface fe-1/2/2.14  
user@PE1# set mpls label-switched-path PE1-to-P1 to 10.9.9.2  
user@PE1# set mpls label-switched-path PE1-to-P2 to 10.9.9.3  
user@PE1# set mpls interface fe-1/2/0.0  
user@PE1# set mpls interface fe-1/2/1.2  
user@PE1# set mpls interface fe-1/2/2.14
```

3. Configure BGP.

```
[edit protocols bgp]  
user@PE1# set group internal type internal  
user@PE1# set group internal local-address 10.9.9.1  
user@PE1# set group internal export SET_EXPORT_ROUTES  
user@PE1# set group internal vpn-apply-export
```

```

user@PE1# set group internal neighbor 10.9.9.4
user@PE1# set group internal neighbor 10.9.9.2
user@PE1# set group internal neighbor 10.9.9.3
user@PE1# set group external type external
user@PE1# set group external export SET_EXPORT_ROUTES
user@PE1# set group external peer-as 7019
user@PE1# set group external neighbor 10.0.0.10

```

4. Enable AIGP.

```

[edit protocols bgp]
user@PE1# set group internal family inet labeled-unicast aigp
user@PE1# set group external family inet labeled-unicast aigp

```

5. Enable the policies.

```

[edit policy-options policy-statement SET_EXPORT_ROUTES term 10]
user@PE1# set from protocol direct
user@PE1# set from protocol bgp
user@PE1# set then next-hop 10.100.1.1
user@PE1# set then accept

```

6. Configure an IGP, such as OSPF, RIP, or IS-IS.

```

[edit protocols ospf area 0.0.0.1]
user@PE1# set interface fe-1/2/0.0 metric 1
user@PE1# set interface fe-1/2/1.2 metric 1
user@PE1# set interface 10.9.9.1 passive
user@PE1# set interface 10.9.9.1 metric 1
user@PE1# set interface 10.100.1.1 passive
user@PE1# set interface 10.100.1.1 metric 1

```

7. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@PE1# set router-id 10.9.9.1
user@PE1# set autonomous-system 13979

```

8. If you are done configuring the device, commit the configuration.

```

user@PE1# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE1# show interfaces
fe-1/2/0 {
  unit 0 {
    description PE1-to-P1;
    family inet {
      address 10.0.0.1/30;
    }
    family mpls;
  }
}
fe-1/2/1 {

```

```
    unit 2 {
      description PE1-to-P2;
      family inet {
        address 10.0.0.5/30;
      }
      family mpls;
    }
  }
  fe-1/2/2 {
    unit 14 {
      description PE1-to-PE7;
      family inet {
        address 10.0.0.9/30;
      }
    }
  }
  lo0 {
    unit 1 {
      family inet {
        address 10.9.9.1/32;
        address 10.100.1.1/32;
      }
    }
  }
}

user@PE1# show policy-options
policy-statement SET_EXPORT_ROUTES {
  term 10 {
    from protocol [ direct bgp ];
    then {
      next-hop 10.100.1.1;
      accept;
    }
  }
}

user@PE1# show protocols
rsvp {
  interface fe-1/2/0.0;
  interface fe-1/2/1.2;
  interface fe-1/2/2.14;
}
mpls {
  label-switched-path PE1-to-P1 {
    to 10.9.9.2;
  }
  label-switched-path PE1-to-P2 {
    to 10.9.9.3;
  }
  interface fe-1/2/0.0;
  interface fe-1/2/1.2;
  interface fe-1/2/2.14;
}
bgp {
  group internal {
    type internal;
    local-address 10.9.9.1;
```

```

    family inet {
        labeled-unicast {
            aigp;
        }
    }
    export SET_EXPORT_ROUTES;
    vpn-apply-export;
    neighbor 10.9.9.4;
    neighbor 10.9.9.2;
    neighbor 10.9.9.3;
}
group external {
    type external;
    family inet {
        labeled-unicast {
            aigp;
        }
    }
    export SET_EXPORT_ROUTES;
    peer-as 7019;
    neighbor 10.0.0.10;
}
}
ospf {
    area 0.0.0.1 {
        interface fe-1/2/0.0 {
            metric 1;
        }
        interface fe-1/2/1.2 {
            metric 1;
        }
        interface 10.9.9.1 {
            passive;
            metric 1;
        }
        interface 10.100.1.1 {
            passive;
            metric 1;
        }
    }
}
}

```

```

user@PE1# show routing-options
router-id 10.9.9.1;
autonomous-system 13979;

```

### Configuring Device PE2

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE2:

1. Configure the interfaces.

```
[edit interfaces]
user@PE2# set fe-1/2/0 unit 11 description PE2-to-PE4
user@PE2# set fe-1/2/0 unit 11 family inet address 10.0.0.22/30
user@PE2# set fe-1/2/0 unit 11 family mpls
user@PE2# set lo0 unit 9 family inet address 10.9.9.5/32 primary
user@PE2# set lo0 unit 9 family inet address 10.100.1.5/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE2# set rsvp interface fe-1/2/0.11
user@PE2# set mpls label-switched-path PE2-to-PE4 to 10.9.9.4
user@PE2# set mpls interface fe-1/2/0.11
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE2# set group external type external
user@PE2# set group external multihop ttl 2
user@PE2# set group external local-address 10.9.9.5
user@PE2# set group external export next-hop
user@PE2# set group external export aigp
user@PE2# set group external export SET_EXPORT_ROUTES
user@PE2# set group external vpn-apply-export
user@PE2# set group external peer-as 13979
user@PE2# set group external neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp]
user@PE2# set group external family inet labeled-unicast aigp
```

5. Originate a prefix, and configure an AIGP distance.

By default, a prefix is originated using the current IGP distance. Optionally, you can configure a distance for the AIGP attribute, using the **distance** option, as shown here.

```
[edit policy-options policy-statement aigp]
user@PE2# set term 10 from route-filter 55.0.0.0/24 exact
user@PE2# set term 10 then aigp-originate distance 20
user@PE2# set term 10 then next-hop 10.100.1.5
user@PE2# set term 10 then accept
user@PE2# set term 20 from route-filter 99.0.0.0/24 exact
user@PE2# set term 20 then aigp-originate distance 30
user@PE2# set term 20 then next-hop 10.100.1.5
user@PE2# set term 20 then accept
```



6. Enable the policies.

```
[edit policy-options]
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
direct
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
static
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
bgp
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 then next-hop
10.100.1.5
user@PE2# set policy-statement SET_EXPORT_ROUTES term 10 then accept
user@PE2# set policy-statement next-hop term 10 from protocol bgp
user@PE2# set policy-statement next-hop term 10 then next-hop 10.100.1.5
user@PE2# set policy-statement next-hop term 10 then accept
user@PE2# set policy-statement next-hop term 20 from protocol direct
user@PE2# set policy-statement next-hop term 20 from route-filter 10.9.9.5/32
exact
user@PE2# set policy-statement next-hop term 20 from route-filter 10.100.1.5/32
exact
user@PE2# set policy-statement next-hop term 20 then next-hop 10.100.1.5
user@PE2# set policy-statement next-hop term 20 then accept
```

7. Enable some static routes.

```
[edit routing-options]
user@PE2# set static route 99.0.0.0/24 discard
user@PE2# set static route 55.0.0.0/24 discard
```

8. Configure an IGP, such as OSPF, RIP, or IS-IS.

```
[edit protocols ospf area 0.0.0.2]
user@PE2# set interface 10.9.9.5 passive
user@PE2# set interface 10.9.9.5 metric 1
user@PE2# set interface 10.100.1.5 passive
user@PE2# set interface 10.100.1.5 metric 1
user@PE2# set interface fe-1/2/0.11 metric 1
```

9. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@PE2# set router-id 10.9.9.5
user@PE2# set autonomous-system 7018
```

10. If you are done configuring the device, commit the configuration.

```
user@PE2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
fe-1/2/0 {
  unit 11 {
    description PE2-to-PE4;
```

```
        family inet {
            address 10.0.0.22/30;
        }
        family mpls;
    }
}
lo0 {
    unit 9 {
        family inet {
            address 10.9.9.5/32 {
                primary;
            }
            address 10.100.1.5/32;
        }
    }
}

user@PE2# show policy-options
policy-statement SET_EXPORT_ROUTES {
    term 10 {
        from protocol [ direct static bgp ];
        then {
            next-hop 10.100.1.5;
            accept;
        }
    }
}

policy-statement aigp {
    term 10 {
        from {
            route-filter 55.0.0.0/24 exact;
        }
        then {
            aigp-originate distance 20;
            next-hop 10.100.1.5;
            accept;
        }
    }
    term 20 {
        from {
            route-filter 99.0.0.0/24 exact;
        }
        then {
            aigp-originate distance 30;
            next-hop 10.100.1.5;
            accept;
        }
    }
}

policy-statement next-hop {
    term 10 {
        from protocol bgp;
        then {
            next-hop 10.100.1.5;
            accept;
        }
    }
}
```

```

    }
    term 20 {
        from {
            protocol direct;
            route-filter 10.9.9.5/32 exact;
            route-filter 10.100.1.5/32 exact;
        }
        then {
            next-hop 10.100.1.5;
            accept;
        }
    }
}

user@PE2# show protocols
rsvp {
    interface fe-1/2/0.11;
}
mpls {
    label-switched-path PE2-to-PE4 {
        to 10.9.9.4;
    }
    interface fe-1/2/0.11;
}
bgp {
    group external {
        type external;
        multihop {
            ttl 2;
        }
        local-address 10.9.9.5;
        family inet {
            labeled-unicast {
                aigp;
            }
        }
        export [ next-hop aigp SET_EXPORT_ROUTES ];
        vpn-apply-export;
        peer-as 13979;
        neighbor 10.9.9.4;
    }
}
ospf {
    area 0.0.0.2 {
        interface 10.9.9.5 {
            passive;
            metric 1;
        }
        interface 10.100.1.5 {
            passive;
            metric 1;
        }
        interface fe-1/2/0.11 {
            metric 1;
        }
    }
}

```

```
}
user@PE2# show routing-options
static {
    route 99.0.0.0/24 discard;
    route 55.0.0.0/24 discard;
}
router-id 10.9.9.5;
autonomous-system 7018;
```

### **Configuring Device PE3**

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE3:

1. Configure the interfaces.

```
[edit interfaces]
user@PE3# set fe-1/2/0 unit 13 description PE3-to-PE4
user@PE3# set fe-1/2/0 unit 13 family inet address 10.0.0.26/30
user@PE3# set fe-1/2/0 unit 13 family mpls
user@PE3# set lo0 unit 11 family inet address 10.9.9.6/32
user@PE3# set lo0 unit 11 family inet address 10.100.1.6/32
```

2. Configure MPLS and a signaling protocol, such as RSVP or LDP.

```
[edit protocols]
user@PE3# set rsvp interface fe-1/2/0.13
user@PE3# set mpls label-switched-path PE3-to-PE4 to 10.9.9.4
user@PE3# set mpls interface fe-1/2/0.13
```

3. Configure BGP.

```
[edit protocols bgp group external]
user@PE3# set type external
user@PE3# set multihop ttl 2
user@PE3# set local-address 10.9.9.6
user@PE3# set export next-hop
user@PE3# set export SET_EXPORT_ROUTES
user@PE3# set vpn-apply-export
user@PE3# set peer-as 13979
user@PE3# set neighbor 10.9.9.4
```

4. Enable AIGP.

```
[edit protocols bgp group external]
user@PE3# set family inet labeled-unicast aigp
```

5. Enable the policies.

```
[edit policy-options]
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
    direct
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
    static
```

```

user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 from protocol
bgp
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 then next-hop
10.100.1.6
user@PE3# set policy-statement SET_EXPORT_ROUTES term 10 then accept
user@PE3# set policy-statement next-hop term 10 from protocol bgp
user@PE3# set policy-statement next-hop term 10 then next-hop 10.100.1.6
user@PE3# set policy-statement next-hop term 10 then accept
user@PE3# set policy-statement next-hop term 20 from protocol direct
user@PE3# set policy-statement next-hop term 20 from route-filter 10.9.9.6/32
exact
user@PE3# set policy-statement next-hop term 20 from route-filter 10.100.1.6/32
exact
user@PE3# set policy-statement next-hop term 20 then next-hop 10.100.1.6
user@PE3# set policy-statement next-hop term 20 then accept

```

6. Configure an IGP, such as OSPF, RIP, or IS-IS.

```

[edit protocols ospf area 0.0.0.3]
user@PE3# set interface 10.9.9.6 passive
user@PE3# set interface 10.9.9.6 metric 1
user@PE3# set interface 10.100.1.6 passive
user@PE3# set interface 10.100.1.6 metric 1
user@PE3# set interface fe-1/2/0.13 metric 1

```

7. Configure the router ID and the autonomous system number.

```

[edit routing-options]
user@PE3# set router-id 10.9.9.6
user@PE3# set autonomous-system 7018

```

8. If you are done configuring the device, commit the configuration.

```

user@PE3# commit

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE3# show interfaces
fe-1/2/0 {
  unit 13 {
    description PE3-to-PE4;
    family inet {
      address 10.0.0.26/30;
    }
    family mpls;
  }
}
lo0 {
  unit 11 {
    family inet {
      address 10.9.9.6/32;
      address 10.100.1.6/32;
    }
  }
}

```

```
    }  
  }  
user@PE3# show policy-options  
policy-statement SET_EXPORT_ROUTES {  
  term 10 {  
    from protocol [ direct static bgp ];  
    then {  
      next-hop 10.100.1.6;  
      accept;  
    }  
  }  
}  
policy-statement next-hop {  
  term 10 {  
    from protocol bgp;  
    then {  
      next-hop 10.100.1.6;  
      accept;  
    }  
  }  
  term 20 {  
    from {  
      protocol direct;  
      route-filter 10.9.9.6/32 exact;  
      route-filter 10.100.1.6/32 exact;  
    }  
    then {  
      next-hop 10.100.1.6;  
      accept;  
    }  
  }  
}  
user@PE3# show protocols  
rsvp {  
  interface fe-1/2/0.13;  
}  
mpls {  
  label-switched-path PE3-to-PE4 {  
    to 10.9.9.4;  
  }  
  interface fe-1/2/0.13;  
}  
bgp {  
  group external {  
    type external;  
    multihop {  
      ttl 2;  
    }  
    local-address 10.9.9.6;  
    family inet {  
      labeled-unicast {  
        aigp;  
      }  
    }  
  }  
  export [ next-hop SET_EXPORT_ROUTES ];
```

```

        vpn-apply-export;
        peer-as 13979;
        neighbor 10.9.9.4;
    }
}
ospf {
    area 0.0.0.3 {
        interface 10.9.9.6 {
            passive;
            metric 1;
        }
        interface 10.100.1.6 {
            passive;
            metric 1;
        }
        interface fe-1/2/0.13 {
            metric 1;
        }
    }
}

user@PE3# show routing-options
router-id 10.9.9.6;
autonomous-system 7018;

```

### Configuring Device PE7

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE7:

1. Configure the interfaces.

```

[edit interfaces]
user@PE7# set fe-1/2/0 unit 15 description PE7-to-PE1
user@PE7# set fe-1/2/0 unit 15 family inet address 10.0.0.10/30
user@PE7# set lo0 unit 13 family inet address 10.9.9.7/32
user@PE7# set lo0 unit 13 family inet address 10.100.1.7/32

```

2. Configure BGP.

```

[edit protocols bgp group external]
user@PE7# set type external
user@PE7# set export SET_EXPORT_ROUTES
user@PE7# set peer-as 13979
user@PE7# set neighbor 10.0.0.9

```

3. Enable AIGP.

```

[edit protocols bgp group external]
user@PE7# set family inet labeled-unicast aigp

```

4. Configure the routing policy.

```

[edit policy-options policy-statement SET_EXPORT_ROUTES term 10]
user@PE7# set from protocol direct

```

```
user@PE7# set from protocol bgp
user@PE7# set then next-hop 10.100.1.7
user@PE7# set then accept
```

5. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@PE7# set router-id 10.9.9.7
user@PE7# set autonomous-system 7019
```

6. If you are done configuring the device, commit the configuration.

```
user@PE7# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE7# show interfaces
interfaces {
  fe-1/2/0 {
    unit 15 {
      description PE7-to-PE1;
      family inet {
        address 10.0.0.10/30;
      }
    }
  }
  lo0 {
    unit 13 {
      family inet {
        address 10.9.9.7/32;
        address 10.100.1.7/32;
      }
    }
  }
}

user@PE7# show policy-options
policy-statement SET_EXPORT_ROUTES {
  term 10 {
    from protocol [ direct bgp ];
    then {
      next-hop 10.100.1.7;
      accept;
    }
  }
}

user@PE7# show protocols
bgp {
  group external {
    type external;
    family inet {
      labeled-unicast {
        aigp;
      }
    }
  }
}
```



```

    }
  }
  export SET_EXPORT_ROUTES;
  peer-as 13979;
  neighbor 10.0.0.9;
}
}

user@PE7# show routing-options
router-id 10.9.9.7;
autonomous-system 7019;

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That Device PE4 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE2 on page 3415](#)
- [Checking the IGP Metric on page 3415](#)
- [Verifying That Device PE4 Adds the IGP Metric to the AIGP Attribute on page 3416](#)
- [Verifying That Device PE7 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE1 on page 3416](#)
- [Verifying the Resolving AIGP Metric on page 3417](#)
- [Verifying the Presence of AIGP Attributes in BGP Updates on page 3420](#)

### *Verifying That Device PE4 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE2*

**Purpose** Make sure that the AIGP policy on Device PE2 is working.

**Action**

```

user@PE4> show route receive-protocol bgp 10.9.9.5 extensive
* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 299888
  Nexthop: 10.100.1.5
  AS path: 7018 I
  AIGP: 20

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 299888
  Nexthop: 10.100.1.5
  AS path: 7018 I
  AIGP: 30

```

**Meaning** On Device PE2, the **aigp-originate** statement is configured with a distance of 20 (**aigp-originate distance 20**). This statement is applied to route 55.0.0.0/24. Likewise, the **aigp-originate distance 30** statement is applied to route 99.0.0.0/24. Thus, when Device PE4 receives these routes, the AIGP attribute is attached with the configured metrics.

### *Checking the IGP Metric*

**Purpose** From Device PE4, check the IGP metric to the BGP next hop 10.100.1.5.

**Action** user@PE4> show route 10.100.1.5  
inet.0: 30 destinations, 40 routes (30 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

10.100.1.5/32      \*[OSPF/10] 05:35:50, metric 2  
                         > to 10.0.0.22 via fe-1/2/2.10  
                         [BGP/170] 03:45:07, localpref 100, from 10.9.9.5  
                         AS path: 7018 I  
                         > to 10.0.0.22 via fe-1/2/2.10

**Meaning** The IGP metric for this route is 2.

***Verifying That Device PE4 Adds the IGP Metric to the AIGP Attribute***

**Purpose** Make sure that Device PE4 adds the IGP metric to the AIGP attribute when it readvertises routes to its IBGP neighbor, Device PE1.

**Action** user@PE4> show route advertising-protocol bgp 10.9.9.1 extensive

\* 55.0.0.0/24 (1 entry, 1 announced)  
BGP group internal type Internal  
Route Label: 300544  
Nexthop: 10.100.1.4  
Flags: Nexthop Change  
Localpref: 100  
AS path: [13979] 7018 I  
AIGP: 22

\* 99.0.0.0/24 (1 entry, 1 announced)  
BGP group internal type Internal  
Route Label: 300544  
Nexthop: 10.100.1.4  
Flags: Nexthop Change  
Localpref: 100  
AS path: [13979] 7018 I  
AIGP: 32

**Meaning** The IGP metric is added to the AIGP metric ( $20 + 2 = 22$  and  $30 + 2 = 32$ ), because the next hop is changed for these routes.

***Verifying That Device PE7 Is Receiving the AIGP Attribute from Its EBGp Neighbor PE1***

**Purpose** Make sure that the AIGP policy on Device PE1 is working.

**Action** user@PE7> show route receive-protocol bgp 10.0.0.9 extensive

\* 44.0.0.0/24 (1 entry, 1 announced)

Accepted  
Route Label: 300096  
Nexthop: 10.0.0.9  
AS path: 13979 I  
AIGP: 203

\* 55.0.0.0/24 (1 entry, 1 announced)

Accepted  
Route Label: 300112  
Nexthop: 10.0.0.9  
AS path: 13979 7018 I  
AIGP: 25

\* 99.0.0.0/24 (1 entry, 1 announced)

Accepted  
Route Label: 300112  
Nexthop: 10.0.0.9  
AS path: 13979 7018 I  
AIGP: 35

**Meaning** The 44.0.0.0/24 route is originated at Device PE4. The 55.0.0.0/24 and 99.0.0.0/24 routes are originated at Device PE2. The IGP distances are added to the configured AIGP distances.

#### *Verifying the Resolving AIGP Metric*

**Purpose** Confirm that if the prefix is resolved through recursion and the recursive next hops have AIGP metrics, the prefix has the sum of the AIGP values that are on the recursive BGP next hops.

**Action** 1. Add a static route to 66.0.0.0/24.

```
[edit routing-options]
user@PE2# set static route 66.0.0.0/24 discard
```

2. Delete the existing terms in the **aigp** policy statement on Device PE2.

```
[edit policy-options policy-statement aigp]
user@PE2# delete term 10
user@PE2# delete term 20
```

3. Configure a recursive route lookup for the route to 66.0.0.0.

The policy shows the AIGP metric for prefix 66.0.0.0/24 (none) and its recursive next hop. Prefix 66.0.0.0/24 is resolved by 55.0.0.1. Prefix 66.0.0.0/24 does not have its own AIGP metric being originated, but its recursive next hop, 55.0.0.1, has an AIGP value.

```
[edit policy-options policy-statement aigp]
user@PE2# set term 10 from route-filter 55.0.0.1/24 exact
user@PE2# set term 10 then aigp-originate distance 20
user@PE2# set term 10 then next-hop 10.100.1.5
user@PE2# set term 10 then accept
user@PE2# set term 20 from route-filter 66.0.0.0/24 exact
user@PE2# set term 20 then next-hop 55.0.0.1
```

user@PE2# set term 20 then accept

4. On Device PE4, run the **show route 55.0.0.0 extensive** command.

The value of Metric2 is the IGP metric to the BGP next hop. When Device PE4 readvertises these routes to its IBGP peer, Device PE1, the AIGP metric is the sum of AIGP + its Resolving AIGP metric + Metric2.

Prefix 55.0.0.0 shows its own IGP metric 20, as defined and advertised by Device PE2. It does not show a resolving AIGP value because it does not have a recursive BGP next hop. The value of Metric2 is 2.

```
user@PE4> show route 55.0.0.0 extensive
inet.0: 31 destinations, 41 routes (31 active, 0 holddown, 0 hidden)
55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 55.0.0.0/24 -> {indirect(262151)}
Page 0 idx 0 Type 1 val 928d1b8
  Flags: Nexthop Change
  Nexthop: 10.100.1.4
  Localpref: 100
  AS path: [13979] 7018 I
  Communities:
  AIGP: 22
Path 55.0.0.0 from 10.9.9.5 Vector len 4. Val: 0
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x925da38
    Next-hop reference count: 4
    Source: 10.9.9.5
    Next hop type: Router, Next hop index: 1004
    Next hop: 10.0.0.22 via fe-1/2/2.10, selected
    Label operation: Push 299888
    Label TTL action: prop-ttl
    Protocol next hop: 10.100.1.5
    Push 299888
    Indirect next hop: 93514d8 262151
    State: <Active Ext>
    Local AS: 13979 Peer AS: 7018
    Age: 22:03:26 Metric2: 2
    AIGP: 20
    Task: BGP_7018.10.9.9.5+58560
    Announcement bits (3): 3-KRT 4-BGP_RT_Background 5-Resolve tree 1
    AS path: 7018 I
    Accepted
    Route Label: 299888
    Localpref: 100
    Router ID: 10.9.9.5
    Indirect next hops: 1
      Protocol next hop: 10.100.1.5 Metric: 2
      Push 299888
      Indirect next hop: 93514d8 262151
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.0.22 via fe-1/2/2.10
      10.100.1.5/32 Originating RIB: inet.0
        Metric: 2 Node path count: 1
        Forwarding nexthops: 1
          Nexthop: 10.0.0.22 via fe-1/2/2.10
```

5. On Device PE4, run the **show route 66.0.0.0 extensive** command.

Prefix 66.0.0.0/24 shows the Resolving AIGP, which is the sum of its own AIGP metric and its recursive BGP next hop:

66.0.0.1 = 0, 55.0.0.1 = 20, 0+20 = 20

```

user@PE4> show route 66.0.0.0 extensive
inet.0: 31 destinations, 41 routes (31 active, 0 holddown, 0 hidden)
66.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT in-kerne1 66.0.0.0/24 -> {indirect(262162)}
Page 0 idx 0 Type 1 val 928cefc
  Flags: Nexthop Change
  Nexthop: 10.100.1.4
  Localpref: 100
  AS path: [13979] 7018 I
  Communities:
Path 66.0.0.0 from 10.9.9.5 Vector len 4. Val: 0
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x925d4e0
    Next-hop reference count: 4
    Source: 10.9.9.5
    Next hop type: Router, Next hop index: 1006
    Next hop: 10.0.0.22 via fe-1/2/2.10, selected
    Label operation: Push 299888, Push 299888(top)
    Label TTL action: prop-ttl, prop-ttl(top)
    Protocol next hop: 55.0.0.1
    Push 299888
    Indirect next hop: 9353e88 262162
    State: <Active Ext>
    Local AS: 13979 Peer AS: 7018
    Age: 31:42 Metric2:2
    Resolving-AIGP: 20
    Task: BGP_7018.10.9.9.5+58560
    Announcement bits (3): 3-KRT 4-BGP_RT_Background 5-Resolve tree 1
    AS path: 7018 I
    Accepted
    Route Label: 299888
    Localpref: 100
    Router ID: 10.9.9.5
    Indirect next hops: 1
      Protocol next hop: 55.0.0.1 Metric: 2 AIGP: 20
      Push 299888
      Indirect next hop: 9353e88 262162
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.0.22 via fe-1/2/2.10
      55.0.0.0/24 Originating RIB: inet.0
        Metric: 2 Node path count: 1
        Indirect nexthops: 1
          Protocol Nexthop: 10.100.1.5 Metric: 2 Push 299888
          Indirect nexthop: 93514d8 262151
          Indirect path forwarding nexthops: 1
            Nexthop: 10.0.0.22 via fe-1/2/2.10
          10.100.1.5/32 Originating RIB: inet.0
            Metric: 2 Node path count: 1
            Forwarding nexthops: 1
              Nexthop: 10.0.0.22 via fe-1/2/2.10

```

### *Verifying the Presence of AIGP Attributes in BGP Updates*

**Purpose** If the AIGP attribute is not enabled under BGP (or the **group** or **neighbor** hierarchies), the AIGP attribute is silently discarded. Enable **traceoptions** and include the **packets** flag in the **detail** option in the configuration to confirm the presence of the AIGP attribute in transmitted or received BGP updates. This is useful when debugging AIGP issues.

**Action** 1. Configure Device PE2 and Device PE4 for **traceoptions**.

```
user@host> show protocols bgp
traceoptions {
  file bgp size 1m files 5;
  flag packets detail;
}
```

2. Check the **traceoptions** file on Device PE2.

The following sample shows Device PE2 advertising prefix 99.0.0.0/24 to Device PE4 (10.9.9.4) with an AIGP metric of 20:

```
user@PE2> show log bgp
Mar 22 09:27:18.982150 BGP SEND 10.9.9.5+49652 -> 10.9.9.4+179
Mar 22 09:27:18.982178 BGP SEND message type 2 (Update) length 70
Mar 22 09:27:18.982198 BGP SEND Update PDU length 70
Mar 22 09:27:18.982248 BGP SEND flags 0x40 code Origin(1): IGP
Mar 22 09:27:18.982273 BGP SEND flags 0x40 code ASPath(2) length 6: 7018
Mar 22 09:27:18.982295 BGP SEND flags 0x80 code AIGP(26): AIGP: 20
Mar 22 09:27:18.982316 BGP SEND flags 0x90 code MP_reach(14): AFI/SAFI 1/4
Mar 22 09:27:18.982341 BGP SEND      nhop 10.100.1.5 len 4
Mar 22 09:27:18.982372 BGP SEND    99.0.0.0/24 (label 301664)
Mar 22 09:27:33.665412 bgp_send: sending 19 bytes to abcd::10:255:170:84
(External AS 13979)
```

3. Verify that the route was received on Device PE4 using the **show route receive-protocol** command.

AIGP is not enabled on Device PE4, so the AIGP attribute is silently discarded for prefix 99.0.0.0/24 and does not appear in the following output:

```
user@PE4> show route receive-protocol bgp 10.9.9.5 extensive | find 55.0.0.0
* 99.0.0.0/24 (2 entries, 1 announced)
  Accepted
  Route Label: 301728
  Nexthop: 10.100.1.5
  AS path: 7018 I
```

4. Check the **traceoptions** file on Device PE4.

The following output from the **traceoptions** log shows that the 99.0.0.0/24 prefix was received with the AIGP attribute attached:

```
user@PE4> show log bgp
Mar 22 09:41:39.650295 BGP RECV 10.9.9.5+64690 -> 10.9.9.4+179
Mar 22 09:41:39.650331 BGP RECV message type 2 (Update) length 70
Mar 22 09:41:39.650350 BGP RECV Update PDU length 70
Mar 22 09:41:39.650370 BGP RECV flags 0x40 code Origin(1): IGP
Mar 22 09:41:39.650394 BGP RECV flags 0x40 code ASPath(2) length 6: 7018
Mar 22 09:41:39.650415 BGP RECV flags 0x80 code AIGP(26): AIGP: 20
Mar 22 09:41:39.650436 BGP RECV flags 0x90 code MP_reach(14): AFI/SAFI 1/4
Mar 22 09:41:39.650459 BGP RECV      nhop 10.100.1.5 len 4
```

```

Mar 22 09:41:39.650495 BGP RECV    99.0.0.0/24 (label 301728)
Mar 22 09:41:39.650574 bgp_rcv_nlri: 99.0.0.0/24
Mar 22 09:41:39.650607 bgp_rcv_nlri: 99.0.0.0/24 belongs to meshgroup
Mar 22 09:41:39.650629 bgp_rcv_nlri: 99.0.0.0/24 qualified bnp->ribact 0x0
12afcb 0x0

```

**Meaning** Performing this verification helps with AIGP troubleshooting and debugging issues. It enables you to verify which devices in your network send and receive AIGP attributes.

- Related Documentation**
- [Understanding BGP Path Selection](#)
  - [Examples: Configuring Internal BGP Peering on page 3284](#)

## BGP Policy Configuration

- [Example: Configuring BGP Interactions with IGP on page 3421](#)
- [Example: Configuring BGP Route Advertisement on page 3425](#)
- [Example: Configuring EBGp Multihop on page 3433](#)
- [Example: Configuring BGP Route Preference \(Administrative Distance\) on page 3442](#)
- [Example: Configuring BGP Path Selection on page 3449](#)
- [Example: Removing Private AS Numbers on page 3456](#)

### Example: Configuring BGP Interactions with IGPs

- [Understanding Routing Policies on page 3421](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table on page 3422](#)

#### Understanding Routing Policies

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration.

Once a policy is created and named, it must be applied before it is active. You apply routing policies using the **import** and **export** statements at the **protocols>protocol-name** level in the configuration hierarchy.

In the **import** statement, you list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

### Example: Injecting OSPF Routes into the BGP Routing Table

---

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

- [Requirements on page 3422](#)
- [Overview on page 3422](#)
- [Configuration on page 3422](#)
- [Verification on page 3424](#)
- [Troubleshooting on page 3424](#)

#### Requirements

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3262](#).
- Configure interior gateway protocol (IGP) sessions between peers.

#### Overview

In this example, you create a routing policy called **injectpolicy1** and a routing term called **injectterm1**. The policy injects OSPF routes into the BGP routing table.

#### Configuration

- [Configuring the Routing Policy on page 3422](#)
- [Configuring Tracing for the Routing Policy on page 3423](#)

#### Configuring the Routing Policy

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

##### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To inject OSPF routes into a BGP routing table:

1. Create the policy term.

```
[edit policy-options policy-statement injectpolicy1]
user@host# set term injectterm1
```



- Specify OSPF as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```

- Specify the routes from an OSPF area as a match condition.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```

- Specify that the route is to be accepted if the previous conditions are matched.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```

- Apply the routing policy to BGP.

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

**Results** Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    from {
      protocol ospf;
      area 0.0.0.1;
    }
    then accept;
  }
}

user@host# show protocols bgp
export injectpolicy1;
```

If you are done configuring the device, enter **commit** from configuration mode.

### *Configuring Tracing for the Routing Policy*

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Include a trace action in the policy.

```
[edit policy-options policy-statement injectpolicy] term injectterm1]
user@host# then trace
```

2. Configure the tracing file for the output.

```
[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
user@host# set file size 5m
user@host# set file files 5
user@host# set flag policy
```

**Results** Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    then {
      trace;
    }
  }
}

user@host# show routing-options
traceoptions {
  file ospf-bgp-policy-log size 5m files 5;
  flag policy;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying That the Expected BGP Routes Are Present**

**Purpose** Verify the effect of the export policy.

**Action** From operational mode, enter the **show route** command.

### **Troubleshooting**

- [Using the show log Command to Examine the Actions of the Routing Policy on page 3424](#)

### **Using the show log Command to Examine the Actions of the Routing Policy**

**Problem** The routing table contains unexpected routes, or routes are missing from the routing table.

**Solution** If you configure policy tracing as shown in this example, you can run the **show log ospf-bgp-policy-log** command to diagnose problems with the routing policy. The **show log ospf-bgp-policy-log** command displays information about the routes that the **injectpolicy1** policy term analyzes and acts upon.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

## Example: Configuring BGP Route Advertisement

- [Understanding Route Advertisement on page 3425](#)
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3429](#)

### Understanding Route Advertisement

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes they should advertise in their protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table. For information about routing policy, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

When configuring BGP routing policy, you can perform the following tasks:

- [Applying Routing Policy on page 3425](#)
- [Setting BGP to Advertise Inactive Routes on page 3426](#)
- [Configuring BGP to Advertise the Best External Route to Internal Peers on page 3426](#)
- [Configuring How Often BGP Exchanges Routes with the Routing Table on page 3428](#)
- [Disabling Suppression of Route Advertisements on page 3429](#)

### Applying Routing Policy

You define routing policy at the **[edit policy-options]** hierarchy level. To apply policies you have defined for BGP, include the **import** and **export** statements within the BGP configuration.

You can apply policies as follows:

- BGP global **import** and **export** statements—Include these statements at the **[edit protocols bgp]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level).
- Group **import** and **export** statements—Include these statements at the **[edit protocols bgp group group-name]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp group group-name]** hierarchy level).
- Peer **import** and **export** statements—Include these statements at the **[edit protocols bgp group group-name neighbor address]** hierarchy level (for routing instances, include these statements at the **[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]** hierarchy level).

A peer-level **import** or **export** statement overrides a group **import** or **export** statement. A group-level **import** or **export** statement overrides a global BGP **import** or **export** statement.

To apply policies, see the following sections:

- [Applying Policies to Routes Being Imported into the Routing Table from BGP on page 3426](#)
- [Applying Policies to Routes Being Exported from the Routing Table into BGP on page 3426](#)

#### ***Applying Policies to Routes Being Imported into the Routing Table from BGP***

To apply policy to routes being imported into the routing table from BGP, include the **import** statement, listing the names of one or more policies to be evaluated:

```
import [ policy-names ];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices.

#### ***Applying Policies to Routes Being Exported from the Routing Table into BGP***

To apply policy to routes being exported from the routing table into BGP, include the **export** statement, listing the names of one or more policies to be evaluated:

```
export [ policy-names ];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify more than one policy, they are evaluated in the order specified, from first to last, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP.

#### ***Setting BGP to Advertise Inactive Routes***

By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. To have the routing table export to BGP the best route learned by BGP even if Junos OS did not select it to be an active route, include the **advertise-inactive** statement:

```
advertise-inactive;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

#### ***Configuring BGP to Advertise the Best External Route to Internal Peers***

In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information

and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.

In Junos OS Release 9.3 and later, you can configure BGP to advertise the best external route into an internal BGP (IBGP) mesh group, a route reflector cluster, or an autonomous system (AS) confederation, even when the best route is an internal route.



**NOTE:** In order to configure the `advertise-external` statement on a route reflector, you must disable intracluster reflection with the `no-client-reflect` statement.

When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.

In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external.

You can also configure BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric is evaluated. As a result, an external route with an AS path worse (that is, longer) than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

To configure BGP to advertise the best external path to internal peers, include the `advertise-external` statement:

```
advertise-external;
```



**NOTE:** The `advertise-external` statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.

For a complete list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To configure BGP to advertise the best external path only if the route selection process reaches the point where the MED value is evaluated, include the `conditional` statement:

```
advertise-external {
  conditional;
}
```

### *Configuring How Often BGP Exchanges Routes with the Routing Table*

BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. To guard against this, you can delay the time between when BGP and the routing table exchange route information.

To configure how often BGP and the routing table exchange route information, include the **out-delay** statement:

**out-delay** *seconds*;

By default, the routing table retains some of the route information learned from BGP. To have the routing table retain all or none of this information, include the **keep** statement:

**keep** (all | none);

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

The routing table can retain the route information learned from BGP in one of the following ways:

- Default (omit the **keep** statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.
- **keep all**—Keep all route information that was learned from BGP.
- **keep none**—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure **keep none** for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.

### *Disabling Suppression of Route Advertisements*

Junos OS does not advertise the routes learned from one EBGp peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGp peers that are in the same AS as the originating peer, regardless of the routing instance. You can modify this behavior by including the **advertise-peer-as** statement in the configuration. To disable the default advertisement suppression, include the **advertise-peer-as** statement:

```
advertise-peer-as;
```



**NOTE:** The route suppression default behavior is disabled if the **as-override** statement is included in the configuration.

If you include the **advertise-peer-as** statement in the configuration, BGP advertises the route regardless of this check.

To restore the default behavior, include the **no-advertise-peer-as** statement in the configuration:

```
no-advertise-peer-as;
```

If you include both the **as-override** and **no-advertise-peer-as** statements in the configuration, the **no-advertise-peer-as** statement is ignored. You can include these statements at multiple hierarchy levels.

For a list of hierarchy levels at which you can include these statements, see the statement summary section for these statements.

### Example: Configuring BGP Prefix-Based Outbound Route Filtering

This example shows how to configure a Juniper Networks router to accept route filters from remote peers and perform outbound route filtering using the received filters.

- [Requirements on page 3429](#)
- [Overview on page 3430](#)
- [Configuration on page 3430](#)
- [Verification on page 3432](#)

#### **Requirements**

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol (IGP).

### Overview

You can configure a BGP peer to accept route filters from remote peers and perform outbound route filtering using the received filters. By filtering out unwanted updates, the sending peer saves resources needed to generate and transmit updates, and the receiving peer saves resources needed to process updates. This feature can be useful, for example, in a virtual private network (VPN) in which subsets of customer edge (CE) devices are not capable of processing all the routes in the VPN. The CE devices can use prefix-based outbound route filtering to communicate to the provider edge (PE) routing device to transmit only a subset of routes, such as routes to the main data centers only.

The maximum number of prefix-based outbound route filters that a BGP peer can accept is 5000. If a remote peer sends more than 5000 outbound route filters to a peer address, the additional filters are discarded, and a system log message is generated.


You can configure interoperability for the routing device as a whole or for specific BGP groups or peers only.

### Topology

In the sample network, Device CE1 is a router from another vendor. The configuration shown in this example is on Juniper Networks Router PE1.

Figure 75 on page 3430 shows the sample network.

Figure 75: BGP Prefix-Based Outbound Route Filtering



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
PE1 set protocols bgp group cisco-peers type external
    set protocols bgp group cisco-peers description "to CE1"
    set protocols bgp group cisco-peers local-address 192.168.165.58
    set protocols bgp group cisco-peers peer-as 35
    set protocols bgp group cisco-peers outbound-route-filter bgp-orf-cisco-mode
    set protocols bgp group cisco-peers outbound-route-filter prefix-based accept inet
    set protocols bgp group cisco-peers neighbor 192.168.165.56
    set routing-options autonomous-system 65500
```



**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1 to accept route filters from Device CE1 and perform outbound route filtering using the received filters:

1. Configure the local autonomous system.

```
[edit routing-options]
user@PE1# set autonomous-system 65500
```

2. Configure external peering with Device CE1.

```
[edit protocols bgp group cisco-peers]
user@PE1# set type external
user@PE1# set description "to CE1"
user@PE1# set local-address 192.168.165.58
user@PE1# set peer-as 35
user@PE1# set neighbor 192.168.165.56
```

3. Configure Router PE1 to accept IPv4 route filters from Device CE1 and perform outbound route filtering using the received filters.

```
[edit protocols bgp group cisco-peers]
user@PE1# set outbound-route-filter prefix-based accept inet
```

4. (Optional) Enable interoperability with routing devices that use the vendor-specific compatibility code of 130 for outbound route filters and the code type of 128.

The IANA standard code is 3, and the standard code type is 64.

```
[edit protocols bgp group cisco-peers]
user@PE1# set outbound-route-filter bgp-orf-cisco-mode
```

**Results** From configuration mode, confirm your configuration by entering the **show protocols** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show protocols
group cisco-peers {
  type external;
  description "to CE1";
  local-address 192.168.165.58;
  peer-as 35;
  outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
      accept {
        inet;
      }
    }
  }
  neighbor 192.168.165.56;
}

user@PE1# show routing-options
autonomous-system 65500;
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the Outbound Route Filter on page 3432](#)
- [Verifying the BGP Neighbor Mode on page 3432](#)

### **Verifying the Outbound Route Filter**

**Purpose** Display information about the prefix-based outbound route filter received from Device CE1.

**Action** From operational mode, enter the [show bgp neighbor orf detail](#) command.

```
user@PE1> show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56 Type: External
Group: cisco-peers

inet-unicast
Filter updates rcv:          4 Immediate:          0
Filter: prefix-based         receive
Updates rcv:                4
Received filter entries:
  seq 10 2.2.0.0/16 deny minlen 0 maxlen 0
  seq 20 3.3.0.0/16 deny minlen 24 maxlen 0
  seq 30 4.4.0.0/16 deny minlen 0 maxlen 28
  seq 40 5.5.0.0/16 deny minlen 24 maxlen 28
```

### **Verifying the BGP Neighbor Mode**

**Purpose** Verify that the **bgp-orf-cisco-mode** setting is enabled for the peer by making sure that the **ORFCiscoMode** option is displayed in the **show bgp neighbor** command output.

**Action** From operational mode, enter the [show bgp neighbor](#) command.

```
user@PE1> show bgp neighbor
Peer: 192.168.165.56 AS 35           Local: 192.168.165.58 AS 65500
Type: External   State: Active      Flags: <>
Last State: Idle   Last Event: Start
Last Error: None
Export: [ adv_stat ]
Options: <Preference LocalAddress AddressFamily PeerAS Refresh>
Options: <ORF ORFCiscoMode>
Address families configured: inet-unicast
Local Address: 192.168.165.58 Holdtime: 90 Preference: 170
Number of flaps: 0
Trace options: detail open detail refresh
Trace file: /var/log/orf size 5242880 files 20
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3261](#)
  - [BGP Configuration Overview](#)
  - [Example: Configuring BGP to Advertise the Best External Route to Internal Peers](#)
  - [Example: Setting BGP to Advertise Inactive Routes](#)

## Example: Configuring EBGP Multihop

- [Understanding BGP Multihop on page 3433](#)
- [Example: Configuring EBGP Multihop Sessions on page 3433](#)

### Understanding BGP Multihop

---

When external BGP (EBGP) peers are not directly connected to each other, they must cross one or more non-BGP routers to reach each other. Configuring multihop EBGP enables the peers to pass through the other routers to form peer relationships and exchange update messages. This type of configuration is typically used when a Juniper Networks routing device needs to run EBGP with a third-party router that does not allow direct connection of the two EBGP peers. EBGP multihop enables a neighbor connection between two EBGP peers that do not have a direct connection.

### Example: Configuring EBGP Multihop Sessions

---

This example shows how to configure an external BGP (EBGP) peer that is more than one hop away from the local router. This type of session is called a *multihop* BGP session.

- [Requirements on page 3433](#)
- [Overview on page 3433](#)
- [Configuration on page 3434](#)
- [Verification on page 3440](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.

#### Overview

The configuration to enable multihop EBGP sessions requires connectivity between the two EBGP peers. This example uses static routes to provide connectivity between the devices.


Unlike directly connected EBGP sessions in which physical address are typically used in the **neighbor** statements, you must use loopback interface addresses for multihop EBGP by specifying the loopback interface address of the indirectly connected peer. In this way, EBGP multihop is similar to internal BGP (IBGP).

Finally, you must add the **multihop** statement. Optionally, you can set a maximum time-to-live (TTL) value with the **ttl** statement. The TTL is carried in the IP header of BGP packets. If you do not specify a TTL value, the system's default maximum TTL value is used. The default TTL value is 64 for multihop EBGP sessions. Another option is to retain the BGP next-hop value for route advertisements by including the **no-nexthop-change** statement.

[Figure 76 on page 3434](#) shows a typical EBGP multihop network.

Device C and Device E have an established EBGP session. Device D is not a BGP-enabled device. All of the devices have connectivity via static routes.

Figure 76: Typical Network with EBGP Multihop Sessions

**Configuration**

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device C**

```

set interfaces fe-1/2/0 unit 9 description to-D
set interfaces fe-1/2/0 unit 9 family inet address 10.10.10.9/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group external-peers type external
set protocols bgp group external-peers multihop ttl 2
set protocols bgp group external-peers local-address 192.168.40.4
set protocols bgp group external-peers export send-static
set protocols bgp group external-peers peer-as 18
set protocols bgp group external-peers neighbor 192.168.6.7
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 10.10.10.14/32 next-hop 10.10.10.10
set routing-options static route 192.168.6.7/32 next-hop 10.10.10.10
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17

```

**Device D**

```

set interfaces fe-1/2/0 unit 10 description to-C
set interfaces fe-1/2/0 unit 10 family inet address 10.10.10.10/30
set interfaces fe-1/2/1 unit 13 description to-E
set interfaces fe-1/2/1 unit 13 family inet address 10.10.10.13/30
set interfaces lo0 unit 4 family inet address 192.168.6.6/32
set routing-options static route 192.168.40.4/32 next-hop 10.10.10.9
set routing-options static route 192.168.6.7/32 next-hop 10.10.10.14
set routing-options router-id 192.168.6.6

```

**Device E**

```

set interfaces fe-1/2/0 unit 14 description to-D
set interfaces fe-1/2/0 unit 14 family inet address 10.10.10.14/30
set interfaces lo0 unit 5 family inet address 192.168.6.7/32
set protocols bgp group external-peers multihop ttl 2
set protocols bgp group external-peers local-address 192.168.6.7
set protocols bgp group external-peers export send-static
set protocols bgp group external-peers peer-as 17
set protocols bgp group external-peers neighbor 192.168.40.4
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept

```

```

set routing-options static route 10.10.10.8/30 next-hop 10.10.10.13
set routing-options static route 192.168.40.4/32 next-hop 10.10.10.13
set routing-options router-id 192.168.6.7
set routing-options autonomous-system 18

```

### Device C

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device C:

1. Configure the interface to the directly connected device (to-D), and configure the loopback interface.

```

[edit interfaces fe-1/2/0 unit 9]
user@C# set description to-D
user@C# set family inet address 10.10.10.9/30

```

```

[edit interfaces lo0 unit 3]
user@C# set family inet address 192.168.40.4/32

```

2. Configure an EBGP session with Device E.

The **neighbor** statement points to the loopback interface on Device E.

```

[edit protocols bgp group external-peers]
user@C# set type external
user@C# set local-address 192.168.40.4
user@C# set export send-static
user@C# set peer-as 18
user@C# set neighbor 192.168.6.7

```

3. Configure the multihop statement to enable Device C and Device E to become EBGP peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```

[edit protocols bgp group external-peers]
user@C# set multihop ttl 2

```

4. Configure connectivity to Device E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```

[edit routing-options]
user@C# set static route 10.10.10.14/32 next-hop 10.10.10.10
user@C# set static route 192.168.6.7/32 next-hop 10.10.10.10

```

5. Configure the local router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@C# set router-id 192.168.40.4
user@C# set autonomous-system 17

```

6. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-static term 1]
user@C# set from protocol static
user@C# set then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@C# show interfaces
fe-1/2/0 {
  unit 9 {
    description to-D;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 192.168.40.4/32;
    }
  }
}

user@C# show protocols
bgp {
  group external-peers {
    type external;
    multihop {
      ttl 2;
    }
    local-address 192.168.40.4;
    export send-static;
    peer-as 18;
    neighbor 192.168.6.7;
  }
}

user@C# show policy-options
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@C# show routing-options
static {
  route 10.10.10.14/32 next-hop 10.10.10.10;
  route 192.168.6.7/32 next-hop 10.10.10.10;
}
```

```
router-id 192.168.40.4;
autonomous-system 17;
```

If you are done configuring the device, enter **commit** from configuration mode.  
Repeat these steps for all BFD sessions in the topology.

### Configuring Device D

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device D:

1. Set the CLI to Device D.

```
user@host> set cli logical-system D
```

2. Configure the interfaces to the directly connected devices, and configure a loopback interface.

```
[edit interfaces fe-1/2/0 unit 10]
user@D# set description to-C
user@D# set family inet address 10.10.10.10/30
```

```
[edit interfaces fe-1/2/1 unit 13]
user@D# set description to-E
user@D# set family inet address 10.10.10.13/30
```

```
[edit interfaces lo0 unit 4]
user@D# set family inet address 192.168.6.6/32
```

3. Configure connectivity to the other devices using static routes to the loopback interface addresses.

On Device D, you do not need static routes to the physical addresses because Device D is directly connected to Device C and Device E.

```
[edit routing-options]
user@D# set static route 192.168.40.4/32 next-hop 10.10.10.9
user@D# set static route 192.168.6.7/32 next-hop 10.10.10.14
```

4. Configure the local router ID.

```
[edit routing-options]
user@D# set router-id 192.168.6.6
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@D# show interfaces
fe-1/2/0 {
  unit 10 {
    description to-C;
    family inet {
      address 10.10.10.10/30;
```

```
    }  
  }  
}  
fe-1/2/1 {  
  unit 13 {  
    description to-E;  
    family inet {  
      address 10.10.10.13/30;  
    }  
  }  
}  
lo0 {  
  unit 4 {  
    family inet {  
      address 192.168.6.6/32;  
    }  
  }  
}  
}  
  
user@D# show protocols  
  
user@D# show routing-options  
static {  
  route 192.168.40.4/32 next-hop 10.10.10.9;  
  route 192.168.6.7/32 next-hop 10.10.10.14;  
}  
router-id 192.168.6.6;
```

If you are done configuring the device, enter **commit** from configuration mode.  
Repeat these steps for all BFD sessions in the topology.

### **Configuring Device E**

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device E:

1. Set the CLI to Device E.
2. Configure the interface to the directly connected device (to-D), and configure the loopback interface.

```
[edit interfaces fe-1/2/0 unit 14]  
user@E# set description to-D  
user@E# set family inet address 10.10.10.14/30
```

```
[edit interfaces lo0 unit 5]  
user@E# set family inet address 192.168.6.7/32
```

3. Configure an EBGP session with Device E.

The **neighbor** statement points to the loopback interface on Device C.

```
[edit protocols bgp group external-peers]
```



```

user@E# set local-address 192.168.6.7
user@E# set export send-static
user@E# set peer-as 17
user@E# set neighbor 192.168.40.4

```

4. Configure the **multihop** statement to enable Device C and Device E to become EBGp peers.

Because the peers are two hops away from each other, the example uses the **ttl 2** statement.

```

[edit protocols bgp group external-peers]
user@E# set multihop ttl 2

```

5. Configure connectivity to Device E, using static routes.

You must configure a route to both the loopback interface address and to the address on the physical interface.

```

[edit routing-options]
user@E# set static route 10.10.10.8/30 next-hop 10.10.10.13
user@E# set static route 192.168.40.4/32 next-hop 10.10.10.13

```

6. Configure the local router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@E# set router-id 192.168.6.7
user@E# set autonomous-system 18

```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```

[edit policy-options policy-statement send-static term 1]
user@E# set from protocol static
user@E# set then accept

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@E# show interfaces
fe-1/2/0 {
  unit 14 {
    description to-D;
    family inet {
      address 10.10.10.14/30;
    }
  }
}
lo0 {
  unit 5 {
    family inet {
      address 192.168.6.7/32;
    }
  }
}

```

```
}
user@E# show protocols
bgp {
  group external-peers {
    multihop {
      ttl 2;
    }
    local-address 192.168.6.7;
    export send-static;
    peer-as 17;
    neighbor 192.168.40.4;
  }
}

user@E# show policy-options
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@E# show routing-options
static {
  route 10.10.10.8/30 next-hop 10.10.10.13;
  route 192.168.40.4/32 next-hop 10.10.10.13;
}
router-id 192.168.6.7;
autonomous-system 18;
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 3440](#)
- [Verifying That BGP Sessions Are Established on page 3441](#)
- [Viewing Advertised Routes on page 3441](#)

### **Verifying Connectivity**

**Purpose** Make sure that Device C can ping Device E, specifying the loopback interface address as the source of the ping request.

The loopback interface address is the source address that BGP will use.

**Action** From operational mode, enter the **ping 10.10.10.14 source 192.168.40.4** command from Device C, and enter the **ping 10.10.10.9 source 192.168.6.7** command from Device E.

```
user@C> ping 10.10.10.14 source 192.168.40.4
```

```
PING 10.10.10.14 (10.10.10.14): 56 data bytes
64 bytes from 10.10.10.14: icmp_seq=0 ttl=63 time=1.262 ms
64 bytes from 10.10.10.14: icmp_seq=1 ttl=63 time=1.202 ms
^C
```

```

--- 10.10.10.14 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.202/1.232/1.262/0.030 ms

```

```
user@E> ping 10.10.10.9 source 192.168.6.7
```

```

PING 10.10.10.9 (10.10.10.9): 56 data bytes
64 bytes from 10.10.10.9: icmp_seq=0 ttl=63 time=1.255 ms
64 bytes from 10.10.10.9: icmp_seq=1 ttl=63 time=1.158 ms
^C
--- 10.10.10.9 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.158/1.206/1.255/0.049 ms

```

**Meaning** The static routes are working if the pings work.

### *Verifying That BGP Sessions Are Established*

**Purpose** Verify that the BGP sessions are up.

**Action** From operational mode, enter the `show bgp summary` command.

```
user@C> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0          2        0          0          0          0          0          0
Peer        AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.6.7      18      147      147        0        1    1:04:27
0/2/2/0          0/0/0/0

```

```
user@E> show bgp summary
```

```

Groups: 1 Peers: 1 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0          2        0          0          0          0          0          0
Peer        AS      InPkt    OutPkt    OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.40.4     17      202      202        0        1    1:02:18
0/2/2/0          0/0/0/0

```

**Meaning** The output shows that both devices have one peer each. No peers are down.

### *Viewing Advertised Routes*

**Purpose** Check to make sure that routes are being advertised by BGP.

**Action** From operational mode, enter the `show route advertising-protocol bgp neighbor` command.

```
user@C> show route advertising-protocol bgp 192.168.6.7
```

```

inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED    Lclpref  AS path
* 10.10.10.14/32    Self              0
* 192.168.6.7/32    Self              0

```

```
user@E> show route advertising-protocol bgp 192.168.40.4
```

```
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED      Lclpref   AS path
* 10.10.10.8/30      Self              0         0         I
* 192.168.40.4/32    Self              0         0         I
```

**Meaning** The **send-static** routing policy is exporting the static routes from the routing table into BGP. BGP is advertising these routes between the peers because the BGP peer session is established.

**Related Documentation**

- [Examples: Configuring External BGP Peering on page 3261](#)
- [BGP Configuration Overview](#)

## Example: Configuring BGP Route Preference (Administrative Distance)

- [Understanding Route Preference Values on page 3442](#)
- [Example: Configuring the Preference Value for BGP Routes on page 3443](#)

### Understanding Route Preference Values

The Junos OS routing protocol process assigns a default preference value (also known as an *administrative distance*) to each route that the routing table receives. The default value depends on the source of the route. The preference value is a value from 0 through 4,294,967,295 ( $2^{32} - 1$ ), with a lower value indicating a more preferred route.

[Table 296 on page 3442](#) lists the default preference values.

**Table 296: Default Route Preference Values**

| How Route Is Learned         | Default Preference | Statement to Modify Default Preference  |
|------------------------------|--------------------|---|
| Directly connected network   | 0                  | —   |
| System routes                | 4                  | —   |
| Static and Static LSPs       | 5                  | <i>static</i>   |
| RSVP-signaled LSPs           | 7                  | RSVP <b>preference</b> as described in the <i>Junos OS MPLS Applications Library for Routing Devices</i>  |
| LDP-signaled LSPs            | 9                  | LDP <b>preference</b> , as described in the <i>Junos OS MPLS Applications Library for Routing Devices</i> |
| OSPF internal route          | 10                 | OSPF <b>preference</b>  |
| IS-IS Level 1 internal route | 15                 | IS-IS <b>preference</b>   |
| IS-IS Level 2 internal route | 18                 | IS-IS <b>preference</b>   |

Table 296: Default Route Preference Values (*continued*)

| How Route Is Learned         | Default Preference | Statement to Modify Default Preference   |
|------------------------------|--------------------|--|
| Redirects                    | 30                 | –  |
| Kernel                       | 40                 | –  |
| SNMP                         | 50                 | –  |
| Router discovery             | 55                 | –  |
| RIP                          | 100                | RIP <a href="#">preference</a>   |
| RIPng                        | 100                | RIPng <a href="#">preference</a>   |
| PIM                          | 105                | <i>Multicast Protocols Feature Guide for Routing Devices</i>                     |
| DVMRP                        | 110                | <i>Multicast Protocols Feature Guide for Routing Devices</i>                     |
| Aggregate                    | 130                | <a href="#">aggregate</a>  |
| OSPF AS external routes      | 150                | OSPF <a href="#">external-preference</a>   |
| IS-IS Level 1 external route | 160                | IS-IS <a href="#">external-preference</a>  |
| IS-IS Level 2 external route | 165                | IS-IS <a href="#">external-preference</a>  |
| BGP                          | 170                | BGP <a href="#">preference</a> , <a href="#">export</a> , <a href="#">import</a> |
| MSDP                         | 175                | <i>Multicast Protocols Feature Guide for Routing Devices</i>                     |

In general, the narrower the scope of the statement, the higher precedence its preference value is given, but the smaller the set of routes it affects. To modify the default preference value for routes learned by routing protocols, you generally apply routing policy when configuring the individual routing protocols. You also can modify some preferences with other configuration statements, which are indicated in the table.

#### Example: Configuring the Preference Value for BGP Routes

This example shows how to specify the preference for routes learned from BGP. Routing information can be learned from multiple sources. To break ties among equally specific routes learned from multiple sources, each source has a preference value. Routes that are learned through explicit administrative action, such as static routes, are preferred over routes learned from a routing protocol, such as BGP or OSPF. This concept is called *administrative distance* by some vendors.

- [Requirements on page 3444](#)
- [Overview on page 3444](#)

- [Configuration on page 3445](#)
- [Verification on page 3447](#)

### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

### **Overview**

Routing information can be learned from multiple sources, such as through static configuration, BGP, or an interior gateway protocol (IGP). When Junos OS determines a route's preference to become the active route, it selects the route with the lowest preference as the active route and installs this route into the forwarding table. By default, the routing software assigns a preference of 170 to routes that originated from BGP. Of all the routing protocols, BGP has the highest default preference value, which means that routes learned by BGP are the least likely to become the active route.

Some vendors have a preference (distance) of 20 for external BGP (EBGP) and a distance of 200 for internal BGP (IBGP). Junos OS uses the same value (170) for both EBGP and IBGP. However, this difference between vendors has no operational impact because Junos OS always prefers EBGP routes over IBGP routes.

Another area in which vendors differ is in regard to IGP distance compared to BGP distance. For example, some vendors assign a distance of 110 to OSPF routes. This is higher than the EBGP distance of 20, and results in the selection of an EBGP route over an equivalent OSPF route. In the same scenario, Junos OS chooses the OSPF route, because of the default preference 10 for an internal OSPF route and 150 for an external OSPF route, which are both lower than the 170 preference assigned to all BGP routes.

In a multivendor environment, you might want to change the preference value for BGP routes so that Junos OS chooses an EBGP route instead of an OSPF route. To accomplish this goal, one option is to include the [preference](#) statement in the EBGP configuration. To modify the default BGP preference value, include the **preference** statement, specifying a value from 0 through 4,294,967,295 ( $2^{32} - 1$ ).



**TIP:** Another way to achieve multivendor compatibility is to include the [advertise-inactive](#) statement in the EBGP configuration. This causes the routing table to export to BGP the best route learned by BGP even if Junos OS did not select it to be an active route. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The `advertise-inactive` statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the `advertise-inactive` statement, the Junos OS device uses the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.

---

### Topology


In the sample network, Device R1 and Device R2 have EBGP routes to each other and also OSPF routes to each other.

This example shows the routing tables in the following cases:

- Accept the default preference values of 170 for BGP and 10 for OSPF.
- Change the BGP preference to 8.

Figure 77 on page 3445 shows the sample network.

**Figure 77: BGP Preference Value Topology**



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 4 family inet address 1.12.0.1/30
set interfaces lo0 unit 2 family inet address 10.255.71.24/32
set protocols bgp export send-direct
set protocols bgp group ext type external
set protocols bgp group ext preference 8
set protocols bgp group ext peer-as 65000
set protocols bgp group ext neighbor 1.12.0.2
set protocols ospf area 0.0.0.0 interface fe-1/2/0.4
set protocols ospf area 0.0.0.0 interface 10.255.71.24
set policy-options policy-statement send-direct term 1 from protocol direct

```

```
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 65500
```

**Device R2**

```
set interfaces fe-1/2/0 unit 6 family inet address 1.12.0.2/30
set interfaces lo0 unit 3 family inet address 10.255.14.177/32
set protocols bgp export send-direct
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65500
set protocols bgp group ext neighbor 1.12.0.1
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface 10.255.14.177
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set routing-options autonomous-system 65000
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.  

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 4 family inet address 1.12.0.1/30
user@R1# set lo0 unit 2 family inet address 10.255.71.24/32
```
2. Configure the local autonomous system.  

```
[edit routing-options]
user@R1# set autonomous-system 65500
```
3. Configure the external peering with Device R2.  

```
[edit protocols bgp]
user@R1# set export send-direct
user@R1# set group ext type external
user@R1# set group ext preference 8
user@R1# set group ext peer-as 65000
user@R1# set group ext neighbor 1.12.0.2
```
4. Configure OSPF.  

```
[edit protocols ospf area 0.0.0.0]
user@R1# set interface fe-1/2/0.4
user@R1# set interface 10.255.71.24
```
5. Configure the routing policy.  

```
[edit policy-options policy-statement send-direct term 1]
user@R1# set from protocol direct
user@R1# set then accept
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.



```

user@R1# show interfaces
fe-1/2/0 {
  unit 4 {
    family inet {
      address 1.12.0.1/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.255.71.24/32;
    }
  }
}

user@R1# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}

user@R1# show protocols
protocols {
  bgp {
    export send-direct;
    group ext {
      type external;
      preference 8;
      peer-as 65000;
      neighbor 1.12.0.2;
    }
  }
  ospf {
    area 0.0.0.0 {
      interface fe-1/2/0.4;
      interface 10.255.71.24;
    }
  }
}

user@R1# show routing-options
autonomous-system 65500;

```

If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps on Device R2.

### Verification

Confirm that the configuration is working properly.

### ***Verifying the Preference***

**Purpose** Make sure that the routing tables on Device R1 and Device R2 reflect the fact that Device R1 is using the configured EBGp preference of 8, and Device R2 is using the default EBGp preference of 170.

**Action** From operational mode, enter the **show route** command.

```
user@R1> show route
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.12.0.0/30      *[Direct/0] 3d 07:03:01
                  > via fe-1/2/0.4
                  [BGP/8] 01:04:49, localpref 100
                  AS path: 65000 I
                  > to 1.12.0.2 via fe-1/2/0.4
1.12.0.1/32      *[Local/0] 3d 07:03:01
                  Local via fe-1/2/0.4
10.255.14.177/32 *[BGP/8] 01:04:49, localpref 100
                  AS path: 65000 I
                  > to 1.12.0.2 via fe-1/2/0.4
                  [OSPF/10] 3d 07:02:16, metric 1
                  > to 1.12.0.2 via fe-1/2/0.4
10.255.71.24/32  *[Direct/0] 3d 07:03:01
                  > via lo0.2
224.0.0.5/32     *[OSPF/10] 5d 03:42:16, metric 1
                  MultiRecv
```

```
user@R2> show route
inet.0: 5 destinations, 7 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.12.0.0/30      *[Direct/0] 3d 07:03:30
                  > via fe-1/2/0.6
                  [BGP/170] 00:45:36, localpref 100
                  AS path: 65500 I
                  > to 1.12.0.1 via fe-1/2/0.6
1.12.0.2/32      *[Local/0] 3d 07:03:30
                  Local via fe-1/2/0.6
10.255.14.177/32 *[Direct/0] 3d 07:03:30
                  > via lo0.3
10.255.71.24/32  *[OSPF/10] 3d 07:02:45, metric 1
                  > to 1.12.0.1 via fe-1/2/0.6
                  [BGP/170] 00:45:36, localpref 100
                  AS path: 65500 I
                  > to 1.12.0.1 via fe-1/2/0.6
224.0.0.5/32     *[OSPF/10] 5d 03:42:45, metric 1
                  MultiRecv
```

**Meaning** The output shows that on Device R1, the active path to Device R2's loopback interface (10.255.14.177/32) is a BGP route. The output also shows that on Device R2, the active path to Device R1's loopback interface (10.255.71.24/32) is an OSPF route.

**Related Documentation**

- [Route Preferences Overview](#)
- [Understanding External BGP Peering Sessions on page 3261](#)

- [BGP Configuration Overview](#)

## Example: Configuring BGP Path Selection

- [Example: Ignoring the AS Path Attribute When Selecting the Best Path on page 3449](#)

### Example: Ignoring the AS Path Attribute When Selecting the Best Path

If multiple BGP routes to the same destination exist, BGP selects the best path based on the route attributes of the paths. One of the route attributes that affects the best-path decision is the length of the AS paths of each route. Routes with shorter AS paths are preferred over those with longer AS paths. Although not typically practical, some scenarios might require that the AS path length be ignored in the route selection process. This example shows how to configure a routing device to ignore the AS path attribute.

- [Requirements on page 3449](#)
- [Overview on page 3449](#)
- [Configuration on page 3450](#)
- [Verification on page 3455](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.

#### Overview

On externally connected routing devices, the purpose of skipping the AS path comparison might be to force an external BGP (EBGP) versus internal BGP (IBGP) decision to remove traffic from your network as soon as possible. On internally connected routing devices, you might want your IBGP-only routers to default to the local externally connected gateway. The local IBGP-only (internal) routers skip the AS path comparison and move down the decision tree to use the closest interior gateway protocol (IGP) gateway (lowest IGP metric). Doing this might be an effective way to force these routers to use a LAN connection instead of their WAN connection.



**CAUTION:** When you include the `as-path-ignore` statement on a routing device in your network, you might need to include it on all other BGP-enabled devices in your network to prevent routing loops and convergence issues. This is especially true for IBGP path comparisons.


In this example, Device R2 is learning about the loopback interface address on Device R4 (4.4.4.4/32) from Device R1 and Device R3. Device R1 is advertising 4.4.4.4/32 with an AS-path of 1 5 4, and Device R3 is advertising 4.4.4.4/32 with an AS-path of 3 4. Device R2 selects the path for 4.4.4.4/32 from Device R3 as the best path because the AS path is shorter than the AS path from Device R1.

This example modifies the BGP configuration on Device R2 so that the AS-path length is not used in the best-path selection.

Device R1 has a lower router ID (1.1.1.1) than Device R3 (3.3.3.3). If all other path selection criteria are equal (or, as in this case, ignored), the route learned from Device R1 is used. Because the AS-path attribute is being ignored, the best path is toward Device R1 because of its lower router ID value.

Figure 78 on page 3450 shows the sample topology.

**Figure 78: Topology for Ignoring the AS-Path Length**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1
set interfaces fe-1/2/0 unit 1 family inet address 192.168.10.1/24
set interfaces fe-1/2/1 unit 10 family inet address 192.168.50.2/24
set interfaces lo0 unit 1 family inet address 1.1.1.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct

```

```

set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.10.2 peer-as 2
set protocols bgp group ext neighbor 192.168.50.1 peer-as 5
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.20.0/24 next-hop 192.168.10.2
set routing-options static route 192.168.30.0/24 next-hop 192.168.10.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.50.1
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1

```

Device R2

```

set interfaces fe-1/2/0 unit 2 family inet address 192.168.10.2/24
set interfaces fe-1/2/1 unit 3 family inet address 192.168.20.2/24
set interfaces lo0 unit 2 family inet address 2.2.2.2/32
set protocols bgp path-selection as-path-ignore
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.10.1 peer-as 1
set protocols bgp group ext neighbor 192.168.20.1 peer-as 3
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.50.0/24 next-hop 192.168.10.1
set routing-options static route 192.168.40.0/24 next-hop 192.168.10.1
set routing-options static route 192.168.30.0/24 next-hop 192.168.20.1
set routing-options router-id 2.2.2.2
set routing-options autonomous-system 2

```

Device R3

```

set interfaces fe-1/2/0 unit 4 family inet address 192.168.20.1/24
set interfaces fe-1/2/1 unit 5 family inet address 192.168.30.1/24
set interfaces lo0 unit 3 family inet address 1.1.1.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.20.2 peer-as 2
set protocols bgp group ext neighbor 192.168.30.2 peer-as 4
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.20.2
set routing-options static route 192.168.50.0/24 next-hop 192.168.20.2

```

```
set routing-options static route 192.168.40.0/24 next-hop 192.168.30.2
set routing-options router-id 3.3.3.3
set routing-options autonomous-system 3
```

**Device R4**

```
set interfaces fe-1/2/0 unit 6 family inet address 192.168.30.2/24
set interfaces fe-1/2/1 unit 7 family inet address 192.168.40.1/24
set interfaces lo0 unit 4 family inet address 4.4.4.4/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.30.1 peer-as 3
set protocols bgp group ext neighbor 192.168.40.2 peer-as 5
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.40.2
set routing-options static route 192.168.50.0/24 next-hop 192.168.40.2
set routing-options static route 192.168.40.0/24 next-hop 192.168.30.1
set routing-options router-id 4.4.4.4
set routing-options autonomous-system 4
```

**Device R5**

```
set interfaces fe-1/2/0 unit 8 family inet address 192.168.40.2/24
set interfaces fe-1/2/1 unit 9 family inet address 192.168.50.1/24
set interfaces lo0 unit 5 family inet address 5.5.5.5/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext export send-local
set protocols bgp group ext neighbor 192.168.40.1 peer-as 4
set protocols bgp group ext neighbor 192.168.50.2 peer-as 1
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-local term 1 from protocol local
set policy-options policy-statement send-local term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.50.2
set routing-options static route 192.168.20.0/24 next-hop 192.168.50.2
set routing-options static route 192.168.30.0/24 next-hop 192.168.40.1
set routing-options router-id 5.5.5.5
set routing-options autonomous-system 5
```

### *Configuring Device R2*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 192.168.10.2/24
user@R2# set fe-1/2/1 unit 3 family inet address 192.168.20.2/24
user@R2# set lo0 unit 2 family inet address 2.2.2.2/32
```

2. Configure EBGp.

```
[edit protocols bgp group ext]
user@R2# set type external
user@R2# set export send-direct
user@R2# set export send-static
user@R2# set export send-local
user@R2# set neighbor 192.168.10.1 peer-as 1
user@R2# set neighbor 192.168.20.1 peer-as 3
```

3. Configure the autonomous system (AS) path attribute to be ignored in the Junos OS path selection algorithm.

```
[edit protocols bgp]
user@R2# set path-selection as-path-ignore
```

4. Configure the routing policy.

```
[edit policy-options]
user@R2# set policy-statement send-direct term 1 from protocol direct
user@R2# set policy-statement send-direct term 1 then accept
user@R2# set policy-statement send-local term 1 from protocol local
user@R2# set policy-statement send-local term 1 then accept
user@R2# set policy-statement send-static term 1 from protocol static
user@R2# set policy-statement send-static term 1 then accept
```

5. Configure some static routes.

```
[edit routing-options static]
user@R2# set route 192.168.50.0/24 next-hop 192.168.10.1
user@R2# set route 192.168.40.0/24 next-hop 192.168.10.1
user@R2# set route 192.168.30.0/24 next-hop 192.168.20.1
```

6. Configure the autonomous system (AS) number and the router ID.

```
[edit routing-options]
user@R2# set router-id 2.2.2.2
user@R2# set autonomous-system 2
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 192.168.10.2/24;
    }
  }
}
fe-1/2/1 {
```

```
    unit 3 {
      family inet {
        address 192.168.20.2/24;
      }
    }
  }
lo0 {
  unit 2 {
    family inet {
      address 2.2.2.2/32;
    }
  }
}

user@R2# show policy-options
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
policy-statement send-local {
  term 1 {
    from protocol local;
    then accept;
  }
}
policy-statement send-static {
  term 1 {
    from protocol static;
    then accept;
  }
}

user@R2# show protocols
bgp {
  path-selection as-path-ignore;
  group ext {
    type external;
    export [ send-direct send-static send-local ];
    neighbor 192.168.10.1 {
      peer-as 1;
    }
    neighbor 192.168.20.1 {
      peer-as 3;
    }
  }
}

user@R2# show routing-options
static {
  route 192.168.50.0/24 next-hop 192.168.10.1;
  route 192.168.40.0/24 next-hop 192.168.10.1;
  route 192.168.30.0/24 next-hop 192.168.20.1;
}
router-id 2.2.2.2;
autonomous-system 2;
```



If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on the other devices in the network, changing the interface names and IP addresses, as needed.

### Verification

Confirm that the configuration is working properly.

- [Checking the Neighbor Status on page 3455](#)

### Checking the Neighbor Status

**Purpose** Make sure that from Device R2, the active path to get to AS 4 is through AS 1 and AS 5, not through AS 3.



**NOTE:** To verify the functionality of the `as-path-ignore` statement, you might need to run the `restart routing` command to force reevaluation of the active path. This is because for BGP, if both paths are external, the Junos OS behavior is to prefer the currently active path. This behavior helps to minimize route-flapping. Use caution when restarting the routing protocol process in a production network.

**Action** From operational mode, enter the `restart routing` command.

```
user@R2> restart routing
Routing protocols process started, pid 49396
```

From operational mode, enter the `show route 4.4.4.4 protocol bgp` command.

```
user@R2> show route 4.4.4.4 protocol bgp
inet.0: 12 destinations, 25 routes (12 active, 0 holddown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

4.4.4.4/32          *[BGP/170] 00:00:12, localpref 100
                    AS path: 154 I
                    > to 192.168.10.1 via fe-1/2/0.2
                    [BGP/170] 00:00:08, localpref 100
                    AS path: 34 I
                    > to 192.168.20.1 via fe-1/2/1.3
```

**Meaning** The asterisk (\*) is next to the path learned from R1, meaning that this is the active path. The AS path for the active path is 1 5 4, which is longer than the AS path (3 4) for the nonactive path learned from Router R3.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

## Example: Removing Private AS Numbers

- [Understanding Private AS Number Removal from AS Paths on page 3456](#)
- [Example: Removing Private AS Numbers from AS Paths on page 3457](#)

### Understanding Private AS Number Removal from AS Paths

---

By default, when BGP advertises AS paths to remote systems, it includes all AS numbers, including private AS numbers. You can configure the software so that it removes private AS numbers from AS paths. Doing this is useful when any of the following circumstances are true:

- A remote AS for which you provide connectivity is multihomed, but only to the local AS.
- The remote AS does not have an officially allocated AS number.
- It is not appropriate to make the remote AS a confederation member AS of the local AS.

Most companies acquire their own AS number. Some companies also use private AS numbers to connect to their public AS network. These companies might use a different private AS number for each region in which their company does business. In any implementation, announcing a private AS number to the Internet must be avoided. Service providers can use the **remove-private** statement to prevent advertising private AS numbers to the Internet.

In an enterprise scenario, suppose that you have multiple AS numbers in your company, some of which are private AS numbers, and one with a public AS number. The one with a public AS number has a direct connection to the service provider. In the AS that connects directly to the service provider, you can use the **remove-private** statement to filter out any private AS numbers in the advertisements that are sent to the service provider.



**CAUTION:** Changing configuration statements that affect BGP peers, such as enabling or disabling **remove-private** or renaming a BGP group, resets the BGP sessions. Changes that affect BGP peers should only be made when resetting a BGP session is acceptable.

The AS numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.



**NOTE:** As of Junos OS 10.0R2 and later, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the **as-override** statement instead of the **remove-private** statement.

The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The software is preconfigured with knowledge of the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document. The set of AS numbers reserved as private are in the range from 64,512 through 65,534, inclusive.

### Example: Removing Private AS Numbers from AS Paths

This example demonstrates the removal of a private AS number from the advertised AS path to avoid announcing the private AS number to the Internet.

- [Requirements on page 3457](#)
- [Overview on page 3457](#)
- [Configuration on page 3458](#)
- [Verification on page 3460](#)

#### Requirements


No special configuration beyond device initialization is required before you configure this example.

#### Overview

Service providers and enterprise networks use the **remove-private** statement to prevent advertising private AS numbers to the Internet. The **remove-private** statement works in the outbound direction. You configure the **remove-private** statement on a device that has a public AS number and that is connected to one or more devices that have private AS numbers. Generally, you would not configure this statement on a device that has a private AS number.

[Figure 79 on page 3457](#) shows the sample topology.

**Figure 79: Topology for Removing a Private AS from the Advertised AS Path**



In this example, Device R1 is connected to its service provider using private AS number 65535. The example shows the **remove-private** statement configured on Device ISP to prevent Device R1's private AS number from being announced to Device R2. Device R2 sees only the AS number of the service provider.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```
set interfaces fe-1/2/0 unit 1 family inet address 192.168.10.1/24
set interfaces lo0 unit 1 family inet address 10.10.10.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 100
set protocols bgp group ext neighbor 192.168.10.10
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.20.0/24 next-hop 192.168.10.10
set routing-options autonomous-system 65535
```

**Device ISP**

```
set interfaces fe-1/2/0 unit 2 family inet address 192.168.10.10/24
set interfaces fe-1/2/1 unit 3 family inet address 192.168.20.20/24
set interfaces lo0 unit 2 family inet address 10.10.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext neighbor 192.168.10.1 peer-as 65535
set protocols bgp group ext neighbor 192.168.20.1 remove-private
set protocols bgp group ext neighbor 192.168.20.1 peer-as 200
set routing-options autonomous-system 100
```

**Device R2**

```
set interfaces fe-1/2/0 unit 4 family inet address 192.168.20.1/24
set interfaces lo0 unit 3 family inet address 10.10.20.1/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct
set protocols bgp group ext export send-static
set protocols bgp group ext peer-as 100
set protocols bgp group ext neighbor 192.168.20.20
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set routing-options static route 192.168.10.0/24 next-hop 192.168.20.20
set routing-options autonomous-system 200
```

### Device ISP

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device ISP:

1. Configure the interfaces.  
[edit interfaces]

```

user@ISP# set fe-1/2/0 unit 2 family inet address 192.168.10.10/24
user@ISP# set fe-1/2/1 unit 3 family inet address 192.168.20.20/24
user@ISP# set lo0 unit 2 family inet address 10.10.0.1/32

```

2. Configure EBGP.

```

[edit protocols bgp group ext]
user@ISP# set type external
user@ISP# set neighbor 192.168.10.1 peer-as 65535
user@ISP# set neighbor 192.168.20.1 peer-as 200

```

3. For the neighbor in autonomous system (AS) 200 (Device R2), remove private AS numbers from the advertised AS paths.

```

[edit protocols bgp group ext]
user@ISP# set neighbor 192.168.20.1 remove-private

```

4. Configure the AS number.

```

[edit routing-options]
user@ISP# set autonomous-system 100

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@ISP# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 192.168.10.10/24;
    }
  }
}
fe-1/2/1 {
  unit 3 {
    family inet {
      address 192.168.20.20/24;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.10.0.1/32;
    }
  }
}
}

user@ISP# show protocols
bgp {
  group ext {
    type external;
    neighbor 192.168.10.1 {
      peer-as 65535;
    }
    neighbor 192.168.20.1 {
      remove-private;
    }
  }
}

```

```
        peer-as 200;
    }
}
}
```

```
user@ISP# show routing-options
autonomous-system 100;
```

If you are done configuring the device, enter **commit** from configuration mode. Repeat the configuration on Device R1 and Device R2, changing the interface names and IP address, as needed, and adding the routing policy configuration.

### **Verification**

Confirm that the configuration is working properly.

- [Checking the Neighbor Status on page 3460](#)
- [Checking the Routing Tables on page 3461](#)
- [Checking the AS Path When the remove-private Statement Is Deactivated on page 3461](#)

### **Checking the Neighbor Status**

**Purpose** Make sure that Device ISP has the **remove-private** setting enabled in its neighbor session with Device R2.

**Action** From operational mode, enter the **show bgp neighbor 192.168.20.1** command.

```
user@ISP> show bgp neighbor 192.168.20.1
Peer: 192.168.20.1+179 AS 200 Local: 192.168.20.20+60216 AS 100
  Type: External State: Established Flags: <ImportEval Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference RemovePrivateAS PeerAS Refresh>
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.10.20.1 Local ID: 10.10.0.1 Active Holdtime: 90
  Keepalive Interval: 30 Peer index: 0
  BFD: disabled, down
  Local Interface: fe-1/2/1.3
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 200)
  Peer does not support Addpath
  Table inet.0 Bit: 10001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes: 1
    Received prefixes: 3
    Accepted prefixes: 2
    Suppressed due to damping: 0
    Advertised prefixes: 1
```

```

Last traffic (seconds): Received 10    Sent 16    Checked 55
Input messages: Total 54    Updates 3    Refreshes 0    Octets 1091
Output messages: Total 54    Updates 1    Refreshes 0    Octets 1118
Output Queue[0]: 0

```

**Meaning** The `RemovePrivateAS` option shows that Device ISP has the expected setting.

### *Checking the Routing Tables*

**Purpose** Make sure that the devices have the expected routes and AS paths.

**Action** From operational mode, enter the `show route protocol bgp` command.

```

user@R1> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.20.1/32    *[BGP/170] 00:28:57, localpref 100
                 AS path: 100 200 I
                 > to 192.168.10.10 via fe-1/2/0.1

user@ISP> show route protocol bgp

inet.0: 7 destinations, 11 routes (7 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32    *[BGP/170] 00:29:40, localpref 100
                 AS path: 65535 I
                 > to 192.168.10.1 via fe-1/2/0.2
10.10.20.1/32    *[BGP/170] 00:29:36, localpref 100
                 AS path: 200 I
                 > to 192.168.20.1 via fe-1/2/1.3
192.168.10.0/24  [BGP/170] 00:29:40, localpref 100
                 AS path: 65535 I
                 > to 192.168.10.1 via fe-1/2/0.2
192.168.20.0/24  [BGP/170] 00:29:36, localpref 100
                 AS path: 200 I
                 > to 192.168.20.1 via fe-1/2/1.3

user@R2> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32    *[BGP/170] 00:29:53, localpref 100
                 AS path: 100 I
                 > to 192.168.20.20 via fe-1/2/0.4

```

**Meaning** Device ISP has the private AS number 65535 in its AS path to Device R1. However, Device ISP does not advertise this private AS number to Device R2. This is shown in the routing table of Device R2. Device R2's path to Device R1 contains only the AS number for Device ISP.

### *Checking the AS Path When the remove-private Statement Is Deactivated*

**Purpose** Verify that without the `remove-private` statement, the private AS number appears in Device R2's routing table.

**Action** From configuration mode on Device ISP, enter the **deactivate remove-private** command and then recheck the routing table on Device R2.

```
[protocols bgp group ext neighbor 192.168.20.1]
user@ISP# deactivate remove-private
user@ISP# commit

user@R2> show route protocol bgp
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.1/32      *[BGP/170] 00:00:54, localpref 100
                   AS path: 100 65535 I
                   > to 192.168.20.20 via fe-1/2/0.4
```

**Meaning** Private AS number 65535 appears in Device R2's AS path to Device R1.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

---

## BGP BFD Configuration

- [Example: Configuring BFD for BGP on page 3462](#)
- [Example: Configuring BFD Authentication for BGP on page 3471](#)

### Example: Configuring BFD for BGP

- [Understanding BFD for BGP on page 3462](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 3463](#)

---

#### Understanding BFD for BGP

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than default failure detection mechanisms for BGP, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



In Junos OS Release 8.3 and later, BFD is supported on internal BGP (IBGP) and multihop external BGP (EBGP) sessions as well as on single-hop EBGP sessions. In Junos OS Release 9.1 through Junos OS Release 11.1, BFD supports IPv6 interfaces in static routes only. In Junos OS Release 11.2 and later, BFD supports IPv6 interfaces with BGP.

### Example: Configuring BFD on Internal BGP Peer Sessions

---

This example shows how to configure internal BGP (IBGP) peer sessions with the Bidirectional Forwarding Detection (BFD) protocol to detect failures in a network.

- [Requirements on page 3463](#)
- [Overview on page 3463](#)
- [Configuration on page 3464](#)
- [Verification on page 3468](#)

#### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

#### **Overview**

The minimum configuration to enable BFD on IBGP sessions is to include the **bfd-liveness-detection minimum-interval** statement in the BGP configuration of all neighbors participating in the BFD session. The **minimum-interval** statement specifies the minimum transmit and receive intervals for failure detection. Specifically, this value represents the minimum interval after which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value from 1 through 255,000 milliseconds.

Optionally, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements. For information about these and other optional BFD configuration statements, see [bfd-liveness-detection](#).



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD less than 100 ms for Routing Engine-based sessions and less than 10 ms for distributed BFD sessions can cause undesired BFD flapping.


Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

BFD is supported on the default routing instance (the main router), routing instances, and logical systems. This example shows BFD on logical systems.

Figure 80 on page 3464 shows a typical network with internal peer sessions.

**Figure 80: Typical Network with IBGP Sessions**



#### Configuration

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

##### Device A

```
set logical-systems A interfaces lt-1/2/0 unit 1 description to-B
set logical-systems A interfaces lt-1/2/0 unit 1 encapsulation ethernet
set logical-systems A interfaces lt-1/2/0 unit 1 peer-unit 2
```

```

set logical-systems A interfaces lt-1/2/0 unit 1 family inet address 10.10.10.1/30
set logical-systems A interfaces lo0 unit 1 family inet address 192.168.6.5/32
set logical-systems A protocols bgp group internal-peers type internal
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-bfd
set logical-systems A protocols bgp group internal-peers traceoptions flag bfd detail
set logical-systems A protocols bgp group internal-peers local-address 192.168.6.5
set logical-systems A protocols bgp group internal-peers export send-direct
set logical-systems A protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems A protocols bgp group internal-peers neighbor 192.163.6.4
set logical-systems A protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems A protocols ospf area 0.0.0.0 interface lo0.1 passive
set logical-systems A protocols ospf area 0.0.0.0 interface lt-1/2/0.1
set logical-systems A policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems A policy-options policy-statement send-direct term 2 then accept
set logical-systems A routing-options router-id 192.168.6.5
set logical-systems A routing-options autonomous-system 17

```

**Device B**

```

set logical-systems B interfaces lt-1/2/0 unit 2 description to-A
set logical-systems B interfaces lt-1/2/0 unit 2 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 2 peer-unit 1
set logical-systems B interfaces lt-1/2/0 unit 2 family inet address 10.10.10.2/30
set logical-systems B interfaces lt-1/2/0 unit 5 description to-C
set logical-systems B interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems B interfaces lt-1/2/0 unit 5 peer-unit 6
set logical-systems B interfaces lt-1/2/0 unit 5 family inet address 10.10.10.5/30
set logical-systems B interfaces lo0 unit 2 family inet address 192.163.6.4/32
set logical-systems B protocols bgp group internal-peers type internal
set logical-systems B protocols bgp group internal-peers local-address 192.163.6.4
set logical-systems B protocols bgp group internal-peers export send-direct
set logical-systems B protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems B protocols bgp group internal-peers neighbor 192.168.40.4
set logical-systems B protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems B protocols ospf area 0.0.0.0 interface lo0.2 passive
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.2
set logical-systems B protocols ospf area 0.0.0.0 interface lt-1/2/0.5
set logical-systems B policy-options policy-statement send-direct term 2 from protocol
    direct
set logical-systems B policy-options policy-statement send-direct term 2 then accept
set logical-systems B routing-options router-id 192.163.6.4
set logical-systems B routing-options autonomous-system 17

```

**Device C**

```

set logical-systems C interfaces lt-1/2/0 unit 6 description to-B
set logical-systems C interfaces lt-1/2/0 unit 6 encapsulation ethernet
set logical-systems C interfaces lt-1/2/0 unit 6 peer-unit 5
set logical-systems C interfaces lt-1/2/0 unit 6 family inet address 10.10.10.6/30
set logical-systems C interfaces lo0 unit 3 family inet address 192.168.40.4/32
set logical-systems C protocols bgp group internal-peers type internal
set logical-systems C protocols bgp group internal-peers local-address 192.168.40.4
set logical-systems C protocols bgp group internal-peers export send-direct
set logical-systems C protocols bgp group internal-peers bfd-liveness-detection
    minimum-interval 1000
set logical-systems C protocols bgp group internal-peers neighbor 192.163.6.4

```

```
set logical-systems C protocols bgp group internal-peers neighbor 192.168.6.5
set logical-systems C protocols ospf area 0.0.0.0 interface lo0.3 passive
set logical-systems C protocols ospf area 0.0.0.0 interface lt-1/2/0.6
set logical-systems C policy-options policy-statement send-direct term 2 from protocol
  direct
set logical-systems C policy-options policy-statement send-direct term 2 then accept
set logical-systems C routing-options router-id 192.168.40.4
set logical-systems C routing-options autonomous-system 17
```

### *Configuring Device A*

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Set the CLI to Logical System A.

```
user@host> set cli logical-system A
```

2. Configure the interfaces.

```
[edit interfaces lt-1/2/0 unit 1]
user@host:A# set description to-B
user@host:A# set encapsulation ethernet
user@host:A# set peer-unit 2
user@host:A# set family inet address 10.10.10.1/30
```

```
[edit interfaces lo0 unit 1]
user@host:A# set family inet address 192.168.6.5/32
```

3. Configure BGP.

The **neighbor** statements are included for both Device B and Device C, even though Device A is not directly connected to Device C.

```
[edit protocols bgp group internal-peers]
user@host:A# set type internal
user@host:A# set local-address 192.168.6.5
user@host:A# set export send-direct
user@host:A# set neighbor 192.163.6.4
user@host:A# set neighbor 192.168.40.4
```

4. Configure BFD.

```
[edit protocols bgp group internal-peers]
user@host:A# set bfd-liveness-detection minimum-interval 1000
```

You must configure the same minimum interval on the connecting peer.

5. (Optional) Configure BFD tracing.

```
[edit protocols bgp group internal-peers]
user@host:A# set traceoptions file bgp-bfd
user@host:A# set traceoptions flag bfd detail
```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@host:A# set interface lo0.1 passive
user@host:A# set interface lt-1/2/0.1
```

7. Configure a policy that accepts direct routes.

Other useful options for this scenario might be to accept routes learned through OSPF or local routes.

```
[edit policy-options policy-statement send-direct term 2]
user@host:A# set from protocol direct
user@host:A# set then accept
```

8. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host:A# set router-id 192.168.6.5
user@host:A# set autonomous-system 17
```

9. If you are done configuring the device, enter **commit** from configuration mode. Repeat these steps to configure Device B and Device C.

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host:A# show interfaces
lt-1/2/0 {
  unit 1 {
    description to-B;
    encapsulation ethernet;
    peer-unit 2;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@host:A# show policy-options
policy-statement send-direct {
  term 2 {
    from protocol direct;
    then accept;
  }
}

user@host:A# show protocols
bgp {
  group internal-peers {
    type internal;
    traceoptions {
```

```
        file bgp-bfd;
        flag bfd detail;
    }
    local-address 192.168.6.5;
    export send-direct;
    bfd-liveness-detection {
        minimum-interval 1000;
    }
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.1 {
            passive;
        }
        interface lt-1/2/0.1;
    }
}

user@host:A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;
```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying That BFD Is Enabled on page 3468](#)
- [Verifying That BFD Sessions Are Up on page 3469](#)
- [Viewing Detailed BFD Events on page 3469](#)
- [Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface on page 3470](#)

### **Verifying That BFD Is Enabled**

**Purpose** Verify that BFD is enabled between the IBGP peers.

**Action** From operational mode, enter the **show bgp neighbor** command. You can use the **| match bfd** filter to narrow the output.

```
user@host:A> show bgp neighbor | match bfd
Options: <BfdEnabled>
  BFD: enabled, up
  Trace file: /var/log/A/bgp-bfd size 131072 files 10
Options: <BfdEnabled>
  BFD: enabled, up
  Trace file: /var/log/A/bgp-bfd size 131072 files 10
```

**Meaning** The output shows that Logical System A has two neighbors with BFD enabled. When BFD is not enabled, the output displays **BFD: disabled, down**, and the **<BfdEnabled>** option is absent. If BFD is enabled and the session is down, the output displays **BFD: enabled, down**. The output also shows that BFD-related events are being written to a log file because trace operations are configured.

**Verifying That BFD Sessions Are Up**

**Purpose** Verify that the BFD sessions are up, and view details about the BFD sessions.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@host:A> show bfd session extensive
```

| Address     | State | Interface | Detect Time | Transmit Interval | Multiplier |
|-------------|-------|-----------|-------------|-------------------|------------|
| 192.163.6.4 | Up    |           | 3.000       | 1.000             | 3          |

Client BGP, TX interval 1.000, RX interval 1.000  
Session up time 00:54:40  
Local diagnostic None, remote diagnostic None  
Remote state Up, version 1  
Logical system 12, routing table index 25  
Min async interval 1.000, min slow interval 1.000  
Adaptive async TX interval 1.000, RX interval 1.000  
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
Local discriminator 10, remote discriminator 9  
Echo mode disabled/inactive  
Multi-hop route table 25, local-address 192.168.6.5

| Address      | State | Interface | Detect Time | Transmit Interval | Multiplier |
|--------------|-------|-----------|-------------|-------------------|------------|
| 192.168.40.4 | Up    |           | 3.000       | 1.000             | 3          |

Client BGP, TX interval 1.000, RX interval 1.000  
Session up time 00:48:03  
Local diagnostic None, remote diagnostic None  
Remote state Up, version 1  
Logical system 12, routing table index 25  
Min async interval 1.000, min slow interval 1.000  
Adaptive async TX interval 1.000, RX interval 1.000  
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3  
Remote min TX interval 1.000, min RX interval 1.000, multiplier 3  
Local discriminator 14, remote discriminator 13  
Echo mode disabled/inactive  
Multi-hop route table 25, local-address 192.168.6.5

2 sessions, 2 clients  
Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps

**Meaning** The TX interval 1.000, RX interval 1.000 output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

**Viewing Detailed BFD Events**

**Purpose** View the contents of the BFD trace file to assist in troubleshooting, if needed.

**Action** From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
Aug 15 17:07:25 trace_on: Tracing to "/var/log/A/bgp-bfd" started
Aug 15 17:07:26.492190 bgp_peer_init: BGP peer 192.163.6.4 (Internal AS 17) local
address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:26.493176 bgp_peer_init: BGP peer 192.168.40.4 (Internal AS 17) local
```

```

address 192.168.6.5 not found. Leaving peer idled
Aug 15 17:07:32.597979 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:07:32.599623 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:07:36.869394 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:07:36.870624 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:04.599220 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host
Aug 15 17:08:04.601135 bgp_connect_start: connect 192.163.6.4 (Internal AS 17):
No route to host
Aug 15 17:08:08.869717 task_connect: task BGP_17.192.168.40.4+179 addr
192.168.40.4+179: No route to host
Aug 15 17:08:08.869934 bgp_connect_start: connect 192.168.40.4 (Internal AS 17):
No route to host
Aug 15 17:08:36.603544 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:08:36.606726 bgp_read_message: 192.163.6.4 (Internal AS 17): 0 bytes
buffered
Aug 15 17:08:36.609119 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:36.734033 advertising receiving-speaker only capability to neighbor
192.168.40.4 (Internal AS 17)
Aug 15 17:08:36.738436 Initiated BFD session to peer 192.168.40.4 (Internal AS
17): address=192.168.40.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:08:40.537552 BFD session to peer 192.163.6.4 (Internal AS 17) up
Aug 15 17:08:40.694410 BFD session to peer 192.168.40.4 (Internal AS 17) up

```

**Meaning** Before the routes are established, the **No route to host** message appears in the output. After the routes are established, the last two lines show that both BFD sessions come up.

#### *Viewing Detailed BFD Events After Deactivating and Reactivating a Loopback Interface*

**Purpose** Check to see what happens after bringing down a router or switch and then bringing it back up. To simulate bringing down a router or switch, deactivate the loopback interface on Logical System B.

**Action** 1. From configuration mode, enter the **deactivate logical-systems B interfaces lo0 unit 2 family inet** command.

```

user@host:A# deactivate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit

```

2. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```

user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:20:55.995648 bgp_read_v4_message:9747: NOTIFICATION received from
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 6 (Other Configuration
Change)
Aug 15 17:20:56.004508 Terminated BFD session to peer 192.163.6.4 (Internal
AS 17)
Aug 15 17:21:28.007755 task_connect: task BGP_17.192.163.6.4+179 addr
192.163.6.4+179: No route to host

```



```
Aug 15 17:21:28.008597 bgp_connect_start: connect 192.163.6.4 (Internal AS 17): No route to host
```

3. From configuration mode, enter the **activate logical-systems B interfaces lo0 unit 2 family inet** command.

```
user@host:A# activate logical-systems B interfaces lo0 unit 2 family inet
user@host:A# commit
```

4. From operational mode, enter the **file show /var/log/A/bgp-bfd** command.

```
user@host:A> file show /var/log/A/bgp-bfd
...
Aug 15 17:25:53.623743 advertising receiving-speaker only capability to neighbor
192.163.6.4 (Internal AS 17)
Aug 15 17:25:53.631314 Initiated BFD session to peer 192.163.6.4 (Internal AS
17): address=192.163.6.4 ifindex=0 ifname=(none) txivl=1000 rxivl=1000 mult=3
ver=255
Aug 15 17:25:57.570932 BFD session to peer 192.163.6.4 (Internal AS 17) up
```

#### Related Documentation

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

### Example: Configuring BFD Authentication for BGP

- [Understanding BFD Authentication for BGP on page 3471](#)
- [Example: Configuring BFD Authentication for BGP on page 3473](#)

#### Understanding BFD Authentication for BGP

Bidirectional Forwarding Detection protocol (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over BGP. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3472](#)
- [Security Authentication Keychains on page 3472](#)
- [Strict Versus Loose Authentication on page 3473](#)

### **BFD Authentication Algorithms**

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

### **Security Authentication Keychains**

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### ***Strict Versus Loose Authentication***

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Example: Configuring BFD Authentication for BGP**

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over BGP. Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the BGP protocol.
2. Associate the authentication keychain with the BGP protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on BGP:

- [Configuring BFD Authentication Parameters on page 3473](#)
- [Viewing Authentication Information for BFD Sessions on page 3475](#)

### ***Configuring BFD Authentication Parameters***

BFD authentication can be configured for the entire BGP protocol, or a specific BGP group, neighbor, or routing instance.

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.

[edit]

```
user@host# set protocols bgp bfd-liveness-detection authentication algorithm
keyed-sha-1
```

```
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
algorithm keyed-sha-1
```

```
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication algorithm keyed-sha-1
```



**NOTE:** Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on BGP with the unique security authentication keychain attributes.

The keychain name you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication keychain bfd-bgp
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
keychain bfd-bgp
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication keychain bfd-bgp
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
user@host# set authentication-key-chains key-chain bfd-bgp key 53 secret
$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host# set protocols bgp bfd-liveness-detection authentication loose-check
user@host# set protocols bgp group bgp-gr1 bfd-liveness-detection authentication
loose-check
user@host# set protocols bgp group bgp-gr1 neighbor 10.10.10.7 bfd-liveness-detection
authentication loose-check
```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat these steps to configure the other end of the BFD session.



**NOTE:** BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

### Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **bgp-gr1** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-bgp**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9L.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols bgp]
group bgp-gr1 {
  bfd-liveness-detection {
    authentication {
      algorithm keyed-sha-1;
      key-chain bfd-bgp;
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-bgp {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREvsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$9$a5jiKW9L.reP38ny.TszF2/9";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

### show bfd session detail

```
user@host# show bfd session detail
```

| Address   | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|---|-------|------------|-------------|-------------------|------------|
| 50.0.0.2  | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |
| Client BGP, TX interval 0.300, RX interval 0.300, <b>Authenticate</b> |       |            |             |                   |            |

```
Session up time 3d 00:34
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
```

#### show bfd session extensive

```
user@host# show bfd session extensive
```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 50.0.0.2 | Up    | ge-0/1/5.0 | 0.900       | 0.300             | 3          |

```
Client BGP, TX interval 0.300, RX interval 0.300, Authenticate
keychain bfd-bgp, algo keyed-sha-1, mode strict
Session up time 00:04:42
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.300, min slow interval 1.000
Adaptive async TX interval 0.300, RX interval 0.300
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
Local discriminator 2, remote discriminator 2
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd-bgp, algo keyed-sha-1, mode strict
```

#### Related Documentation

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

---

## BGP Load Balancing Configuration

- [Examples: Configuring BGP Multipath on page 3476](#)
- [Example: Advertising Multiple BGP Paths to a Destination on page 3494](#)
- [Example: Advertising Multiple Paths in BGP on page 3520](#)
- [Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing on page 3545](#)

### Examples: Configuring BGP Multipath

- [Understanding BGP Multipath on page 3476](#)
- [Example: Load Balancing BGP Traffic on page 3477](#)
- [Example: Configuring Single-Hop EBGP Peers to Accept Remote Next Hops on page 3482](#)

---

#### Understanding BGP Multipath

BGP multipath allows you to select multiple internal or external BGP peers as active paths. Selecting multiple paths enables BGP peering to load-balance traffic across an Autonomous System (AS) confederation boundary.

A path is considered a BGP equal-cost path (and is used for forwarding) if a tie-break is performed. The tie-break is performed after the BGP route path selection step that chooses the next-hop path that is resolved through the IGP route with the lowest metric. All paths with the same neighboring AS, learned by a multipath-enabled BGP neighbor are considered in the path selection process.

BGP, typically selects only one best path for each prefix and installs that route in the routing table. When BGP multipath is enabled, the device selects multiple equal-cost EBGp paths as the best paths to reach a given destination, and all these paths are installed in the routing table. BGP advertises only the active path to its neighbors. However, you can configure BGP to advertise multiple paths to the same destination for redundancy and load balancing.

The Junos OS BGP multipath feature supports the following applications:

- Load balancing across multiple links between two routing devices belonging to different autonomous systems (ASs)
- Load balancing across a common subnet or multiple subnets to different routing devices belonging to the same peer AS
- Load balancing across multiple links between two routing devices belonging to different external confederation peers
- Load balancing across a common subnet or multiple subnets to different routing devices belonging to external confederation peers

In a common scenario for load balancing, a customer is multihomed to multiple routers in a point of presence (POP). The default behavior is to send all traffic across only one of the available links. Load balancing causes traffic to use two or more of the links.



**NOTE:** BGP multipath does not apply to paths that share the same MED-plus-IGP cost, yet differ in IGP cost. Multipath path selection is based on the IGP cost metric, even if two paths have the same MED-plus-IGP cost.

### Example: Load Balancing BGP Traffic

This example shows how to configure BGP to select multiple equal-cost external BGP (EBGP) or internal BGP (IBGP) paths as active paths.

- [Requirements on page 3477](#)
- [Overview on page 3478](#)
- [Configuration on page 3479](#)
- [Verification on page 3481](#)

#### Requirements

Before you begin:

- Configure the device interfaces.
- Configure an interior gateway protocol (IGP).
- Configure BGP.
- Configure a routing policy that exports routes (such as direct routes or IGP routes) from the routing table into BGP.

## Overview

The following steps show how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {  
  from {  
    match-conditions;  
    route-filter destination-prefix match-type <actions>;  
    prefix-list name;  
  }  
  then {  
    load-balance per-packet;  
  }  
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {  
  export policy-name;  
}
```

You cannot apply the export policy to VRF routing instances.

3. Specify all next hops of that route, if more than one exists, when allocating a label corresponding to a route that is being advertised.
4. Configure the forwarding-options hash key for MPLS to include the IP payload.



**NOTE:** On some platforms, you can increase the number of paths that are load balanced by using the **chassis maximum-ecmp** statement. With this statement, you can change the maximum number of equal-cost load-balanced paths to 32 or 64.

---


In this example, Device R1 is in AS 65000 and is connected to both Device R2 and Device R3, which are in AS 65001. This example shows the configuration on Device R1.

## Topology

[Figure 81 on page 3479](#) shows the topology used in this example.



Figure 81: BGP Load Balancing



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols bgp group external type external
set protocols bgp group external peer-as 65001
set protocols bgp group external multipath
set protocols bgp group external neighbor 10.0.1.1
set protocols bgp group external neighbor 10.0.0.2
set policy-options policy-statement loadbal from route-filter 10.0.0.0/16 orlonger
set policy-options policy-statement loadbal then load-balance per-packet
set routing-options forwarding-table export loadbal
set routing-options autonomous-system 65000
```

#### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the BGP peer sessions:

1. Configure the BGP group.
 

```
[edit protocols bgp group external]
user@R1# set type external
user@R1# set peer-as 65001
user@R1# set neighbor 10.0.1.1
user@R1# set neighbor 10.0.0.2
```
2. Enable the BGP group to use multiple paths.



**NOTE:** To disable the default check requiring that paths accepted by BGP multipath must have the same neighboring autonomous system (AS), include the `multiple-as` option.

```
[edit protocols bgp group external]
user@R1# set multipath
```

3. Configure the load-balancing policy.

```
[edit policy-options policy-statement loadbal]
user@R1# set from route-filter 10.0.0.0/16 orlonger
user@R1# set then load-balance per-packet
```

4. Apply the load-balancing policy.

```
[edit routing-options]
user@R1# set forwarding-table export loadbal
```

5. Configure the local autonomous system (AS) number.

```
[edit routing-options]
user@R1# set autonomous-system 65000
```

**Results** From configuration mode, confirm your configuration by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show protocols
bgp {
  group external {
    type external;
    peer-as 65001;
    multipath;
    neighbor 10.0.1.1;
    neighbor 10.0.0.2;
  }
}
```

```
[edit]
user@R1# show policy-options
policy-statement loadbal {
  from {
    route-filter 10.0.0.0/16 orlonger;
  }
  then {
    load-balance per-packet;
  }
}
```

```
[edit]
user@R1# show routing-options
autonomous-system 65000;
forwarding-table {
```

```
export loadbal;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly:

- [Verifying Routes on page 3481](#)
- [Verifying Forwarding on page 3482](#)

### Verifying Routes

**Purpose** Verify that routes are learned from both routers in the neighboring AS.

**Action** From operational mode, run the **show route** command.

```
user@R1> show route 10.0.2.0
inet.0: 12 destinations, 15 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.2.0/30          *[BGP/170] 03:12:32, localpref 100
                    AS path: 65001 I
                    to 10.0.1.1 via ge-1/2/0.0
                    > to 10.0.0.2 via ge-1/2/1.0
                    [BGP/170] 03:12:32, localpref 100
                    AS path: 65001 I
                    > to 10.0.1.1 via ge-1/2/0.0

user@R1> show route 10.0.2.0 detail
inet.0: 12 destinations, 15 routes (12 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
   *BGP      Preference: 170/-101
             Next hop type: Router, Next hop index: 262142
             Next-hop reference count: 3
             Source: 10.0.0.2
             Next hop: 10.0.1.1 via ge-1/2/0.0
             Next hop: 10.0.0.2 via ge-1/2/1.0, selected
             State: <Active Ext>
             Local AS: 65000 Peer AS: 65001
             Age: 3:18:30
             Task: BGP_65001.10.0.0.2+55402
             Announcement bits (1): 2-KRT
             AS path: 65001 I
             Accepted Multipath
             Localpref: 100
             Router ID: 192.168.2.1
   BGP      Preference: 170/-101
             Next hop type: Router, Next hop index: 602
             Next-hop reference count: 5
             Source: 10.0.1.1
             Next hop: 10.0.1.1 via ge-1/2/0.0, selected
             State: <NotBest Ext>
             Inactive reason: Not Best in its group - Active preferred
             Local AS: 65000 Peer AS: 65001
             Age: 3:18:30
             Task: BGP_65001.10.0.1.1+53135
             AS path: 65001 I
```

Accepted  
Localpref: 100  
Router ID: 192.168.3.1

**Meaning** The active path, denoted with an asterisk (\*), has two next hops: 10.0.1.1 and 10.0.0.2 to the 10.0.2.0 destination. The 10.0.1.1 next hop is copied from the inactive path to the active path.



**NOTE:** The `show route detail` command output designates one gateway as selected. This output is potentially confusing in the context of load balancing. The selected gateway is used for many purposes in addition to deciding which gateway to install into the kernel when Junos OS is not performing per-packet load-balancing. For instance, the `ping mpls` command uses the selected gateway when sending packets. Multicast protocols use the selected gateway in some cases to determine the upstream interface. Therefore, even when Junos OS is performing per-packet load-balancing by way of a forwarding-table policy, the selected gateway information is still required for other purposes. It is useful to display the selected gateway for troubleshooting purposes. Additionally, it is possible to use forwarding-table policy to override what is installed into the kernel (for example, by using the `install-nexthop` action). In this case, the next-hop gateway installed in the forwarding table might be a subset of the total gateways displayed in the `show route` command.

---

### Verifying Forwarding

**Purpose** Verify that both next hops are installed in the forwarding table.

**Action** From operational mode, run the `show route forwarding-table` command.

```
user@R1> show route forwarding-table destination 10.0.2.0
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
10.0.2.0/30      user  0         10.0.1.1             ucst  602   5 ge-1/2/0.0
                  10.0.0.2             ucst  522   6 ge-1/2/1.0
```

### Example: Configuring Single-Hop EBGPeers to Accept Remote Next Hops

---

This example shows how to configure a single-hop external BGP (EBGP) peer to accept a remote next hop with which it does not share a common subnet.

- [Requirements on page 3483](#)
- [Overview on page 3483](#)
- [Configuration on page 3484](#)
- [Verification on page 3491](#)

### Requirements

No special configuration beyond device initialization is required before you configure this example.

### Overview

In some situations, it is necessary to configure a single-hop EBGP peer to accept a remote next hop with which it does not share a common subnet. The default behavior is for any next-hop address received from a single-hop EBGP peer that is not recognized as sharing a common subnet to be discarded. The ability to have a single-hop EBGP peer accept a remote next hop to which it is not directly connected also prevents you from having to configure the single-hop EBGP neighbor as a multihop session. When you configure a multihop session in this situation, all next-hop routes learned through this EBGP peer are labeled indirect even when they do share a common subnet. This situation breaks multipath functionality for routes that are recursively resolved over routes that include these next-hop addresses. Configuring the `accept-remote-nexthop` statement allows a single-hop EBGP peer to accept a remote next hop, which restores multipath functionality for routes that are resolved over these next-hop addresses. You can configure this statement at the global, group, and neighbor hierarchy levels for BGP. The statement is also supported on logical systems and the VPN routing and forwarding (VRF) routing instance type. Both the remote next-hop and the EBGP peer must support BGP route refresh as defined in RFC 2918, *Route Refresh Capability in BGP-4*. If the remote peer does not support BGP route refresh, the session is reset.



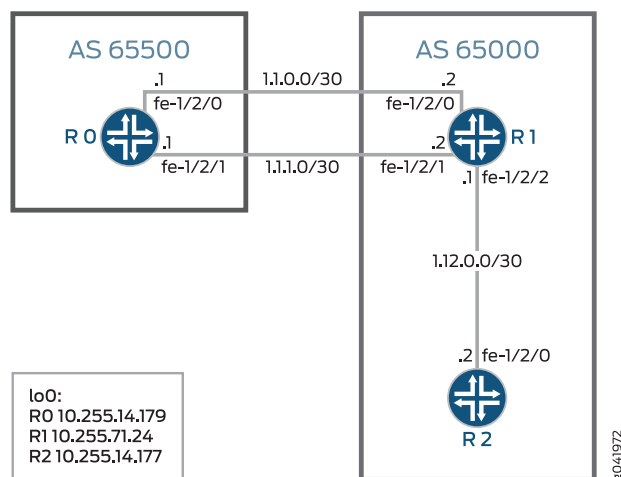
**NOTE:** You cannot configure both the `multihop` and `accept-remote-nexthop` statements for the same EBGP peer.

When you enable a single-hop EBGP peer to accept a remote next hop, you must also configure an import routing policy on the EBGP peer that specifies the remote next-hop address.

This example includes an import routing policy, `agg_route`, that enables a single-hop external BGP peer (Device R1) to accept the remote next-hop 1.1.10.10 for the route to the 1.1.230.0/23 network. At the `[edit protocols bgp]` hierarchy level, the example includes the `import agg_route` statement to apply the policy to the external BGP peer and includes the `accept-remote-nexthop` statement to enable the single-hop EBGP peer to accept the remote next hop.

Figure 82 on page 3484 shows the sample topology.

Figure 82: Topology for Accepting a Remote Next Hop

**Configuration**

- [Device R0 on page 3485](#)
- [Configuring Device R1 on page 3487](#)
- [Configuring Device R2 on page 3490](#)

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R0**

```

set interfaces fe-1/2/0 unit 1 family inet address 1.1.0.1/30
set interfaces fe-1/2/1 unit 2 family inet address 1.1.1.1/30
set interfaces lo0 unit 1 family inet address 10.255.14.179/32
set protocols bgp group ext type external
set protocols bgp group ext export test_route
set protocols bgp group ext export agg_route
set protocols bgp group ext peer-as 65000
set protocols bgp group ext multipath
set protocols bgp group ext neighbor 1.1.0.2
set protocols bgp group ext neighbor 1.1.1.2
set policy-options policy-statement agg_route term 1 from protocol static
set policy-options policy-statement agg_route term 1 from route-filter 1.1.230.0/23 exact
set policy-options policy-statement agg_route term 1 then accept
set policy-options policy-statement test_route term 1 from protocol static
set policy-options policy-statement test_route term 1 from route-filter 1.1.10.10/32 exact
set policy-options policy-statement test_route term 1 then accept
set routing-options static route 1.1.10.10/32 reject
set routing-options static route 1.1.230.0/23 reject
set routing-options autonomous-system 65500

```

**Device R1**

```

set interfaces fe-1/2/0 unit 3 family inet address 1.1.0.2/30
set interfaces fe-1/2/1 unit 4 family inet address 1.12.0.1/30
set interfaces fe-1/2/2 unit 5 family inet address 1.1.1.2/30
set interfaces lo0 unit 2 family inet address 10.255.71.24/32
set protocols bgp accept-remote-nexthop

```

```

set protocols bgp group ext type external
set protocols bgp group ext import agg_route
set protocols bgp group ext peer-as 65500
set protocols bgp group ext multipath
set protocols bgp group ext neighbor 1.1.0.1
set protocols bgp group ext neighbor 1.1.1.1
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.71.24
set protocols bgp group int neighbor 10.255.14.177
set protocols ospf area 0.0.0.0 interface fe-1/2/1.4
set protocols ospf area 0.0.0.0 interface 10.255.71.24
set policy-options policy-statement agg_route term 1 from protocol bgp
set policy-options policy-statement agg_route term 1 from route-filter 1.1.230.0/23 exact
set policy-options policy-statement agg_route term 1 then next-hop 1.1.10.10
set policy-options policy-statement agg_route term 1 then accept
set routing-options autonomous-system 65000

```

**Device R2**

```

set interfaces fe-1/2/0 unit 6 family inet address 1.12.0.2/30
set interfaces lo0 unit 3 family inet address 10.255.14.177/32
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.14.177
set protocols bgp group int neighbor 10.255.71.24
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface 10.255.14.177
set routing-options autonomous-system 65000

```

### Device R0

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R0:

1. Configure the interfaces.
 

```

[edit interfaces fe-1/2/0 unit 1]
user@R0# set family inet address 1.1.0.1/30

[edit interfaces fe-1/2/1 unit 2]
user@R0# set family inet address 1.1.1.1/30

[edit interfaces lo0 unit 1]
user@R0# set family inet address 10.255.14.179/32

```
2. Configure EBGp.
 

```

[edit protocols bgp group ext]
user@R0# set type external
user@R0# set peer-as 65000
user@R0# set neighbor 1.1.0.2
user@R0# set neighbor 1.1.1.2

```
3. Enable multipath BGP between Device R0 and Device R1.
 

```

[edit protocols bgp group ext]
user@R0# set multipath

```

4. Configure static routes to remote networks.  
These routes are not part of the topology. The purpose of these routes is to demonstrate the functionality in this example.

```
[edit routing-options]
user@R0# set static route 1.1.10.10/32 reject
user@R0# set static route 1.1.230.0/23 reject
```

5. Configure routing policies that accept the static routes.

```
[edit policy-options policy-statement agg_route term 1]
user@R0# set from protocol static
user@R0# set from route-filter 1.1.230.0/23 exact
user@R0# set then accept
```

```
[edit policy-options policy-statement test_route term 1]
user@R0# set from protocol static
user@R0# set from route-filter 1.1.10.10/32 exact
user@R0# set then accept
```

6. Export the **agg\_route** and **test\_route** policies from the routing table into BGP.

```
[edit protocols bgp group ext]
user@R0# set export test_route
user@R0# set export agg_route
```

7. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R0# set autonomous-system 65500
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 1.1.0.1/30;
    }
  }
}
fe-1/2/1 {
  unit 2 {
    family inet {
      address 1.1.1.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.255.14.179/32;
    }
  }
}
```



```

}

user@R0# show policy-options
policy-statement agg_route {
  term 1 {
    from {
      protocol static;
      route-filter 1.1.230.0/23 exact;
    }
    then accept;
  }
}
policy-statement test_route {
  term 1 {
    from {
      protocol static;
      route-filter 1.1.10.10/32 exact;
    }
    then accept;
  }
}

user@R0# show protocols
bgp {
  group ext {
    type external;
    export [ test_route agg_route ];
    peer-as 65000;
    multipath;
    neighbor 1.1.0.2;
    neighbor 1.1.1.2;
  }
}

user@R0# show routing-options
static {
  route 1.1.10.10/32 reject;
  route 1.1.230.0/23 reject;
}
autonomous-system 65500;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Configuring Device R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.
 

```

[edit interfaces fe-1/2/0 unit 3]
user@R1# set family inet address 1.1.0.2/30

[edit interfaces fe-1/2/1 unit 4]

```

```
user@R1# set family inet address 1.12.0.1/30
```

```
[edit interfaces fe-1/2/2 unit 5]
```

```
user@R1# set family inet address 1.1.1.2/30
```

```
[edit interfaces lo0 unit 2]
```

```
user@R1# set family inet address 10.255.71.24/32
```

2. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@R1# set interface fe-1/2/1.4
```

```
user@R1# set interface 10.255.71.24
```

3. Enable Device R1 to accept the remote next hop.

```
[edit protocols bgp]
```

```
user@R1# set accept-remote-nexthop
```

4. Configure IBGP.

```
[edit protocols bgp group int]
```

```
user@R1# set type internal
```

```
user@R1# set local-address 10.255.71.24
```

```
user@R1# set neighbor 10.255.14.177
```

5. Configure EBGP.

```
[edit protocols bgp group ext]
```

```
user@R1# set type external
```

```
user@R1# set peer-as 65500
```

```
user@R1# set neighbor 1.1.0.1
```

```
user@R1# set neighbor 1.1.1.1
```

6. Enable multipath BGP between Device R0 and Device R1.

```
[edit protocols bgp group ext]
```

```
user@R1# set multipath
```

7. Configure a routing policy that enables a single-hop external BGP peer (Device R1) to accept the remote next-hop 1.1.10.10 for the route to the 1.1.230.0/23 network.

```
[edit policy-options policy-statement agg_route term 1]
```

```
user@R1# set from protocol bgp
```

```
user@R1# set from route-filter 1.1.230.0/23 exact
```

```
user@R1# set then next-hop 1.1.10.10
```

```
user@R1# set then accept
```

8. Import the **agg\_route** policy into the routing table on Device R1.

```
[edit protocols bgp group ext]
```

```
user@R1# set import agg_route
```

9. Configure the autonomous system (AS) number.

```
[edit routing-options]
```

```
user@R1# set autonomous-system 65000
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 3 {
    family inet {
      address 1.1.0.2/30;
    }
  }
}
fe-1/2/1 {
  unit 4 {
    family inet {
      address 1.12.0.1/30;
    }
  }
}
fe-1/2/2 {
  unit 5 {
    family inet {
      address 1.1.1.2/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 10.255.71.24/32;
    }
  }
}

user@R1# show policy-options
policy-statement agg_route {
  term 1 {
    from {
      protocol bgp;
      route-filter 1.1.230.0/23 exact;
    }
    then {
      next-hop 1.1.10.10;
      accept;
    }
  }
}

user@R1# show protocols
bgp {
  accept-remote-nexthop;
  group ext {
    type external;
    import agg_route;
    peer-as 65500;
    multipath;
    neighbor 1.1.0.1;
  }
}

```

```
        neighbor 1.1.1.1;
    }
    group int {
        type internal;
        local-address 10.255.71.24;
        neighbor 10.255.14.177;
    }
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/1.4;
        interface 10.255.71.24;
    }
}
```

```
user@R1# show routing-options
autonomous-system 65000;
```

If you are done configuring the device, enter **commit** from configuration mode.

### *Configuring Device R2*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.  

```
[edit interfaces fe-1/2/0 unit 6]
user@R2# set family inet address 1.12.0.2/30

[edit interfaces lo0 unit 3]
user@R2# set family inet address 10.255.14.177/32
```
2. Configure OSPF.  

```
[edit protocols ospf area 0.0.0.0]
user@R2# set interface fe-1/2/0.6
user@R2# set interface 10.255.14.177
```
3. Configure IBGP.  

```
[edit protocols bgp group int]
user@R2# set type internal
user@R2# set local-address 10.255.14.177
user@R2# set neighbor 10.255.71.24
```
4. Configure the autonomous system (AS) number.  

```
[edit routing-options]
user@R1# set autonomous-system 65000
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R2# show interfaces
fe-1/2/0 {
  unit 6 {
    family inet {
      address 1.12.0.2/30;
    }
  }
}
lo0 {
  unit 3 {
    family inet {
      address 10.255.14.177/32;
    }
  }
}

user@R2# show protocols
bgp {
  group int {
    type internal;
    local-address 10.255.14.177;
    neighbor 10.255.71.24;
  }
}
ospf {
  area 0.0.0.0 {
    interface fe-1/2/0.6;
    interface 10.255.14.177;
  }
}

user@R2# show routing-options
autonomous-system 65000;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Multipath Route with the Indirect Next Hop Is in the Routing Table on page 3491](#)
- [Deactivating and Reactivating the accept-remote-nexthop Statement on page 3493](#)

### Verifying That the Multipath Route with the Indirect Next Hop Is in the Routing Table

**Purpose** Verify that Device R1 has a route to the 1.1.230.0/23 network.

**Action** From operational mode, enter the **show route 1.1.230.0 extensive** command.

```

user@R1> show route 1.1.230.0 extensive
inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
1.1.230.0/23 (2 entries, 1 announced)
TSI:
KRT in-kernel 1.1.230.0/23 -> {indirect(262142)}
Page 0 idx 1 Type 1 val 9168f6c

```

```

Nexthop: 1.1.10.10
Localpref: 100
AS path: [65000] 65500 I
Communities:
Path 1.1.230.0 from 1.1.0.1 Vector len 4. Val: 1
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x90c44d8
    Next-hop reference count: 4
    Source: 1.1.0.1
    Next hop type: Router, Next hop index: 262143
    Next hop: 1.1.0.1 via fe-1/2/0.3, selected
    Next hop: 1.1.1.1 via fe-1/2/2.5
    Protocol next hop: 1.1.10.10
    Indirect next hop: 91c0000 262142
    State: <Active Ext>
    Local AS: 65000 Peer AS: 65500
    Age: 2:55:31 Metric2: 0
    Task: BGP_65500.1.1.0.1+64631
    Announcement bits (3): 2-KRT 3-BGP_RT_Background 4-Resolve tree
1
  AS path: 65500 I
  Accepted Multipath
  Localpref: 100
  Router ID: 10.255.14.179
  Indirect next hops: 1
    Protocol next hop: 1.1.10.10
    Indirect next hop: 91c0000 262142
    Indirect path forwarding next hops: 2
      Next hop type: Router
      Next hop: 1.1.0.1 via fe-1/2/0.3
      Next hop: 1.1.1.1 via fe-1/2/2.5
    1.1.10.10/32 Originating RIB: inet.0
    Node path count: 1
    Forwarding nexthops: 2
      Nexthop: 1.1.0.1 via fe-1/2/0.3
      Nexthop: 1.1.1.1 via fe-1/2/2.5
  BGP Preference: 170/-101
    Next hop type: Indirect
    Address: 0x90c44d8
    Next-hop reference count: 4
    Source: 1.1.1.1
    Next hop type: Router, Next hop index: 262143
    Next hop: 1.1.0.1 via fe-1/2/0.3, selected
    Next hop: 1.1.1.1 via fe-1/2/2.5
    Protocol next hop: 1.1.10.10
    Indirect next hop: 91c0000 262142
    State: <NotBest Ext>
    Inactive reason: Not Best in its group - Update source
    Local AS: 65000 Peer AS: 65500
    Age: 2:55:27 Metric2: 0
    Task: BGP_65500.1.1.1.1+53260
    AS path: 65500 I
    Accepted
    Localpref: 100
    Router ID: 10.255.14.179
    Indirect next hops: 1
      Protocol next hop: 1.1.10.10
      Indirect next hop: 91c0000 262142
      Indirect path forwarding next hops: 2
        Next hop type: Router

```

```

Next hop: 1.1.0.1 via fe-1/2/0.3
Next hop: 1.1.1.1 via fe-1/2/2.5
1.1.10.10/32 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 2
  Nexthop: 1.1.0.1 via fe-1/2/0.3
  Nexthop: 1.1.1.1 via fe-1/2/2.5

```

**Meaning** The output shows that Device R1 has a route to the 1.1.230.0 network with the multipath feature enabled (**Accepted Multipath**). The output also shows that the route has an indirect next hop of 1.1.10.10.

### *Deactivating and Reactivating the accept-remote-nexthop Statement*

**Purpose** Make sure that the multipath route with the indirect next hop is removed from the routing table when you deactivate the **accept-remote-nexthop** statement.

**Action** 1. From configuration mode, enter the **deactivate protocols bgp accept-remote-nexthop** command.

```

user@R1# deactivate protocols bgp accept-remote-nexthop
user@R1# commit

```

2. From operational mode, enter the **show route 1.1.230.0** command.

```

user@R1> show route 1.1.230.0

```

3. From configuration mode, reactivate the statement by entering the **activate protocols bgp accept-remote-nexthop** command.

```

user@R1# activate protocols bgp accept-remote-nexthop
user@R1# commit

```

4. From operational mode, reenter the **show route 1.1.230.0** command.

```

user@R1> show route 1.1.230.0

```

```

inet.0: 11 destinations, 13 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

```

```

1.1.230.0/23      *[BGP/170] 03:13:19, localpref 100
                  AS path: 65500 I
                  > to 1.1.0.1 via fe-1/2/0.3
                  to 1.1.1.1 via fe-1/2/2.5
                  [BGP/170] 03:13:15, localpref 100, from 1.1.1.1
                  AS path: 65500 I
                  > to 1.1.0.1 via fe-1/2/0.3
                  to 1.1.1.1 via fe-1/2/2.5

```

**Meaning** When the **accept-remote-nexthop** statement is deactivated, the multipath route to the 1.1.230.0 network is removed from the routing table .

**Related Documentation**

- *Example: Overriding the Default BGP Routing Policy on PTX Series Packet Transport Routers*
- *Example: Load Balancing BGP Traffic with Unequal Bandwidth Allocated to the Paths*

## Example: Advertising Multiple BGP Paths to a Destination

- [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP on page 3494](#)
- [Example: Advertising Multiple Paths in BGP on page 3495](#)

### Understanding the Advertisement of Multiple Paths to a Single Destination in BGP

---

BGP peers advertise routes to each other in update messages. BGP stores its routes in the Junos OS routing table (**inet.0**). For each prefix in the routing table, the routing protocol process selects a single best path, called the active path. Unless you configure BGP to advertise multiple paths to the same destination, BGP advertises only the active path.

Instead of advertising only the active path to a destination, you can configure BGP to advertise multiple paths to the destination. Within an autonomous system (AS), the availability of multiple exit points to reach a destination provides the following benefits:

- **Fault tolerance**—Path diversity leads to reduction in restoration time after failure. For instance, a border after receiving multiple paths to the same destination can precompute a backup path and have it ready so that when the primary path becomes invalid, the border routing device can use the backup to quickly restore connectivity. Without a backup path, the restoration time depends on BGP reconvergence, which includes withdraw and advertisement messages in the network before a new best path can be learned.
- **Load balancing**—The availability of multiple paths to reach the same destination enables load balancing of traffic, if the routing within the AS meets certain constraints.
- **Maintenance**—The availability of alternate exit points allows for graceful maintenance operation of routers.

The following limitations apply to advertising multiple routes in BGP:

- Address families supported:
  - IPv4 unicast (**family inet unicast**)
  - IPv6 unicast (**family inet6 unicast**)
  - IPv4 labeled unicast (**family inet labeled-unicast**)
  - IPv6 labeled unicast (**family inet6 labeled-unicast**)
- Internal BGP (IBGP) peers only. No support on external BGP (EBGP) peers.
- Master instance only. No support for routing instances.
- Graceful restart and nonstop active routing (NSR) are supported.
- No BGP Monitoring Protocol (BMP) support.



- No support for EBGP sessions between confederations.
- Prefix policies enable you to filter routes on a router that is configured to advertise multiple paths to a destination. Prefix policies can only match prefixes. They cannot match route attributes, and they cannot change the attributes of routes.

### Example: Advertising Multiple Paths in BGP

In this example, BGP routers are configured to advertise multiple paths instead of advertising only the active path. Advertising multiple paths in BGP is specified in Internet draft draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*.

- [Requirements on page 3495](#)
- [Overview on page 3495](#)
- [Configuration on page 3496](#)
- [Verification on page 3515](#)

#### Requirements

This example uses the following hardware and software components:

- Eight BGP-enabled devices.
- Five of the BGP-enabled devices do not necessarily need to be routers. For example, they can be EX Series Ethernet Switches.
- Three of the BGP-enabled devices are configured to send multiple paths or receive multiple paths (or both send and receive multiple paths). These three BGP-enabled devices must be M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- The three routers must be running Junos OS Release 11.4 or later.

#### Overview

The following statements are used for configuring multiple paths to a destination:

```
[edit protocols bgp group group-name family family]
add-path {
  receive;
  send {
    path-count number;
    prefix-policy [ policy-names ];
  }
}
```

In this example, Router R5, Router R6, and Router R7 redistribute static routes into BGP. Router R1 and Router R4 are route reflectors. Router R2 and Router R3 are clients to Route Reflector R1. Router R8 is a client to Route Reflector R4.

Route reflection is optional when multiple-path advertisement is enabled in BGP.

With the **add-path send path-count 6** configuration, Router R1 is configured to send up to six paths (per destination) to Router R4.

With the **add-path receive** configuration, Router R4 is configured to receive multiple paths from Router R1.

With the **add-path send path-count 6** configuration, Router R4 is configured to send up to six paths to Router R8.

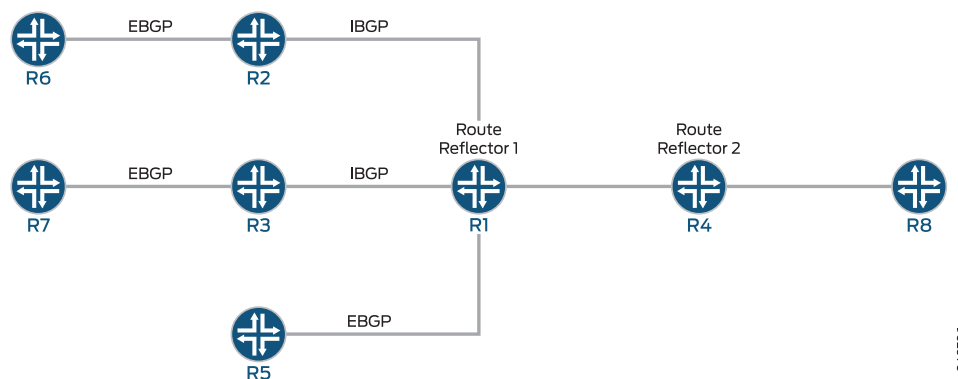
With the **add-path receive** configuration, Router R8 is configured to receive multiple paths from Router R4.

The **add-path send prefix-policy allow\_199** policy configuration (along with the corresponding route filter) limits Router R4 to sending multiple paths for only the 199.1.1.1/32 route.

### Topology Diagram

Figure 83 on page 3496 shows the topology used in this example.

Figure 83: Advertisement of Multiple Paths in BGP



### Configuration

- [Configuring Router R1 on page 3499](#)
- [Configuring Router R2 on page 3502](#)
- [Configuring Router R3 on page 3504](#)
- [Configuring Router R4 on page 3506](#)
- [Configuring Router R5 on page 3508](#)
- [Configuring Router R6 on page 3510](#)
- [Configuring Router R7 on page 3512](#)
- [Configuring Router R8 on page 3513](#)
- [Results on page 3514](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Router R1**

```
set interfaces fe-0/0/0 unit 12 family inet address 10.0.12.1/24
set interfaces fe-0/0/1 unit 13 family inet address 10.0.13.1/24
set interfaces fe-1/0/0 unit 14 family inet address 10.0.14.1/24
```

```

set interfaces fe-1/2/0 unit 15 family inet address 10.0.15.1/24
set interfaces lo0 unit 10 family inet address 10.0.0.10/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.10
set protocols bgp group rr cluster 10.0.0.10
set protocols bgp group rr neighbor 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.30
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.2 local-address 10.0.15.1
set protocols bgp group e1 neighbor 10.0.15.2 peer-as 2
set protocols bgp group rr_rr type internal
set protocols bgp group rr_rr local-address 10.0.0.10
set protocols bgp group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
  path-count 6
set protocols ospf area 0.0.0.0 interface lo0.10 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.12
set protocols ospf area 0.0.0.0 interface fe-0/0/1.13
set protocols ospf area 0.0.0.0 interface fe-1/0/0.14
set protocols ospf area 0.0.0.0 interface fe-1/2/0.15
set routing-options router-id 10.0.0.10
set routing-options autonomous-system 1

```

Router R2

```

set interfaces fe-1/2/0 unit 21 family inet address 10.0.12.2/24
set interfaces fe-1/2/1 unit 26 family inet address 10.0.26.1/24
set interfaces lo0 unit 20 family inet address 10.0.0.20/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.20 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.21
set protocols ospf area 0.0.0.0 interface fe-1/2/1.28
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1

```

Router R3

```

set interfaces fe-1/0/1 unit 31 family inet address 10.0.13.2/24
set interfaces fe-1/0/2 unit 37 family inet address 10.0.37.1/24
set interfaces lo0 unit 30 family inet address 10.0.0.30/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.30
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.30 passive
set protocols ospf area 0.0.0.0 interface fe-1/0/1.31
set protocols ospf area 0.0.0.0 interface fe-1/0/2.37
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1

```

Router R4

```

set interfaces fe-1/2/0 unit 41 family inet address 10.0.14.2/24
set interfaces fe-1/2/1 unit 48 family inet address 10.0.48.1/24
set interfaces lo0 unit 40 family inet address 10.0.0.40/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.40

```

```
set protocols bgp group rr family inet unicast add-path receive
set protocols bgp group rr neighbor 10.0.0.10
set protocols bgp group rr_client type internal
set protocols bgp group rr_client local-address 10.0.0.40
set protocols bgp group rr_client cluster 10.0.0.40
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
  path-count 6
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
  prefix-policy allow_199
set protocols ospf area 0.0.0.0 interface fe-1/2/0.41
set protocols ospf area 0.0.0.0 interface lo0.40 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.48
set routing-options autonomous-system 1
set policy-options policy-statement allow_199 from route-filter 199.1.1.1/32 exact
set policy-options policy-statement allow_199 term match_199 from prefix-list match_199
set policy-options policy-statement allow_199 then add-path send-count 20
set policy-options policy-statement allow_199 then accept
```

**Router R5**

```
set interfaces fe-1/2/0 unit 51 family inet address 10.0.15.2/24
set interfaces lo0 unit 50 family inet address 10.0.0.50/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.1 export s2b
set protocols bgp group e1 neighbor 10.0.15.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then as-path-expand 2
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject
```

**Router R6**

```
set interfaces fe-1/2/0 unit 62 family inet address 10.0.26.2/24
set interfaces lo0 unit 60 family inet address 10.0.0.60/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.1 export s2b
set protocols bgp group e1 neighbor 10.0.26.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject
```

**Router R7**

```
set interfaces fe-1/2/0 unit 73 family inet address 10.0.37.2/24
set interfaces lo0 unit 70 family inet address 10.0.0.70/32
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.1 export s2b
set protocols bgp group e1 neighbor 10.0.37.1 peer-as 1
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
```

**Router R8**

```
set interfaces fe-1/2/0 unit 84 family inet address 10.0.48.2/24
```

```

set interfaces lo0 unit 80 family inet address 10.0.0.80/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.80
set protocols bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive
set protocols ospf area 0.0.0.0 interface lo0.80 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.84
set routing-options autonomous-system 1

```

### Configuring Router R1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1:

1. Configure the interfaces to Router R2, Router R3, Router R4, and Router R5, and configure the loopback (lo0) interface.

```

[edit interfaces]
user@R1# set fe-0/0/0 unit 12 family inet address 10.0.12.1/24

user@R1# set fe-0/0/1 unit 13 family inet address 10.0.13.1/24

user@R1# set fe-1/0/0 unit 14 family inet address 10.0.14.1/24

user@R1# set fe-1/2/0 unit 15 family inet address 10.0.15.1/24

user@R1# set lo0 unit 10 family inet address 10.0.0.10/32

```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```

[edit protocols bgp]
user@R1# set group rr type internal
user@R1# set group rr local-address 10.0.0.10
user@R1# set group rr cluster 10.0.0.10
user@R1# set group rr neighbor 10.0.0.20
user@R1# set group rr neighbor 10.0.0.30

user@R1# set group rr_rr type internal
user@R1# set group rr_rr local-address 10.0.0.10

user@R1# set group e1 type external
user@R1# set group e1 neighbor 10.0.15.2 local-address 10.0.15.1
user@R1# set group e1 neighbor 10.0.15.2 peer-as 2

```

3. Configure Router R1 to send up to six paths to its neighbor, Router R4.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```

[edit protocols bgp]
user@R1# set group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
path-count 6

```

4. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@R1# set area 0.0.0.0 interface lo0.10 passive
user@R1# set area 0.0.0.0 interface fe-0/0/0.12
user@R1# set area 0.0.0.0 interface fe-0/0/1.13
user@R1# set area 0.0.0.0 interface fe-1/0/0.14
user@R1# set area 0.0.0.0 interface fe-1/2/0.15
```

5. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@R1# set router-id 10.0.0.10
user@R1# set autonomous-system 1
```

6. If you are done configuring the device, commit the configuration.

```
user@R1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-0/0/0 {
  unit 12 {
    family inet {
      address 10.0.12.1/24;
    }
  }
}
fe-0/0/1 {
  unit 13 {
    family inet {
      address 10.0.13.1/24;
    }
  }
}
fe-1/0/0 {
  unit 14 {
    family inet {
      address 10.0.14.1/24;
    }
  }
}
fe-1/2/0 {
  unit 15 {
    family inet {
      address 10.0.15.1/24;
    }
  }
}
lo0 {
  unit 10 {
    family inet {
```

```

        address 10.0.0.10/32;
    }
}
}
user@R1# show protocols
bgp {
    group rr {
        type internal;
        local-address 10.0.0.10;
        cluster 10.0.0.10;
        neighbor 10.0.0.20;
        neighbor 10.0.0.30;
    }
    group e1 {
        type external;
        neighbor 10.0.15.2 {
            local-address 10.0.15.1;
            peer-as 2;
        }
    }
    group rr_rr {
        type internal;
        local-address 10.0.0.10;
        neighbor 10.0.0.40 {
            family inet {
                unicast {
                    add-path {
                        send {
                            path-count 6;
                        }
                    }
                }
            }
        }
    }
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.10 {
            passive;
        }
        interface fe-0/0/0.12;
        interface fe-0/0/1.13;
        interface fe-1/0/0.14;
        interface fe-1/2/0.15;
    }
}
}
user@R1# show routing-options
router-id 10.0.0.10;
autonomous-system 1;

```

### *Configuring Router R2*

#### **Step-by-Step Procedure**

To configure Router R2:

1. Configure the loopback (lo0) interface and the interfaces to Router R6 and Router R1.

```
[edit interfaces]
```

```
user@R2# set fe-1/2/0 unit 21 family inet address 10.0.12.2/24
```

```
user@R2# set fe-1/2/1 unit 26 family inet address 10.0.26.1/24
```

```
user@R2# set lo0 unit 20 family inet address 10.0.0.20/32
```

2. Configure BGP and OSPF on Router R2's interfaces.

```
[edit protocols]
```

```
user@R2# set bgp group rr type internal
```

```
user@R2# set bgp group rr local-address 10.0.0.20
```

```
user@R2# set bgp group e1 type external
```

```
user@R2# set bgp group e1 neighbor 10.0.26.2 peer-as 2
```

```
user@R2# set ospf area 0.0.0.0 interface lo0.20 passive
```

```
user@R2# set ospf area 0.0.0.0 interface fe-1/2/0.21
```

```
user@R2# set ospf area 0.0.0.0 interface fe-1/2/1.28
```

3. For routes sent from Router R2 to Router R1, advertise Router R2 as the next hop, because Router R1 does not have a route to Router R6's address on the 10.0.26.0/24 network.

```
[edit]
```

```
user@R2# set policy-options policy-statement set_nh_self then next-hop self
```

```
user@R2# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
```

4. Configure the autonomous system number.

```
[edit]
```

```
user@R2# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 21 {
    family inet {
      address 10.0.12.2/24;
```



```

    }
  }
}
fe-1/2/1 {
  unit 26 {
    family inet {
      address 10.0.26.1/24;
    }
  }
}
lo0 {
  unit 20 {
    family inet {
      address 10.0.0.20/32;
    }
  }
}
}

user@R2# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.20;
    neighbor 10.0.0.10 {
      export set_nh_self;
    }
  }
  group e1 {
    type external;
    neighbor 10.0.26.2 {
      peer-as 2;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.20 {
      passive;
    }
    interface fe-1/2/0.21;
    interface fe-1/2/1.28;
  }
}

user@R2# show policy-options
policy-statement set_nh_self {
  then {
    next-hop self;
  }
}

user@R2# show routing-options
autonomous-system 1;

```

### Configuring Router R3

#### Step-by-Step Procedure

To configure Router R3:

1. Configure the loopback (lo0) interface and the interfaces to Router R7 and Router R1.

```
[edit interfaces]
```

```
user@R3# set fe-1/0/1 unit 31 family inet address 10.0.13.2/24
```

```
user@R3# set fe-1/0/2 unit 37 family inet address 10.0.37.1/24
```

```
user@R3# set lo0 unit 30 family inet address 10.0.0.30/32
```

2. Configure BGP and OSPF on Router R3's interfaces.

```
[edit protocols]
```

```
user@R3# set bgp group rr type internal
```

```
user@R3# set bgp group rr local-address 10.0.0.30
```

```
user@R3# set bgp group e1 type external
```

```
user@R3# set bgp group e1 neighbor 10.0.37.2 peer-as 2
```

```
user@R3# set ospf area 0.0.0.0 interface lo0.30 passive
```

```
user@R3# set ospf area 0.0.0.0 interface fe-1/0/1.31
```

```
user@R3# set ospf area 0.0.0.0 interface fe-1/0/2.37
```

3. For routes sent from Router R3 to Router R1, advertise Router R3 as the next hop, because Router R1 does not have a route to Router R7's address on the 10.0.37.0/24 network.

```
[edit]
```

```
user@R3# set policy-options policy-statement set_nh_self then next-hop self
```

```
user@R3# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
```

4. Configure the autonomous system number.

```
[edit]
```

```
user@R3# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R3# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/0/1 {
  unit 31 {
    family inet {
      address 10.0.13.2/24;
```

```

    }
  }
}
fe-1/0/2 {
  unit 37 {
    family inet {
      address 10.0.37.1/24;
    }
  }
}
lo0 {
  unit 30 {
    family inet {
      address 10.0.0.30/32;
    }
  }
}
}

user@R3# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.30;
    neighbor 10.0.0.10 {
      export set_nh_self;
    }
  }
  group e1 {
    type external;
    neighbor 10.0.37.2 {
      peer-as 2;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.30 {
      passive;
    }
    interface fe-1/0/1.31;
    interface fe-1/0/2.37;
  }
}

user@R3# show policy-options
policy-statement set_nh_self {
  then {
    next-hop self;
  }
}

user@R3# show routing-options
autonomous-system 1;

```

### ***Configuring Router R4***

#### **Step-by-Step Procedure**

To configure Router R4:

1. Configure the interfaces to Router R1 and Router R8, and configure the loopback (lo0) interface.

```
[edit interfaces]
```

```
user@R4# set fe-1/2/0 unit 41 family inet address 10.0.14.2/24
```

```
user@R4# set fe-1/2/1 unit 48 family inet address 10.0.48.1/24
```

```
user@R4# set lo0 unit 40 family inet address 10.0.0.40/32
```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
```

```
user@R4# set group rr type internal
```

```
user@R4# set group rr local-address 10.0.0.40
```

```
user@R4# set group rr neighbor 10.0.0.10
```

```
user@R4# set group rr_client type internal
```

```
user@R4# set group rr_client local-address 10.0.0.40
```

```
user@R4# set group rr_client cluster 10.0.0.40
```

3. Configure Router R4 to send up to six paths to its neighbor, Router R8.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

```
[edit protocols bgp]
```

```
user@R4# set group rr_client neighbor 10.0.0.80 family inet unicast add-path send  
path-count 6
```

4. Configure Router R4 to receive multiple paths from its neighbor, Router R1.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```
[edit protocols bgp group rr family inet unicast]
```

```
user@R4# set add-path receive
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@R4# set interface fe-1/2/0.41
```

```
user@R4# set interface lo0.40 passive
```

```
user@R4# set interface fe-1/2/1.48
```

6. Configure a policy that allows Router R4 to send Router R8 multiple paths to the 199.1.1.1/32 route.

- Router R4 receives multiple paths for the 198.1.1.1/32 route and the 199.1.1.1/32 route. However, because of this policy, Router R4 only sends multiple paths for the 199.1.1.1/32 route.

```
[edit protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast]
```

```
user@R4# set add-path send prefix-policy allow_199
```

```
[edit policy-options policy-statement allow_199]
user@R4# set from route-filter 199.1.1/32 exact
user@R4# set then accept
```

- Router R4 can also be configured to send up-to 20 BGP **add-path** routes for a subset of *add-path advertised prefixes*.

```
[edit policy-options policy-statement allow_199]
user@R4# set term match_199 from prefix-list match_199
user@R4# set then add-path send-count 20
```

7. Configure the autonomous system number.

```
[edit routing-options]
user@R4# set autonomous-system 1
```

8. If you are done configuring the device, commit the configuration.

```
user@R4# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 41 {
    family inet {
      address 10.0.14.2/24;
    }
  }
}
fe-1/2/1 {
  unit 48 {
    family inet {
      address 10.0.48.1/24;
    }
  }
}
lo0 {
  unit 40 {
    family inet {
      address 10.0.0.40/32;
    }
  }
}
```

```
user@R4# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.40;
    family inet {
      unicast {
        add-path {
          receive;
        }
      }
    }
  }
}
```

```

    }
}
neighbor 10.0.0.10;
}
group rr_client {
    type internal;
    local-address 10.0.0.40;
    cluster 10.0.0.40;
    neighbor 10.0.0.80 {
        family inet {
            unicast {
                add-path {
                    send {
                        path-count 6;
                        prefix-policy allow_199;
                    }
                }
            }
        }
    }
}
}
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.40 {
            passive;
        }
        interface fe-1/2/0.41;
        interface fe-1/2/1.48;
    }
}
}
}
user@R4# show policy-options
policy-statement allow_199 {
    from {
        route-filter 199.1.1.1/32 exact;
    }
    from term match_199 {
        prefix-list match_199;
    }
    then add-path send-count 20;
    then accept;
}
}
user@R4# show routing-options
autonomous-system 1;
```

### Configuring Router R5

## Step-by-Step Procedure

To configure Router R5:

1. Configure the loopback (lo0) interface and the interface to Router R1.

[edit interfaces]

```
user@R5# set fe-1/2/0 unit 51 family inet address 10.0.15.2/24
```

```
user@R5# set lo0 unit 50 family inet address 10.0.0.50/32
```

2. Configure BGP on Router R5's interface.

```
[edit protocols bgp group e1]
user@R5# set type external
user@R5# set neighbor 10.0.15.1 peer-as 1
```

3. Create static routes for redistribution into BGP.

```
[edit routing-options]
user@R5# set static route 199.1.1.1/32 reject
user@R5# set static route 198.1.1.1/32 reject
```

4. Redistribute static and direct routes into BGP.

```
[edit protocols bgp group e1 neighbor 10.0.15.1]
user@R5# set export s2b
```

```
[edit policy-options policy-statement s2b]
user@R5# set from protocol static
user@R5# set from protocol direct
user@R5# set then as-path-expand 2
user@R5# set then accept
```

5. Configure the autonomous system number.

```
[edit routing-options]
user@R5# set autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R5# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R5# show interfaces
fe-1/2/0 {
  unit 51 {
    family inet {
      address 10.0.15.2/24;
    }
  }
}
lo0 {
  unit 50 {
    family inet {
      address 10.0.0.50/32;
    }
  }
}

user@R5# show protocols
bgp {
  group e1 {
```

```
type external;
neighbor 10.0.15.1 {
    export s2b;
    peer-as 1;
}
}
}

user@R5# show policy-options
policy-statement s2b {
    from protocol [ static direct ];
    then {
        as-path-expand 2;
        accept;
    }
}

user@R5# show routing-options
static {
    route 198.1.1.1/32 reject;
    route 199.1.1.1/32 reject;
}
autonomous-system 2;
```

### ***Configuring Router R6***

#### **Step-by-Step Procedure**

To configure Router R6:

1. Configure the loopback (lo0) interface and the interface to Router R2.  
  
[edit interfaces]  
user@R6# set fe-1/2/0 unit 62 family inet address 10.0.26.2/24  
  
user@R6# set lo0 unit 60 family inet address 10.0.0.60/32
2. Configure BGP on Router R6's interface.  
  
[edit protocols]  
user@R6# set bgp group e1 type external  
user@R6# set bgp group e1 neighbor 10.0.26.1 peer-as 1
3. Create static routes for redistribution into BGP.  
  
[edit]  
user@R6# set routing-options static route 199.1.1.1/32 reject  
user@R6# set routing-options static route 198.1.1.1/32 reject
4. Redistribute static and direct routes from Router R6's routing table into BGP.  
  
[edit protocols bgp group e1 neighbor 10.0.26.1]  
user@R6# set export s2b  
  
[edit policy-options policy-statement s2b]  
user@R6# set from protocol static  
user@R6# set from protocol direct  
user@R6# set then accept



5. Configure the autonomous system number.

```
[edit routing-options]
user@R6# set autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R6# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R6# show interfaces
fe-1/2/0 {
  unit 62 {
    family inet {
      address 10.0.26.2/24;
    }
  }
}
lo0 {
  unit 60 {
    family inet {
      address 10.0.0.60/32;
    }
  }
}

user@R6# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.26.1 {
      export s2b;
      peer-as 1;
    }
  }
}

user@R6# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then accept;
}

user@R6# show routing-options
static {
  route 198.1.1.1/32 reject;
  route 199.1.1.1/32 reject;
}
autonomous-system 2;
```

### *Configuring Router R7*

#### **Step-by-Step Procedure**

To configure Router R7:

1. Configure the loopback (lo0) interface and the interface to Router R3.  
  
[edit interfaces]  
user@R7# set fe-1/2/0 unit 73 family inet address 10.0.37.2/24  
  
user@R7# set lo0 unit 70 family inet address 10.0.0.70/32
2. Configure BGP on Router R7's interface.  
  
[edit protocols bgp group e1]  
user@R7# set type external  
user@R7# set neighbor 10.0.37.1 peer-as 1
3. Create a static route for redistribution into BGP.  
  
[edit]  
user@R7# set routing-options static route 199.1.1.1/32 reject
4. Redistribute static and direct routes from Router R7's routing table into BGP.  
  
[edit protocols bgp group e1 neighbor 10.0.37.1]  
user@R7# set export s2b  
  
[edit policy-options policy-statement s2b]  
user@R7# set from protocol static  
user@R7# set from protocol direct  
user@R7# set then accept
5. Configure the autonomous system number.  
  
[edit routing-options]  
user@R7# set autonomous-system 2
6. If you are done configuring the device, commit the configuration.  
  
user@R7# commit

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R7# show interfaces
fe-1/2/0 {
  unit 73 {
    family inet {
      address 10.0.37.2/24;
    }
  }
}
lo0 {
  unit 70 {
    family inet {
```

```

        address 10.0.0.70/32;
    }
}
}

user@R7# show protocols
bgp {
    group e1 {
        type external;
        neighbor 10.0.37.1 {
            export s2b;
            peer-as 1;
        }
    }
}

user@R7# show policy-options
policy-statement s2b {
    from protocol [ static direct ];
    then accept;
}

user@R7# show routing-options
static {
    route 199.1.1.1/32 reject;
}
autonomous-system 2;

```

### Configuring Router R8

#### Step-by-Step Procedure

To configure Router R8:

1. Configure the loopback (lo0) interface and the interface to Router R4.

```

[edit interfaces]
user@R8# set fe-1/2/0 unit 84 family inet address 10.0.48.2/24

user@R8# set lo0 unit 80 family inet address 10.0.0.80/32

```

2. Configure BGP and OSPF on Router R8's interface.

```

[edit protocols]
user@R8# set bgp group rr type internal
user@R8# set bgp group rr local-address 10.0.0.80

user@R8# set ospf area 0.0.0.0 interface lo0.80 passive
user@R8# set ospf area 0.0.0.0 interface fe-1/2/0.84

```

3. Configure Router R8 to receive multiple paths from its neighbor, Router R4.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

```

[edit protocols]
user@R8# set bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive

```

4. Configure the autonomous system number.

```

[edit]

```

```
user@R8# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R8# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R8# show interfaces
fe-1/2/0 {
  unit 84 {
    family inet {
      address 10.0.48.2/24;
    }
  }
}
lo0 {
  unit 80 {
    family inet {
      address 10.0.0.80/32;
    }
  }
}

user@R8# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.80;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            receive;
          }
        }
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.80 {
      passive;
    }
    interface fe-1/2/0.84;
  }
}

user@R8# show routing-options
autonomous-system 1;
```

**Verification**

Confirm that the configuration is working properly.

- [Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths on page 3515](#)
- [Verifying That Router R1 Is Advertising Multiple Paths on page 3515](#)
- [Verifying That Router R4 Is Receiving and Advertising Multiple Paths on page 3516](#)
- [Verifying That Router R8 Is Receiving Multiple Paths on page 3517](#)
- [Checking the Path ID on page 3517](#)

**Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths**

**Purpose** Make sure that one or both of the following strings appear in the output of the **show bgp neighbor** command:

- NLRI's for which peer can receive multiple paths: inet-unicast
- NLRI's for which peer can send multiple paths: inet-unicast

**Action**

```

user@R1> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1      Local: 10.0.0.10+65237 AS 1
  Type: Internal  State: Established  Flags: <Sync>
... NLRI's for which peer can receive multiple paths: inet-unicast
...

user@R4> show bgp neighbor 10.0.0.10
Peer: 10.0.0.10+65237 AS 1      Local: 10.0.0.40+179 AS 1
  Type: Internal  State: Established  Flags: <Sync>
...
  NLRI's for which peer can send multiple paths: inet-unicast
...

user@R4> show bgp neighbor 10.0.0.80
Peer: 10.0.0.80+55416 AS 1      Local: 10.0.0.40+179 AS 1
  Type: Internal  State: Established (route reflector client)Flags: <Sync>
...
  NLRI's for which peer can receive multiple paths: inet-unicast
...

user@R8> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1      Local: 10.0.0.80+55416 AS 1
  Type: Internal  State: Established  Flags: <Sync>
...
  NLRI's for which peer can send multiple paths: inet-unicast
...

```

**Verifying That Router R1 Is Advertising Multiple Paths**

**Purpose** Make sure that multiple paths to the 198.1.1/32 destination and multiple paths to the 199.1.1/32 destination are advertised to Router R4.

**Action** user@R1> show route advertising-protocol bgp 10.0.0.40  
 inet.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

**Meaning** When you see one prefix and more than one next hop, it means that multiple paths are advertised to Router R4.

#### *Verifying That Router R4 Is Receiving and Advertising Multiple Paths*

**Purpose** Make sure that multiple paths to the 199.1.1.1/32 destination are received from Router R1 and advertised to Router R8. Make sure that multiple paths to the 198.1.1.1/32 destination are received from Router R1, but only one path to this destination is advertised to Router R8.

**Action** user@R4> show route receive-protocol bgp 10.0.0.10  
 inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

user@R4> show route advertising-protocol bgp 10.0.0.80  
 inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

**Meaning** The **show route receive-protocol** command shows that Router R4 receives two paths to the 198.1.1.1/32 destination and three paths to the 199.1.1.1/32 destination. The **show route advertising-protocol** command shows that Router R4 advertises only one path to the 198.1.1.1/32 destination and advertises all three paths to the 199.1.1.1/32 destination.

Because of the prefix policy that is applied to Router R4, Router R4 does not advertise multiple paths to the 198.1.1.1/32 destination. Router R4 advertises only one path to the 198.1.1.1/32 destination even though it receives multiple paths to this destination.

### *Verifying That Router R8 Is Receiving Multiple Paths*

**Purpose** Make sure that Router R8 receives multiple paths to the 199.1.1.1/32 destination through Router R4. Make sure that Router R8 receives only one path to the 198.1.1.1/32 destination through Router R4.

**Action** user@R8> show route receive-protocol bgp 10.0.0.40  
 inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

### *Checking the Path ID*

**Purpose** On the downstream devices, Router R4 and Router R8, verify that a path ID uniquely identifies the path. Look for the **Addpath Path ID:** string.

**Action** user@R4> show route 199.1.1.1/32 detail

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 3 announced)
  *BGP    Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 9
          Source: 10.0.0.10
          Next hop type: Router, Next hop index: 676
          Next hop: 10.0.14.1 via lt-1/2/0.41, selected
          Protocol next hop: 10.0.0.20
          Indirect next hop: 92041c8 262146
          State: <Active Int Ext>
          Local AS:      1 Peer AS:      1
          Age: 1:44:37    Metric2: 2
          Task: BGP_1.10.0.0.10+65237
          Announcement bits (3): 2-KRT 3-BGP RT Background 4-Resolve tree

1
  AS path: 2 I (Originator) Cluster list: 10.0.0.10
  AS path: Originator ID: 10.0.0.20
  Accepted
  Localpref: 100
  Router ID: 10.0.0.10
  Addpath Path ID: 1
  BGP    Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 4
          Source: 10.0.0.10
          Next hop type: Router, Next hop index: 676
          Next hop: 10.0.14.1 via lt-1/2/0.41, selected
          Protocol next hop: 10.0.0.30
          Indirect next hop: 92042ac 262151
          State: <NotBest Int Ext>
          Inactive reason: Not Best in its group - Router ID
          Local AS:      1 Peer AS:      1
          Age: 1:44:37    Metric2: 2
          Task: BGP_1.10.0.0.10+65237
          Announcement bits (1): 3-BGP RT Background
          AS path: 2 I (Originator) Cluster list: 10.0.0.10
          AS path: Originator ID: 10.0.0.30
          Accepted
          Localpref: 100
          Router ID: 10.0.0.10
          Addpath Path ID: 2
  BGP    Preference: 170/-101
          Next hop type: Indirect
          Next-hop reference count: 4
          Source: 10.0.0.10
          Next hop type: Router, Next hop index: 676
          Next hop: 10.0.14.1 via lt-1/2/0.41, selected
          Protocol next hop: 10.0.15.2
          Indirect next hop: 92040e4 262150
          State: <Int Ext>
          Inactive reason: AS path
          Local AS:      1 Peer AS:      1
          Age: 1:44:37    Metric2: 2
          Task: BGP_1.10.0.0.10+65237
          Announcement bits (1): 3-BGP RT Background
          AS path: 2 2 I
          Accepted
```



```

Localpref: 100
Router ID: 10.0.0.10
Addpath Path ID: 3

```

```
user@R8> show route 199.1.1.1/32 detail
```

```

inet.0: 17 destinations, 19 routes (17 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 1 announced)
*BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 9
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.20
      Indirect next hop: 91fc0e4 262148
      State: <Active Int Ext>
      Local AS:      1 Peer AS:      1
      Age: 1:56:51    Metric2: 3
      Task: BGP_1.10.0.0.40+179
      Announcement bits (2): 2-KRT 4-Resolve tree 1
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.20
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 1
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.30
      Indirect next hop: 91fc1c8 262152
      State: <NotBest Int Ext>
      Inactive reason: Not Best in its group - Router ID
      Local AS:      1 Peer AS:      1
      Age: 1:56:51    Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.30
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 2
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.15.2
      Indirect next hop: 91fc2ac 262153
      State: <Int Ext>
      Inactive reason: AS path
      Local AS:      1 Peer AS:      1
      Age: 1:56:51    Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 2 I (Originator) Cluster list: 10.0.0.40
      AS path: Originator ID: 10.0.0.10

```

Accepted  
Localpref: 100  
Router ID: 10.0.0.40  
Addpath Path ID: 3

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3261](#)
  - [BGP Configuration Overview](#)

## Example: Advertising Multiple Paths in BGP

In this example, BGP routers are configured to advertise multiple paths instead of advertising only the active path. Advertising multiple paths in BGP is specified in Internet draft draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*.

- [Requirements on page 3520](#)
- [Overview on page 3520](#)
- [Configuration on page 3521](#)
- [Verification on page 3540](#)

### Requirements

---

This example uses the following hardware and software components:

- Eight BGP-enabled devices.
- Five of the BGP-enabled devices do not necessarily need to be routers. For example, they can be EX Series Ethernet Switches.
- Three of the BGP-enabled devices are configured to send multiple paths or receive multiple paths (or both send and receive multiple paths). These three BGP-enabled devices must be M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, or T Series Core Routers.
- The three routers must be running Junos OS Release 11.4 or later.

### Overview

---

The following statements are used for configuring multiple paths to a destination:

```
[edit protocols bgp group group-name family family]  
add-path {  
  receive;  
  send {  
    path-count number;  
    prefix-policy [ policy-names ];  
  }  
}
```

In this example, Router R5, Router R6, and Router R7 redistribute static routes into BGP. Router R1 and Router R4 are route reflectors. Router R2 and Router R3 are clients to Route Reflector R1. Router R8 is a client to Route Reflector R4.

Route reflection is optional when multiple-path advertisement is enabled in BGP.

With the **add-path send path-count 6** configuration, Router R1 is configured to send up to six paths (per destination) to Router R4.

With the **add-path receive** configuration, Router R4 is configured to receive multiple paths from Router R1.

With the **add-path send path-count 6** configuration, Router R4 is configured to send up to six paths to Router R8.


With the **add-path receive** configuration, Router R8 is configured to receive multiple paths from Router R4.

The **add-path send prefix-policy allow\_199** policy configuration (along with the corresponding route filter) limits Router R4 to sending multiple paths for only the 199.1.1.1/32 route.

### Topology Diagram

Figure 83 on page 3496 shows the topology used in this example.

Figure 84: Advertisement of Multiple Paths in BGP



### Configuration

- [Configuring Router R1 on page 3524](#)
- [Configuring Router R2 on page 3527](#)
- [Configuring Router R3 on page 3529](#)
- [Configuring Router R4 on page 3531](#)
- [Configuring Router R5 on page 3533](#)
- [Configuring Router R6 on page 3535](#)
- [Configuring Router R7 on page 3537](#)
- [Configuring Router R8 on page 3538](#)
- [Results on page 3539](#)

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Router R1**

```
set interfaces fe-0/0/0 unit 12 family inet address 10.0.12.1/24
set interfaces fe-0/0/1 unit 13 family inet address 10.0.13.1/24
set interfaces fe-1/0/0 unit 14 family inet address 10.0.14.1/24
set interfaces fe-1/2/0 unit 15 family inet address 10.0.15.1/24
set interfaces lo0 unit 10 family inet address 10.0.0.10/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.10
set protocols bgp group rr cluster 10.0.0.10
set protocols bgp group rr neighbor 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.30
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.2 local-address 10.0.15.1
set protocols bgp group e1 neighbor 10.0.15.2 peer-as 2
set protocols bgp group rr_rr type internal
set protocols bgp group rr_rr local-address 10.0.0.10
set protocols bgp group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
  path-count 6
set protocols ospf area 0.0.0.0 interface lo0.10 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.12
set protocols ospf area 0.0.0.0 interface fe-0/0/1.13
set protocols ospf area 0.0.0.0 interface fe-1/0/0.14
set protocols ospf area 0.0.0.0 interface fe-1/2/0.15
set routing-options router-id 10.0.0.10
set routing-options autonomous-system 1
```

**Router R2**

```
set interfaces fe-1/2/0 unit 21 family inet address 10.0.12.2/24
set interfaces fe-1/2/1 unit 26 family inet address 10.0.26.1/24
set interfaces lo0 unit 20 family inet address 10.0.0.20/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.20
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.20 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.21
set protocols ospf area 0.0.0.0 interface fe-1/2/1.28
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1
```

**Router R3**

```
set interfaces fe-1/0/1 unit 31 family inet address 10.0.13.2/24
set interfaces fe-1/0/2 unit 37 family inet address 10.0.37.1/24
set interfaces lo0 unit 30 family inet address 10.0.0.30/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.30
set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.2 peer-as 2
set protocols ospf area 0.0.0.0 interface lo0.30 passive
set protocols ospf area 0.0.0.0 interface fe-1/0/1.31
set protocols ospf area 0.0.0.0 interface fe-1/0/2.37
set policy-options policy-statement set_nh_self then next-hop self
set routing-options autonomous-system 1
```

**Router R4**

```
set interfaces fe-1/2/0 unit 41 family inet address 10.0.14.2/24
set interfaces fe-1/2/1 unit 48 family inet address 10.0.48.1/24
```

```

set interfaces lo0 unit 40 family inet address 10.0.0.40/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.40
set protocols bgp group rr family inet unicast add-path receive
set protocols bgp group rr neighbor 10.0.0.10
set protocols bgp group rr_client type internal
set protocols bgp group rr_client local-address 10.0.0.40
set protocols bgp group rr_client cluster 10.0.0.40
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
  path-count 6
set protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast add-path send
  prefix-policy allow_199
set protocols ospf area 0.0.0.0 interface fe-1/2/0.41
set protocols ospf area 0.0.0.0 interface lo0.40 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.48
set routing-options autonomous-system 1
set policy-options policy-statement allow_199 from route-filter 199.1.1.1/32 exact
set policy-options policy-statement allow_199 term match_199 from prefix-list match_199
set policy-options policy-statement allow_199 then add-path send-count 20
set policy-options policy-statement allow_199 then accept

```

**Router R5**

```

set interfaces fe-1/2/0 unit 51 family inet address 10.0.15.2/24
set interfaces lo0 unit 50 family inet address 10.0.0.50/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.15.1 export s2b
set protocols bgp group e1 neighbor 10.0.15.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then as-path-expand 2
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject

```

**Router R6**

```

set interfaces fe-1/2/0 unit 62 family inet address 10.0.26.2/24
set interfaces lo0 unit 60 family inet address 10.0.0.60/32
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.26.1 export s2b
set protocols bgp group e1 neighbor 10.0.26.1 peer-as 1
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set routing-options autonomous-system 2
set routing-options static route 199.1.1.1/32 reject
set routing-options static route 198.1.1.1/32 reject

```

**Router R7**

```

set interfaces fe-1/2/0 unit 73 family inet address 10.0.37.2/24
set interfaces lo0 unit 70 family inet address 10.0.0.70/32
set policy-options policy-statement s2b from protocol static
set policy-options policy-statement s2b from protocol direct
set policy-options policy-statement s2b then accept
set protocols bgp group e1 type external
set protocols bgp group e1 neighbor 10.0.37.1 export s2b
set protocols bgp group e1 neighbor 10.0.37.1 peer-as 1
set routing-options autonomous-system 2

```

```
set routing-options static route 199.1.1.1/32 reject
```

**Router R8**

```
set interfaces fe-1/2/0 unit 84 family inet address 10.0.48.2/24
set interfaces lo0 unit 80 family inet address 10.0.0.80/32
set protocols bgp group rr type internal
set protocols bgp group rr local-address 10.0.0.80
set protocols bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive
set protocols ospf area 0.0.0.0 interface lo0.80 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/0.84
set routing-options autonomous-system 1
```

### *Configuring Router R1*

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1:

1. Configure the interfaces to Router R2, Router R3, Router R4, and Router R5, and configure the loopback (lo0) interface.

```
[edit interfaces]
```

```
user@R1# set fe-0/0/0 unit 12 family inet address 10.0.12.1/24
```

```
user@R1# set fe-0/0/1 unit 13 family inet address 10.0.13.1/24
```

```
user@R1# set fe-1/0/0 unit 14 family inet address 10.0.14.1/24
```

```
user@R1# set fe-1/2/0 unit 15 family inet address 10.0.15.1/24
```

```
user@R1# set lo0 unit 10 family inet address 10.0.0.10/32
```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
```

```
user@R1# set group rr type internal
```

```
user@R1# set group rr local-address 10.0.0.10
```

```
user@R1# set group rr cluster 10.0.0.10
```

```
user@R1# set group rr neighbor 10.0.0.20
```

```
user@R1# set group rr neighbor 10.0.0.30
```

```
user@R1# set group rr_rr type internal
```

```
user@R1# set group rr_rr local-address 10.0.0.10
```

```
user@R1# set group e1 type external
```

```
user@R1# set group e1 neighbor 10.0.15.2 local-address 10.0.15.1
```

```
user@R1# set group e1 neighbor 10.0.15.2 peer-as 2
```

3. Configure Router R1 to send up to six paths to its neighbor, Router R4.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```
[edit protocols bgp]
```

```
user@R1# set group rr_rr neighbor 10.0.0.40 family inet unicast add-path send
path-count 6
```

4. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@R1# set area 0.0.0.0 interface lo0.10 passive
user@R1# set area 0.0.0.0 interface fe-0/0/0.12
user@R1# set area 0.0.0.0 interface fe-0/0/1.13
user@R1# set area 0.0.0.0 interface fe-1/0/0.14
user@R1# set area 0.0.0.0 interface fe-1/2/0.15
```

5. Configure the router ID and the autonomous system number.

```
[edit routing-options]
user@R1# set router-id 10.0.0.10
user@R1# set autonomous-system 1
```

6. If you are done configuring the device, commit the configuration.

```
user@R1# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-0/0/0 {
  unit 12 {
    family inet {
      address 10.0.12.1/24;
    }
  }
}
fe-0/0/1 {
  unit 13 {
    family inet {
      address 10.0.13.1/24;
    }
  }
}
fe-1/0/0 {
  unit 14 {
    family inet {
      address 10.0.14.1/24;
    }
  }
}
fe-1/2/0 {
  unit 15 {
    family inet {
      address 10.0.15.1/24;
    }
  }
}
```

```
lo0 {
  unit 10 {
    family inet {
      address 10.0.0.10/32;
    }
  }
}

user@R1# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.10;
    cluster 10.0.0.10;
    neighbor 10.0.0.20;
    neighbor 10.0.0.30;
  }
  group e1 {
    type external;
    neighbor 10.0.15.2 {
      local-address 10.0.15.1;
      peer-as 2;
    }
  }
  group rr_rr {
    type internal;
    local-address 10.0.0.10;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            send {
              path-count 6;
            }
          }
        }
      }
    }
  }
}

ospf {
  area 0.0.0.0 {
    interface lo0.10 {
      passive;
    }
    interface fe-0/0/0.12;
    interface fe-0/0/1.13;
    interface fe-1/0/0.14;
    interface fe-1/2/0.15;
  }
}

user@R1# show routing-options
router-id 10.0.0.10;
autonomous-system 1;
```



**Configuring Router R2****Step-by-Step Procedure**

To configure Router R2:

1. Configure the loopback (lo0) interface and the interfaces to Router R6 and Router R1.

```
[edit interfaces]
```

```
user@R2# set fe-1/2/0 unit 21 family inet address 10.0.12.2/24
```

```
user@R2# set fe-1/2/1 unit 26 family inet address 10.0.26.1/24
```

```
user@R2# set lo0 unit 20 family inet address 10.0.0.20/32
```

2. Configure BGP and OSPF on Router R2's interfaces.

```
[edit protocols]
```

```
user@R2# set bgp group rr type internal
```

```
user@R2# set bgp group rr local-address 10.0.0.20
```

```
user@R2# set bgp group e1 type external
```

```
user@R2# set bgp group e1 neighbor 10.0.26.2 peer-as 2
```

```
user@R2# set ospf area 0.0.0.0 interface lo0.20 passive
```

```
user@R2# set ospf area 0.0.0.0 interface fe-1/2/0.21
```

```
user@R2# set ospf area 0.0.0.0 interface fe-1/2/1.28
```

3. For routes sent from Router R2 to Router R1, advertise Router R2 as the next hop, because Router R1 does not have a route to Router R6's address on the 10.0.26.0/24 network.

```
[edit]
```

```
user@R2# set policy-options policy-statement set_nh_self then next-hop self
```

```
user@R2# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
```

4. Configure the autonomous system number.

```
[edit]
```

```
user@R2# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R2# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 21 {
    family inet {
      address 10.0.12.2/24;
```

```
    }
  }
}
fe-1/2/1 {
  unit 26 {
    family inet {
      address 10.0.26.1/24;
    }
  }
}
lo0 {
  unit 20 {
    family inet {
      address 10.0.0.20/32;
    }
  }
}

user@R2# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.20;
    neighbor 10.0.0.10 {
      export set_nh_self;
    }
  }
  group e1 {
    type external;
    neighbor 10.0.26.2 {
      peer-as 2;
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.20 {
      passive;
    }
    interface fe-1/2/0.21;
    interface fe-1/2/1.28;
  }
}

user@R2# show policy-options
policy-statement set_nh_self {
  then {
    next-hop self;
  }
}

user@R2# show routing-options
autonomous-system 1;
```

**Configuring Router R3****Step-by-Step Procedure**

To configure Router R3:

1. Configure the loopback (lo0) interface and the interfaces to Router R7 and Router R1.

```
[edit interfaces]
```

```
user@R3# set fe-1/0/1 unit 31 family inet address 10.0.13.2/24
```

```
user@R3# set fe-1/0/2 unit 37 family inet address 10.0.37.1/24
```

```
user@R3# set lo0 unit 30 family inet address 10.0.0.30/32
```

2. Configure BGP and OSPF on Router R3's interfaces.

```
[edit protocols]
```

```
user@R3# set bgp group rr type internal
```

```
user@R3# set bgp group rr local-address 10.0.0.30
```

```
user@R3# set bgp group e1 type external
```

```
user@R3# set bgp group e1 neighbor 10.0.37.2 peer-as 2
```

```
user@R3# set ospf area 0.0.0.0 interface lo0.30 passive
```

```
user@R3# set ospf area 0.0.0.0 interface fe-1/0/1.31
```

```
user@R3# set ospf area 0.0.0.0 interface fe-1/0/2.37
```

3. For routes sent from Router R3 to Router R1, advertise Router R3 as the next hop, because Router R1 does not have a route to Router R7's address on the 10.0.37.0/24 network.

```
[edit]
```

```
user@R3# set policy-options policy-statement set_nh_self then next-hop self
```

```
user@R3# set protocols bgp group rr neighbor 10.0.0.10 export set_nh_self
```

4. Configure the autonomous system number.

```
[edit]
```

```
user@R3# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R3# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show interfaces
fe-1/0/1 {
  unit 31 {
    family inet {
      address 10.0.13.2/24;
```

```
    }  
  }  
}  
fe-1/0/2 {  
  unit 37 {  
    family inet {  
      address 10.0.37.1/24;  
    }  
  }  
}  
lo0 {  
  unit 30 {  
    family inet {  
      address 10.0.0.30/32;  
    }  
  }  
}  
}  
  
user@R3# show protocols  
bgp {  
  group rr {  
    type internal;  
    local-address 10.0.0.30;  
    neighbor 10.0.0.10 {  
      export set_nh_self;  
    }  
  }  
  group e1 {  
    type external;  
    neighbor 10.0.37.2 {  
      peer-as 2;  
    }  
  }  
}  
ospf {  
  area 0.0.0.0 {  
    interface lo0.30 {  
      passive;  
    }  
    interface fe-1/0/1.31;  
    interface fe-1/0/2.37;  
  }  
}  
  
user@R3# show policy-options  
policy-statement set_nh_self {  
  then {  
    next-hop self;  
  }  
}  
  
user@R3# show routing-options  
autonomous-system 1;
```

**Configuring Router R4****Step-by-Step Procedure**

To configure Router R4:

1. Configure the interfaces to Router R1 and Router R8, and configure the loopback (lo0) interface.

```
[edit interfaces]
```

```
user@R4# set fe-1/2/0 unit 41 family inet address 10.0.14.2/24
```

```
user@R4# set fe-1/2/1 unit 48 family inet address 10.0.48.1/24
```

```
user@R4# set lo0 unit 40 family inet address 10.0.0.40/32
```

2. Configure BGP on the interfaces, and configure IBGP route reflection.

```
[edit protocols bgp]
```

```
user@R4# set group rr type internal
```

```
user@R4# set group rr local-address 10.0.0.40
```

```
user@R4# set group rr neighbor 10.0.0.10
```

```
user@R4# set group rr_client type internal
```

```
user@R4# set group rr_client local-address 10.0.0.40
```

```
user@R4# set group rr_client cluster 10.0.0.40
```

3. Configure Router R4 to send up to six paths to its neighbor, Router R8.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

```
[edit protocols bgp]
```

```
user@R4# set group rr_client neighbor 10.0.0.80 family inet unicast add-path send path-count 6
```

4. Configure Router R4 to receive multiple paths from its neighbor, Router R1.

The destination of the paths can be any destination that Router R1 can reach through multiple paths.

```
[edit protocols bgp group rr family inet unicast]
```

```
user@R4# set add-path receive
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf area 0.0.0.0]
```

```
user@R4# set interface fe-1/2/0.41
```

```
user@R4# set interface lo0.40 passive
```

```
user@R4# set interface fe-1/2/1.48
```

6. Configure a policy that allows Router R4 to send Router R8 multiple paths to the 199.1.1.1/32 route.

- Router R4 receives multiple paths for the 198.1.1.1/32 route and the 199.1.1.1/32 route. However, because of this policy, Router R4 only sends multiple paths for the 199.1.1.1/32 route.

```
[edit protocols bgp group rr_client neighbor 10.0.0.80 family inet unicast]
```

```
user@R4# set add-path send prefix-policy allow_199
```

```
[edit policy-options policy-statement allow_199]
user@R4# set from route-filter 199.1.1/32 exact
user@R4# set then accept
```

- Router R4 can also be configured to send up-to 20 BGP **add-path** routes for a subset of *add-path advertised prefixes*.

```
[edit policy-options policy-statement allow_199]
user@R4# set term match_199 from prefix-list match_199
user@R4# set then add-path send-count 20
```

7. Configure the autonomous system number.

```
[edit routing-options]
user@R4# set autonomous-system 1
```

8. If you are done configuring the device, commit the configuration.

```
user@R4# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R4# show interfaces
fe-1/2/0 {
  unit 41 {
    family inet {
      address 10.0.14.2/24;
    }
  }
}
fe-1/2/1 {
  unit 48 {
    family inet {
      address 10.0.48.1/24;
    }
  }
}
lo0 {
  unit 40 {
    family inet {
      address 10.0.0.40/32;
    }
  }
}
```

```
user@R4# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.40;
    family inet {
      unicast {
        add-path {
          receive;
        }
      }
    }
  }
}
```

### Configuring Router R5

To configure Router R5:

- 
- Copyright © 2014, Juniper Networks, Inc. 3533

```
user@R5# set lo0 unit 50 family inet address 10.0.0.50/32
```

2. Configure BGP on Router R5's interface.

```
[edit protocols bgp group e1]
user@R5# set type external
user@R5# set neighbor 10.0.15.1 peer-as 1
```

3. Create static routes for redistribution into BGP.

```
[edit routing-options]
user@R5# set static route 199.1.1.1/32 reject
user@R5# set static route 198.1.1.1/32 reject
```

4. Redistribute static and direct routes into BGP.

```
[edit protocols bgp group e1 neighbor 10.0.15.1]
user@R5# set export s2b
```

```
[edit policy-options policy-statement s2b]
user@R5# set from protocol static
user@R5# set from protocol direct
user@R5# set then as-path-expand 2
user@R5# set then accept
```

5. Configure the autonomous system number.

```
[edit routing-options]
user@R5# set autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R5# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R5# show interfaces
fe-1/2/0 {
  unit 51 {
    family inet {
      address 10.0.15.2/24;
    }
  }
}
lo0 {
  unit 50 {
    family inet {
      address 10.0.0.50/32;
    }
  }
}

user@R5# show protocols
bgp {
  group e1 {
```



```

        type external;
        neighbor 10.0.15.1 {
            export s2b;
            peer-as 1;
        }
    }
}

user@R5# show policy-options
policy-statement s2b {
    from protocol [ static direct ];
    then {
        as-path-expand 2;
        accept;
    }
}

user@R5# show routing-options
static {
    route 198.1.1.1/32 reject;
    route 199.1.1.1/32 reject;
}
autonomous-system 2;

```

### Configuring Router R6

#### Step-by-Step Procedure

To configure Router R6:

1. Configure the loopback (lo0) interface and the interface to Router R2.
 

```

[edit interfaces]
user@R6# set fe-1/2/0 unit 62 family inet address 10.0.26.2/24

user@R6# set lo0 unit 60 family inet address 10.0.0.60/32

```
2. Configure BGP on Router R6's interface.
 

```

[edit protocols]
user@R6# set bgp group e1 type external
user@R6# set bgp group e1 neighbor 10.0.26.1 peer-as 1

```
3. Create static routes for redistribution into BGP.
 

```

[edit]
user@R6# set routing-options static route 199.1.1.1/32 reject
user@R6# set routing-options static route 198.1.1.1/32 reject

```
4. Redistribute static and direct routes from Router R6's routing table into BGP.
 

```

[edit protocols bgp group e1 neighbor 10.0.26.1]
user@R6# set export s2b

[edit policy-options policy-statement s2b]
user@R6# set from protocol static
user@R6# set from protocol direct
user@R6# set then accept

```

5. Configure the autonomous system number.

```
[edit routing-options]
user@R6# set autonomous-system 2
```

6. If you are done configuring the device, commit the configuration.

```
user@R6# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R6# show interfaces
fe-1/2/0 {
  unit 62 {
    family inet {
      address 10.0.26.2/24;
    }
  }
}
lo0 {
  unit 60 {
    family inet {
      address 10.0.0.60/32;
    }
  }
}

user@R6# show protocols
bgp {
  group e1 {
    type external;
    neighbor 10.0.26.1 {
      export s2b;
      peer-as 1;
    }
  }
}

user@R6# show policy-options
policy-statement s2b {
  from protocol [ static direct ];
  then accept;
}

user@R6# show routing-options
static {
  route 198.1.1.1/32 reject;
  route 199.1.1.1/32 reject;
}
autonomous-system 2;
```

*Configuring Router R7***Step-by-Step Procedure**

To configure Router R7:

1. Configure the loopback (lo0) interface and the interface to Router R3.  
  

```
[edit interfaces]
user@R7# set fe-1/2/0 unit 73 family inet address 10.0.37.2/24

user@R7# set lo0 unit 70 family inet address 10.0.0.70/32
```
2. Configure BGP on Router R7's interface.  
  

```
[edit protocols bgp group e1]
user@R7# set type external
user@R7# set neighbor 10.0.37.1 peer-as 1
```
3. Create a static route for redistribution into BGP.  
  

```
[edit]
user@R7# set routing-options static route 199.1.1.1/32 reject
```
4. Redistribute static and direct routes from Router R7's routing table into BGP.  
  

```
[edit protocols bgp group e1 neighbor 10.0.37.1]
user@R7# set export s2b

[edit policy-options policy-statement s2b]
user@R7# set from protocol static
user@R7# set from protocol direct
user@R7# set then accept
```
5. Configure the autonomous system number.  
  

```
[edit routing-options]
user@R7# set autonomous-system 2
```
6. If you are done configuring the device, commit the configuration.  
  

```
user@R7# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R7# show interfaces
fe-1/2/0 {
  unit 73 {
    family inet {
      address 10.0.37.2/24;
    }
  }
}
lo0 {
  unit 70 {
    family inet {
```

```
        address 10.0.0.70/32;
    }
}

user@R7# show protocols
bgp {
    group e1 {
        type external;
        neighbor 10.0.37.1 {
            export s2b;
            peer-as 1;
        }
    }
}

user@R7# show policy-options
policy-statement s2b {
    from protocol [ static direct ];
    then accept;
}

user@R7# show routing-options
static {
    route 199.1.1.1/32 reject;
}
autonomous-system 2;
```

### ***Configuring Router R8***

#### **Step-by-Step Procedure**

To configure Router R8:

1. Configure the loopback (lo0) interface and the interface to Router R4.

[edit interfaces]

```
user@R8# set fe-1/2/0 unit 84 family inet address 10.0.48.2/24
```

```
user@R8# set lo0 unit 80 family inet address 10.0.0.80/32
```

2. Configure BGP and OSPF on Router R8's interface.

[edit protocols]

```
user@R8# set bgp group rr type internal
```

```
user@R8# set bgp group rr local-address 10.0.0.80
```

```
user@R8# set ospf area 0.0.0.0 interface lo0.80 passive
```

```
user@R8# set ospf area 0.0.0.0 interface fe-1/2/0.84
```

3. Configure Router R8 to receive multiple paths from its neighbor, Router R4.

The destination of the paths can be any destination that Router R4 can reach through multiple paths.

[edit protocols]

```
user@R8# set bgp group rr neighbor 10.0.0.40 family inet unicast add-path receive
```

4. Configure the autonomous system number.

[edit]

```
user@R8# set routing-options autonomous-system 1
```

5. If you are done configuring the device, commit the configuration.

```
user@R8# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R8# show interfaces
fe-1/2/0 {
  unit 84 {
    family inet {
      address 10.0.48.2/24;
    }
  }
}
lo0 {
  unit 80 {
    family inet {
      address 10.0.0.80/32;
    }
  }
}

user@R8# show protocols
bgp {
  group rr {
    type internal;
    local-address 10.0.0.80;
    neighbor 10.0.0.40 {
      family inet {
        unicast {
          add-path {
            receive;
          }
        }
      }
    }
  }
}

ospf {
  area 0.0.0.0 {
    interface lo0.80 {
      passive;
    }
    interface fe-1/2/0.84;
  }
}

user@R8# show routing-options
autonomous-system 1;
```

## Verification

---

Confirm that the configuration is working properly.

- [Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths on page 3540](#)
- [Verifying That Router R1 Is Advertising Multiple Paths on page 3540](#)
- [Verifying That Router R4 Is Receiving and Advertising Multiple Paths on page 3541](#)
- [Verifying That Router R8 Is Receiving Multiple Paths on page 3542](#)
- [Checking the Path ID on page 3542](#)

### *Verifying That the BGP Peers Have the Ability to Send and Receive Multiple Paths*

**Purpose** Make sure that one or both of the following strings appear in the output of the **show bgp neighbor** command:

- NLRI's for which peer can receive multiple paths: inet-unicast
- NLRI's for which peer can send multiple paths: inet-unicast

**Action**

```
user@R1> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1      Local: 10.0.0.10+65237 AS 1
Type: Internal    State: Established    Flags: <Sync>
... NLRI's for which peer can receive multiple paths: inet-unicast
...

user@R4> show bgp neighbor 10.0.0.10
Peer: 10.0.0.10+65237 AS 1    Local: 10.0.0.40+179 AS 1
Type: Internal    State: Established    Flags: <Sync>
...
NLRI's for which peer can send multiple paths: inet-unicast
...

user@R4> show bgp neighbor 10.0.0.80
Peer: 10.0.0.80+55416 AS 1    Local: 10.0.0.40+179 AS 1
Type: Internal    State: Established (route reflector client)Flags: <Sync>
...
NLRI's for which peer can receive multiple paths: inet-unicast
...

user@R8> show bgp neighbor 10.0.0.40
Peer: 10.0.0.40+179 AS 1      Local: 10.0.0.80+55416 AS 1
Type: Internal    State: Established    Flags: <Sync>
...
NLRI's for which peer can send multiple paths: inet-unicast
...
```

### *Verifying That Router R1 Is Advertising Multiple Paths*

**Purpose** Make sure that multiple paths to the 198.1.1.1/32 destination and multiple paths to the 199.1.1.1/32 destination are advertised to Router R4.

**Action** user@R1> show route advertising-protocol bgp 10.0.0.40  
 inet.0: 21 destinations, 25 routes (21 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

**Meaning** When you see one prefix and more than one next hop, it means that multiple paths are advertised to Router R4.

### *Verifying That Router R4 Is Receiving and Advertising Multiple Paths*

**Purpose** Make sure that multiple paths to the 199.1.1.1/32 destination are received from Router R1 and advertised to Router R8. Make sure that multiple paths to the 198.1.1.1/32 destination are received from Router R1, but only one path to this destination is advertised to Router R8.

**Action** user@R4> show route receive-protocol bgp 10.0.0.10  
 inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

user@R4> show route advertising-protocol bgp 10.0.0.80  
 inet.0: 19 destinations, 22 routes (19 active, 0 holddown, 0 hidden)

| Prefix         | Nexthop   | MED | Lc1pref | AS path |
|----------------|-----------|-----|---------|---------|
| * 10.0.0.50/32 | 10.0.15.2 |     | 100     | 2 2 I   |
| * 10.0.0.60/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 10.0.0.70/32 | 10.0.0.30 |     | 100     | 2 I     |
| * 198.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
| * 199.1.1.1/32 | 10.0.0.20 |     | 100     | 2 I     |
|                | 10.0.0.30 |     | 100     | 2 I     |
|                | 10.0.15.2 |     | 100     | 2 2 I   |
| * 200.1.1.0/30 | 10.0.0.20 |     | 100     | 2 I     |

**Meaning** The **show route receive-protocol** command shows that Router R4 receives two paths to the 198.1.1.1/32 destination and three paths to the 199.1.1.1/32 destination. The **show route advertising-protocol** command shows that Router R4 advertises only one path to the 198.1.1.1/32 destination and advertises all three paths to the 199.1.1.1/32 destination.

Because of the prefix policy that is applied to Router R4, Router R4 does not advertise multiple paths to the 198.1.1.1/32 destination. Router R4 advertises only one path to the 198.1.1.1/32 destination even though it receives multiple paths to this destination.

### ***Verifying That Router R8 Is Receiving Multiple Paths***

**Purpose** Make sure that Router R8 receives multiple paths to the 199.1.1.1/32 destination through Router R4. Make sure that Router R8 receives only one path to the 198.1.1.1/32 destination through Router R4.

**Action** user@R8> **show route receive-protocol bgp 10.0.0.40**  
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)  
Prefix Nexthop MED Lc1pref AS path  
\* 10.0.0.50/32 10.0.15.2 100 2 2 I  
\* 10.0.0.60/32 10.0.0.20 100 2 I  
\* 10.0.0.70/32 10.0.0.30 100 2 I  
\* 198.1.1.1/32 10.0.0.20 100 2 I  
\* 199.1.1.1/32 10.0.0.20 100 2 I  
10.0.0.30 100 2 I  
10.0.15.2 100 2 2 I  
\* 200.1.1.0/30 10.0.0.20 100 2 I

### ***Checking the Path ID***

**Purpose** On the downstream devices, Router R4 and Router R8, verify that a path ID uniquely identifies the path. Look for the **Addpath Path ID:** string.



**Action** user@R4> show route 199.1.1.1/32 detail

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 3 announced)
  *BGP
    Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 9
    Source: 10.0.0.10
    Next hop type: Router, Next hop index: 676
    Next hop: 10.0.14.1 via lt-1/2/0.41, selected
    Protocol next hop: 10.0.0.20
    Indirect next hop: 92041c8 262146
    State: <Active Int Ext>
    Local AS: 1 Peer AS: 1
    Age: 1:44:37 Metric2: 2
    Task: BGP_1.10.0.0.10+65237
    Announcement bits (3): 2-KRT 3-BGP RT Background 4-Resolve tree

  1
    AS path: 2 I (Originator) Cluster list: 10.0.0.10
    AS path: Originator ID: 10.0.0.20
    Accepted
    Localpref: 100
    Router ID: 10.0.0.10
    Addpath Path ID: 1
  BGP
    Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 4
    Source: 10.0.0.10
    Next hop type: Router, Next hop index: 676
    Next hop: 10.0.14.1 via lt-1/2/0.41, selected
    Protocol next hop: 10.0.0.30
    Indirect next hop: 92042ac 262151
    State: <NotBest Int Ext>
    Inactive reason: Not Best in its group - Router ID
    Local AS: 1 Peer AS: 1
    Age: 1:44:37 Metric2: 2
    Task: BGP_1.10.0.0.10+65237
    Announcement bits (1): 3-BGP RT Background
    AS path: 2 I (Originator) Cluster list: 10.0.0.10
    AS path: Originator ID: 10.0.0.30
    Accepted
    Localpref: 100
    Router ID: 10.0.0.10
    Addpath Path ID: 2
  BGP
    Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 4
    Source: 10.0.0.10
    Next hop type: Router, Next hop index: 676
    Next hop: 10.0.14.1 via lt-1/2/0.41, selected
    Protocol next hop: 10.0.15.2
    Indirect next hop: 92040e4 262150
    State: <Int Ext>
    Inactive reason: AS path
    Local AS: 1 Peer AS: 1
    Age: 1:44:37 Metric2: 2
    Task: BGP_1.10.0.0.10+65237
    Announcement bits (1): 3-BGP RT Background
    AS path: 2 2 I
    Accepted
```

```

Localpref: 100
Router ID: 10.0.0.10
Addpath Path ID: 3

```

```
user@R8> show route 199.1.1.1/32 detail
```

```

inet.0: 17 destinations, 19 routes (17 active, 0 holddown, 0 hidden)
199.1.1.1/32 (3 entries, 1 announced)
*BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 9
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.20
      Indirect next hop: 91fc0e4 262148
      State: <Active Int Ext>
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      Announcement bits (2): 2-KRT 4-Resolve tree 1
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.20
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 1
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.0.30
      Indirect next hop: 91fc1c8 262152
      State: <NotBest Int Ext>
      Inactive reason: Not Best in its group - Router ID
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 I (Originator) Cluster list: 10.0.0.40 10.0.0.10
      AS path: Originator ID: 10.0.0.30
      Accepted
      Localpref: 100
      Router ID: 10.0.0.40
      Addpath Path ID: 2
BGP   Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 4
      Source: 10.0.0.40
      Next hop type: Router, Next hop index: 1045
      Next hop: 10.0.48.1 via lt-1/2/0.84, selected
      Protocol next hop: 10.0.15.2
      Indirect next hop: 91fc2ac 262153
      State: <Int Ext>
      Inactive reason: AS path
      Local AS:      1 Peer AS:      1
      Age: 1:56:51   Metric2: 3
      Task: BGP_1.10.0.0.40+179
      AS path: 2 2 I (Originator) Cluster list: 10.0.0.40
      AS path: Originator ID: 10.0.0.10

```

```
Accepted
Localpref: 100
Router ID: 10.0.0.40
Addpath Path ID: 3
```

- Related Documentation**
- [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP on page 3258](#)
  - [Understanding Adding AS Numbers to BGP AS Paths](#)

## Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing

The Junos OS supports configurations of 16, 32, or 64 equal-cost multipath (ECMP) next hops for RSVP and LDP LSPs on M10i routers with an Enhanced CFEB, M320, M120, MX Series, and T Series routers, and routing devices. For networks with high-volume traffic, this provides more flexibility to load-balance the traffic over as many as 64 LSPs.

To configure the maximum limit for ECMP next hops, include the **maximum-ecmp next-hops** statement at the **[edit chassis]** hierarchy level:

```
[edit chassis]
maximum-ecmp next-hops;
```

You can configure a maximum ECMP next-hop limit of 16, 32, or 64 using this statement. The default limit is 16.



**NOTE:** MX Series routers with one or more Modular Port Concentrator (MPC) cards and with Junos OS 11.4 or earlier installed, support the configuration of the **maximum-ecmp** statement with only 16 next hops. You should *not* configure the **maximum-ecmp** statement with 32 or 64 next hops. When you commit the configuration with 32 or 64 next hops, the following warning message appears:

**Error: Number of members in Unilist NH exceeds the maximum supported 16 on Trio.**

The following types of routes support the ECMP maximum next-hop configuration for as many as 64 ECMP gateways:

- Static IPv4 and IPv6 routes with direct and indirect next-hop ECMPs
- LDP ingress and transit routes learned through associated IGP routes
- RSVP ECMP next hops created for LSPs
- OSPF IPv4 and IPv6 route ECMPs
- ISIS IPv4 and IPv6 route ECMPs
- EBGp IPv4 and IPv6 route ECMPs
- IBGP (resolving over IGP routes) IPv4 and IPv6 route ECMPs

The enhanced ECMP limit of up to 64 ECMP next hops is also applicable for Layer 3 VPNs, Layer 2 VPNs, Layer 2 circuits, and VPLS services that resolve over an MPLS route, because the available ECMP paths in the MPLS route can also be used by such traffic.



**NOTE:**

The following FPCs on M320, T640, and T1600 routers only support 16 ECMP next hops:

- (M320, T640, and T1600 routers only) Enhanced II FPC1
- (M320, T640, and T1600 routers only) Enhanced II FPC2
- (M320 and T640 routers only) Enhanced II FPC3
- (T640 and T1600 routers only) FPC2
- (T640 and T1600 routers only) FPC3

If a maximum ECMP next-hop limit of 32 or 64 is configured on an M320, T640, or T1600 router with any of these FPCs installed, the Packet Forwarding Engines on these FPCs use only the first 16 ECMP next hops. For Packet Forwarding Engines on FPCs that support only 16 ECMP next hops, the Junos OS generates a system log message if a maximum ECMP next-hop limit of 32 or 64 is configured. However, for Packet Forwarding Engines on other FPCs installed on the router, a maximum configured ECMP limit of 32 or 64 ECMP next hops is applicable.



**NOTE:** If RSVP LSPs are configured with bandwidth allocation, for ECMP next hops with more than 16 LSPs, traffic is not distributed optimally based on bandwidths configured. Some LSPs with smaller allocated bandwidths receive more traffic than the ones configured with higher bandwidths. Traffic distribution does not strictly comply with the configured bandwidth allocation. This caveat is applicable to the following routers:

- T1600 and T640 routers with Enhanced Scaling FPC1, Enhanced Scaling FPC2, Enhanced Scaling FPC3, Enhanced Scaling FPC 4, and all Type 4 FPCs
- M320 routers with Enhanced III FPC1, Enhanced III FPC2, and Enhanced III FPC3
- MX Series routers with all types of FPCs and DPCs, excluding MPCs. This caveat is not applicable to MX Series routers with line cards based on the Junos Trio chipset.
- M120 routers with Type 1, Type 2, and Type 3 FPCs
- M10i routers with Enhanced CFEB

Next-hop cloning and permutations are disabled on T Series routers with Enhanced Scaling FPCs (Enhanced Scaling FPC1, Enhanced Scaling FPC2, Enhanced Scaling FPC3,

and Enhanced Scaling FPC 4) that support enhanced load-balancing capability. As a result, memory utilization is reduced for a highly scaled system with a high number of next hops on ECMP or aggregated interfaces. Next-hop cloning and permutations are also disabled on T Series routers with Type-4 FPCs.

To view the details of the ECMP next hops, issue the **show route** command. The **show route summary** command also shows the current configuration for the maximum ECMP limit. To view details of the ECMP LDP paths, issue the **traceroute mpls ldp** command.

**Related Documentation**

- [maximum-ecmp](#)

## IBGP Scaling Configuration

- [Example: Configuring BGP Route Reflectors on page 3547](#)
- [Example: Configuring BGP Confederations on page 3564](#)

### Example: Configuring BGP Route Reflectors

- [Understanding BGP Route Reflectors on page 3547](#)
- [Example: Configuring a Route Reflector on page 3549](#)

#### Understanding BGP Route Reflectors

Because of the internal BGP (IBGP) full-mesh requirement, most networks use route reflectors to simplify configuration. The formula to compute the number of sessions required for a full mesh is  $v * (v - 1) / 2$ , where  $v$  is the number of BGP-enabled devices. The full-mesh model does not scale well. Using a route reflector, you group routers into clusters, which are identified by numeric identifiers unique to the autonomous system (AS). Within the cluster, you must configure a BGP session from a single router (the route reflector) to each internal peer. With this configuration, the IBGP full-mesh requirement is met.

To use route reflection in an AS, you designate one or more routers as a route reflector—typically, one per point of presence (POP). Route reflectors have the special BGP ability to readvertise routes learned from an internal peer to other internal peers. So rather than requiring all internal peers to be fully meshed with each other, route reflection requires only that the route reflector be fully meshed with all internal peers. The route reflector and all of its internal peers form a cluster, as shown in [Figure 85 on page 3548](#).



**NOTE:** For some Juniper Networks devices, you must have an Advanced BGP Feature license installed on each device that uses a route reflector. For license details, see the *Junos OS Initial Configuration Guide for Security Devices*.


**Figure 85: Simple Route Reflector Topology (One Cluster)**

Figure 85 on page 3548 shows Router RR configured as the route reflector for Cluster 127. The other routers are designated internal peers within the cluster. BGP routes are advertised to Router RR by any of the internal peers. RR then readvertises those routes to all other peers within the cluster.

You can configure multiple clusters and link them by configuring a full mesh of route reflectors (see Figure 86 on page 3548).



**Figure 86: Basic Route Reflection (Multiple Clusters)**

Figure 86 on page 3548 shows Route Reflectors RR 1, RR 2, RR 3, and RR 4 as fully meshed internal peers. When a router advertises a route to RR 1, RR 1 readvertises the route to the other route reflectors, which, in turn, readvertise the route to the remaining routers within the AS. Route reflection allows the route to be propagated throughout the AS without the scaling problems created by the full mesh requirement.

However, as clusters become large, a full mesh with a route reflector becomes difficult to scale, as does a full mesh between route reflectors. To help offset this problem, you can group clusters of routers together into clusters of clusters for hierarchical route reflection (see [Figure 87 on page 3549](#)).

**Figure 87: Hierarchical Route Reflection (Clusters of Clusters)**



[Figure 87 on page 3549](#) shows RR 2, RR 3, and RR 4 as the route reflectors for Clusters 127, 19, and 45, respectively. Rather than fully mesh those route reflectors, the network administrator has configured them as part of another cluster (Cluster 6) for which RR 1 is the route reflector. When a router advertises a route to RR 2, RR 2 readvertises the route to all the routers within its own cluster, and then readvertises the route to RR 1. RR 1 readvertises the route to the routers in its cluster, and those routers propagate the route down through their clusters.

### Example: Configuring a Route Reflector

This example shows how to configure a route reflector.

- [Requirements on page 3549](#)
- [Overview on page 3549](#)
- [Configuration on page 3551](#)
- [Verification on page 3559](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.

#### Overview

Generally, internal BGP (IBGP)-enabled devices need to be fully meshed, because IBGP does not readvertise updates to other IBGP-enabled devices. The full mesh is a logical mesh achieved through configuration of multiple **neighbor** statements on each IBGP-enabled device. The full mesh is not necessarily a physical full mesh. Maintaining a full mesh (logical or physical) does not scale well in large deployments.

Figure 88 on page 3551 shows an IBGP network with Device A acting as a route reflector. Device B and Device C are clients of the route reflector. Device D and Device E are outside the cluster, so they are nonclients of the route reflector.

On Device A (the route reflector), you must form peer relationships with all of the IBGP-enabled devices by including the **neighbor** statement for the clients (Device B and Device C) and the nonclients (Device D and Device E). You must also include the **cluster** statement and a cluster identifier. The cluster identifier can be any 32-bit value. This example uses the loopback interface IP address of the route reflector.

On Device B and Device C, the route reflector clients, you only need one **neighbor** statement that forms a peer relationship with the route reflector, Device A.

On Device D and Device E, the nonclients, you need a **neighbor** statement for each nonclient device (D-to-E and E-to-D). You also need a **neighbor** statement for the route reflector (D-to-A and E-to-A). Device D and Device E do not need **neighbor** statements for the client devices (Device B and Device C).




**TIP:** Device D and Device E are considered to be nonclients because they have explicitly configured peer relationships with each other. To make them RRroute reflector clients, remove the **neighbor 192.168.5.5** statement from the configuration on Device D, and remove the **neighbor 192.168.0.1** statement from the configuration on Device E.

---



Figure 88: IBGP Network Using a Route Reflector

**Configuration**

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device A**

```

set interfaces fe-0/0/0 unit 1 description to-B
set interfaces fe-0/0/0 unit 1 family inet address 10.10.10.1/30
set interfaces fe-0/0/1 unit 3 description to-D
set interfaces fe-0/0/1 unit 3 family inet address 10.10.10.9/30
set interfaces lo0 unit 1 family inet address 192.168.6.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.6.5
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers cluster 192.168.6.5
set protocols bgp group internal-peers neighbor 192.163.6.4
set protocols bgp group internal-peers neighbor 192.168.40.4
set protocols bgp group internal-peers neighbor 192.168.0.1
set protocols bgp group internal-peers neighbor 192.168.5.5
set protocols ospf area 0.0.0.0 interface lo0.1 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.1
set protocols ospf area 0.0.0.0 interface fe-0/0/1.3
set policy-options policy-statement send-ospf term 2 from protocol ospf

```

```
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.6.5
set routing-options autonomous-system 17
```

**Device B**

```
set interfaces fe-0/0/0 unit 2 description to-A
set interfaces fe-0/0/0 unit 2 family inet address 10.10.10.2/30
set interfaces fe-0/0/1 unit 5 description to-C
set interfaces fe-0/0/1 unit 5 family inet address 10.10.10.5/30
set interfaces lo0 unit 2 family inet address 192.163.6.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.163.6.4
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.2
set protocols ospf area 0.0.0.0 interface fe-0/0/1.5
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.163.6.4
set routing-options autonomous-system 17
```

**Device C**

```
set interfaces fe-0/0/0 unit 6 description to-B
set interfaces fe-0/0/0 unit 6 family inet address 10.10.10.6/30
set interfaces lo0 unit 3 family inet address 192.168.40.4/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.40.4
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.6
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.40.4
set routing-options autonomous-system 17
```

**Device D**

```
set interfaces fe-0/0/0 unit 4 description to-A
set interfaces fe-0/0/0 unit 4 family inet address 10.10.10.10/30
set interfaces fe-0/0/1 unit 7 description to-E
set interfaces fe-0/0/1 unit 7 family inet address 10.10.10.13/30
set interfaces lo0 unit 4 family inet address 192.168.0.1/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.0.1
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols bgp group internal-peers neighbor 192.168.5.5
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.4
set protocols ospf area 0.0.0.0 interface fe-0/0/1.7
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 17
```

**Device E**

```
set interfaces fe-0/0/0 unit 8 description to-D
set interfaces fe-0/0/0 unit 8 family inet address 10.10.10.14/30
```

```

set interfaces lo0 unit 5 family inet address 192.168.5.5/32
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.5.5
set protocols bgp group internal-peers export send-ospf
set protocols bgp group internal-peers neighbor 192.168.0.1
set protocols bgp group internal-peers neighbor 192.168.6.5
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/0.8
set policy-options policy-statement send-ospf term 2 from protocol ospf
set policy-options policy-statement send-ospf term 2 then accept
set routing-options router-id 192.168.5.5
set routing-options autonomous-system 17

```

### Configuring the Route Reflector

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IBGP in the network using Juniper Networks Device A as a route reflector:

1. Configure the interfaces.

```

[edit interfaces]
user@A# set fe-0/0/0 unit 1 description to-B
user@A# set fe-0/0/0 unit 1 family inet address 10.10.1/30
user@A# set fe-0/0/1 unit 3 description to-D
user@A# set fe-0/0/1 unit 3 family inet address 10.10.9/30
user@A# set lo0 unit 1 family inet address 192.168.6.5/32

```

2. Configure BGP, including the cluster identifier and neighbor relationships with all IBGP-enabled devices in the autonomous system (AS).

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@A# set type internal
user@A# set local-address 192.168.6.5
user@A# set export send-ospf
user@A# set cluster 192.168.6.5
user@A# set neighbor 192.163.6.4
user@A# set neighbor 192.168.40.4
user@A# set neighbor 192.168.0.1
user@A# set neighbor 192.168.5.5

```

3. Configure static routing or an interior gateway protocol (IGP).

This example uses OSPF.

```

[edit protocols ospf area 0.0.0.0]
user@A# set interface lo0.1 passive
user@A# set interface fe-0/0/0.1
user@A# set interface fe-0/0/1.3

```

4. Configure the policy that redistributes OSPF routes into BGP.

```

[edit policy-options policy-statement send-ospf term 2]
user@A# set from protocol ospf
user@A# set then accept

```

5. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@A# set router-id 192.168.6.5
user@A# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@A# show interfaces
fe-0/0/0 {
  unit 1 {
    description to-B;
    family inet {
      address 10.10.10.1/30;
    }
  }
}
fe-0/0/1 {
  unit 3 {
    description to-D;
    family inet {
      address 10.10.10.9/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 192.168.6.5/32;
    }
  }
}

user@A# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.168.6.5;
    export send-ospf;
    cluster 192.168.6.5;
    neighbor 192.163.6.4;
    neighbor 192.168.40.4;
    neighbor 192.168.0.1;
    neighbor 192.168.5.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.1 {
      passive;
    }
    interface fe-0/0/0.1;
    interface fe-0/0/1.3;
```

```

    }
  }

user@A# show policy-options
policy-statement send-ospf {
  term 2 {
    from protocol ospf;
    then accept;
  }
}

user@A# show routing-options
router-id 192.168.6.5;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



**NOTE:** Repeat these steps for each nonclient BGP peer within the cluster that you are configuring, if the other nonclient devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

### Configuring Client Peers

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure client peers:

1. Configure the interfaces.
 

```

[edit interfaces]
user@B# set fe-0/0/0 unit 2 description to-A
user@B# set fe-0/0/0 unit 2 family inet address 10.10.10.2/30
user@B# set fe-0/0/1 unit 5 description to-C
user@B# set fe-0/0/1 unit 5 family inet address 10.10.10.5/30
user@B# set lo0 unit 2 family inet address 192.163.6.4/32
      
```

2. Configure the BGP neighbor relationship with the route reflector.

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@B# set type internal
user@B# set local-address 192.163.6.4
user@B# set export send-ospf
user@B# set neighbor 192.168.6.5

```

3. Configure OSPF.
 

```

[edit protocols ospf area 0.0.0.0]
user@B# set interface lo0.2 passive
user@B# set interface fe-0/0/0.2
user@B# set interface fe-0/0/1.5
      
```
4. Configure the policy that redistributes OSPF routes into BGP.

```
[edit policy-options policy-statement send-ospf term 2]
user@B# set from protocol ospf
user@B# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@B# set router-id 192.163.6.4
user@B# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/0/0 {
  unit 2 {
    description to-A;
    family inet {
      address 10.10.10.2/30;
    }
  }
}
fe-0/0/1 {
  unit 5 {
    description to-C;
    family inet {
      address 10.10.10.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.163.6.4/32;
    }
  }
}

user@B# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.163.6.4;
    export send-ospf;
    neighbor 192.168.6.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.2 {
      passive;
    }
    interface fe-0/0/0.2;
    interface fe-0/0/1.5;
```

```

    }
  }

user@B# show policy-options
policy-statement send-ospf {
  term 2 {
    from protocol ospf;
    then accept;
  }
}

user@B# show routing-options
router-id 192.163.6.4;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



**NOTE:** Repeat these steps for each client BGP peer within the cluster that you are configuring if the other client devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

### Configuring Nonclient Peers

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure nonclient peers:

1. Configure the interfaces.
 

```

[edit interfaces]
user@D# set fe-0/0/0 unit 4 description to-A
user@D# set fe-0/0/0 unit 4 family inet address 10.10.10.10/30
user@D# set fe-0/0/1 unit 7 description to-E
user@D# set fe-0/0/1 unit 7 family inet address 10.10.10.13/30
user@D# set lo0 unit 4 family inet address 192.168.0.1/32
      
```
2. Configure the BGP neighbor relationships with the RRroute reflector and with the other nonclient peers.

Also apply the policy that redistributes OSPF routes into BGP.

```

[edit protocols bgp group internal-peers]
user@D# set type internal
user@D# set local-address 192.168.0.1
user@D# set export send-ospf
user@D# set neighbor 192.168.6.5
user@D# set neighbor 192.168.5.5

```

3. Configure OSPF.
 

```

[edit protocols ospf area 0.0.0.0]
user@D# set interface lo0.4 passive
user@D# set interface fe-0/0/0.4
      
```

```
user@D# set interface fe-0/0/1.7
```

4. Configure the policy that redistributes OSPF routes into BGP.

```
[edit policy-options policy-statement send-ospf term 2]
user@D# set from protocol ospf
user@D# set then accept
```

5. Configure the router ID and the AS number.

```
[edit routing-options]
user@D# set router-id 192.168.0.1
user@D# set autonomous-system 17
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@D# show interfaces
fe-0/0/0 {
  unit 4 {
    description to-A;
    family inet {
      address 10.10.10.10/30;
    }
  }
}
fe-0/0/1 {
  unit 7 {
    description to-E;
    family inet {
      address 10.10.10.13/30;
    }
  }
}
lo0 {
  unit 4 {
    family inet {
      address 192.168.0.1/32;
    }
  }
}
```

```
user@D# show protocols
bgp {
  group internal-peers {
    type internal;
    local-address 192.168.0.1;
    export send-ospf;
    neighbor 192.168.6.5;
    neighbor 192.168.5.5;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.4 {
```



```

        passive;
    }
    interface fe-0/0/0.4;
    interface fe-0/0/1.7;
}
}

user@D# show policy-options
policy-statement send-ospf {
    term 2 {
        from protocol ospf;
        then accept;
    }
}

user@D# show routing-options
router-id 192.168.0.1;
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.



**NOTE:** Repeat these steps for each nonclient BGP peer within the cluster that you are configuring if the other nonclient devices are from Juniper Networks. Otherwise, consult the device's documentation for instructions.

### Verification

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3559](#)
- [Verifying BGP Groups on page 3562](#)
- [Verifying BGP Summary Information on page 3562](#)
- [Verifying Routing Table Information on page 3562](#)

### Verifying BGP Neighbors

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is established for each neighbor address.

**Action** From operational mode, enter the **show bgp neighbor** command.

```

user@A> show bgp neighbor
Peer: 192.163.6.4+179 AS 17    Local: 192.168.6.5+62857 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.163.6.4    Local ID: 192.168.6.5    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  BFD: disabled, down

```

```

NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        6
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 5    Sent 3    Checked 19
Input messages: Total 2961    Updates 7    Refreshes 0    Octets 56480
Output messages: Total 2945    Updates 6    Refreshes 0    Octets 56235
Output Queue[0]: 0

Peer: 192.168.0.1+179 AS 17    Local: 192.168.6.5+60068 AS 17
Type: Internal    State: Established (route reflector client)Flags: <Sync>
Last State: OpenConfirm    Last Event: RecvKeepAlive
Last Error: None
Export: [ send-ospf ]
Options: <Preference LocalAddress Cluster Refresh>
Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 192.168.0.1    Local ID: 192.168.6.5    Active Holdtime: 90
Keepalive Interval: 30    Peer index: 3
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 120
Stale routes from peer are kept for: 300
Restart time requested by this peer: 120
NLRI that peer supports restart for: inet-unicast
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        6
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 18    Sent 20    Checked 12
Input messages: Total 15    Updates 5    Refreshes 0    Octets 447
Output messages: Total 554    Updates 4    Refreshes 0    Octets 32307

```

Output Queue[0]: 0

```

Peer: 192.168.5.5+57458 AS 17 Local: 192.168.6.5+179 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.168.5.5      Local ID: 192.168.6.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 2
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 17)
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          0
    Received prefixes:        7
    Accepted prefixes:        7
    Suppressed due to damping: 0
    Advertised prefixes:      6
  Last traffic (seconds): Received 17 Sent 3 Checked 9
  Input messages: Total 2967 Updates 7 Refreshes 0 Octets 56629
  Output messages: Total 2943 Updates 6 Refreshes 0 Octets 56197
  Output Queue[0]: 0

```

```

Peer: 192.168.40.4+53990 AS 17 Local: 192.168.6.5+179 AS 17
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ send-ospf ]
  Options: <Preference LocalAddress Cluster Refresh>
  Local Address: 192.168.6.5 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.168.40.4     Local ID: 192.168.6.5      Active Holdtime: 90
  Keepalive Interval: 30      Peer index: 1
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-unicast
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast

```

```

Peer supports 4 byte AS extension (peer-as 17)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:        7
  Accepted prefixes:        7
  Suppressed due to damping: 0
  Advertised prefixes:      6
Last traffic (seconds): Received 5   Sent 23   Checked 52
Input messages: Total 2960   Updates 7   Refreshes 0   Octets 56496
Output messages: Total 2943   Updates 6   Refreshes 0   Octets 56197
Output Queue[0]: 0

```

### Verifying BGP Groups

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From operational mode, enter the **show bgp group** command.

```

user@A> show bgp group
Group Type: Internal      AS: 17                Local AS: 17
Name: internal-peers     Index: 0              Flags: <>
Export: [ send-ospf ]
Options: <Cluster>
Holdtime: 0
Total peers: 4           Established: 4
192.163.6.4+179
192.168.40.4+53990
192.168.0.1+179
192.168.5.5+57458
inet.0: 0/26/16/0

Groups: 1  Peers: 4   External: 0   Internal: 4   Down peers: 0   Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0          26         0         0         0         0         0         0

```

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From operational mode, enter the **show bgp summary** command.

```

user@A> show bgp summary

Groups: 1 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State   Pending
inet.0          26         0         0         0         0         0         0
Peer      AS      InPkt    OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.163.6.4      17      2981      2965        0        0  22:19:15 0/6/1/0      0/0/0/0
192.168.0.1      17        36        575        0        0   13:43 0/6/1/0      0/0/0/0
192.168.5.5      17      2988      2964        0        0  22:19:10 0/7/7/0      0/0/0/0
192.168.40.4     17      2980      2964        0        0  22:19:14 0/7/7/0      0/0/0/0

```

### Verifying Routing Table Information

**Purpose** Verify that the routing table contains the IBGP routes.

**Action** From operational mode, enter the **show route** command.

```

user@A> show route
inet.0: 12 destinations, 38 routes (12 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.0/30    * [Direct/0] 22:22:03
                > via fe-0/0/0.1
                [BGP/170] 22:20:55, MED 2, localpref 100, from 192.168.40.4
                AS path: I
                > to 10.10.10.2 via fe-0/0/0.1
                [BGP/170] 22:20:51, MED 3, localpref 100, from 192.168.5.5
                AS path: I
                > to 10.10.10.10 via fe-0/0/1.3
10.10.10.1/32    * [Local/0] 22:22:03
                Local via fe-0/0/0.1
10.10.10.4/30    * [OSPF/10] 22:21:13, metric 2
                > to 10.10.10.2 via fe-0/0/0.1
                [BGP/170] 22:20:51, MED 4, localpref 100, from 192.168.5.5
                AS path: I
                > to 10.10.10.10 via fe-0/0/1.3
10.10.10.8/30    * [Direct/0] 22:22:03
                > via fe-0/0/1.3
                [BGP/170] 22:20:51, MED 2, localpref 100, from 192.168.5.5
                AS path: I
                > to 10.10.10.10 via fe-0/0/1.3
                [BGP/170] 22:20:55, MED 3, localpref 100, from 192.168.40.4
                AS path: I
                > to 10.10.10.2 via fe-0/0/0.1
10.10.10.9/32    * [Local/0] 22:22:03
                Local via fe-0/0/1.3
10.10.10.12/30   * [OSPF/10] 22:21:08, metric 2
                > to 10.10.10.10 via fe-0/0/1.3
                [BGP/170] 22:20:55, MED 4, localpref 100, from 192.168.40.4
                AS path: I
                > to 10.10.10.2 via fe-0/0/0.1
192.163.6.4/32   * [OSPF/10] 22:21:13, metric 1
                > to 10.10.10.2 via fe-0/0/0.1
                [BGP/170] 22:20:55, MED 1, localpref 100, from 192.168.40.4
                AS path: I
                > to 10.10.10.2 via fe-0/0/0.1
                [BGP/170] 22:20:51, MED 3, localpref 100, from 192.168.5.5
                AS path: I
                > to 10.10.10.10 via fe-0/0/1.3
192.168.0.1/32   * [OSPF/10] 22:21:08, metric 1
                > to 10.10.10.10 via fe-0/0/1.3
                [BGP/170] 22:20:51, MED 1, localpref 100, from 192.168.5.5
                AS path: I
                > to 10.10.10.10 via fe-0/0/1.3
                [BGP/170] 22:20:55, MED 3, localpref 100, from 192.168.40.4
                AS path: I
                > to 10.10.10.2 via fe-0/0/0.1
192.168.5.5/32   * [OSPF/10] 22:21:08, metric 2
                > to 10.10.10.10 via fe-0/0/1.3
                [BGP/170] 00:15:24, MED 1, localpref 100, from 192.168.0.1
                AS path: I
                > to 10.10.10.10 via fe-0/0/1.3
                [BGP/170] 22:20:55, MED 4, localpref 100, from 192.168.40.4
                AS path: I
                > to 10.10.10.2 via fe-0/0/0.1
192.168.6.5/32   * [Direct/0] 22:22:04

```

```
> via lo0.1
[BGP/170] 22:20:51, MED 2, localpref 100, from 192.168.5.5
AS path: I
> to 10.10.10.10 via fe-0/0/1.3
[BGP/170] 22:20:55, MED 2, localpref 100, from 192.168.40.4
AS path: I
192.168.40.4/32 > to 10.10.10.2 via fe-0/0/0.1
*[OSPF/10] 22:21:13, metric 2
> to 10.10.10.2 via fe-0/0/0.1
[BGP/170] 22:20:55, MED 1, localpref 100, from 192.163.6.4
AS path: I
> to 10.10.10.2 via fe-0/0/0.1
[BGP/170] 22:20:51, MED 4, localpref 100, from 192.168.5.5
AS path: I
224.0.0.5/32 > to 10.10.10.10 via fe-0/0/1.3
*[OSPF/10] 22:22:07, metric 1
MultiRecv
```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3261](#)
  - [BGP Configuration Overview](#)

## Example: Configuring BGP Confederations

- [Understanding BGP Confederations on page 3564](#)
- [Example: Configuring BGP Confederations on page 3565](#)

### Understanding BGP Confederations

---

BGP confederations are another way to solve the scaling problems created by the BGP full mesh requirement. BGP confederations effectively break up a large autonomous system (AS) into subautonomous systems (sub-ASs). Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64,512 and 65,535.

Within a sub-AS, the same internal BGP (IBGP) full mesh requirement exists. Connections to other confederations are made with standard external BGP (EBGP), and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

The confederation AS appears whole to other confederation ASs. The AS path received by other ASs shows only the globally assigned AS number. It does not include the confederation sequence or the privately assigned sub-AS numbers. The sub-AS numbers are removed when the route is advertised out of the confederation AS.

[Figure 89 on page 3565](#) shows an AS divided into four confederations.

Figure 89: BGP Confederations




Figure 89 on page 3565 shows AS 3 divided into four sub-ASs, 64517, 64550, 65300, and 65410, which are linked through EBGP sessions. Because the confederations are connected by EBGP, they do not need to be fully meshed. EBGP routes are readvertised to other sub-ASs.

### Example: Configuring BGP Confederations

This example shows how to configure BGP confederations.

- [Requirements on page 3565](#)
- [Overview on page 3565](#)
- [Configuration on page 3566](#)
- [Verification on page 3568](#)

#### Requirements

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3262](#).
- Configure interior gateway protocol (IGP) sessions between peers.
- Configure a routing policy to advertise the BGP routes.

#### Overview


Within a BGP confederation, the links between the confederation member autonomous systems (ASs) must be external BGP (EBGP) links, not internal BGP (IBGP) links.

Similar to route reflectors, BGP confederations reduce the number of peer sessions and TCP sessions to maintain connections between IBGP routing devices. BGP confederation is one method used to solve the scaling problems created by the IBGP full mesh requirement. BGP confederations effectively break up a large AS into subautonomous

systems. Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64512 and 65535. Within a sub-AS, the same IBGP full mesh requirement exists. Connections to other confederations are made with standard EBGP, and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

Figure 90 on page 3566 shows a sample network in which AS 17 has two separate confederations: sub-AS 64512 and sub-AS 64513, each of which has multiple routers. Within a sub-AS, an IGP is used to establish network connectivity with internal peers. Between sub-ASs, an EBGP peer session is established.

**Figure 90: Typical Network Using BGP Confederations**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

|                               |   |
|-------------------------------|---|
| All Devices in Sub-AS 64512   | <pre> set routing-options autonomous-system 64512 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64512 type internal set protocols bgp group sub-AS-64512 local-address 192.168.5.1 set protocols bgp group sub-AS-64512 neighbor 192.168.8.1 set protocols bgp group sub-AS-64512 neighbor 192.168.15.1 </pre> |
| Border Device in Sub-AS 64512 | <pre> set protocols bgp group to-sub-AS-64513 type external set protocols bgp group to-sub-AS-64513 peer-as 64513 set protocols bgp group to-sub-AS-64513 neighbor 192.168.5.2 </pre>   |
| All Devices in Sub-AS 64513   | <pre> set routing-options autonomous-system 64513 set routing-options confederation 17 members 64512 set routing-options confederation 17 members 64513 set protocols bgp group sub-AS-64513 type internal set protocols bgp group sub-AS-64513 local-address 192.168.5.2 set protocols bgp group sub-AS-64513 neighbor 192.168.9.1 </pre>  |



|                               |   |
|-------------------------------|---|
|                               | <pre>set protocols bgp group sub-AS-64513 neighbor 192.168.16.1</pre>   |
| Border Device in Sub-AS 64513 | <pre>set protocols bgp group to-sub-AS-64512 type external set protocols bgp group to-sub-AS-64512 peer-as 64512 set protocols bgp group to-sub-AS-64512 neighbor 192.168.5.1</pre>   |
| Step-by-Step Procedure        | <p>This procedure shows the steps for the devices that are in sub-AS 64512.</p> <p>The <b>autonomous-system</b> statement sets the sub-AS number of the device.</p> <p>The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see <i>Using the CLI Editor in Configuration Mode</i> in the <i>CLI User Guide</i>.</p> <p>To configure BGP confederations:</p> <ol style="list-style-type: none"> <li>Set the sub-AS number for the device. <pre>[edit routing-options] user@host# set autonomous-system 64512</pre> </li> <li>In the confederation, include all sub-ASs in the main AS. <p>The number 17 represents the main AS. The <b>members</b> statement lists all the sub-ASs in the main AS.</p> <pre>[edit routing-options confederation] user@host# set 17 members 64512 user@host# set 17 members 64513</pre> </li> <li>On the border device in sub-AS 64512, configure an EBGP connection to the border device in AS 64513. <pre>[edit protocols bgp group to-sub-AS-64513] user@host# set type external user@host# set neighbor 192.168.5.2 user@host# set peer-as 64513</pre> </li> <li>Configure an IBGP group for peering with the devices within sub-AS 64512. <pre>[edit protocols bgp group sub-AS-64512] user@host# set type internal user@host# set local-address 192.168.5.1 user@host# neighbor 192.168.8.1 user@host# neighbor 192.168.15.1</pre> </li> </ol> |
| Results                       | <p>From configuration mode, confirm your configuration by entering the <b>show routing-options</b> and <b>show protocols</b> commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.</p> <pre>user@host# show routing-options autonomous-system 64512; confederation 17 members [ 64512 64513 ];  user@host# show protocols bgp {   group to-sub-AS-64513 { # On the border devices only     type external;</pre>  |

```
    peer-as 64513;
    neighbor 192.168.5.2;
  }
  group sub-AS-64512 {
    type internal;
    local-address 192.168.5.1;
    neighbor 192.168.8.1;
    neighbor 192.168.15.1;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.  
Repeat these steps for sSub-AS 64513.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying BGP Neighbors on page 3568](#)
- [Verifying BGP Groups on page 3569](#)
- [Verifying BGP Summary Information on page 3570](#)

### **Verifying BGP Neighbors**

**Purpose** Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

**Action** From the CLI, enter the **show bgp neighbor** command.

### **Sample Output**

```
user@host> show bgp neighbor
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: Sync
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh

  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
```

```

Active prefixes: 4
Received prefixes: 6
Suppressed due to damping: 0
Table inet6.0 Bit: 20000
RIB State: restart is complete
Send state: in sync
Active prefixes: 0
Received prefixes: 2
Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

**Meaning** The output shows a list of the BGP neighbors with detailed session information. Verify the following information:

- Each configured peering neighbor is listed.
- For **State**, each BGP session is **Established**.
- For **Type**, each peer is configured as the correct type (either internal or external).
- For **AS**, the AS number of the BGP neighbor is correct.

#### *Verifying BGP Groups*

**Purpose** Verify that the BGP groups are configured correctly.

**Action** From the CLI, enter the **show bgp group** command.

### Sample Output

```

user@host> show bgp group
Group Type: Internal  AS: 10045      Local AS: 10045
Name: pe-to-asbr2
Export: [ match-all ]
Total peers: 1      Established: 1
10.0.0.4+179
bgp.l3vpn.0: 1/1/0
vpn-green.inet.0: 1/1/0

Groups: 1  Peers: 1  External: 0  Internal: 1  Down peers: 0  Flaps: 0
Table      Tot Paths  Act Paths Suppressed  History Damp State  Pending
bgp.l3vpn.0      1          1          0          0          0          0

```

**Meaning** The output shows a list of the BGP groups with detailed group information. Verify the following information:

- Each configured group is listed.
- For **AS**, each group's remote AS is configured correctly.
- For **Local AS**, each group's local AS is configured correctly.
- For **Group Type**, each group has the correct type (either internal or external).

- For **Total peers**, the expected number of peers within the group is shown.
- For **Established**, the expected number of peers within the group have BGP sessions in the **Established** state.
- The IP addresses of all the peers within the group are present.

### Verifying BGP Summary Information

**Purpose** Verify that the BGP configuration is correct.

**Action** From the CLI, enter the **show bgp summary** command.

### Sample Output

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table      Tot Paths  Act Paths  Suppressed    History Damp State    Pending
inet.0           6           4           0           0      0      0      0
Peer      AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2    65002    88675    88652      0        2      42:38 2/4/0
           0/0/0
10.0.0.3    65002    54528    54532      0        1     2w4d22h 0/0/0
           0/0/0
10.0.0.4    65002    51597    51584      0        0     2w3d22h 2/2/0
           0/0/0
```

**Meaning** The output shows a summary of BGP session information. Verify the following information:

- For **Groups**, the total number of configured groups is shown.
- For **Peers**, the total number of BGP peers is shown.
- For **Down Peers**, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under **Peer**, the IP address for each configured peer is shown.
- Under **AS**, the peer AS for each configured peer is correct.
- Under **Up/Dwn State**, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being damped (such as 0/0/0). If the field is **Active**, it indicates a problem in the establishment of the BGP session.

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

## BGP Security Configuration

- [Example: Configuring BGP Route Authentication on page 3571](#)
- [Examples: Configuring TCP and BGP Security on page 3577](#)

## Example: Configuring BGP Route Authentication

- [Understanding Route Authentication on page 3571](#)
- [Example: Configuring Route Authentication for BGP on page 3572](#)

### Understanding Route Authentication

The use of router and route authentication and route integrity greatly mitigates the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. In this kind of attack, the attacked router can be tricked into creating a routing loop, or the attacked router's routing table can be greatly increased thus impacting performance, or routing information can be redirected to a place in the network for the attacker to analyze it. Bogus route advertisements can be sent out on a segment. These updates can be accepted into the routing tables of neighbor routers unless an authentication mechanism is in place to verify the source of the routes.

Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, it accepts the route. By using a hashing algorithm, the key is not sent over the wire in plain text. Instead, a hash is calculated using the configured key. The routing update is used as the input text, along with the key, into the hashing function. This hash is sent along with the route update to the receiving router. The receiving router compares the received hash with a hash it generates on the route update using the preshared key configured on it. If the two hashes are the same, the route is assumed to be from a trusted source. The key is known only to the sending and receiving routers.

To further strengthen security, you can configure a series of authentication keys (a *keychain*). Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as *hitless* because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.

The sending peer uses the following rules to identify the active authentication key:

- The start time is less than or equal to the current time (in other words, not in the future).
- The start time is greater than that of all other keys in the chain whose start time is less than the current time (in other words, closest to the current time).

The receiving peer determines the key with which it authenticates based on the incoming key identifier.

The sending peer identifies the current authentication key based on a configured start time and then generates a hash value using the current key. The sending peer then inserts a TCP-enhanced authentication option object into the BGP update message. The object contains an object ID (assigned by IANA), the object length, the current key, and a hash value.

The receiving peer examines the incoming TCP-enhanced authentication option, looks up the received authentication key, and determines whether the key is acceptable based on the start time, the system time, and the tolerance parameter. If the key is accepted, the receiving peer calculates a hash and authenticates the update message.

Initial application of a keychain to a TCP session causes the session to reset. However, once the keychain is applied, the addition or removal of a password from the keychain does not cause the TCP session to reset. Also, the TCP session does not reset when the keychain changes from one authentication algorithm to another.

### Example: Configuring Route Authentication for BGP

---

All BGP protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in autonomous system (AS) routing updates. By default, authentication is disabled.

- [Requirements on page 3572](#)
- [Overview on page 3572](#)
- [Configuration on page 3573](#)
- [Verification on page 3575](#)

#### **Requirements**

Before you begin:

- Configure the router interfaces.
- Configure an interior gateway protocol (IGP).

#### **Overview**

When you configure authentication, the algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

This example includes the following statements for configuring and applying the keychain:

- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.  
  
The key can be up to 126 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
- **tolerance**—(Optional) For each keychain, you can configure a clock-skew tolerance value in seconds. The clock-skew tolerance is applicable to the receiver accepting keys for BGP updates. The configurable range is 0 through 999,999,999 seconds. During the tolerance period, either the current or previous password is acceptable.
- **key-chain**—For each keychain, you must specify a name. This example defines one keychain: **bgp-auth**. You can have multiple keychains on a routing device. For example, you can have a keychain for BGP, a keychain for OSPF, and a keychain for LDP.

- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the keychain.
- **authentication-key-chain**—Enables you to apply a keychain at the global BGP level for all peers, for a group, or for a neighbor. This example applies the keychain to the peers defined in the external BGP (EBGP) group called **ext**.
- **authentication-algorithm**—For each keychain, you can specify a hashing algorithm. The algorithm can be AES-128, MD5, or SHA-1.


You associate a keychain and an authentication algorithm with a BGP neighboring session.

This example configures a keychain named **bgp-auth**. Key 0 will be sent and accepted starting at 2011-6-23.20:19:33 -0700, and will stop being sent and accepted when the next key in the keychain (key 1) becomes active. Key 1 becomes active one year later at 2012-6-23.20:19:33 -0700, and will not stop being sent and accepted unless another key is configured with a start time that is later than the start time of key 1. A clock-skew tolerance of 30 seconds applies to the receiver accepting the keys. During the tolerance period, either the current or previous key is acceptable. The keys are shared-secret passwords. This means that the neighbors receiving the authenticated routing updates must have the same authentication keychain configuration, including the same keys (passwords). So Router R0 and Router R1 must have the same authentication-key-chain configuration if they are configured as peers. This example shows the configuration on only one of the routing devices.

### Topology Diagram

Figure 91 on page 3573 shows the topology used in this example.

Figure 91: Authentication for BGP



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols bgp group ext type external
set protocols bgp group ext peer-as 65530
set protocols bgp group ext neighbor 172.16.2.1
set routing-options autonomous-system 65533
```

```
set protocols bgp group ext authentication-key-chain bgp-auth
set protocols bgp group ext authentication-algorithm md5
set security authentication-key-chains key-chain bgp-auth tolerance 30
set security authentication-key-chains key-chain bgp-auth key 0 secret
  this-is-the-secret-password
set security authentication-key-chains key-chain bgp-auth key 0 start-time
  2011-6-23.20:19:33-0700
set security authentication-key-chains key-chain bgp-auth key 1 secret
  this-is-another-secret-password
set security authentication-key-chains key-chain bgp-auth key 1 start-time
  2012-6-23.20:19:33-0700
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1 to accept route filters from Device CE1 and perform outbound route filtering using the received filters:

1. Configure the local autonomous system.

```
[edit routing-options]
user@R1# set autonomous-system 65533
```

2. Configure one or more BGP groups.

```
[edit protocols bgp group ext]
user@R1# set type external
user@R1# set peer-as 65530
user@R1# set neighbor 172.16.2.1
```

3. Configure authentication with multiple keys.

```
[edit security authentication-key-chains key-chain bgp-auth]
user@R1# set key 0 secret this-is-the-secret-password
user@R1# set key 0 start-time 2011-6-23.20:19:33-0700
user@R1# set key 1 secret this-is-another-secret-password
user@R1# set key 1 start-time 2012-6-23.20:19:33-0700
```

The start time of each key must be unique within the keychain.

4. Apply the authentication keychain to BGP, and set the hashing algorithm.

```
[edit protocols bgp group ext]
user@R1# set authentication-key-chain bgp-auth
user@R1# set authentication-algorithm md5
```

5. (Optional) Apply a clock-skew tolerance value in seconds.

```
[edit security authentication-key-chains key-chain bgp-auth]
user@R1# set tolerance 30
```

**Results** From configuration mode, confirm your configuration by entering the **show protocols**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
bgp {
```



```

group ext {
  type external;
  peer-as 65530;
  neighbor 172.16.2.1;
  authentication-key-chain bgp-auth;
  authentication-algorithm md5;
}
}

user@R1# show routing-options
autonomous-system 65533;

user@R1# show security
authentication-key-chains {
  key-chain bgp-auth {
    tolerance 30;
    key 0 {
      secret
      "$9$ST6AREyK8RhXNdwaJn/Ct0IcykWWx9AyIMWdVgoJDqP5FCA0z3IEhcMWLxNbgJDiF6A";
      ## SECRET-DATA
      start-time "2011-6-23.20:19:33 -0700";
    }
    key 1 {
      secret "$9$UyD.59Cu0Ih9AylKW-dqmfT369CuRhSP5hrvMN-JGDiqfu0lleWpuh.";
      ## SECRET-DATA
      start-time "2012-6-23.20:19:33 -0700";
    }
  }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure for every BGP-enabled device in the network, using the appropriate interface names and addresses for each BGP-enabled device.

### Verification

Confirm that the configuration is working properly.

- [Verifying Authentication for the Neighbor on page 3575](#)
- [Verifying That Authorization Messages Are Sent on page 3576](#)
- [Checking Authentication Errors on page 3577](#)
- [Verifying the Operation of the Keychain on page 3577](#)

### Verifying Authentication for the Neighbor

**Purpose** Make sure that the **AuthKeyChain** option appears in the output of the **show bgp neighbor** command.

**Action** From operational mode, enter the **show bgp neighbor** command.

```

user@R1> show bgp neighbor
Peer: 172.16.2.1+179 AS 65530  Local: 172.16.2.2+1222 AS 65533
  Type: External  State: Established  Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None

```

```

Export: [ direct-lo0 ]
Options: <Preference PeerAS Refresh>
Options: <AuthKeyChain>
Authentication key is configured
Authentication key chain: jni
Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 172.16.2.1      Local ID: 10.255.124.35   Active Holdtime: 90
Keepalive Interval: 30      Peer index: 0
Local Interface: fe-0/0/1.0
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          2
  Received prefixes:        2
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 2    Sent 2    Checked 2
Input messages: Total 21    Updates 2    Refreshes 0    Octets 477
Output messages: Total 22    Updates 1    Refreshes 0    Octets 471
Output Queue[0]: 0

```

### Verifying That Authorization Messages Are Sent

**Purpose** Confirm that BGP has the enhanced authorization option.

**Action** From operational mode, enter the `monitor traffic interface fe-0/0/1` command.

```

user@R1> monitor traffic interface fe-0/0/1
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Listening on fe-0/0/1, capture size 96 bytes

13:08:00.618402 In arp who-has 172.16.2.66 tell 172.16.2.69
13:08:02.408249 Out IP 172.16.2.2.1122 > 172.16.2.1.646: P
1889289217:1889289235(18) ack 2215740969 win 58486 <nop,nop,timestamp 167557
1465469,nop,Enhanced Auth keyid 0 diglen 12 digest: fe3366001f45767165f17037>:
13:08:02.418396 In IP 172.16.2.1.646 > 172.16.2.2.1122: P 1:19(18) ack 18 win
57100 <nop,nop,timestamp 1466460 167557,nop,Enhanced Auth keyid 0 diglen 12
digest: a18c31eda1b14b2900921675>:
13:08:02.518146 Out IP 172.16.2.2.1122 > 172.16.2.1.646: . ack 19 win 58468
<nop,nop,timestamp 167568 1466460,nop,Enhanced Auth keyid 0 diglen 12 digest:
c3b6422eb6bd3fd9cf79742b>
13:08:28.199557 Out IP 172.16.2.2.nerv > 172.16.2.1.bgp: P
286842489:286842508(19) ack 931203976 win 57200 <nop,Enhanced Auth keyid 0
diglen 12 digest: fc0e42900a73736bcc07c1a4>: BGP, length: 19
13:08:28.209661 In IP 172.16.2.1.bgp > 172.16.2.2.nerv: P 1:20(19) ack 19 win
56835 <nop,Enhanced Auth keyid 0 diglen 12 digest: 0fc8578c489fabce63aeb2c3>:
BGP, length: 19
13:08:28.309525 Out IP 172.16.2.2.nerv > 172.16.2.1.bgp: . ack 20 win 57181
<nop,Enhanced Auth keyid 0 diglen 12 digest: ef03f282fb2ece0039491df8>
13:08:32.439708 Out IP 172.16.2.2.1122 > 172.16.2.1.646: P 54:72(18) ack 55 win
58432 <nop,nop,timestamp 170560 1468472,nop,Enhanced Auth keyid 0 diglen 12
digest: 76e0cf926f348b726c631944>:
13:08:32.449795 In IP 172.16.2.1.646 > 172.16.2.2.1122: P 55:73(18) ack 72 win
57046 <nop,nop,timestamp 1469463 170560,nop,Enhanced Auth keyid 0 diglen 12
digest: dae3eec390d18a114431f4d8>:
13:08:32.549726 Out IP 172.16.2.2.1122 > 172.16.2.1.646: . ack 73 win 58414
<nop,nop,timestamp 170571 1469463,nop,Enhanced Auth keyid 0 diglen 12 digest:

```

```

851df771aee2ea7a43a0c46c>
13:08:33.719880 In arp who-has 172.16.2.66 tell 172.16.2.69
^C
35 packets received by filter
0 packets dropped by kernel

```

### Checking Authentication Errors

**Purpose** Check the number of packets dropped by TCP because of authentication errors.

**Action** From operational mode, enter the **show system statistics tcp | match auth** command.

```

user@R1> show system statistics tcp | match auth
0 send packets dropped by TCP due to auth errors
58 rcv packets dropped by TCP due to auth errors

```

### Verifying the Operation of the Keychain

**Purpose** Check the number of packets dropped by TCP because of authentication errors.

**Action** From operational mode, enter the **show security keychain detail** command.

```

user@R1> show security keychain detail
keychain          Active-ID      Next-ID      Transition  Tolerance
                  Send Receive   Send Receive
bgp-auth          3      3      1      1      1d 23:58    30
Id 3, Algorithm hmac-md5, State send-receive, Option basic
Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
Id 1, Algorithm hmac-md5, State inactive, Option basic
Start-time Fri Aug 20 11:30:57 2010, Mode send-receive

```

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

## Examples: Configuring TCP and BGP Security

- [Understanding Security Options for BGP with TCP on page 3577](#)
- [Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers on page 3578](#)
- [Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List on page 3583](#)
- [Example: Limiting TCP Segment Size for BGP on page 3586](#)

### Understanding Security Options for BGP with TCP

Among routing protocols, BGP is unique in using TCP as its transport protocol. BGP peers are established by manual configuration between routing devices to create a TCP session on port 179. A BGP-enabled device periodically sends keepalive messages to maintain the connection.

Over time, BGP has become the dominant interdomain routing protocol on the Internet. However, it has limited guarantees of stability and security. Configuring security options for BGP must balance suitable security measures with acceptable costs. No one method

has emerged as superior to other methods. Each network administrator must configure security measures that meet the needs of the network being used.

For detailed information about the security issues associated with BGP's use of TCP as a transport protocol, see RFC 4272, *BGP Security Vulnerabilities Analysis*.

### Example: Configuring a Filter to Block TCP Access to a Port Except from Specified BGP Peers

This example shows how to configure a standard stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except from specified BGP peers.

- [Requirements on page 3578](#)
- [Overview on page 3578](#)
- [Configuration on page 3578](#)
- [Verification on page 3581](#)

#### Requirements

No special configuration beyond device initialization is required before you configure this example.


#### Overview

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except the specified BGP peers.

The stateless firewall filter **filter\_bgp179** matches all packets from the directly connected interfaces on Device A and Device B to the destination port number 179.

Figure 92 on page 3578 shows the topology used in this example. Device C attempts to make a TCP connection to Device E. Device E blocks the connection attempt. This example shows the configuration on Device E.

Figure 92: Typical Network with BGP Peer Sessions



#### Configuration

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device C    set interfaces ge-1/2/0 unit 10 description to-E
            set interfaces ge-1/2/0 unit 10 family inet address 10.10.10.10/30
            set protocols bgp group external-peers type external
            set protocols bgp group external-peers peer-as 17
            set protocols bgp group external-peers neighbor 10.10.10.9
            set routing-options autonomous-system 22

Device E    set interfaces ge-1/2/0 unit 0 description to-A
            set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30
            set interfaces ge-1/2/1 unit 5 description to-B
            set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30
            set interfaces ge-1/0/0 unit 9 description to-C
            set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30
            set interfaces lo0 unit 2 family inet filter input filter_bgp179
            set interfaces lo0 unit 2 family inet address 192.168.0.1/32
            set protocols bgp group external-peers type external
            set protocols bgp group external-peers peer-as 22
            set protocols bgp group external-peers neighbor 10.10.10.2
            set protocols bgp group external-peers neighbor 10.10.10.6
            set protocols bgp group external-peers neighbor 10.10.10.10
            set routing-options autonomous-system 17
            set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.2/32
            set firewall family inet filter filter_bgp179 term 1 from source-address 10.10.10.6/32
            set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
            set firewall family inet filter filter_bgp179 term 1 then accept
            set firewall family inet filter filter_bgp179 term 2 then reject

```

### Configuring Device E

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device E with a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requestors except specified BGP peers:

1. Configure the interfaces.
 

```

user@E# set interfaces ge-1/2/0 unit 0 description to-A
user@E# set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/30

user@E# set interfaces ge-1/2/1 unit 5 description to-B
user@E# set interfaces ge-1/2/1 unit 5 family inet address 10.10.10.5/30

user@E# set interfaces ge-1/0/0 unit 9 description to-C
user@E# set interfaces ge-1/0/0 unit 9 family inet address 10.10.10.9/30

```
2. Configure BGP.
 

```

[edit protocols bgp group external-peers]
user@E# set type external
user@E# set peer-as 22
user@E# set neighbor 10.10.10.2
user@E# set neighbor 10.10.10.6
user@E# set neighbor 10.10.10.10

```

3. Configure the autonomous system number.

```
[edit routing-options]
user@E# set autonomous-system 17
```

4. Define the filter term that accepts TCP connection attempts to port 179 from the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 1 from source-address 10.10.10.2/32
user@E# set term 1 from source-address 10.10.10.6/32
user@E# set term 1 from destination-port bgp
user@E# set term 1 then accept
```

5. Define the other filter term to reject packets from other sources.

```
[edit firewall family inet filter filter_bgp179]
user@E# set term 2 then reject
```

6. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 2 family inet]
user@E# set filter input filter_bgp179
user@E# set address 192.168.0.1/32
```

**Results** From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@E# show firewall
family inet {
  filter filter_bgp179 {
    term 1 {
      from {
        source-address {
          10.10.10.2/32;
          10.10.10.6/32;
        }
        destination-port bgp;
      }
      then accept;
    }
    term 2 {
      then {
        reject;
      }
    }
  }
}

user@E# show interfaces
lo0 {
  unit 2 {
    family inet {
      filter {
        input filter_bgp179;
```

```

        }
        address 192.168.0.1/32;
    }
}
ge-1/2/0 {
    unit 0 {
        description to-A;
        family inet {
            address 10.10.10.1/30;
        }
    }
}
ge-1/2/1 {
    unit 5 {
        description to-B;
        family inet {
            address 10.10.10.5/30;
        }
    }
}
ge-1/0/0 {
    unit 9 {
        description to-C;
        family inet {
            address 10.10.10.9/30;
        }
    }
}
}

user@E# show protocols
bgp {
    group external-peers {
        type external;
        peer-as 22;
        neighbor 10.10.10.2;
        neighbor 10.10.10.6;
        neighbor 10.10.10.10;
    }
}

user@E# show routing-options
autonomous-system 17;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Filter Is Configured on page 3582](#)
- [Verifying the TCP Connections on page 3582](#)
- [Monitoring Traffic on the Interfaces on page 3582](#)

*Verifying That the Filter Is Configured*

**Purpose** Make sure that the filter is listed in output of the **show firewall filter** command.

**Action** user@E> show firewall filter filter\_bgp179  
Filter: filter\_bgp179

*Verifying the TCP Connections*

**Purpose** Verify the TCP connections.

**Action** From operational mode, run the **show system connections extensive** command on Device C and Device E.

The output on Device C shows the attempt to establish a TCP connection. The output on Device E shows that connections are established with Device A and Device B only.

user@C> show system connections extensive | match 10.10.10

|      |   |   |                  |                 |          |
|------|---|---|------------------|-----------------|----------|
| tcp4 | 0 | 0 | 10.10.10.9.51872 | 10.10.10.10.179 | SYN_SENT |
|------|---|---|------------------|-----------------|----------|

user@E> show system connections extensive | match 10.10.10

|      |   |   |                  |                  |             |
|------|---|---|------------------|------------------|-------------|
| tcp4 | 0 | 0 | 10.10.10.5.179   | 10.10.10.6.62096 | ESTABLISHED |
| tcp4 | 0 | 0 | 10.10.10.6.62096 | 10.10.10.5.179   | ESTABLISHED |
| tcp4 | 0 | 0 | 10.10.10.1.179   | 10.10.10.2.61506 | ESTABLISHED |
| tcp4 | 0 | 0 | 10.10.10.2.61506 | 10.10.10.1.179   | ESTABLISHED |

*Monitoring Traffic on the Interfaces*

**Purpose** Use the **monitor traffic** command to compare the traffic on an interface that establishes a TCP connection with the traffic on an interface that does not establish a TCP connection.

**Action** From operational mode, run the **monitor traffic** command on the Device E interface to Device B and on the Device E interface to Device C. The following sample output verifies that in the first example, acknowledgment (**ack**) messages are received. In the second example, **ack** messages are not received.

```
user@E> monitor traffic size 1500 interface ge-1/2/1.5
19:02:49.700912 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P
3330573561:3330573580(19) ack 915601686 win 16384 <nop,nop,timestamp 1869518816
1869504850>: BGP, length: 19
19:02:49.801244 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 19 win 16384
<nop,nop,timestamp 1869518916 1869518816>
19:03:03.323018 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: P 1:20(19) ack 19 win
16384 <nop,nop,timestamp 1869532439 1869518816>: BGP, length: 19
19:03:03.422418 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: . ack 20 win 16384
<nop,nop,timestamp 1869532539 1869532439>
19:03:17.220162 Out IP 10.10.10.5.bgp > 10.10.10.6.62096: P 19:38(19) ack 20 win
16384 <nop,nop,timestamp 1869546338 1869532439>: BGP, length: 19
19:03:17.320501 In IP 10.10.10.6.62096 > 10.10.10.5.bgp: . ack 38 win 16384
<nop,nop,timestamp 1869546438 1869546338>
```

user@E> monitor traffic size 1500 interface ge-1/0/0.9

```
18:54:20.175471 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
```



```

win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869009240 0,sackOK,eol>
18:54:23.174422 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869012240 0,sackOK,eol>
18:54:26.374118 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,nop,wscale 0,nop,nop,timestamp 1869015440 0,sackOK,eol>
18:54:29.573799 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:32.773493 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>
18:54:35.973185 Out IP 10.10.10.9.61335 > 10.10.10.10.bgp: S 573929123:573929123(0)
win 16384 <mss 1460,sackOK,eol>

```

### Example: Configuring a Filter to Limit TCP Access to a Port Based On a Prefix List

This example shows how to configure a standard stateless firewall filter that limits certain TCP and Internet Control Message Protocol (ICMP) traffic destined for the Routing Engine by specifying a list of prefix sources that contain allowed BGP peers.

- [Requirements on page 3583](#)
- [Overview on page 3583](#)
- [Configuration on page 3583](#)
- [Verification on page 3585](#)

#### Requirements

No special configuration beyond device initialization is required before configuring this example.

#### Overview

In this example, you create a stateless firewall filter that blocks all TCP connection attempts to port 179 from all requesters except BGP peers that have a specified prefix.

A source prefix list, **plist\_bgp179**, is created that specifies the list of source prefixes that contain allowed BGP peers.

The stateless firewall filter **filter\_bgp179** matches all packets from the source prefix list **plist\_bgp179** to the destination port number 179.

#### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

- [Configure the Filter on page 3584](#)
- [Results on page 3584](#)

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set policy-options prefix-list plist_bgp179 apply-path "protocols bgp group <*> neighbor
<*>"
set firewall family inet filter filter_bgp179 term 1 from source-address 0.0.0.0/0

```

```
set firewall family inet filter filter_bgp179 term 1 from source-prefix-list plist_bgp179 except
set firewall family inet filter filter_bgp179 term 1 from destination-port bgp
set firewall family inet filter filter_bgp179 term 1 then reject
set firewall family inet filter filter_bgp179 term 2 then accept
set interfaces lo0 unit 0 family inet filter input filter_bgp179
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
```

### Configure the Filter

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the filter:

1. Expand the prefix list **bgp179** to include all prefixes pointed to by the BGP peer group defined by **protocols bgp group <\*> neighbor <\*>**.

```
[edit policy-options prefix-list plist_bgp179]
user@host# set apply-path "protocols bgp group <*> neighbor <*>"
```

2. Define the filter term that rejects TCP connection attempts to port 179 from all requesters except the specified BGP peers.

```
[edit firewall family inet filter filter_bgp179]
user@host# set term term1 from source-address 0.0.0.0/0
user@host# set term term1 from source-prefix-list bgp179 except
user@host# set term term1 from destination-port bgp
user@host# set term term1 then reject
```

3. Define the other filter term to accept all packets.

```
[edit firewall family inet filter filter_bgp179]
user@host# set term term2 then accept
```

4. Apply the firewall filter to the loopback interface.

```
[edit interfaces lo0 unit 0 family inet]
user@host# set filter input filter_bgp179
user@host# set address 127.0.0.1/32
```

### Results

From configuration mode, confirm your configuration by entering the **show firewall**, **show interfaces**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show firewall
family inet {
  filter filter_bgp179 {
    term 1 {
      from {
        source-address {
          0.0.0.0/0;
        }
        source-prefix-list {
          plist_bgp179 except;
        }
      }
    }
  }
}
```

```

    }
    destination-port bgp;
  }
  then {
    reject;
  }
}
term 2 {
  then {
    accept;
  }
}
}
}

user@host# show interfaces
lo0 {
  unit 0 {
    family inet {
      filter {
        input filter_bgp179;
      }
      address 127.0.0.1/32;
    }
  }
}

user@host# show policy-options
prefix-list plist_bgp179 {
  apply-path "protocols bgp group <*> neighbor <*>";
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Repeat the procedure, where appropriate, for every BGP-enabled device in the network, using the appropriate interface names and addresses for each BGP-enabled device.

### Verification

Confirm that the configuration is working properly.

### Displaying the Firewall Filter Applied to the Loopback Interface

**Purpose** Verify that the firewall filter **filter\_bgp179** is applied to the IPv4 input traffic at logical interface **lo0.0**.

**Action** Use the **show interfaces statistics** operational mode command for logical interface **lo0.0**, and include the **detail** option. Under the **Protocol inet** section of the command output section, the **Input Filters** field displays the name of the stateless firewall filter applied to the logical interface in the input direction:

```

[edit]
user@host> show interfaces statistics lo0.0 detail
Logical interface lo0.0 (Index 321) (SNMP ifIndex 16) (Generation 130)
Flags: SNMP-Traps Encapsulation: Unspecified
Traffic statistics:
Input bytes : 0

```

```
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: Unlimited, Generation: 145, Route table: 0
Flags: Sendbroadcast-pkt-to-re
Input Filters: filter_bgp179
Addresses, Flags: Primary
Destination: Unspecified, Local: 127.0.0.1, Broadcast: Unspecified,
Generation: 138
```

---

### Example: Limiting TCP Segment Size for BGP

This example shows how to avoid Internet Control Message Protocol (ICMP) vulnerability issues by limiting TCP segment size when you are using maximum transmission unit (MTU) discovery. Using MTU discovery on TCP paths is one method of avoiding BGP packet fragmentation.

- [Requirements on page 3586](#)
- [Overview on page 3586](#)
- [Configuration on page 3587](#)
- [Verification on page 3589](#)
- [Troubleshooting on page 3589](#)

#### **Requirements**

No special configuration beyond device initialization is required before you configure this example.

#### **Overview**

TCP negotiates a maximum segment size (MSS) value during session connection establishment between two peers. The MSS value negotiated is primarily based on the maximum transmission unit (MTU) of the interfaces to which the communicating peers are directly connected. However, due to variations in link MTU on the path taken by the TCP packets, some packets in the network that are well within the MSS value might be fragmented when the packet size exceeds the link's MTU.

To configure the TCP MSS value, include the **tcp-mss** statement with a segment size from 1 through 4096.

If the router receives a TCP packet with the SYN bit and the MSS option set, and the MSS option specified in the packet is larger than the MSS value specified by the **tcp-mss** statement, the router replaces the MSS value in the packet with the lower value specified by the **tcp-mss** statement.


The configured MSS value is used as the maximum segment size for the sender. The assumption is that the TCP MSS value used by the sender to communicate with the BGP neighbor is the same as the TCP MSS value that the sender can accept from the BGP neighbor. If the MSS value from the BGP neighbor is less than the MSS value configured, the MSS value from the BGP neighbor is used as the maximum segment size for the sender.

This feature is supported with TCP over IPv4 and TCP over IPv6.

### Topology Diagram

Figure 93 on page 3587 shows the topology used in this example.

**Figure 93: TCP Maximum Segment Size for BGP**



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
R0
set interfaces fe-1/2/0 unit 1 family inet address 1.1.0.1/30
set interfaces lo0 unit 1 family inet address 10.255.14.179/32
set protocols bgp group-int tcp-mss 2020
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.14.179
set protocols bgp group int mtu-discovery
set protocols bgp group int neighbor 10.255.71.24 tcp-mss 2000
set protocols bgp group int neighbor 10.255.14.177
set protocols bgp group int neighbor 10.0.14.4 tcp-mss 4000
set protocols ospf area 0.0.0.0 interface fe-1/2/0.1
set protocols ospf area 0.0.0.0 interface 10.255.14.179
set routing-options autonomous-system 65000
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R0:

1. Configure the interfaces.  

```
[edit interfaces]
user@R0# set fe-1/2/0 unit 1 family inet address 1.1.0.1/30
user@R0# set lo0 unit 1 family inet address 10.255.14.179/32
```
2. Configure an interior gateway protocol (IGP), OSPF in this example.  

```
[edit protocols ospf area 0.0.0.0]
```

```
user@R0# set interface fe-1/2/0.1
user@R0# set interface 10.255.14.179
```

3. Configure one or more BGP groups.

```
[edit protocols bgp group int]
user@R0# set type internal
user@R0# set local-address 10.255.14.179
```

4. Configure MTU discovery to prevent packet fragmentation.

```
[edit protocols bgp group int]
user@R0# set mtu-discovery
```

5. Configure the BGP neighbors, with the TCP MSS set globally for the group or specifically for the various neighbors.

```
[edit protocols bgo group int]
user@R0# set tcp-mss 2020
user@R0# set neighbor 10.255.14.177
user@R0# set neighbor 10.255.71.24 tcp-mss 2000
user@R0# set neighbor 10.0.14.4 tcp-mss 4000
```



**NOTE:** The TCP MSS neighbor setting overrides the group setting.

---

6. Configure the local autonomous system.

```
[edit routing-options]
user@R0# set autonomous-system 65000
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 1.1.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 10.255.14.179/32;
    }
  }
}

user@R0# show protocols
bgp {
  group int {
    type internal;
    local-address 10.255.14.179;
```

```

mtu-discovery;
tcp-mss 2020;
neighbor 10.255.71.24 {
    tcp-mss 2000;
}
neighbor 10.255.14.177;
neighbor 10.0.14.4 {
    tcp-mss 4000;
}
}
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/0.1;
        interface 10.255.14.179;
    }
}
}

user@R0# show routing-options
autonomous-system 65000;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

To confirm that the configuration is working properly, run the following commands:

- **show system connections extensive | find <neighbor-address>**, to check the negotiated TCP MSS value.
- **monitor traffic interface**, to monitor BGP traffic and to make sure that the configured TCP MSS value is used as the MSS option in the TCP SYN packet.

### Troubleshooting

- [MSS Calculation with MTU Discovery on page 3589](#)

### MSS Calculation with MTU Discovery

**Problem** Consider an example in which two routing devices (R1 and R2) have an internal BGP (IBGP) connection. On both of the routers, the connected interfaces have 4034 as the IPv4 MTU.

```

user@R1# show protocols bgp | display set
[edit]
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 45.45.45.2
set protocols bgp group ibgp mtu-discovery
set protocols bgp group ibgp neighbor 45.45.45.1

```

```

user@R1# run show interfaces xe-0/0/3 extensive | match mtu

```

```

Link-level type: Ethernet, MTU: 4048, LAN-PHY mode, Speed: 10Gbps,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Protocol inet, MTU: 4034, Generation: 180, Route table: 0
Protocol multiservice, MTU: Unlimited, Generation: 181, Route table: 0

```

In the following packet capture on Device R1, the negotiated MSS is 3994. In the **show system connections extensive** information for MSS, it is set to 2048.

```
05:50:01.575218 Out
  Juniper PCAP Flags [Ext], PCAP Extension(s) total length 16
    Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
    Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
    Device Interface Index Extension TLV #1, length 2, value: 137
    Logical Interface Index Extension TLV #4, length 4, value: 69
----original packet-----
00:21:59:e1:e8:03 > 00:19:e2:20:79:01, ethertype IPv4 (0x0800), length
78: (tos 0xc0, ttl 64, id 53193, offset 0, flags [DF], proto: TCP (6), length:
64) 45.45.45.2.62840 > 45.45.45.1.bgp: S 2939345813:2939345813(0) win 16384 **mss
3994,nop,wscale 0,nop,nop,timestamp 70559970 0,sackOK,eol>
05:50:01.575875 In
  Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 16
    Device Media Type Extension TLV #3, length 1, value: Ethernet (1)
    Logical Interface Encapsulation Extension TLV #6, length 1, value:
Ethernet (14)
    Device Interface Index Extension TLV #1, length 2, value: 137
    Logical Interface Index Extension TLV #4, length 4, value: 69
----original packet-----
PFE proto 2 (ipv4): (tos 0xc0, ttl 255, id 37709, offset 0, flags [DF], proto:
TCP (6), length: 64) 45.45.45.1.bgp > 45.45.45.2.62840: S 2634967984:2634967984(0)
ack 2939345814 win 16384 **mss 3994,nop,wscale 0,nop,nop,timestamp 174167273
70559970,sackOK,eol>
```

user@R1# run show system connections extensive | find 45.45

```
tcp4      0      0 45.45.45.2.62840          45.45.45.1.179
ESTABLISHED
  sndsbcc:      0 sndsbmbcnt:      0 sndsbmbmax:    131072
  sndsblowat:   2048 sndsbhiwat:    16384
  rcvsbcc:      0 rcvsbmbcnt:      0 rcvsbmbmax:    131072
  rcvsblowat:   1 rcvsbhiwat:    16384
  proc id:     19725 proc name:    rpd
    iss: 2939345813    sndup: 2939345972
  snduna: 2939345991    sndnxt: 2939345991    sndwnd:    16384
  sndmax: 2939345991    sndcwnd:    10240 sndssthresh: 1073725440
  irs: 2634967984    rcvup: 2634968162
  rcvnxt: 2634968162    rcvadv: 2634984546    rcvwnd:    16384
  rtt: 0    srtt: 1538    rttv: 1040
  rxcur: 1200    rxtshift: 0    rtseq: 2939345972
  rttmin: 1000    mss: 2048
```

**Solution** This is expected behavior with Junos OS. The MSS value is equal to the MTU value minus the IP or IPv6 and TCP headers. This means that the MSS value is generally 40 bytes less than the MTU (for IPv4) and 60 bytes less than the MTU (for IPv6). This value is negotiated between the peers. In this example, it is  $4034 - 40 = 3994$ . Junos OS then rounds this value to a multiple of 2 KB. The value is  $3994 / 2048 * 2048 = 2048$ . So it is not necessary to see same MSS value with in the **show system connections** output.

$3994 / 2048 = 1.95$

1.95 is rounded to 1.

$1 * 2048 = 2048$



- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3261](#)
  - [BGP Configuration Overview](#)

## BGP Flap Configuration

- [Example: Preventing BGP Session Resets on page 3591](#)
- [Examples: Configuring BGP Flap Damping on page 3599](#)

### Example: Preventing BGP Session Resets

- [Understanding BGP Session Resets on page 3591](#)
- [Example: Preventing BGP Session Flaps When VPN Families Are Configured on page 3591](#)

#### Understanding BGP Session Resets

Certain configuration actions and events cause BGP sessions to be reset (dropped and then reestablished).

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same autonomous system (AS) number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an internal BGP (IBGP) group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.
- Changing configuration statements that affect BGP peers, such as renaming a BGP group, resets the BGP sessions.
- If you change the address family specified in the **[edit protocols bgp family]** hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

#### Example: Preventing BGP Session Flaps When VPN Families Are Configured

This example shows a workaround for a known issue in which BGP sessions sometimes go down and then come back up (in other words, flap) when virtual private network (VPN) families are configured. If any VPN family (for example, **inet-vpn**, **inet6-vpn**, **inet-mpvn**, **inet-mdt**, **inet6-mpvn**, **l2vpn**, **iso-vpn**, and so on) is configured on a BGP master instance, a flap of either a route reflector (RR) internal BGP (IBGP) session or an external BGP (EBGP) session causes flaps of other BGP sessions configured with the same VPN family.

- [Requirements on page 3592](#)
- [Overview on page 3593](#)
- [Configuration on page 3595](#)
- [Verification on page 3598](#)

### ***Requirements***

Before you begin:

- Configure router interfaces.
- Configure an interior gateway protocol (IGP).
- Configure BGP.
- Configure VPNs.

## *Overview*

When a router or switch is configured as either a route reflector (RR) or an AS boundary router (an external BGP peer) and a VPN family (for example, the **family inet-vpn unicast** statement) is configured, a flap of either the RR IBGP session or the EBGP session causes flaps of all other BGP sessions that are configured with the **family inet-vpn unicast** statement. This example shows how to prevent these unnecessary session flaps.

The reason for the flapping behavior is related to BGP operation in Junos OS when originating VPN routes.

BGP has the following two modes of operation with respect to originating VPN routes:

- If BGP does not need to propagate VPN routes because the session has no EBGP peer and no RR clients, BGP exports VPN routes directly from the *instance.inet.0* routing table to other PE routers. This behavior is efficient in that it avoids the creation of two copies of many routes (one in the *instance.inet.0* table and one in the *bgp.l3vpn.0* table).
- If BGP does need to propagate VPN routes because the session has an EBGP peer or RR clients, BGP first exports the VPN routes from the *instance.inet.0* table to the *bgp.l3vpn.0* table. Then BGP exports the routes to other PE routers. In this scenario, two copies of the route are needed to enable best-route selection. A PE router might receive the same VPN route from a CE device and also from an RR client or EBGP peer.



.....

**NOTE:** The route export is not performed if the route in *instance.inet.0* is a secondary route. In Junos OS, a route is only exported one time from one routing table as a primary route to another routing table as a secondary route. Because the route in *instance.inet.0* is already a secondary route, it is not allowed to be moved again to the *bgp.l3vpn.0* table, as needed to be advertised. The route does not reach the *bgp.l3vpn.0* table and thus is not advertised. One workaround is to send the routes that should be advertised to *inet.0* so that they are advertised.

.....

When, because of a configuration change, BGP transitions from needing two copies of a route to not needing two copies of a route (or the reverse), all sessions over which VPN routes are exchanged go down and then come back up. Although this example focuses on the **family inet-vpn unicast** statement, the concept applies to all VPN network layer reachability information (NLRI) families. This issue impacts logical systems as well. All BGP sessions in the master instance related to the VPN NLRI family are brought down to implement the table advertisement change for the VPN NLRI family. Changing an RR to a non-RR or the reverse (by adding or removing the **cluster** statement) causes the table advertisement change. Also, configuring the first EBGP session or removing the EBGP session from the configuration in the master instance for a VPN NLRI family causes the table advertisement change.

The way to prevent these unnecessary session flaps is to configure an extra RR client or EBGP session as a passive session with a neighbor address that does not exist. This example focuses on the EBGP case, but the same workaround works for the RR case.

When a session is passive, the routing device does not send Open requests to a peer. Once you configure the routing device to be passive, the routing device does not originate the TCP connection. However, when the routing device receives a connection from the peer and an Open message, it replies with another BGP Open message. Each routing device declares its own capabilities.

Figure 94 on page 3595 shows the topology for the EBGp case. Router R1 has an IBGP session with Routers R2 and R3 and an EBGp session with Router R4. All sessions have the **family inet-vpn unicast** statement configured. If the R1-R4 EBGp session flaps, the R1-R2 and R1-R3 BGP sessions flap also.

Figure 94: Topology for the EBGp Case





Figure 95 on page 3595 shows the topology for the RR case. Router R1 is the RR, and Router R3 is the client. Router R1 has IBGP sessions with Routers R2 and R3. All sessions have the **family inet-vpn unicast** statement configured. If the R1-R3 session flaps, the R1-R2 and R1-R4 sessions flap also.

Figure 95: Topology for the RR Case



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols bgp family inet-vpn unicast
set protocols bgp family l2vpn signaling
set protocols bgp group R1-R4 type external
set protocols bgp group R1-R4 local-address 4.4.4.2
set protocols bgp group R1-R4 neighbor 4.4.4.1 peer-as 200
set protocols bgp group R1-R2-R3 type internal
set protocols bgp group R1-R2-R3 log-updown
set protocols bgp group R1-R2-R3 local-address 15.15.15.15
set protocols bgp group R1-R2-R3 neighbor 12.12.12.12
set protocols bgp group R1-R2-R3 neighbor 13.13.13.13
set protocols bgp group Fake type external
set protocols bgp group Fake passive
set protocols bgp group Fake neighbor 100.100.100.100 peer-as 500
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the EBGp scenario:

1. Configure one or more VPN families.

```
[edit protocols bgp]
user@R1# set family inet-vpn unicast
user@R1# set family l2vpn signaling
```

2. Configure the EBGp session.

```
[edit protocols bgp]
user@R1# set group R1-R4 type external
user@R1# set group R1-R4 local-address 4.4.4.2
user@R1# set group R1-R4 neighbor 4.4.4.1 peer-as 200
```

3. Configure the IBGP sessions.

```
[edit protocols bgp]
user@R1# set group R1-R2-R3 type internal
user@R1# set group R1-R2-R3 local-address 15.15.15.15
user@R1# set group R1-R2-R3 neighbor 12.12.12.12
user@R1# set group R1-R2-R3 neighbor 13.13.13.13
```

4. (Optional) Configure BGP so that it generates a **syslog** message whenever a BGP peer makes a state transition.

```
[edit protocols bgp]
user@R1# set group R1-R2-R3 log-updown
```

Enabling the **log-updown** statement causes BGP state transitions to be logged at **warning** level.

**Step-by-Step Procedure** To verify that unnecessary session flaps are occurring:

1. Run the **show bgp summary** command to verify that the sessions have been established.

```
user@R1> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table      Tot Paths Act Paths Suppressed History Damp State Pending
```

```

bgp.13vpn.0 0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0
inet.0      0      0      0      0      0      0
Peer        AS  InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1     200 6      5      0      0      1:08 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 3      7      0      0      1:18 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
13.13.13.13 100 3      6      0      0      1:14 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0

```

2. Deactivate the EBGp session.

```

user@R1# deactivate group R1-R4
user@R1# commit

```

```

Mar 10 18:27:40 R1: rpd[1464]: bgp_peer_delete:6589: NOTIFICATION sent to 4.4.4.1 (External AS 200): code
 6 (Cease) subcode 3 (Peer Unconfigured), Reason: Peer Deletion
Mar 10 18:27:40 R1: rpd[1464]: bgp_adv_main_update:7253: NOTIFICATION sent to 12.12.12.12 (Internal AS
100): code 6 (Cease) subcode 6 (Other Configuration Change), Reason: Configuration change - VPN table
advertise
Mar 10 18:27:40 R1: rpd[1464]: bgp_adv_main_update:7253: NOTIFICATION sent to 13.13.13.13 (Internal AS
100): code 6 (Cease) subcode 6 (Other Configuration Change), Reason: Configuration change - VPN table
advertise

```

3. Run the **show bgp summary** command to view the session flaps.

```

user@R1> show bgp summary
Groups: 1 Peers: 2 Down peers: 2
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0
inet.0     0      0      0      0      0      0
Peer       AS   InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 4      9      0      1      19 Active
13.13.13.13 100 4      8      0      1      19 Active

user@R1> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0      0      0      0      0      0
bgp.12vpn.0 0      0      0      0      0      0
inet.0     0      0      0      0      0      0
Peer       AS   InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 2      3      0      1      0 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
13.13.13.13 100 2      3      0      1      0 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To prevent unnecessary BGP session flaps:

1. Add a passive EBGP session with a neighbor address that does not exist in the peer autonomous system (AS).

```
[edit protocols bgp]
user@R1# set group Fake type external
user@R1# set group Fake passive
user@R1# set neighbor 100.100.100.100 peer-as 500
```

2. Run the **show bgp summary** command to verify that the real sessions have been established and the passive session is idle.

```
user@R1> show bgp summary
Groups: 3 Peers: 4 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0 0 0 0 0 0
bgp.12vpn.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1 200 9500 9439 0 0 2d 23:14:23 Estab1
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 10309 10239 0 0 3d 5:17:49 Estab1
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10306 10241 0 0 3d 5:18:25 Estab1
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 0 2d 23:38:52 Idle
```

### Verification

Confirm that the configuration is working properly.

- [Bringing Down the EBGP Session on page 3598](#)
- [Verifying That the IBGP Sessions Remain Up on page 3598](#)

### Bringing Down the EBGP Session

**Purpose** Try to cause the flap issue after the workaround is configured.

**Action** user@R1# deactivate group R1-R4  
user@R1# commit

### Verifying That the IBGP Sessions Remain Up

**Purpose** Make sure that the IBGP sessions do not flap after the EBGP session is deactivated.



```

Action user@R1> show bgp summary
Groups: 2 Peers: 3 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0 0 0 0 0 0
bgp.12vpn.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
12.12.12.12 100 10312 10242 0 0 3d 5:19:01 Establ
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10309 10244 0 0 3d 5:19:37 Establ
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 0 2d 23:40:04 Idle

user@R1> show bgp summary
Groups: 3 Peers: 4 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 0 0 0 0 0 0
bgp.12vpn.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
4.4.4.1 200 5 4 0 0 28 Establ
bgp.13vpn.0: 0/0/0/0
bgp.12vpn.0: 0/0/0/0
12.12.12.12 100 10314 10244 0 0 3d 5:19:55 Establ
bgp.13vpn.0: 0/0/0/0
13.13.13.13 100 10311 10246 0 0 3d 5:20:31 Establ
bgp.13vpn.0: 0/0/0/0
100.100.100.100 500 0 0 0 0 2d 23:40:58 Idle

```

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3261](#)
  - [BGP Configuration Overview](#)

## Examples: Configuring BGP Flap Damping

- [Understanding BGP Route Flap Damping Parameters on page 3599](#)
- [Example: Configuring BGP Route Flap Damping Parameters on page 3600](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family on page 3609](#)

### Understanding BGP Route Flap Damping Parameters

BGP *route flapping* describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. BGP *flap damping* is a method of reducing the number of update messages sent between BGP peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Marking routes in this way leads to some delay, or *suppression*, in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (routes in different ASs). You can also apply flap damping within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do

not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level, which is supported in Junos OS Release 12.2 and later. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

By default, route flap damping is not enabled. Damping is applied to external peers and to peers at confederation boundaries.

When you enable damping, default parameters are applied, as summarized in [Table 297 on page 3600](#).

**Table 297: Damping Parameters**

| Damping Parameter                  | Description   | Default Value | Possible Values  |
|------------------------------------|---|---------------|------------------|
| <b>half-life <i>minutes</i></b>    | Decay half-life—Number of minutes after which an arbitrary value is halved if a route stays stable.                     | 15 (minutes)  | 1 through 45     |
| <b>max-suppress <i>minutes</i></b> | Maximum hold-down time for a route, in minutes.   | 60 (minutes)  | 1 through 720    |
| <b>reuse</b>                       | Reuse threshold—Arbitrary value below which a suppressed route can be used again.                                       | 750           | 1 through 20,000 |
| <b>suppress</b>                    | Cutoff (suppression) threshold—Arbitrary value above which a route can no longer be used or included in advertisements. | 3000          | 1 through 20,000 |

To change the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the damping action. For the damping routing policy to work, you also must enable BGP route flap damping.

#### **Example: Configuring BGP Route Flap Damping Parameters**

This example shows how to configure damping parameters.

- [Requirements on page 3600](#)
- [Overview on page 3600](#)
- [Configuration on page 3601](#)
- [Verification on page 3605](#)

#### **Requirements**

Before you begin, configure router interfaces and configure routing protocols.

#### **Overview**

This example has three routing devices. Device R2 has external BGP (EBGP) connections with Device R1 and Device R3.

Device R1 and Device R3 have some static routes configured for testing purposes, and these static routes are advertised through BGP to Device R2.


Device R2 damps routes received from Device R1 and Device R3 according to these criteria:

- Damp all prefixes with a mask length equal to or greater than 17 more aggressively than routes with a mask length between 9 and 16.
- Damp routes with a mask length between 0 and 8, inclusive, less than routes with a mask length greater than 8.
- Do not damp the 10.128.0.0/9 prefix at all.

The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only the active routes are exported from the routing table.

Figure 96 on page 3601 shows the sample network.

**Figure 96: BGP Flap Damping Topology**



"CLI Quick Configuration" on page 3601 shows the configuration for all of the devices in Figure 96 on page 3601.

The section "Step-by-Step Procedure" on page 3602 describes the steps on Device R2.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device R1  set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
            set interfaces lo0 unit 0 family inet address 192.168.0.1/32
            set protocols bgp group ext type external
            set protocols bgp group ext export send-direct-and-static
            set protocols bgp group ext peer-as 200
            set protocols bgp group ext neighbor 10.0.0.2
            set policy-options policy-statement send-direct-and-static term 1 from protocol direct
            set policy-options policy-statement send-direct-and-static term 1 from protocol static
            set policy-options policy-statement send-direct-and-static term 1 then accept
            set routing-options static route 172.16.0.0/16 reject
            set routing-options static route 172.16.128.0/17 reject
            set routing-options static route 172.16.192.0/20 reject
            set routing-options static route 10.0.0.0/9 reject
            set routing-options static route 224.0.0.0/7 reject
            set routing-options static route 10.224.0.0/11 reject
            set routing-options static route 0.0.0.0/0 reject
            set routing-options autonomous-system 100

Device R2  set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
            set interfaces fe-1/2/1 unit 0 family inet address 10.1.0.1/30
            set interfaces lo0 unit 0 family inet address 192.168.0.2/32
  
```

```
set protocols bgp damping
set protocols bgp group ext type external
set protocols bgp group ext import damp
set protocols bgp group ext export send-direct
set protocols bgp group ext neighbor 10.0.0.1 peer-as 100
set protocols bgp group ext neighbor 10.1.0.2 peer-as 300
set policy-options policy-statement damp term 1 from route-filter 10.128.0.0/9 exact
  damping dry
set policy-options policy-statement damp term 1 from route-filter 0.0.0.0/0
  prefix-length-range /0-/8 damping timid
set policy-options policy-statement damp term 1 from route-filter 0.0.0.0/0
  prefix-length-range /17-/32 damping aggressive
set policy-options policy-statement send-direct term 1 from protocol direct
set policy-options policy-statement send-direct term 1 then accept
set policy-options damping aggressive half-life 30
set policy-options damping aggressive suppress 2500
set policy-options damping timid half-life 5
set policy-options damping dry disable
set routing-options autonomous-system 200
```

Device R3

```
set interfaces fe-1/2/1 unit 0 family inet address 10.1.0.2/30
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set protocols bgp group ext type external
set protocols bgp group ext export send-direct-and-static
set protocols bgp group ext peer-as 200
set protocols bgp group ext neighbor 10.1.0.1
set policy-options policy-statement send-direct-and-static term 1 from protocol direct
set policy-options policy-statement send-direct-and-static term 1 from protocol static
set policy-options policy-statement send-direct-and-static term 1 then accept
set routing-options static route 10.128.0.0/9 reject
set routing-options autonomous-system 300
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure damping parameters:

1. Configure the interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 0 family inet address 10.1.0.1/30

user@R2# set lo0 unit 0 family inet address 192.168.0.2/32
```

2. Configure the BGP neighbors.

```
[edit protocols bgp group ext]
user@R2# set type external
user@R2# set neighbor 10.0.0.1 peer-as 100
user@R2# set neighbor 10.1.0.2 peer-as 300
```

3. Create and configure the damping parameter groups.

```
[edit policy-options]
user@R2# set damping aggressive half-life 30
user@R2# set damping aggressive suppress 2500
user@R2# set damping timid half-life 5
user@R2# set damping dry disable
```

4. Configure the damping policy.

```
[edit policy-options policy-statement damp term 1]
user@R2# set from route-filter 10.128.0.0/9 exact damping dry
user@R2# set from route-filter 0.0.0.0/0 prefix-length-range /0-/8 damping timid
user@R2# set from route-filter 0.0.0.0/0 prefix-length-range /17-/32 damping
    aggressive
```

5. Enable damping for BGP.

```
[edit protocols bgp]
user@R2# set damping
```

6. Apply the policy as an import policy for the BGP neighbor.

```
[edit protocols bgp group ext]
user@R2# set import damp
```



**NOTE:** You can refer to the same routing policy one or more times in the same or different import statements.

7. Configure an export policy.

```
[edit policy-options policy-statement send-direct term 1]
user@R2# set from protocol direct
user@R2# set then accept
```

8. Apply the export policy.

```
[edit protocols bgp group ext]
user@R2# set export send-direct
```

9. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@R2# set autonomous-system 200
```

**Results** From configuration mode, confirm your configuration by issuing the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
```

```
    }
  }
  fe-1/2/1 {
    unit 0 {
      family inet {
        address 10.1.0.1/30;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.0.2/32;
      }
    }
  }
}

user@R2# show protocols
bgp {
  damping;
  group ext {
    type external;
    import damp;
    export send-direct;
    neighbor 10.0.0.1 {
      peer-as 100;
    }
    neighbor 10.1.0.2 {
      peer-as 300;
    }
  }
}

user@R2# show policy-options
policy-statement damp {
  term 1 {
    from {
      route-filter 10.128.0.0/9 exact damping dry;
      route-filter 0.0.0.0/0 prefix-length-range /0-/8 damping timid;
      route-filter 0.0.0.0/0 prefix-length-range /17-/32 damping aggressive;
    }
  }
}
policy-statement send-direct {
  term 1 {
    from protocol direct;
    then accept;
  }
}
damping aggressive {
  half-life 30;
  suppress 2500;
}
damping timid {
  half-life 5;
}
damping dry {
```

```

    disable;
}

user@R2# show routing-options
autonomous-system 200;

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Causing Some Routes to Flap on page 3605](#)
- [Checking the Route Flaps on page 3605](#)
- [Verifying Route Flap Damping on page 3606](#)
- [Displaying the Details of a Damped Route on page 3607](#)
- [Verifying That Default Damping Parameters Are in Effect on page 3607](#)
- [Filtering the Damping Information on page 3608](#)

### Causing Some Routes to Flap

**Purpose** To verify your route flap damping policy, some routes must flap. Having a live Internet feed almost guarantees that a certain number of route flaps will be present. If you have control over a remote system that is advertising the routes, you can modify the advertising router's policy to effect the advertisement and withdrawal of all routes or of a given prefix. In a test environment, you can cause routes to flap by clearing the BGP neighbors or by restarting the routing process on the BGP neighbors, as shown here.

**Action** From operational mode on Device R1 and Device R3, enter the **restart routing** command.



**CAUTION:** Use this command cautiously in a production network.

```
user@R1> restart routing
```

```
R1 started, pid 10474
```

```
user@R3> restart routing
```

```
R3 started, pid 10478
```

**Meaning** On Device R2, all of the routes from the neighbors are withdrawn and re-advertised.

### Checking the Route Flaps

**Purpose** View the number of neighbor flaps.

**Action** From operational mode, enter the **show bgp summary** command.

```
user@R2> show bgp summary
```

```

Groups: 1 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State  Pending
inet.0
          12         1         11         0         11         0
Peer      AS      InPkt   OutPkt   OutQ   Flaps  Last  Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1   100        10       10       0       4      2:50
0/9/0/9    0/0/0/0
10.1.0.2   300        10       10       0       4      2:53
1/3/1/2    0/0/0/0

```

**Meaning** This output was captured after the routing process was restarted on Device R2's neighbors four times.

### *Verifying Route Flap Damping*

**Purpose** Verify that routes are being hidden due to damping.

**Action** From operational mode, enter the **show route damping suppressed** command.

```
user@R2> show route damping suppressed
```

```
inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

0.0.0.0/0      [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.0.0.0/9     [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.0.0.0/30    [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
10.1.0.0/30    [BGP ] 00:00:15, localpref 100
                AS path: 300 I, validation-state: unverified
                > to 10.1.0.2 via fe-1/2/1.0
10.224.0.0/11  [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.0.0/16  [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.128.0/17 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
172.16.192.0/20 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
192.168.0.1/32 [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0
192.168.0.3/32 [BGP ] 00:00:15, localpref 100
                AS path: 300 I, validation-state: unverified
                > to 10.1.0.2 via fe-1/2/1.0
224.0.0.0/7    [BGP ] 00:00:12, localpref 100
                AS path: 100 I, validation-state: unverified
                > to 10.0.0.1 via fe-1/2/0.0

```



**Meaning** The output shows some routing instability. Eleven routes are hidden due to damping.

### *Displaying the Details of a Damped Route*

**Purpose** Display the details of damped routes.

**Action** From operational mode, enter the **show route damping suppressed 172.16.192.0/20 detail** command.

```
user@R2> show route damping suppressed 172.16.192.0/20 detail

inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
172.16.192.0/20 (1 entry, 0 announced)
    BGP                /-101
        Next hop type: Router, Next hop index: 758
        Address: 0x9414484
        Next-hop reference count: 9
        Source: 10.0.0.1
        Next hop: 10.0.0.1 via fe-1/2/0.0, selected
        Session Id: 0x100201
        State: <Hidden Ext>
        Local AS: 200 Peer AS: 100
        Age: 52
        Validation State: unverified
        Task: BGP_100.10.0.0.1+55922
        AS path: 100 I
        Localpref: 100
        Router ID: 192.168.0.1
        Merit (last update/now): 4278/4196
        damping-parameters: aggressive
        Last update: 00:00:52 First update: 01:01:55
        Flaps: 8
        Suppressed. Reusable in: 01:14:40
        Preference will be: 170
```

**Meaning** This output indicates that the displayed route has a mask length that is equal to or greater than /17, and confirms that it has been correctly mapped to the aggressive damping profile. You can also see the route's current (and last) figure of merit value, and when the route is expected to become active if it remains stable.

### *Verifying That Default Damping Parameters Are in Effect*

**Purpose** Locating a damped route with a /16 mask confirms that the default parameters are in effect.

**Action** From operational mode, enter the **show route damping suppressed detail | match 0/16** command.

```
user@R2> show route damping suppressed detail | match 0/16

172.16.0.0/16 (1 entry, 0 announced)

user@R2> show route damping suppressed 172.16.0.0/16 detail

inet.0: 15 destinations, 17 routes (6 active, 0 holddown, 11 hidden)
172.16.0.0/16 (1 entry, 0 announced)
```

```
BGP                               /-101
Next hop type: Router, Next hop index: 758
Address: 0x9414484
Next-hop reference count: 9
Source: 10.0.0.1
Next hop: 10.0.0.1 via fe-1/2/0.0, selected
Session Id: 0x100201
State: <Hidden Ext>
Local AS: 200 Peer AS: 100
Age: 1:58
Validation State: unverified
Task: BGP_100.10.0.0.1+55922
AS path: 100 I
Localpref: 100
Router ID: 192.168.0.1
Merit (last update/now): 3486/3202
Default damping parameters used
Last update: 00:01:58 First update: 01:03:01
Flaps: 8
Suppressed. Reusable in: 00:31:40
Preference will be: 170
```

**Meaning** Routes with a /16 mask are not impacted by the custom damping rules. Therefore, the default damping rules are in effect.

To repeat, the custom rules are as follows:

- Damp all prefixes with a mask length equal to or greater than 17 more aggressively than routes with a mask length between 9 and 16.
- Damp routes with a mask length between 0 and 8, inclusive, less than routes with a mask length greater than 8.
- Do not damp the 10.128.0.0/9 prefix at all.

#### *Filtering the Damping Information*

**Purpose** Use OR groupings or cascaded piping to simplify the determination of what damping profile is being used for routes with a given mask length.

**Action** From operational mode, enter the **show route damping suppressed** command.

```
user@R2> show route damping suppressed detail | match "0 announced | damp"
```

```
0.0.0.0/0 (1 entry, 0 announced)
    damping-parameters: timid
10.0.0.0/9 (1 entry, 0 announced)
    Default damping parameters used
    damping-parameters: aggressive
    damping-parameters: aggressive
10.224.0.0/11 (1 entry, 0 announced)
    Default damping parameters used
172.16.0.0/16 (1 entry, 0 announced)
    Default damping parameters used
172.16.128.0/17 (1 entry, 0 announced)
    damping-parameters: aggressive
172.16.192.0/20 (1 entry, 0 announced)
    damping-parameters: aggressive
192.168.0.1/32 (1 entry, 0 announced)
```

```

damping-parameters: aggressive
192.168.0.3/32 (1 entry, 0 announced)
damping-parameters: aggressive
224.0.0.0/7 (1 entry, 0 announced)
damping-parameters: timid

```

**Meaning** When you are satisfied that your EBGp routes are correctly associated with a damping profile, you can issue the **clear bgp damping** operational mode command to restore an active status to your damped routes, which will return your connectivity to normal operation.

### Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family

This example shows how to configure an multiprotocol BGP multicast VPN (also called Next-Generation MVPN) with BGP route flap damping.

- [Requirements on page 3609](#)
- [Overview on page 3609](#)
- [Configuration on page 3610](#)
- [Verification on page 3617](#)

#### Requirements

This example uses Junos OS Release 12.2. BGP route flap damping support for MBGP MVPN, specifically, and on an address family basis, in general, is introduced in Junos OS Release 12.2.


#### Overview

BGP route flap damping helps to diminish route instability caused by routes being repeatedly withdrawn and readvertised when a link is intermittently failing.

This example uses the default damping parameters and demonstrates an MBGP MVPN scenario with three provider edge (PE) routing devices, three customer edge (CE) routing devices, and one provider (P) routing device.

[Figure 97 on page 3609](#) shows the topology used in this example.

**Figure 97: MBGP MVPN with BGP Route Flap Damping**



On PE Device R4, BGP route flap damping is configured for address family **inet-mvpn**. A routing policy called **dampPolicy** uses the **nlri-route-type** match condition to damp only MVPN route types 3, 4, and 5. All other MVPN route types are not damped.

This example shows the full configuration on all devices in the “[CLI Quick Configuration](#)” on page 3610 section. The “[Configuring Device R4](#)” on page 3613 section shows the step-by-step configuration for PE Device R4.

### *Configuration*

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1    set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30
              set interfaces ge-1/2/0 unit 1 family mpls
              set interfaces lo0 unit 1 family inet address 1.1.1.1/32
              set protocols ospf area 0.0.0.0 interface lo0.1 passive
              set protocols ospf area 0.0.0.0 interface ge-1/2/0.1
              set protocols pim rp static address 100.1.1.2
              set protocols pim interface all
              set routing-options router-id 1.1.1.1

Device R2    set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30
              set interfaces ge-1/2/0 unit 2 family mpls
              set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30
              set interfaces ge-1/2/1 unit 5 family mpls
              set interfaces vt-1/2/0 unit 2 family inet
              set interfaces lo0 unit 2 family inet address 1.1.1.2/32
              set interfaces lo0 unit 102 family inet address 100.1.1.2/32
              set protocols mpls interface ge-1/2/1.5
              set protocols bgp group ibgp type internal
              set protocols bgp group ibgp local-address 1.1.1.2
              set protocols bgp group ibgp family inet-vpn any
              set protocols bgp group ibgp family inet-mvpn signaling
              set protocols bgp group ibgp neighbor 1.1.1.4
              set protocols bgp group ibgp neighbor 1.1.1.5
              set protocols ospf area 0.0.0.0 interface lo0.2 passive
              set protocols ospf area 0.0.0.0 interface ge-1/2/1.5
              set protocols ldp interface ge-1/2/1.5
              set protocols ldp p2mp
              set policy-options policy-statement parent_vpn_routes from protocol bgp
              set policy-options policy-statement parent_vpn_routes then accept
              set routing-instances vpn-1 instance-type vrf
              set routing-instances vpn-1 interface ge-1/2/0.2
              set routing-instances vpn-1 interface vt-1/2/0.2
              set routing-instances vpn-1 interface lo0.102
              set routing-instances vpn-1 route-distinguisher 100:100
              set routing-instances vpn-1 provider-tunnel ldp-p2mp
              set routing-instances vpn-1 vrf-target target:1:1
              set routing-instances vpn-1 protocols ospf export parent_vpn_routes
              set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.102 passive
              set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/0.2
              set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
```

```

set routing-instances vpn-1 protocols pim interface ge-1/2/0.2 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 1001

```

Device R3

```

set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 6 family mpls
set interfaces ge-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces ge-1/2/1 unit 9 family mpls
set interfaces ge-1/2/2 unit 13 family inet address 10.1.1.13/30
set interfaces ge-1/2/2 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/1.9
set protocols mpls interface ge-1/2/2.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/1.9
set protocols ospf area 0.0.0.0 interface ge-1/2/2.13
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/1.9
set protocols ldp interface ge-1/2/2.13
set protocols ldp p2mp
set routing-options router-id 1.1.1.3

```

Device R4

```

set interfaces ge-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/0 unit 10 family mpls
set interfaces ge-1/2/1 unit 17 family inet address 10.1.1.17/30
set interfaces ge-1/2/1 unit 17 family mpls
set interfaces vt-1/2/0 unit 4 family inet
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set interfaces lo0 unit 104 family inet address 100.1.1.4/32
set protocols rsvp interface all aggregate
set protocols mpls interface all
set protocols mpls interface ge-1/2/0.10
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp family inet-vpn unicast
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling damping
set protocols bgp group ibgp neighbor 1.1.1.2 import dampPolicy
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.10
set protocols ldp interface ge-1/2/0.10
set protocols ldp p2mp
set policy-options policy-statement dampPolicy term term1 from family inet-mvpn
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 3
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 4
set policy-options policy-statement dampPolicy term term1 from nlri-route-type 5
set policy-options policy-statement dampPolicy term term1 then accept
set policy-options policy-statement dampPolicy then damping no-damp
set policy-options policy-statement dampPolicy then accept

```

```
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set policy-options damping no-damp disable
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.4
set routing-instances vpn-1 interface ge-1/2/1.17
set routing-instances vpn-1 interface lo0.104
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.104 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.17
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.17 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 1001
```

```
Device R5  set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces ge-1/2/1 unit 21 family inet address 10.1.1.21/30
set interfaces ge-1/2/1 unit 21 family mpls
set interfaces vt-1/2/0 unit 5 family inet
set interfaces lo0 unit 5 family inet address 1.1.1.5/32
set interfaces lo0 unit 105 family inet address 100.1.1.5/32
set protocols mpls interface ge-1/2/0.14
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp family inet-mvpn signaling
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
set protocols ldp interface ge-1/2/0.14
set protocols ldp p2mp
set policy-options policy-statement parent_vpn_routes from protocol bgp
set policy-options policy-statement parent_vpn_routes then accept
set routing-instances vpn-1 instance-type vrf
set routing-instances vpn-1 interface vt-1/2/0.5
set routing-instances vpn-1 interface ge-1/2/1.21
set routing-instances vpn-1 interface lo0.105
set routing-instances vpn-1 route-distinguisher 100:100
set routing-instances vpn-1 vrf-target target:1:1
set routing-instances vpn-1 protocols ospf export parent_vpn_routes
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface lo0.105 passive
set routing-instances vpn-1 protocols ospf area 0.0.0.0 interface ge-1/2/1.21
set routing-instances vpn-1 protocols pim rp static address 100.1.1.2
set routing-instances vpn-1 protocols pim interface ge-1/2/1.21 mode sparse
set routing-instances vpn-1 protocols mvpn
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 1001
```

```
Device R6  set interfaces ge-1/2/0 unit 18 family inet address 10.1.1.18/30
set interfaces ge-1/2/0 unit 18 family mpls
```

```

set interfaces lo0 unit 6 family inet address 1.1.1.6/32
set protocols sap listen 224.1.1.1
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.18
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.6

```

Device R7

```

set interfaces ge-1/2/0 unit 22 family inet address 10.1.1.22/30
set interfaces ge-1/2/0 unit 22 family mpls
set interfaces lo0 unit 7 family inet address 1.1.1.7/32
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-1/2/0.22
set protocols pim rp static address 100.1.1.2
set protocols pim interface all
set routing-options router-id 1.1.1.7

```

### Configuring Device R4

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R4:

1. Configure the interfaces.

```

[edit interfaces]
user@R4# set ge-1/2/0 unit 10 family inet address 10.1.1.10/30
user@R4# set ge-1/2/0 unit 10 family mpls

user@R4# set ge-1/2/1 unit 17 family inet address 10.1.1.17/30
user@R4# set ge-1/2/1 unit 17 family mpls

user@R4# set vt-1/2/0 unit 4 family inet

user@R4# set lo0 unit 4 family inet address 1.1.1.4/32
user@R4# set lo0 unit 104 family inet address 100.1.1.4/32

```

2. Configure MPLS and the signaling protocols on the interfaces.

```

[edit protocols]
user@R4# set mpls interface all
user@R4# set mpls interface ge-1/2/0.10
user@R4# set rsvp interface all aggregate
user@R4# set ldp interface ge-1/2/0.10
user@R4# set ldp p2mp

```

3. Configure BGP.

The BGP configuration enables BGP route flap damping for the **inet-mvpn** address family. The BGP configuration also imports into the routing table the routing policy called **dampPolicy**. This policy is applied to neighbor PE Device R2.

```

[edit protocols bgp group ibgp]
user@R4# set type internal

```

```
user@R4# set local-address 1.1.1.4
user@R4# set family inet-vpn unicast
user@R4# set family inet-vpn any
user@R4# set family inet-mvpn signaling damping
user@R4# set neighbor 1.1.1.2 import dampPolicy
user@R4# set neighbor 1.1.1.5
```

4. Configure an interior gateway protocol.

```
[edit protocols ospf]
user@R4# set traffic-engineering
```

```
[edit protocols ospf area 0.0.0.0]
user@R4# set interface all
user@R4# set interface lo0.4 passive
user@R4# set interface ge-1/2/0.10
```

5. Configure a damping policy that uses the **nlri-route-type** match condition to damp only MVPN route types 3, 4, and 5.

```
[edit policy-options policy-statement dampPolicy term term1]
user@R4# set from family inet-mvpn
user@R4# set from nlri-route-type 3
user@R4# set from nlri-route-type 4
user@R4# set from nlri-route-type 5
user@R4# set then accept
```

6. Configure the **damping** policy to disable BGP route flap damping.

The **no-damp** policy (**damping no-damp disable**) causes any damping state that is present in the routing table to be deleted. The **then damping no-damp** statement applies the **no-damp** policy as an action and has no **from** match conditions. Therefore, all routes that are not matched by **term1** are matched by this term, with the result that all other MVPN route types are not damped.

```
[edit policy-options policy-statement dampPolicy]
user@R4# set then damping no-damp
user@R4# set then accept
```

```
[edit policy-options]
user@R4# set damping no-damp disable
```

7. Configure the **parent\_vpn\_routes** to accept all other BGP routes that are not from the **inet-mvpn** address family.

This policy is applied as an OSPF export policy in the routing instance.

```
[edit policy-options policy-statement parent_vpn_routes]
user@R4# set from protocol bgp
user@R4# set then accept
```

8. Configure the VPN routing and forwarding (VRF) instance.

```
[edit routing-instances vpn-1]
user@R4# set instance-type vrf
user@R4# set interface vt-1/2/0.4
user@R4# set interface ge-1/2/1.17
user@R4# set interface lo0.104
```



```

user@R4# set route-distinguisher 100:100
user@R4# set vrf-target target:1:1
user@R4# set protocols ospf export parent_vpn_routes
user@R4# set protocols ospf area 0.0.0.0 interface lo0.104 passive
user@R4# set protocols ospf area 0.0.0.0 interface ge-1/2/1.17
user@R4# set protocols pim rp static address 100.1.1.2
user@R4# set protocols pim interface ge-1/2/1.17 mode sparse
user@R4# set protocols mvpn

```

9. Configure the router ID and the autonomous system (AS) number.

```

[edit routing-options]
user@R4# set router-id 1.1.1.4
user@R4# set autonomous-system 1001

```

10. If you are done configuring the device, commit the configuration.

```

user@R4# commit

```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R4# show interfaces
ge-1/2/0 {
  unit 10 {
    family inet {
      address 10.1.1.10/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 17 {
    family inet {
      address 10.1.1.17/30;
    }
    family mpls;
  }
}
vt-1/2/0 {
  unit 4 {
    family inet;
  }
}
lo0 {
  unit 4 {
    family inet {
      address 1.1.1.4/32;
    }
  }
}
unit 104 {
  family inet {

```

```
        address 100.1.1.4/32;
    }
}
}
user@R4# show protocols
rsvp {
    interface all {
        aggregate;
    }
}
mpls {
    interface all;
    interface ge-1/2/0.10;
}
bgp {
    group ibgp {
        type internal;
        local-address 1.1.1.4;
        family inet-vpn {
            unicast;
            any;
        }
        family inet-mvpn {
            signaling {
                damping;
            }
        }
        neighbor 1.1.1.2 {
            import dampPolicy;
        }
        neighbor 1.1.1.5;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface lo0.4 {
            passive;
        }
        interface ge-1/2/0.10;
    }
}
ldp {
    interface ge-1/2/0.10;
    p2mp;
}
user@R4# show policy-options
policy-statement dampPolicy {
    term term1 {
        from {
            family inet-mvpn;
            nlri-route-type [ 3 4 5 ];
        }
        then accept;
    }
}
```

```

    }
    then {
        damping no-damp;
        accept;
    }
}
policy-statement parent_vpn_routes {
    from protocol bgp;
    then accept;
}
damping no-damp {
    disable;
}

```

```

user@R4# show routing-instances

```

```

vpn-1 {
    instance-type vrf;
    interface vt-1/2/0.4;
    interface ge-1/2/1.17;
    interface lo0.104;
    route-distinguisher 100:100;
    vrf-target target:1:1;
    protocols {
        ospf {
            export parent_vpn_routes;
            area 0.0.0.0 {
                interface lo0.104 {
                    passive;
                }
                interface ge-1/2/1.17;
            }
        }
        pim {
            rp {
                static {
                    address 100.1.1.2;
                }
            }
            interface ge-1/2/1.17 {
                mode sparse;
            }
        }
        mvpn;
    }
}

```

```

user@R4# show routing-options
router-id 1.1.1.4;
autonomous-system 1001;

```

### Verification

Confirm that the configuration is working properly.

- [Verifying That Route Flap Damping Is Disabled on page 3618](#)
- [Verifying Route Flap Damping on page 3618](#)

**Verifying That Route Flap Damping Is Disabled**

**Purpose** Verify the presence of the **no-damp** policy, which disables damping for MVPN route types other than 3, 4, and 5.

**Action** From operational mode, enter the **show policy damping** command.

```
user@R4> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "no-damp":
  Damping disabled
```

**Meaning** The output shows that the default damping parameters are in effect and that the **no-damp** policy is also in effect for the specified route types.

**Verifying Route Flap Damping**

**Purpose** Check whether BGP routes have been damped.

**Action** From operational mode, enter the **show bgp summary** command.

```
user@R4> show bgp summary
Groups: 1 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0
      6      6      0      0      0      0
bgp.13vpn.2
      0      0      0      0      0      0
bgp.mvpn.0
      2      2      0      0      0      0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2 1001 3159 3155 0 0 23:43:47
Establ
  bgp.13vpn.0: 3/3/3/0
  bgp.13vpn.2: 0/0/0/0
  bgp.mvpn.0: 1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0
1.1.1.5 1001 3157 3154 0 0 23:43:40
Establ
  bgp.13vpn.0: 3/3/3/0
  bgp.13vpn.2: 0/0/0/0
  bgp.mvpn.0: 1/1/1/0
  vpn-1.inet.0: 3/3/3/0
  vpn-1.mvpn.0: 1/1/1/0
```

**Meaning** The Damp State field shows that zero routes in the bgp.mvpn.0 routing table have been damped. Further down, the last number in the State field shows that zero routes have been damped for BGP peer 1.1.1.2.

- Related Documentation**
- [Understanding External BGP Peering Sessions on page 3261](#)
  - [BGP Configuration Overview](#)

## BGP Monitoring Configuration

---

- [Example: Configuring BGP Trace Operations on page 3619](#)
- [Tracing BMP Operations on page 3625](#)

### Example: Configuring BGP Trace Operations

- [Understanding Trace Operations for BGP Protocol Traffic on page 3619](#)
- [Example: Viewing BGP Trace Files on Logical Systems on page 3621](#)

#### Understanding Trace Operations for BGP Protocol Traffic

---

You can trace various BGP protocol traffic to help you debug BGP protocol issues. To trace BGP protocol traffic, include the **traceoptions** statement at the **[edit protocols bgp]** hierarchy level. For routing instances, include the **traceoptions** statement at the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

You can specify the following BGP protocol-specific trace options using the **flag** statement:

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages.
- **nsr-synchronization**—Nonstop active routing synchronization events.
- **open**—BGP open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—BGP update packets. These packets provide routing updates to BGP systems.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the BGP protocol using the **traceoptions flag** statement included at the **[edit protocols bgp]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information.
- **filter**—Filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Packets being received.
- **send**—Packets being transmitted.



**NOTE:** Use the **all** trace flag and the **detail** flag modifier with caution because these might cause the CPU to become very busy.

---



**NOTE:** If you only enable the **update** flag, received keepalive messages do not generate a trace message.

---

You can filter trace statements and display only the statement information that passes through the filter by specifying the **filter** flag modifier. The **filter** modifier is only supported for the **route** and **damping** tracing flags.

The **match-on** statement specifies filter matches based on prefixes. It is used to match on route filters.



**NOTE:** Per-neighbor trace filtering is not supported on a BGP per-neighbor level for **route** and **damping** flags. Trace option filtering support is on a peer group level.

---

### Example: Viewing BGP Trace Files on Logical Systems

This example shows how to list and view files that are stored on a logical system.

- [Requirements on page 3621](#)
- [Overview on page 3621](#)
- [Configuration on page 3622](#)
- [Verification on page 3625](#)

#### Requirements

- You must have the **view** privilege for the logical system.
- Configure a network, such as the BGP network shown in “[Example: Configuring Internal BGP Peering Sessions on Logical Systems](#)” on page 3296.

#### Overview

Logical systems have their individual directory structure created in the `/var/logical-systems/logical-system-name` directory. It contains the following subdirectories:

- **/config**—Contains the active configuration specific to the logical system.
- **/log**—Contains system log and tracing files specific to the logical system.

To maintain backward compatibility for the log files with previous versions of Junos OS, a symbolic link (symlink) from the `/var/logs/logical-system-name` directory to the `/var/logical-systems/logical-system-name` directory is created when a logical system is configured.

- **/tmp**—Contains temporary files specific to the logical system.

The file system for each logical system enables logical system users to view trace logs and modify logical system files. Logical system administrators have full access to view and modify all files specific to the logical system.

Logical system users and administrators can save and load configuration files at the logical-system level using the **save** and **load** configuration mode commands. In addition, they can also issue the **show log**, **monitor**, and **file** operational mode commands at the logical-system level.

This example shows how to configure and view a BGP trace file on a logical system. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.



**TIP:** To view a list of hierarchy levels that support tracing operations, enter the **help apropos traceoptions** command in configuration mode.

### **Configuration**

- [Configuring Trace Operations on page 3622](#)
- [Viewing the Trace File on page 3622](#)
- [Deactivating and Reactivating Trace Logging on page 3624](#)
- [Results on page 3625](#)

#### **CLI Quick Configuration**

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems A protocols bgp group internal-peers traceoptions file bgp-log
set logical-systems A protocols bgp group internal-peers traceoptions file size 10k
set logical-systems A protocols bgp group internal-peers traceoptions file files 2
set logical-systems A protocols bgp group internal-peers traceoptions flag update detail
```

### **Configuring Trace Operations**

#### **Step-by-Step Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the trace operations:

1. Configure trace operations on the logical system.

```
[edit logical-systems A protocols bgp group internal-peers]
user@host# set traceoptions file bgp-log
user@host# set traceoptions file size 10k
user@host# set traceoptions file files 2
user@host# set traceoptions flag update detail
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

### **Viewing the Trace File**

#### **Step-by-Step Procedure**

To view the trace file:

1. In operational mode on the main router, list the directories on the logical system.

```
user@host> file list /var/logical-systems/A
/var/logical-systems/A:
config/
log/
tmp/
```

2. In operational mode on the main router, list the log files on the logical system.

```
user@host> file list /var/logical-systems/A/log/
/var/logical-systems/A/log:
bgp-log
```

3. View the contents of the **bgp-log** file.



```

user@host> file show /var/logical-systems/A/log/bgp-log
Aug 10 17:12:01 trace_on: Tracing to "/var/log/A/bgp-log" started
Aug 10 17:14:22.826182 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.163.6.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.826445 bgp_send: sending 21 bytes to 192.163.6.4 (Internal
AS 17)
Aug 10 17:14:22.826499
Aug 10 17:14:22.826499 BGP SEND 192.168.6.5+64965 -> 192.163.6.4+179
Aug 10 17:14:22.826559 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.826598 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
Aug 10 17:14:22.831756 bgp_peer_mgmt_clear:5829: NOTIFICATION sent to
192.168.40.4 (Internal AS 17): code 6 (Cease) subcode 4 (Administratively
Reset), Reason: Management session cleared BGP neighbor
Aug 10 17:14:22.831851 bgp_send: sending 21 bytes to 192.168.40.4 (Internal
AS 17)
Aug 10 17:14:22.831901
Aug 10 17:14:22.831901 BGP SEND 192.168.6.5+53889 -> 192.168.40.4+179
Aug 10 17:14:22.831959 BGP SEND message type 3 (Notification) length 21
Aug 10 17:14:22.831999 BGP SEND Notification code 6 (Cease) subcode 4
(Administratively Reset)
...

```

4. Filter the output of the log file.

```

user@host> file show /var/logical-systems/A/log/bgp-log | match "flags 0x40"
Aug 10 17:14:54.867460 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.867595 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.867650 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.867692 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.884529 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.884581 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.884628 BGP RECV flags 0x40 code NextHop(3): 192.163.6.4
Aug 10 17:14:54.884667 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.911377 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.911422 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.911466 BGP RECV flags 0x40 code NextHop(3): 192.168.40.4
Aug 10 17:14:54.911507 BGP RECV flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.916008 BGP SEND flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.916054 BGP SEND flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.916100 BGP SEND flags 0x40 code NextHop(3): 192.168.6.5
Aug 10 17:14:54.916143 BGP SEND flags 0x40 code LocalPref(5): 100
Aug 10 17:14:54.920304 BGP RECV flags 0x40 code Origin(1): IGP
Aug 10 17:14:54.920348 BGP RECV flags 0x40 code ASPath(2) length 0: <null>
Aug 10 17:14:54.920393 BGP RECV flags 0x40 code NextHop(3): 10.0.0.10
Aug 10 17:14:54.920434 BGP RECV flags 0x40 code LocalPref(5): 100

```

5. View the tracing operations in real time.

```

user@host> clear bgp neighbor logical-system A
Cleared 2 connections

```



**CAUTION:** Clearing the BGP neighbor table is disruptive in a production environment.

6. Run the **monitor start** command with an optional **match** condition.

```

user@host> monitor start A/bgp-log | match 0.0.0.0/0

```

```
Aug 10 19:21:40.773467 BGP RECV          0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlri: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlri: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlri: 0.0.0.0/0 qualified bnp->ribact 0x0
12afcb 0x0
```

7. Pause the **monitor** command by pressing Esc-Q.  
To unpause the output, press Esc-Q again.
8. Halt the **monitor** command by pressing Enter and typing **monitor stop**.  
[Enter]  
user@host> **monitor stop**
9. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

```
[edit protocols bgp group internal-peers]
user@host:A# deactivate traceoptions
user@host:A# commit
```

When configuration is deactivated, it appears in the configuration with the **inactive** tag. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit protocols bgp group internal-peers]
user@host:A# show
```

```
type internal;
inactive: traceoptions {
  file bgp-log size 10k files 2;
  flag update detail;
  flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;
```

10. To reactivate trace operations, use the **activate** configuration-mode statement.  
[edit protocols bgp group internal-peers]  
user@host:A# **activate traceoptions**  
user@host:A# **commit**

### ***Deactivating and Reactivating Trace Logging***

#### **Step-by-Step Procedure**

To deactivate and reactivate the trace file:

1. When you are finished troubleshooting, consider deactivating trace logging to avoid an unnecessary impact to system resources.  
[edit protocols bgp group internal-peers]  
user@host:A# **deactivate traceoptions**  
user@host:A# **commit**

When configuration is deactivated, the statement appears in the configuration with the **inactive** tag.

```
[edit protocols bgp group internal-peers]
user@host:A# show
```

```

type internal;
inactive: traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
    flag all;
}
local-address 192.168.6.5;
export send-direct;
neighbor 192.163.6.4;
neighbor 192.168.40.4;

```

2. To reactivate logging, use the **activate** configuration-mode statement.

```

[edit protocols bgp group internal-peers]
user@host:A# activate traceoptions
user@host:A# commit

```

### Results

From configuration mode, confirm your configuration by entering the **show logical-systems A protocols bgp group internal-peers** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show logical-systems A protocols bgp group internal-peers
traceoptions {
    file bgp-log size 10k files 2;
    flag update detail;
}

```

### Verification

Confirm that the configuration is working properly.

#### Verifying That the Trace Log File Is Operating

**Purpose** Make sure that events are being written to the log file.

**Action** user@host:A> **show log bgp-log**  
Aug 12 11:20:57 trace\_on: Tracing to "/var/log/A/bgp-log" started

**Related Documentation**

- [Understanding External BGP Peering Sessions on page 3261](#)
- [BGP Configuration Overview](#)

## Tracing BMP Operations

You can trace BMP operations for all BMP stations by configuring the **traceoptions** statement at the **[edit routing-options bmp]** hierarchy level or for specific BMP stations at the **[edit routing-options bmp station station-name]** hierarchy level.

To trace BMP operations, complete the following steps:

1. Configure the **traceoptions** statement:

```

traceoptions {

```

```
file filename <files number> <size size> <world-readable | no-world-readable>;  
flag flag <flag-modifier> <disable>;  
}
```

2. Specify the name of the file to receive the output of the tracing operation using the **file** option. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place BMP tracing output in the file **bmp-log**.
3. (Optional) Specify the maximum number of trace files using the **files** option. When a trace file named **trace-file.0** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.
4. (Optional) Specify the maximum size of each trace file using the **size** option in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.
5. (Optional) You can specify that the log files are either **world-readable** (accessible to all users on the device) or **no-world-readable** (not accessible to all users on the device).
6. You can specify the following BMP-specific trace options using the **flag** statement:
  - **all**—Trace all BMP monitoring operations.
  - **down**—Down messages.
  - **error**—Error conditions.
  - **event**—Major events, station establishment, errors, and events.
  - **general**—General events.
  - **normal**—Normal events.
  - **packets**—All messages.
  - **policy**—Policy processing.
  - **route**—Routing information.
  - **route-monitoring**—Route monitoring messages.
  - **state**—State transitions.
  - **statistics**—Statistics messages.
  - **task**—Routing protocol task processing.
  - **timer**—Routing protocol timer processing.
  - **up**—Up messages.
  - **write**—Writing of messages.

You can optionally specify one or more of the following flag modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable the tracing flag.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.



**NOTE:** Use the all trace flag and the detail flag modifier with caution due to the increased computer processing power required.

#### Related Documentation

- [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

## Configuration Statements

- [accept-remote-nexthop on page 3630](#)
- [advertise-external on page 3631](#)
- [advertise-inactive on page 3633](#)
- [advertise-peer-as on page 3634](#)
- [algorithm \(BGP BFD Authentication\) on page 3635](#)
- [apply-groups on page 3637](#)
- [apply-groups-except on page 3637](#)
- [authentication \(BGP BFD Liveness Detection\) on page 3638](#)
- [authentication-algorithm on page 3640](#)
- [authentication-key \(Protocols BGP and BMP\) on page 3641](#)
- [authentication-key-chain \(Protocols BGP and BMP\) on page 3642](#)
- [bfd-liveness-detection \(Protocols BGP\) on page 3643](#)
- [bgp on page 3647](#)
- [bgp-orf-cisco-mode on page 3648](#)
- [cluster on page 3650](#)
- [connection-mode on page 3651](#)
- [damping \(Protocols BGP\) on page 3652](#)
- [description \(Protocols BGP\) on page 3654](#)
- [detection-time \(BFD Liveness Detection\) on page 3655](#)
- [disable \(Protocols BGP\) on page 3656](#)
- [disable \(BGP Graceful Restart\) on page 3657](#)
- [export \(Protocols BGP\) on page 3658](#)
- [family \(Protocols BGP\) on page 3659](#)

- graceful-restart (Protocols BGP) on page 3663
- group (Protocols BGP) on page 3664
- hold-down on page 3667
- hold-down-interval (BGP BFD Liveness Detection) on page 3669
- hold-time (Protocols BGP) on page 3671
- import (Protocols BGP) on page 3673
- include-mp-next-hop on page 3675
- initiation-message on page 3676
- keep on page 3677
- key-chain (BGP BFD Authentication) on page 3679
- local-address (Protocols BGP) on page 3681
- local-address (Protocols BMP) on page 3683
- local-as on page 3684
- local-port on page 3686
- local-preference on page 3687
- log-updown (Protocols BGP) on page 3688
- loops on page 3689
- loose-check (BGP BFD Authentication) on page 3691
- maximum-ecmp on page 3692
- metric-out (Protocols BGP) on page 3693
- minimum-interval (BFD Liveness Detection) on page 3695
- minimum-interval (transmit-interval) on page 3697
- minimum-receive-interval (BFD Liveness Detection) on page 3699
- monitor (Protocols BMP) on page 3700
- mtu-discovery on page 3701
- multihop on page 3703
- multiplier (BFD Liveness Detection) on page 3705
- neighbor (Protocols BGP) on page 3707
- no-adaptation (BFD Liveness Detection) on page 3710
- no-advertise-peer-as on page 3711
- no-aggregator-id on page 3712
- no-client-reflect on page 3713
- out-delay on page 3714
- outbound-route-filter on page 3716
- passive (Protocols BGP) on page 3717
- path-selection on page 3718
- peer-as (Protocols BGP) on page 3720

- [post-policy](#) on page 3721
- [pre-policy](#) on page 3722
- [preference \(Protocols BGP\)](#) on page 3723
- [priority \(Protocols BMP\)](#) on page 3724
- [remove-private](#) on page 3725
- [restart-time \(BGP Graceful Restart\)](#) on page 3727
- [route-monitoring](#) on page 3728
- [session-mode](#) on page 3729
- [stale-routes-time](#) on page 3730
- [station](#) on page 3731
- [station-address](#) on page 3732
- [station-port](#) on page 3733
- [statistics-timeout](#) on page 3734
- [tcp-mss \(Protocols BGP\)](#) on page 3735
- [threshold \(detection-time\)](#) on page 3736
- [threshold \(transmit-interval\)](#) on page 3738
- [traceoptions \(Protocols BGP\)](#) on page 3740
- [traceoptions \(Protocols BMP\)](#) on page 3743
- [transmit-interval \(BFD Liveness Detection\)](#) on page 3745
- [version \(BFD Liveness Detection\)](#) on page 3747

## accept-remote-nexthop

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | accept-remote-nexthop;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>[edit protocols bgp],<br/>[edit protocols bgp group <i>group-name</i>],<br/>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Specify that a single-hop EBGp peer accepts a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGp peer to specify the remote next hop. You cannot configure <b>multihop</b> and <b>accept-remote-nexthop</b> statements for the same EBGp peer.</p> <p>For Junos OS Release 13.3 and later releases, specify that a multihop EBGp peer accepts a remote next hop with which it does not share a common subnet. This allows working around current resolver limitations to realize multipath forwarding in recursive next-hop resolution scenarios.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Single-Hop EBGp Peers to Accept Remote Next Hops on page 3482</a></li><li>• <a href="#">Understanding Route Advertisement on page 3425</a></li><li>• <i>multipath</i></li></ul>  |



## advertise-external

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>advertise-external {<i>conditional</i>};</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit protocols bgp group <i>group-name</i>],<br/>         [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-address</i>]</p>   |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>         | <p>Specify BGP to advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route.</p> <p>In general, deployed BGP implementations do not advertise the external route with the highest local preference value to internal peers unless it is the best route. Although this behavior was required by an earlier version of the BGP version 4 specification, RFC 1771, it was typically not followed in order to minimize the amount of advertised information and to prevent routing loops. However, there are scenarios in which advertising the best external route is beneficial, in particular, situations that can result in IBGP route oscillation.</p> <p>The <b>advertise-external</b> statement is supported at both the group and neighbor level. If you configure the statement at the neighbor level, you must configure it for all neighbors in a group. Otherwise, the group is automatically split into different groups.</p> <p>In a confederation, when advertising a route to a confederation border router, any route from a different confederation sub-AS is considered external. When configuring the <b>advertise-external</b> statement for an AS confederation, it is recommended that EBGp peers belonging to different autonomous systems are configured in a separate EBGp peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.</p> <p>To configure the <b>advertise-external</b> statement on a route reflector, you must disable intracluster reflection with the <b>no-client-reflect</b> statement.</p> <p>When a routing device is configured as a route reflector for a cluster, a route advertised by the route reflector is considered internal if it is received from an internal peer with the same cluster identifier or if both peers have no cluster identifier configured. A route received from an internal peer that belongs to another cluster, that is, with a different cluster identifier, is considered external.</p> <p>The <b>conditional</b> option causes BGP to advertise the external route only if the route selection process reaches the point where the multiple exit discriminator (MED) metric</p> |

is evaluated. As a result, an external route with an AS path longer than that of the active path is not advertised.

Junos OS also provides support for configuring a BGP export policy that matches on the state of an advertised route. You can match on either active or inactive routes.

**Default** BGP does not advertise the external route with the highest local preference value to internal peers unless it is the best route.

**Options** **conditional**—(Optional) Advertise the best external path only if the route selection process reaches the point at which the multiple exit discriminator (MED) metric is evaluated. The **conditional** option restricts advertisement to when the best external path and the active path are equal until the MED step of the route selection process. This implies that external routes with a longer AS path length than the active path, for instance, are not advertised. The criteria used for selecting the best external path is the same whether or not the **conditional** option is configured.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring BGP to Advertise the Best External Route to Internal Peers*
- [advertise-inactive on page 3633](#)

## advertise-inactive

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | advertise-inactive;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure the routing table to export to BGP the best route learned by BGP even if Junos OS did not select this route to be an active route.</p> <p>One way to achieve multivendor compatibility is to include the <b>advertise-inactive</b> statement in the external BGP (EBGP) configuration. By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers. The <b>advertise-inactive</b> statement causes Junos OS to advertise the best BGP route that is inactive because of IGP preference. When you use the <b>advertise-inactive</b> statement, the Junos OS device uses, for example, the OSPF route for forwarding, and the other vendor's device uses the EBGP route for forwarding. However, from the perspective of an EBGP peer in a neighboring AS, both vendors' devices appear to behave the same way.</p>   |
| <b>Default</b>                  | By default, BGP stores the route information it receives from update messages in the Junos OS routing table, and the routing table exports only active routes into BGP, which BGP then advertises to its peers.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Setting BGP to Advertise Inactive Routes</i></li> <li>• <a href="#">Example: Configuring the Preference Value for BGP Routes on page 3443</a></li> <li>• <a href="#">Example: Configuring BGP Route Preference (Administrative Distance) on page 3442</a></li> </ul>   |

- [advertise-external on page 3631](#)

## advertise-peer-as

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | advertise-peer-as;  |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit protocols bgp], [edit protocols bgp <b>group</b> <i>group-name</i>], [edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>   <b>neighbor</b> <i>address</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Disable the default behavior of suppressing AS routes.</p> <p>If you include the <b>advertise-peer-as</b> statement in the configuration, BGP advertises routes learned from one external BGP (EBGP) peer back to another EBGP peer in the same autonomous system (AS) but not back to the originating peer.</p> <p>Another way to disable the route suppression default behavior is with the <b>as-override</b> statement. If you include both the <b>as-override</b> and <b>no-advertise-peer-as</b> statements in the configuration, the <b>no-advertise-peer-as</b> statement is ignored.</p>  |
| <b>Default</b>                  | By default, Junos OS does not advertise the routes learned from one EBGP peer back to the same external BGP (EBGP) peer. In addition, the software does not advertise those routes back to any EBGP peers that are in the same AS as the originating peer, regardless of the routing instance.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Disabling Suppression of Route Advertisements</i></li><li>• <i>Example: Configuring a Layer 3 VPN with Route Reflection and AS Override</i></li><li>• <a href="#">no-advertise-peer-as on page 3711</a></li></ul>   |

## algorithm (BGP BFD Authentication)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>algorithm <i>algorithm-name</i>;</code>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>         | Configure the algorithm used to authenticate the specified BFD session.  |
| <b>Options</b>             | <p><b><i>algorithm-name</i></b>—Authentication algorithm name: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, <b>meticulous-keyed-sha-1</b>.</p> <p><b>simple-password</b>—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.</p> <p><b>keyed-md5</b>—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.</p>   |

**meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method can take additional time to authenticate the session.

**keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.

**meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method can take additional time to authenticate the session.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Example: Configuring BGP Route Authentication on page 3571</a></li><li>• <a href="#">Example: Configuring EBGp Multihop Sessions on page 3433</a></li><li>• <a href="#">authentication on page 3638</a></li><li>• <a href="#">bfd-liveness-detection on page 3643</a></li><li>• <a href="#">key-chain on page 3679</a></li><li>• <a href="#">loose-check on page 3691</a></li></ul> |
|------------------------------|--|

## apply-groups

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>apply-groups [ <i>group-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | All hierarchy levels   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.  |
| <b>Description</b>              | <p>Apply a configuration group to a specific hierarchy level in a configuration, to have a configuration inherit the statements in the configuration group.</p> <p>You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.</p> |
| <b>Options</b>                  | <i>group-names</i> —One or more names specified in the <b>groups</b> statement.  |
| <b>Required Privilege Level</b> | configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Applying a Junos OS Configuration Group</i></li> <li>• <i>groups</i></li> </ul>  |

## apply-groups-except

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>apply-groups-except [ <i>group-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | All hierarchy levels except the top level   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.   |
| <b>Description</b>              | Disable inheritance of a configuration group.   |
| <b>Options</b>                  | <i>group-names</i> —One or more names specified in the <b>groups</b> statement.   |
| <b>Required Privilege Level</b> | configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy. |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>groups</i></li> <li>• <i>Disabling Inheritance of a Junos OS Configuration Group</i></li> </ul>         |

## authentication (BGP BFD Liveness Detection)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>authentication {<br/>    algorithm <i>algorithm-name</i>;<br/>    key-chain <i>key-chain-name</i>;<br/>    loose-check ;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i><br/>    bfd-liveness-detection],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i><br/>    bfd-liveness-detection],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    bgp bfd-liveness-detection],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    bgp group <i>group-name</i> bfd-liveness-detection],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>    bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],<br/>[edit protocols bgp bgp bfd-liveness-detection],<br/>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],<br/>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i><br/>    bfd-liveness-detection],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor<br/>    <i>address</i> bfd-liveness-detection]</pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Specify the router and route authentication to mitigate the risk of being attacked by a machine or router that has been configured to share incorrect routing information with another router. Router and route authentication enables routers to share information only if they can verify that they are talking to a trusted source, based on a password (key). In this method, a hashed key is sent along with the route being sent to another router. The receiving router compares the sent key to its own configured key. If they are the same, the receiving router accepts the route.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Example: Configuring BGP Route Authentication on page 3571</a></li><li>• <a href="#">algorithm on page 3635</a></li></ul>  |



- [bfd-liveness-detection on page 3643](#)
- [key-chain on page 3679](#)
- [loose-check on page 3691](#)

## authentication-algorithm

**Syntax** authentication-algorithm *algorithm*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],  
 [edit logical-systems *logical-system-name* protocols ldp session *session-address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *session-address*],  
 [edit logical-systems *logical-system-name* routing-options **bmp**],  
 [edit logical-systems *logical-system-name* routing-options bmp **station** *station-name*],  
 [edit protocols bgp],  
 [edit protocols bgp **group** *group-name*],  
 [edit protocols bgp group *group-name* **neighbor** *address*],  
 [edit protocols ldp session *session-address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],  
 [edit routing-instances *routing-instance-name* protocols ldp session *session-address*],  
 [edit routing-options **bmp**],  
 [edit routing-options bmp **station** *station-name*]

**Release Information** Statement introduced in Junos OS Release 7.6.  
 Statement introduced for BGP in Junos OS Release 8.0.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.  
 Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.  
 Statement introduced for BMP in Junos OS Release 13.3.

**Description** Configure an authentication algorithm type.

**Options** *algorithm*—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

**Default:** hmac-sha-1-96



**NOTE:** The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Route Authentication for BGP on page 3572](#)
- [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

## authentication-key (Protocols BGP and BMP)

**Syntax** authentication-key *key*;

**Hierarchy Level**

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name
neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name
neighbor address],
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

**Release Information**

Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.  
Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.  
Statement introduced for BMP version 3 in Junos OS Release 13.3.

**Description** Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.

**Options** *key*—Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Route Authentication for BGP on page 3572](#)
- [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

## authentication-key-chain (Protocols BGP and BMP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>authentication-key-chain <i>key-chain</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>bmp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp <b>station</b> <i>station-name</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-options <b>bmp</b>],</p> <p>[edit routing-options bmp <b>station</b> <i>station-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.3.</p>  |
| <b>Description</b>              | Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update feature for BGP, you cannot commit the <b>0.0.0.0/allow</b> statement with authentication keys or key chains. The CLI issues a warning and fails to commit the configuration.  |
| <b>Options</b>                  | <b>key-chain</b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Route Authentication for BGP on page 3572</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li> <li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li> <li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li> </ul>  |

## bfd-liveness-detection (Protocols BGP)

```
Syntax  bfd-liveness-detection {
        authentication {
            algorithm algorithm-name;
            key-chain key-chain-name;
            loose-check;
        }
        detection-time {
            threshold milliseconds;
        }
        hold-down-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        session-mode (automatic | multihop | single-hop);
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (1 | automatic);
    }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 bgp group *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 bgp group *group-name* neighbor *address*],  
 [edit protocols bgp],  
 [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor  
*address*]

**Release Information** Statement introduced in Junos OS Release 8.1.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
**detection-time threshold** and **transmit-interval threshold** options introduced in Junos OS Release 8.2  
 Support for logical routers introduced in Junos OS Release 8.3.  
 Support for IBGP and multihop EBGP sessions introduced in Junos OS Release 8.3.  
**holddown-interval** statement introduced in Junos OS Release 8.5. You can configure this statement only for EBGP peers at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.  
**no-adaptation** statement introduced in Junos OS Release 9.0.  
 Support for BFD authentication introduced in Junos OS Release 9.6.

Support for BFD on IPv6 interfaces with BGP introduced in Junos OS Release 11.2.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure bidirectional failure detection (BFD) timers and authentication for BGP.

For IBGP and multihop EBGp support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

**Options** **authentication algorithm** *algorithm-name* (Optional)—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

**authentication key-chain** *key-chain-name* (Optional)—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold** *milliseconds* (Optional)—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**holddown-interval** *milliseconds* (Optional)—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes it is down and does not send a state change notification. The **holddown-interval** statement is supported only for EBGp peers at the **[edit protocols bgp group group-name neighbor address]** hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the **local-address** statement at the **[edit protocols bgp group group-name]** hierarchy level.

**Range:** 0 through 255,000

**Default:** 0

**minimum-interval** *milliseconds* (Required)—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval that the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately (using the **minimum-receive-interval** and **transmit-interval** statements).

**Range:** 1 through 255,000

**minimum-receive-interval** *milliseconds* (Optional)—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

**Range:** 1 through 255,000

**multiplier *number*** (Optional)—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation** (Optional)—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds*** (Optional)—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds*** (Optional)—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

**Range:** 1 through 255,000

**version** (Optional)—Configure the BFD version to detect.

**Range:** 1 or **automatic** (autodetect the BFD version)

**Default:** **automatic**

The remaining statements are explained separately.

|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Example: Configuring BFD on Internal BGP Peer Sessions on page 3463</a></li><li>• <a href="#">Example: Configuring BFD Authentication for BGP on page 3473</a></li><li>• <a href="#">Understanding BFD for BGP on page 3462</a></li></ul> |
|------------------------------|--|




## bgp

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>bgp { ... }</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols bgp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Enable BGP on the routing device or for a routing instance.   |
| <b>Default</b>                  | BGP is disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>BGP Feature Guide for Routing Devices</i></li> </ul>  |

## bgp-orf-cisco-mode

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | bgp-orf-cisco-mode;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> <b>outbound-route-filter</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options <b>outbound-route-filter</b>],</p> <p>[edit routing-options <b>outbound-route-filter</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.</p> <p>Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>  |
| <b>Description</b>              | Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.  |
|                                 | <p> <b>NOTE:</b> To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.</p>  |
| <b>Default</b>                  | Disabled   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3429](#)

## cluster

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>cluster <i>cluster-identifier</i>;</code>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>         | Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.   |



### CAUTION:

If you configure both route reflection and VPNs on the same routing device, the following modifications to the route reflection configuration cause current BGP sessions to be reset:

- Adding a cluster ID—If a BGP session shares the same AS number with the group where you add the cluster ID, all BGP sessions are reset regardless of whether the BGP sessions are contained in the same group.
- Creating a new route reflector—If you have an IBGP group with an AS number and create a new route reflector group with the same AS number, all BGP sessions in the IBGP group and the new route reflector group are reset.



**NOTE:** If you change the address family specified in the [edit protocols bgp family] hierarchy level, all current BGP sessions on the routing device are dropped and then reestablished.

|                                 |  |
|---------------------------------|--|
| <b>Options</b>                  | <i>cluster-identifier</i> —4-byte identifier (such as an IPv4 address).  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BGP Route Reflectors on page 3547</a></li> <li>• <a href="#">Understanding External BGP Peering Sessions on page 3261</a></li> <li>• <a href="#">no-client-reflect on page 3713</a></li> </ul> |

## connection-mode

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | connection-mode (active   passive);   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station station-name</a> ],<br>[edit routing-options <a href="#">bmp</a> ],<br>[edit routing-options bmp <a href="#">station station-name</a> ] |
| <b>Release Information</b>      | Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced for BMP in Junos OS Release 13.3.  |
| <b>Description</b>              | Specifies whether the BMP station connection is <b>active</b> or <b>passive</b> .   |
| <b>Options</b>                  | <p><b>active</b>—BMP initiates the connection to the BMP station.</p> <p><b>passive</b>—BMP does not initiate a connection the BMP station. However, it does listen for a connection request from active BMP stations and will connect if a station is available.</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li> </ul>  |

## damping (Protocols BGP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | damping;  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> family <i>family</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> family <i>family</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for flap damping at the address family level introduced in Junos OS Release 12.2.</p>  |
| <b>Description</b>         | <p>Enable route flap damping. BGP route flapping describes the situation in which BGP systems send an excessive number of update messages to advertise network reachability information. Flap damping reduces the number of update messages sent between BGP</p>  |

peers, thereby reducing the load on these peers, without adversely affecting the route convergence time for stable routes.

You typically apply flap damping to external BGP (EBGP) routes (that is, to routes in different ASs). You can also apply it within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.) The exception to this rule is when flap damping is applied at the address family level. When you apply flap damping at the address family level, it works for both IBGP and EBGP.

**Default** Flap damping is disabled on the routing device.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- [Examples: Configuring BGP Flap Damping on page 3599](#)
- [Example: Configuring BGP Route Flap Damping Based on the MBGP MVPN Address Family on page 3609](#)

## description (Protocols BGP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>description text-description;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems logical-system-name protocols bgp],</code><br><code>[edit logical-systems logical-system-name protocols bgp group group-name],</code><br><code>[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],</code><br><code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name neighbor address],</code><br><code>[edit protocols bgp],</code><br><code>[edit protocols bgp group group-name],</code><br><code>[edit protocols bgp group group-name neighbor address],</code><br><code>[edit routing-instances routing-instance-name protocols bgp],</code><br><code>[edit routing-instances routing-instance-name protocols bgp group group-name],</code><br><code>[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Provide a description of the global, group, or neighbor configuration. If the text includes one or more spaces, enclose it in quotation marks (" "). The text is displayed in the output of the <b>show</b> command and has no effect on the configuration.   |
| <b>Options</b>                  | <i>text-description</i> —Text description of the configuration. It is limited to 255 characters.  |
| <b>Required Privilege Level</b> | <code>routing</code> —To view this statement in the configuration.<br><code>routing-control</code> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>BGP Feature Guide for Routing Devices</i></li></ul>  |



## detection-time (BFD Liveness Detection)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> detection-time {     threshold milliseconds; } </pre>   |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>   bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>   bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   l2vpn oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   vpls oam bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>   bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor   <i>address</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam   bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>   neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection] </pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPNs and VPLS.</p>   |
| <b>Description</b>         | <p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance</p>   |

is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

The remaining statement is explained separately.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for Layer 2 VPN and VPLS</a></li><li>• <a href="#">Example: Configuring BFD for BGP on page 3462</a></li><li>• <a href="#">bfd-liveness-detection on page 3643</a></li><li>• <a href="#">threshold on page 3736</a></li></ul> |

---

## disable (Protocols BGP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | disable;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols bgp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br>[edit protocols bgp],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | Disable BGP on the system.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |

## disable (BGP Graceful Restart)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | disable;   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols bgp graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> graceful-restart],<br>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> graceful-restart],<br>[edit protocols bgp graceful-restart],<br>[edit protocols bgp <b>group</b> <i>group-name</i> graceful-restart],<br>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i> graceful-restart] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>         | Disable graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition.   |



**NOTE:** When you disable graceful restart at one level in the configuration statement hierarchy, it is also disabled at lower levels in the same hierarchy. For example, if you disable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level, it is disabled for all the peers in the group. Therefore, if you want to enable graceful restart for some peers in a group and disable it for others, enable graceful restart at the [edit protocols bgp group *group-name*] hierarchy level and disable graceful restart for each peer at the [edit protocols bgp group *group-name* neighbor *address*] hierarchy level.

|                              |  |
|------------------------------|--|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.   |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li> <li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li> <li>• <a href="#">graceful-restart on page 2299</a></li> </ul> |

## export (Protocols BGP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>export [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit protocols bgp],<br/>[edit protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Apply one or more policies to routes being exported from the routing table into BGP.</p> <p>If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of <b>metric 500</b> defined in the next policy.</p>  |
| <b>Options</b>                  | <p><b><i>policy-names</i></b>—Name of one or more policies.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Route Advertisement on page 3425</a></li><li>• <a href="#">Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</a></li><li>• <a href="#">import on page 3673</a></li></ul>   |

## family (Protocols BGP)

```
Syntax  family {
    (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
        (any | flow | labeled-unicast | multicast | unicast) {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage-threshold> idle-timeout (forever | minutes);
            }
            add-path {
                send {
                    path-count number;
                    prefix-policy [ policy-names ];
                }
                receive;
            }
            algp [disable];
            loops number;
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            protection;
            rib-group group-name;
            topology name {
                community {
                    target identifier;
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib (inet.3 | inet6.3);
            rib-group group-name;
            traffic-statistics {
                file filename <world-readable | no-world-readable>;
                interval seconds;
            }
        }
    }
}
```

```
    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      proxy-generate <route-target-policy route-target-policy-name>;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  (evpn | inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
    signaling {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage-threshold> idle-timeout (forever | minutes);
      }
      add-path {
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
        receive;
      }
      aigp [disable];
      damping;
      loops number;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
  }
}
```

|                            |   |
|----------------------------|---|
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit protocols bgp], [edit protocols bgp <b>group</b> <i>group-name</i>], [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>   <b>neighbor</b> <i>address</i>] </pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>inet-mvpn</b> and <b>inet6-mvpn</b> statements introduced in Junos OS Release 8.4.</p> <p><b>inet-mdt</b> statement introduced in Junos OS Release 9.4.</p> <p>Support for the <b>loops</b> statement introduced in Junos OS Release 9.6.</p> <p><b>evpn</b> statement introduced in Junos OS Release 13.2.</p>   |
| <b>Description</b>         | <p>Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.</p>  |

- Options**
- any**—Configure the family type to be both unicast and multicast.
  - evpn**—Configure NLRI parameters for Ethernet VPNs (EVPNs).
  - inet**—Configure NLRI parameters for IPv4.
  - inet6**—Configure NLRI parameters for IPv6.
  - inet-mdt**—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.
  - inet-mvpn**—Configure NLRI parameters for IPv4 for multicast VPNs.
  - inet6-mvpn**—Configure NLRI parameters for IPv6 for multicast VPNs.
  - inet-vpn**—Configure NLRI parameters for IPv4 for Layer 3 VPNs.
  - inet6-vpn**—Configure NLRI parameters for IPv6 for Layer 3 VPNs.
  - iso-vpn**—Configure NLRI parameters for IS-IS for Layer 3 VPNs.
  - l2vpn**—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.
  - labeled-unicast**—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.
  - multicast**—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.
  - unicast**—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes. The default family type is **unicast**.


The remaining statements are explained separately.

- Required Privilege Level**
- routing—To view this statement in the configuration.
  - routing-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring IBGP Sessions Between PE Routers in VPNs*
  - *Understanding Multiprotocol BGP*
  - [autonomous-system on page 2945](#)
  - [local-as on page 3684](#)



## graceful-restart (Protocols BGP)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> graceful-restart {   disable;   restart-time seconds;   stale-routes-time seconds; } </pre>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/> [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/> [edit protocols bgp],<br/> [edit protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/> Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/> Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure graceful restart for BGP. Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. Graceful restart is disabled by default.</p> <p>To configure the duration of the BGP graceful restart period, include the <b>restart-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level. To set the length of time the router waits to receive messages from restarting neighbors before declaring them down, include the <b>stale-routes-time</b> statement at the [edit protocols bgp graceful-restart] hierarchy level.</p> <hr/> <div>  <p><b>NOTE:</b> If you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.</p> </div> <hr/> <p>Configure graceful restart globally at the [edit routing-options] or [edit routing-instances <i>instance-name</i> routing-options] hierarchy level to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/> routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li> <li>• <i>Configuring Graceful Restart for QFabric Systems</i></li> <li>• <i>Junos OS High Availability Library for Routing Devices</i></li> </ul>   |

## group (Protocols BGP)

---

```
Syntax  group group-name {  
    advertise-inactive;  
    allow [ network/mask-length ];  
    authentication-key key;  
    cluster cluster-identifier;  
    damping;  
    description text-description;  
    export [ policy-names ];  
    family {  
        (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {  
            (any | multicast | unicast | signaling) {  
                accepted-prefix-limit {  
                    maximum number;  
                    teardown <percentage> <idle-timeout (forever | minutes)>;  
                }  
            }  
            add-path {  
                send {  
                    path-count number;  
                    prefix-policy [ policy-names ];  
                }  
                receive;  
            }  
            aigp [disable];  
            damping;  
            prefix-limit {  
                maximum number;  
                teardown <percentage> <idle-timeout (forever | minutes)>;  
            }  
            rib-group group-name;  
            topology name {  
                community {  
                    target identifier;  
                }  
            }  
        }  
    }  
    flow {  
        no-validate policy-name;  
    }  
    labeled-unicast {  
        accepted-prefix-limit {  
            maximum number;  
            teardown <percentage> <idle-timeout (forever | minutes)>;  
        }  
        explicit-null {  
            connected-only;  
        }  
        prefix-limit {  
            maximum number;  
            teardown <percentage> <idle-timeout (forever | minutes)>;  
        }  
        resolve-vpn;  
        rib inet.3;  
    }  
}
```

```

        rib-group group-name;
    }
}
route-target {
    accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
    }
}
}
hold-time seconds;
import [ policy-names ];
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <private>;
local-preference local-preference;
log-updown;
metric-out metric;
multihop <ttl-value>;
multipath {
    multiple-as;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
passive;
peer-as autonomous-system;
preference preference;
remove-private;
tcp-aggressive-transmission;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
type type;
neighbor address {
    ... peer-specific-options ...
}
}

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
bgp],  
[edit protocols bgp],  
[edit routing-instances *routing-instance-name* protocols bgp]

|                                 |   |
|---------------------------------|---|
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>              | <p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple <b>group</b> statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the <b>group</b> statement.</p> <p>The <b>group</b> statement is one of the statements you must include in the configuration to run BGP on the routing device.</p> <p>Each group must contain at least one peer.</p> |
| <b>Options</b>                  | <p><b>group-name</b>—Name of the BGP group.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>BGP Feature Guide for Routing Devices</i></li></ul>  |

## hold-down

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>hold-down {     seconds;     flaps number;     period seconds; }</pre>  |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems logical-system-name routing-options bmp], [edit logical-systems logical-system-name routing-options bmp station station-name], [edit routing-options bmp], [edit routing-options bmp station station-name]</pre>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.3.</p>  |
| <b>Description</b>              | <p>If the connection to a BMP station flaps and the <b>hold-down</b> statement is configured, the station is prevented from reconnecting to the device for the specified period of time. A flap is when the TCP session unexpectedly switches from established to non-established. If you alter the configuration of the <b>hold-down</b> statement, the hold down timer and flap counter are reset.</p> <p>You can effectively disable the <b>hold-down</b> statement by setting the <b>flaps</b> option to 10 and the <b>period</b> option to 30 seconds.</p>  |
| <b>Options</b>                  | <p><b>seconds</b>—Specify the time in seconds to wait before allowing the BMP station to reconnect to the device.</p> <p><b>Default:</b> 600 seconds</p> <p><b>Range:</b> 30 through 65,535 seconds</p> <p><b>flaps number</b>—Specify the number of BMP station flaps allowed before terminating the connection to the BMP station and triggering the hold down timer.</p> <p><b>Default:</b> 3 flaps</p> <p><b>Range:</b> 2 to 10 flaps</p> <p><b>period seconds</b>—Specify the time in seconds for the BGP station flaps (specified using the <b>flaps</b> option) to occur before triggering the hold down timer. Every time a flap occurs, the number of flaps in the last time period is checked to see if the criteria is met.</p> <p>For example, if you defined the <b>period</b> as 60 seconds and the <b>flaps</b> as 4 and the BGP station flaps just 2 times in a 60 second period, the hold down timer would not be triggered. However, if the BGP station flaps 4 times in a 60 second period, the hold down timer would be triggered.</p> <p><b>Default:</b> 300 seconds</p> <p><b>Range:</b> 30 through 65,535 seconds</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

## hold-down-interval (BGP BFD Liveness Detection)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>holddown-interval <i>milliseconds</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure an interval specifying how long a BFD session must remain up before a state change notification is sent.</p> <p>When you configure the hold-down interval for the BFD protocol for EBGp, the BFD session is unaware of the BGP session during this time. In this case, if the BGP session goes down during the configured hold-down interval, BFD already assumes the BGP session is down and does not send a state change notification. The <b>holddown-interval</b> statement is supported only for EBGp peers at the <b>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]</b> hierarchy level. If the BFD session goes down and then comes back up during the configured hold-down interval, the timer is restarted. You must configure the hold-down interval on both EBGp peers. If you configure the hold-down interval for a multihop EBGp session, you must also configure a local IP address by including the <b>local-address</b> statement at the <b>[edit protocols bgp group <i>group-name</i>]</b> hierarchy level.</p>  |
| <b>Options</b>                  | <p><b><i>milliseconds</i></b>—Specify the hold-down interval value.</p> <p><b>Range:</b> 0 through 255,000</p> <p><b>Default:</b> 0</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- [Example: Configuring BFD for Static Routes on page 2918](#)
  - [bfd-liveness-detection on page 3643](#)



## hold-time (Protocols BGP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>hold-time seconds;</code>   |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>] </pre>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>         | <p>Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.</p> <p>The hold time is three times the interval at which keepalive messages are sent.</p> <p>BGP on the local routing device uses the smaller of either the local hold-time value or the peer's hold-time value received in the open message as the hold time for the BGP connection between the two peers.</p> <p>Starting in Junos OS Release 12.3, the BGP hold-time value can be zero (0). This implies that the speaker does not expect keepalive messages from its peer to maintain the BGP session. When negotiating between two peers, if one side requests a nonzero hold time and the other requests a zero hold time, the negotiation settles on the nonzero value and keepalive intervals are determined accordingly. Both sides must be set to zero for keepalive messages to stop being sent.</p> |
| <b>Options</b>             | <p><b>seconds</b>—Hold time.</p> <p><b>Range:</b> 10 through 65,535 seconds (or 0 for infinite hold time)</p> <p><b>Default:</b> 90 seconds</p>   |



**TIP:** When you set a hold-time value of 1 through 19 seconds, we recommend that you also configure the BGP `precision-timers` statement. The `precision-timers` statement ensures that if scheduler slip messages occur, the routing device continues to send keepalive messages. When the `precision-timers` statement is included, keepalive message generation is performed in a dedicated kernel thread, which helps to prevent BGP session flaps.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">BGP Messages Overview on page 3257</a></li><li>• <i>precision-timers</i></li></ul> |

## import (Protocols BGP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>import [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>         | <p>Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.</p> <p>If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices. The policy framework software evaluates the routing policies in a chain sequentially. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of <b>metric 500</b> defined in the next policy.</p> <p>It is also important to understand that in Junos OS, although an import policy (inbound route filter) might reject a route, not use it for traffic forwarding, and not include it in an advertisement to other peers, the router retains these routes as hidden routes. These hidden routes are not available for policy or routing purposes. However, they do occupy memory space on the router. A service provider filtering routes to control the amount of information being kept in memory and processed by a router might want the router to entirely drop the routes being rejected by the import policy.</p> <p>Hidden routes can be viewed by using the <b>show route receive-protocol bgp neighbor-address hidden</b> command. The hidden routes can then be retained or dropped from the routing table by configuring the <b>keep all   none</b> statement at the <b>[edit protocols bgp]</b> or <b>[edit protocols bgp group group-name]</b> hierarchy level.</p> |

The rules of BGP route retention are as follows:

- By default, all routes learned from BGP are retained, except those where the AS path is looped. (The AS path includes the local AS.)
- By configuring the **keep all** statement, all routes learned from BGP are retained, even those with the local AS in the AS path.
- By configuring the **keep none** statement, all routes received are discarded. When this statement is configured and the inbound policy changes, Junos OS re-advertises all the routes advertised by the peer.

**Options**    *policy-names*—Name of one or more policies.

**Required Privilege**    routing—To view this statement in the configuration.  
**Level**    routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring BGP Interactions with IGP's on page 3421](#)
- [Understanding Route Advertisement on page 3425](#)
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
- [export on page 3658](#)

## include-mp-next-hop

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | include-mp-next-hop;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | Enable multiprotocol updates to contain next-hop reachability information.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Examples: Configuring Multiprotocol BGP</i></li> </ul>   |

## initiation-message

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>initiation-message text;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station station-name</a> ],<br>[edit routing-options <a href="#">bmp</a> ],<br>[edit routing-options bmp <a href="#">station station-name</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.   |
| <b>Description</b>              | <p>(Optional) Allows you to specify an initiation message for a type 0 TLV to be sent to the BMP station. The message is transmitted when a BMP station establishes a connection to the device. You can provide some information to the BMP station system administrator (for example, a contact phone number). The initiation message includes a type 1 TLV containing the SNMP sysDescr value specified in RFC 1213 <i>Management Information Base for Network Management of TCP/IP-based internets: MIB-II</i> and a type 2 TLV containing the SNMP sysName value also from RFC 1213. The string in the initiation-message message is UTF-8.</p> <p>The normal time for sending an initiation message is when the BMP session is first established. However, an initiation message change also triggers the transmission of an initiation message to current BMP sessions.</p> <p>Another event that triggers the transmission of an initiation message is when you change in the sysName or sysDescr values in the SNMP configuration. The initiation message is sent to current BMP sessions.</p> |
| <b>Options</b>                  | <b>text</b> —Specify a character string for a type 0 TLV to send with the initiation message.<br><b>Range:</b> 1 through 255 characters  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>   |

## keep

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | keep (all   none);   |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>], [edit protocols bgp], [edit protocols bgp <i>group group-name</i>], [edit protocols bgp <i>group group-name neighbor address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i> <i>neighbor address</i>]</pre>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>         | <p>Control whether or not Junos OS keeps in memory and hides certain routes.</p> <p>If the <b>keep none</b> statement is used, Junos OS does not retain in memory and hide routes that are rejected because of a BGP import policy. Nor does BGP keep in memory and hide routes that are declared unfeasible due to BGP sanity checks. The <b>keep none</b> statement causes Junos OS to discard from memory the routes that are rejected due to BGP-specific logic or BGP evaluation. When a route is rejected because of some non-BGP-specific reason, the <b>keep none</b> statement has no effect on this route. This rejected route is retained in memory and hidden even though <b>keep none</b> is configured. An example of this type of hidden route is a route for which the protocol nexthop is unresolved.</p> <p>The routing table can retain the route information learned from BGP in one of the following ways:</p> <ul style="list-style-type: none"> <li>• Default (omit the <b>keep</b> statement)—Keep all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.</li> <li>• <b>keep all</b>—Keep all route information that was learned from BGP.</li> <li>• <b>keep none</b>—Discard routes that were received from a peer and that were rejected by import policy or other sanity checking, such as AS path or next hop. When you configure <b>keep none</b> for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer.</li> </ul> |

In an AS path healing situation, routes with looped paths theoretically could become usable during a soft reconfiguration when the AS path loop limit is changed. However, there is a significant memory usage difference between the default and **keep all**.

Consider the following scenarios:

- A peer readvertises routes back to the peer from which it learned them.

This can happen in the following cases:

- Another vendor's routing device advertises the routes back to the sending peer.
- The Junos OS peer's default behavior of not readvertising routes back to the sending peer is overridden by configuring **advertise-peer-as**.
- A provider edge (PE) routing device discards any VPN route that does not have any of the expected route targets.

When **keep all** is configured, the behavior of discarding routes received in the above scenarios is overridden.



**CAUTION:** If you add or remove **keep all** or **keep none** and the peer does not support session restart, the associated BGP sessions are restarted (flapped). To determine if a peer supports refresh, check for **Peer supports Refresh capability** in the output of the **show bgp neighbor** command.

|                                 |  |
|---------------------------------|--|
| <b>Default</b>                  | By default, BGP retains incoming rejected routes in memory and hides them. If you do not include the <b>keep</b> statement, most routes are retained in the routing table. BGP keeps all route information that was learned from BGP, except for routes whose AS path is looped and whose loop includes the local AS.                        |
| <b>Options</b>                  | <b>all</b> —Retain all routes.<br><br><b>none</b> —Discard routes that were received from a peer and that were rejected by import policy or other sanity checking. When <b>keep none</b> is configured for the BGP session and the inbound policy changes, Junos OS forces readvertisement of the full set of routes advertised by the peer. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">out-delay on page 3714</a></li><li>• <i>Interprovider VPN Example—MP-EBGP Between ISP Peer Routers</i></li><li>• <i>Example: Configuring Conditional Installation of Prefixes in a Routing Table</i></li></ul>   |



## key-chain (BGP BFD Authentication)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>key-chain <i>key-chain-name</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Associate a security key with the specified BFD session using the name of the security keychain. Each key has a unique start time within the keychain. Keychain authentication allows you to change the password information periodically without bringing down peering sessions. This keychain authentication method is referred to as <i>hitless</i> because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol.  |
| <b>Options</b>                  | <b><i>key-chain-name</i></b> —Name of the authentication keychain. The keychain name must match one of the keychains configured with the <b>key-chain <i>key-chain-name</i></b> statement at the [edit security authentication-key-chain] hierarchy level.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li> <li>• <a href="#">Example: Configuring BFD on Internal BGP Peer Sessions on page 3463</a></li> <li>• <a href="#">Example: Configuring BGP Route Authentication on page 3571</a></li> </ul>   |

- [Example: Configuring EBGp Multihop Sessions on page 3433](#)

## local-address (Protocols BGP)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>local-address address;</code>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>         | <p>Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.</p> <p>You generally configure a local address to explicitly configure the system's IP address from BGP's point of view. This IP address can be either an IPv6 or IPv4 address. Typically, an IP address is assigned to a loopback interface, and that IP address is configured here.</p> <p>For internal BGP (IBGP) peering sessions, generally the loopback interface (lo0) is used to establish connections between the IBGP peers. The loopback interface is always up as long as the device is operating. If there is a route to the loopback address, the IBGP peering session stays up. If a physical interface address is used instead and that interface goes up and down, the IBGP peering session also goes up and down. Thus, the loopback interface provides fault tolerance in case the physical interface or the link goes down, if the device has link redundancy.</p> <p>When a device peers with a remote device's loopback interface address, the local device expects BGP update messages to come from (be sourced by) the remote device's loopback interface address. The <b>local-address</b> statement enables you to specify the source information in BGP update messages. If you omit the <b>local-address</b> statement, the expected source of BGP update messages is based on the device's source address selection rules, which normally result in the egress interface address being the expected source of update messages. When this happens, the peering session is not established because a mismatch exists between the expected source address (the egress interface</p> |

of the peer) and the actual source (the loopback interface of the peer). To ensure that the expected source address matches the actual source address, specify the loopback interface address in the **local-address** statement.



**NOTE:** Although a BGP session can be established when only one of the paired routing devices has **local-address** configured, we strongly recommend that you configure **local-address** on both paired routing devices for IBGP and multihop EBGP sessions. The **local-address** statement ensures that deterministic fixed addresses are used for the BGP session end-points.

If you include the **default-address-selection** statement in the configuration, the software chooses the system default address as the source for most locally generated IP packets. For protocols in which the local address is unconstrained by the protocol specification, for example IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same methods as other locally generated IP packets.

**Default** If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address.

**Options** **address**—IPv6 or IPv4 address of the local end of the connection.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Internal BGP Peering Sessions on Logical Systems on page 3296](#)
- [Example: Configuring Internal BGP Peer Sessions on page 3285](#)
- [Understanding Internal BGP Peering Sessions on page 3284](#)
- [router-id on page 3036](#)

## local-address (Protocols BMP)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>local-address <i>address</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station</a> <i>station-name</i>],</p> <p>[edit routing-options <a href="#">bmp</a>],</p> <p>[edit routing-options bmp <a href="#">station</a> <i>station-name</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.3.</p>   |
| <b>Description</b>              | <p>(Optional) Specifies the IPv4 or IPv6 address for the BMP connection on the device. We recommend that you configure a local address. For both active and passive modes, configure a loopback local address. This provides a consistent local endpoint, is useful for debugging, and assures greater reliability for the BMP connection since it is not tied to a single router interface.</p> <p>For passive mode, specifying a local address is required. It also provides some security against a malicious BMP connection. For active mode, we also recommend configuring a local address to help ensure reliability.</p> <p>If you change the local address, the BMP station connection flaps when you commit the configuration.</p> |
| <b>Options</b>                  | <b><i>address</i></b> —Specify the IPv4 or IPv6 address for the BMP connection on the local device.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li> </ul>  |

## local-as

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>local-as <i>autonomous-system</i> &lt;loops <i>number</i>&gt; &lt;private   alias&gt; &lt;no-prepend-global-as&gt;;</code>  |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor   <i>address</i>]</pre> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>alias</b> option introduced in Junos OS Release 9.5.</p> <p><b>no-prepend-global-as</b> option introduced in Junos OS Release 9.6.</p>  |
| <b>Description</b>         | <p>Specify the local autonomous system (AS) number. An AS is a set of routing devices that are under a single technical administration and generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices.</p> <p>Internet service providers (ISPs) sometimes acquire networks that belong to a different AS. When this occurs, there is no seamless method for moving the BGP peers of the acquired network to the AS of the acquiring ISP. The process of configuring the BGP peers with the new AS number can be time-consuming and cumbersome. In this case, it might not be desirable to modify peer arrangements or configuration. During this kind of transition period, it can be useful to configure BGP-enabled devices in the new AS to use the former AS number in BGP updates. This former AS number is called a <i>local</i> AS.</p>   |



**NOTE:** If you are using BGP on the routing device, you must configure an AS number before you specify the local as number.

In Junos OS Release 9.1 and later, the AS numeric range in plain-number format is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For

example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

**Options** **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

**autonomous-system**—AS number.

**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format

**Range:** 0.0 through 65535.65535 in AS-dot notation format

**loops number**—(Optional) Specify the number of times detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.



**NOTE:** If you configure the local AS values for any BGP group, the detection of routing loops is performed using both the AS and the local AS values for all BGP groups.

If the local AS for the EBGP or IBGP peer is the same as the current AS, do not use the **local-as** statement to specify the local AS number.

When you configure the local AS within a VRF, this impacts the AS path loop-detection mechanism. All of the **local-as** statements configured on the device are part of a single AS domain. The AS path loop-detection mechanism is based on looking for a matching AS present in the domain.

**Range:** 1 through 10

**Default:** 1

**no-prepend-global-as**—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

**private**—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.



**NOTE:** The **private** and **alias** options are mutually exclusive. You cannot configure both options with the same **local-as** statement.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Examples: Configuring BGP Local AS on page 3362</a></li><li>• <a href="#">Example: Configuring a Local AS for EBGp Sessions on page 3367</a></li><li>• <a href="#">autonomous-system on page 2945</a></li><li>• <a href="#">family on page 3659</a></li></ul> |

---

## local-port

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>local-port port;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station station-name</a> ],<br>[edit routing-options <a href="#">bmp</a> ],<br>[edit routing-options bmp <a href="#">station station-name</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.   |
| <b>Description</b>              | <p>Specifies the listening port for the BMP station connection.</p> <p>If you configure the <a href="#">connection-mode</a> statement as <b>active</b>, do not configure the <b>local-port</b> statement. If you configure the <a href="#">connection-mode</a> statement as <b>passive</b>, you must configure <b>local-port</b> statement.</p> <p>If you change the local port, the BMP station connection flaps when you commit the configuration.</p> |
| <b>Options</b>                  | <p><b>port</b>—Specify the local port for the BMP station connection.</p> <p><b>Range:</b> 1 through 65,535</p>  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>   |



## local-preference

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>local-preference local-preference;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit protocols bgp],<br/>         [edit protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Modify the value of the <b>LOCAL_PREF</b> path attribute, which is a metric used by BGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.</p> <p>The <b>LOCAL_PREF</b> path attribute always is used in inbound routing policy and is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.</p>  |
| <b>Default</b>                  | If you omit this statement, the <b>LOCAL_PREF</b> path attribute, if present, is not modified.   |
| <b>Options</b>                  | <p><b>local-preference</b>—Preference to assign to routes learned from BGP or from the group or peer.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> If the <b>LOCAL_PREF</b> path attribute is present, do not modify its value. If a BGP route is received without a <b>LOCAL_PREF</b> attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a <b>LOCAL_PREF</b> value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a <b>LOCAL_PREF</b> value of 100.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Local Preference Value for BGP Routes on page 3310</a></li> <li>• <a href="#">Understanding Internal BGP Peering Sessions on page 3284</a></li> </ul>  |

- [preference on page 3723](#)

## log-updown (Protocols BGP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | log-updown;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit protocols bgp],<br/>[edit protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit protocols bgp <i>group</i> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Specify to generate a log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.<br/>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Preventing BGP Session Resets on page 3591</a></li><li>• <i>Junos OS Administration Library for Routing Devices</i></li><li>• <a href="#">traceoptions on page 3740</a></li></ul>  |

## loops

|                        |   |
|------------------------|---|
| <b>Syntax</b>          | <code>loops <i>number</i>;</code>   |
| <b>Hierarchy Level</b> | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp family <i>address-family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family <i>address-family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> local-as],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>address-family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> local-as]</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp local-as],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options autonomous-system <i>as-number</i>],</p> <p>[edit protocols bgp family <i>address-family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> family <i>address-family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> local-as],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family <i>address-family</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> local-as]</p> <p>[edit protocols bgp local-as],</p> <p>[edit routing-options autonomous-system <i>as-number</i>]</p> |

**Release Information** Statement introduced in Junos OS Release 9.6.

**Description** Globally, for the local-AS BGP attribute, or the specified address family, allow the local device's AS number to be in the received AS paths, and specify the number of times detection of the local device's AS number in the AS\_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the local device's AS number is detected in the path one or more times. This prevents routing loops and is the default behavior. If you configure **loops 2**, the route is hidden if the local device's AS number is detected in the path two or more times.

Some examples of BGP address families are as follows:

- **inet unicast**
- **inet-vpn multicast**
- **inet6 any**
- **l2vpn auto-discovery-only**
- ...

This list is truncated for brevity. For a complete list of protocol families for which you can specify the **loops** statement, enter the **help apropos loops** configuration command at the **[edit protocols bgp]** hierarchy level on your device.

```
[edit protocols bgp]
user@host# help apropos loops
set family inet unicast loops
    Allow local AS in received AS paths
set family inet unicast loops <loops>
    AS-Path loop count
set family inet multicast loops
```

```
    Allow local AS in received AS paths
set family inet multicast loops <loops>
    AS-Path loop count
set family inet flow loops
    Allow local AS in received AS paths
set family inet flow loops <loops>
    AS-Path loop count
set family inet any loops
    Allow local AS in received AS paths
set family inet any loops <loops>
    AS-Path loop count
set family inet labeled-unicast loops
    Allow local AS in received AS paths
...
```



**NOTE:** When you configure the `loops` statement for a specific BGP address family, that value is used to evaluate the AS path for routes received by a BGP peer for the specified address family, rather than the `loops` value configured for the global AS number with the `loops` statement at the `[edit routing-options autonomous-system as-number]` hierarchy level.

**Options** *number*—Number of times detection of the AS number in the AS\_PATH attribute causes the route to be discarded or hidden.

**Range:** 1 through 10

**Default:** 1

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Disabling Suppression of Route Advertisements*
- [autonomous-system on page 2945](#)
- [family on page 3659](#)
- [local-as on page 3684](#)

## loose-check (BGP BFD Authentication)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | loose-check ;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp bgp bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection authentication],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection authentication]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li> <li>• <a href="#">Example: Configuring BFD on Internal BGP Peer Sessions on page 3463</a></li> <li>• <a href="#">Example: Configuring BGP Route Authentication on page 3571</a></li> <li>• <a href="#">Example: Configuring EBGp Multihop Sessions on page 3433</a></li> </ul>   |

## maximum-ecmp

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>maximum-ecmp <i>next-hops</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit chassis]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for QFX switches.   |
| <b>Description</b>              | Configure 16, 32, or 64 ECMP next hops for RSVP or LDP LSPs, or MPLS static LSPs that are configured using <code>set protocols mpls static-label-switched-path</code> . |
| <b>Default</b>                  | 16  |
| <b>Options</b>                  | <b>next-hops</b> —Specify the number of next hops (16, 32, or 64) for RSVP or LDP LSPs, or MPLS static LSPs   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>•</li></ul>   |

## metric-out (Protocols BGP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>metric-out (<i>metric</i>   minimum-igp <i>offset</i>   igp (delay-med-update   <i>offset</i>);</code>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <i>group group-name</i>],</p> <p>[edit protocols bgp <i>group group-name neighbor address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group group-name neighbor address</i>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Option <b>delay-med-update</b> introduced in Junos OS Release 9.0.</p>   |
| <b>Description</b>         | <p>Specify the metric for all routes sent using the multiple exit discriminator (MED, or <b>MULTI_EXIT_DISC</b>) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the <b>metric</b> option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the <b>multihop</b> command—you can specify a variable metric by including the <b>minimum-igp</b> or <b>igp</b> option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the <b>igp</b> or <b>minimum-igp</b> statement) by specifying a value for <b>offset</b>. The metric is increased by specifying a positive value for <b>offset</b>, and decreased by specifying a negative value for <b>offset</b>.</p> <p>In Junos OS Release 9.0 and later, you can specify that a BGP group or peer not advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the <b>med-igp-update-interval minutes</b> statement at the [edit routing-options] hierarchy level.</p> |
| <b>Options</b>             | <p><b>delay-med-update</b>—Specify that a BGP group or peer configured with the <b>metric-out igp</b> statement not advertise MED updates unless the current MED value is lower than</p>  |

the previously advertised MED value, or another attribute associated with the route has changed, or the BGP peer is responding to a refresh route request.



**NOTE:** You cannot configure the `delay-med-update` statement at the global BGP level.

**igp**—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop. Routes learned from an EBGP peer usually have a next hop on a directly connected interface and thus the IGP value is equal to zero. This is the value advertised.

**metric**—Primary metric on all routes sent to peers.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**Default:** No metric is sent.

**minimum-igp**—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value. When you change a neighbor's export policy from any configuration to a configuration that sets the minimum IGP offset on an exported route, the advertised MED is not updated if the value would increase as a result, even if the previous configuration does not use a minimum IGP-based MED value. This behavior helps to prevent unnecessary route flapping when an IGP cost changes, by not forcing a route update if the metric value increases past the previous lowest known value.

**offset**—Increases or decreases the metric by this value.

**Range:**  $-2^{31}$  through  $2^{31} - 1$

**Default:** None

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates on page 3352</a></li><li>• <a href="#">Examples: Configuring BGP MED on page 3323</a></li><li>• <a href="#">Example: Configuring the MED Attribute Directly on page 3325</a></li><li>• <a href="#">Understanding the MED Attribute on page 3323</a></li><li>• <a href="#">med-igp-update-interval on page 2996</a></li></ul> |
|------------------------------|---|



## minimum-interval (BFD Liveness Detection)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>minimum-interval <i>milliseconds</i>;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>   |
| <b>Description</b>         | <p>Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <a href="#">minimum-interval</a> (specified under the <a href="#">transmit-interval</a> statement) and <a href="#">minimum-receive-interval</a> statements.</p>   |
| <b>Options</b>             | <p><i>milliseconds</i>—Specify the minimum interval value for BFD liveliness detection.</p> <p><b>Range:</b> 1 through 255,000</p>   |

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BFD for Layer 2 VPN and VPLS</i></li><li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li><li>• <a href="#">bfd-liveness-detection on page 3643</a></li><li>• <a href="#">minimum-receive-interval on page 3699</a></li><li>• <a href="#">transmit-interval on page 3745</a></li></ul> |

## minimum-interval (transmit-interval)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>minimum-interval <i>milliseconds</i>;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>   |
| <b>Description</b>         | Configure the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using   |

this statement at this hierarchy level, you can configure the minimum transmit interval using the [minimum-interval](#) statement at the **bfd-liveness-detection** hierarchy level.

**Options** *milliseconds*—Minimum transmit interval value.

**Range:** 1 through 255,000



**NOTE:** The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

---

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes on page 2918](#)
- [bfd-liveness-detection on page 3643](#)
- [minimum-interval on page 3695](#)
- [threshold on page 3738](#)

## minimum-receive-interval (BFD Liveness Detection)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>minimum-receive-interval <i>milliseconds</i>;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>   |
| <b>Description</b>         | Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the <b>minimum-interval</b> statement.   |
| <b>Options</b>             | <p><b><i>milliseconds</i></b>—Specify the minimum receive interval value.</p> <p><b>Range:</b> 1 through 255,000</p>   |

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes on page 2918](#)
- [bfd-liveness-detection on page 3643](#)
- [minimum-interval on page 3695](#)
- [transmit-interval on page 3745](#)

---

## monitor (Protocols BMP)

---

**Syntax** monitor (enable | disable);

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp [bmp](#)],  
[edit logical-systems *logical-system-name* protocols bgp group *group-name* bmp],  
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* bmp],  
[edit logical-systems *logical-system-name* routing-options [bmp](#)],  
[edit logical-systems *logical-system-name* routing-options bmp [station](#) *station-name*],  
[edit protocols bgp [bmp](#)],  
[edit protocols bgp group *group-name* bmp],  
[edit protocols bgp group *group-name* neighbor *address* bmp],  
[edit routing-options [bmp](#)],  
[edit routing-options bmp [station](#) *station-name*]

**Release Information** Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.  
Statement introduced in Junos OS Release 13.3.

**Description** BMP monitoring is enabled by default. You can explicitly enable BMP monitoring or disable it. You can also selectively enable or disable BMP monitoring at various hierarchy levels (for example, [edit protocols bgp group *group-name*] or [edit protocols bgp group *group-name* neighbor *address*]). If you disable BMP monitoring, withdrawal messages are sent for any previously advertised routes. These are followed by a down message. If you enable BMP monitoring, an up message is sent first and then the route advertisements follow.

**Options** **enable**—Enable BMP monitoring.  
**Default:** BMP monitoring is enabled by default.  
**disable**—Disable BMP monitoring.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## mtu-discovery


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | mtu-discovery;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/> [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/> [edit protocols bgp],<br/> [edit protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/> [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure TCP path maximum transmission unit (MTU) discovery.</p> <p>TCP path MTU discovery enables BGP to automatically discover the best TCP path MTU for each BGP session. In Junos OS, TCP path MTU discovery is disabled by default for all BGP neighbor sessions.</p> <p>When MTU discovery is disabled, TCP sessions that are not directly connected transmit packets of 512-byte maximum segment size (MSS). These small packets minimize the chances of packet fragmentation at a device along the path to the destination. However, because most links use an MTU of at least 1500 bytes, 512-byte packets do not result in the most efficient use of link bandwidth. For directly connected EBGP sessions, MTU mismatches prevent the BGP session from being established. As a workaround, enable path MTU discovery within the EBGP group.</p> <p>Path MTU discovery dynamically determines the MTU size on the network path between the source and the destination, with the goal of avoiding IP fragmentation. Path MTU discovery works by setting the Don't Fragment (DF) bit in the IP headers of outgoing packets. When a device along the path has an MTU that is smaller than the packet, the device drops the packet. The device also sends back an ICMP Fragmentation Needed (Type 3, Code 4) message that contains the device's MTU, thus allowing the source to reduce its path MTU appropriately. The process repeats until the MTU is small enough to traverse the entire path without fragmentation.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |

**Related  
Documentation**

- [Example: Limiting TCP Segment Size for BGP on page 3586](#)
- *Configuring Junos OS for IPv6 Path MTU Discovery*
- *Configuring Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections*



## multihop

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>multihop {     no-nexthop-change;     ttl <i>ttl-value</i>; }</pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit protocols bgp],<br/>         [edit protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>         | <p>Configure an EBGp multihop session.</p> <p>For Layer 3 VPNs, you configure the EBGp multihop session between the PE and CE routing devices. This allows you to configure one or more routing devices between the PE and CE routing devices.</p> <p>An external confederation peer is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case because multihop behavior is implied.</p> <p>If you have external BGP confederation peer-to-loopback addresses, you still need the multihop configuration.</p>  |
|                            | <div>  <p><b>NOTE:</b> You cannot configure the <code>accept-remote-nexthop</code> statement at the same time.</p> </div>  |
| <b>Default</b>             | <p>If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.</p>   |

The remaining statements are explained separately.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li data-bbox="466 432 1373 464">• <a href="#">Example: Configuring EBGp Multihop Sessions on page 3433</a></li><li data-bbox="466 478 1373 510">• <i>Configuring EBGp Multihop Sessions Between PE and CE Routers in Layer 3 VPNs</i></li><li data-bbox="466 525 1373 556">• <a href="#">accept-remote-nextthop on page 3630</a></li><li data-bbox="466 571 1373 602">• <i>no-nextthop-change</i></li><li data-bbox="466 617 1373 648">• <i>tth</i></li></ul> |

## multiplier (BFD Liveness Detection)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>multiplier <i>number</i>;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>   |
| <b>Description</b>         | Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.  |
| <b>Options</b>             | <p><i>number</i>—Number of hello packets.</p> <p><b>Range:</b> 1 through 255</p> <p><b>Default:</b> 3</p>  |

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring BFD for Layer 2 VPN and VPLS</i></li><li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li><li>• <a href="#">bfd-liveness-detection on page 3643</a></li></ul> |

## neighbor (Protocols BGP)

```
Syntax  neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                damping;
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
                topology name {
                    community {
                        target identifier;
                    }
                }
            }
        }
        flow {
            no-validate policy-name;
        }
        labeled-unicast {
            accepted-prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            resolve-vpn;
            rib inet.3;
        }
    }
}
```

```
        rib-group group-name;  
        topology name {  
            community {  
                target identifier;  
            }  
        }  
    }  
}  
route-target {  
    advertise-default;  
    external-paths number;  
    accepted-prefix-limit {  
        maximum number;  
        teardown <percentage> <idle-timeout (forever | minutes)>;  
    }  
    prefix-limit {  
        maximum number;  
        teardown <percentage> <idle-timeout (forever | minutes)>;  
    }  
}  
signaling {  
    prefix-limit {  
        maximum number;  
        teardown <percentage> <idle-timeout (forever | minutes)>;  
    }  
}  
}  
graceful-restart {  
    disable;  
    restart-time seconds;  
    stale-routes-time seconds;  
}  
hold-time seconds;  
import [ policy-names ];  
ipsec-sa ipsec-sa;  
keep (all | none);  
local-address address;  
local-as autonomous-system <private>;  
local-interface interface-name;  
local-preference preference;  
log-updown;  
metric-out (metric | minimum-igp <offset> | igp <offset>);  
mtu-discovery;  
multihop <ttl-value>;  
multipath {  
    multiple-as;  
}  
no-aggregator-id;  
no-client-reflect;  
out-delay seconds;  
passive;  
peer-as autonomous-system;  
preference preference;  
tcp-aggressive-transmission;  
tcp-mss segment-size;  
traceoptions {
```

```

    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  vpn-apply-export;
}

```

|                          |  |
|--------------------------|--|
| Hierarchy Level          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>]</p>   |
| Release Information      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| Description              | <p>Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple <b>neighbor</b> statements.</p> <p>By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the <b>neighbor</b> statement.</p> <p>The <b>neighbor</b> statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an <b>allow all</b> statement in place of a <b>neighbor</b> statement.)</p> |
| Options                  | <p><b>address</b>—IPv6 or IPv4 address of a single peer.</p> <p>The remaining statements are explained separately.</p>   |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| Related Documentation    | <ul style="list-style-type: none"> <li>• <i>BGP Feature Guide for Routing Devices</i></li> </ul>   |

## no-adaptation (BFD Liveness Detection)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-adaptation;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>  |
| <b>Description</b>              | Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring BFD for Layer 2 VPN and VPLS</li> </ul>   |



- [Example: Configuring BFD for Static Routes on page 2918](#)
- [bfd-liveness-detection on page 3643](#)

## no advertise-peer-as

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-advertise-peer-as;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | Enable the default behavior of suppressing AS routes.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BGP Route Advertisement on page 3425</a></li> <li>• <a href="#">Understanding Route Advertisement on page 3425</a></li> <li>• <a href="#">advertise-peer-as on page 3634</a></li> </ul>   |

## no-aggregator-id

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-aggregator-id;  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit protocols bgp],<br/>[edit protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Prevent different routing devices within an AS from creating aggregate routes that contain different AS paths.</p> <p>Junos OS performs route aggregation, which is the process of combining the characteristics of different routes so that only a single route is advertised. Aggregation reduces the amount of information that BGP must store and exchange with other BGP systems. When aggregation occurs, the local routing device adds the local AS number and the router ID to the aggregator path attribute. The <b>no-aggregator-id</b> statement causes Junos OS to place a 0 in the router ID field and thus eliminate the possibility of having multiple aggregate advertisements in the network, each with different path information.</p>  |
| <b>Default</b>                  | If you omit this statement, the router ID is included in the BGP aggregator path attribute.  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Update Messages on page 3258</a></li></ul>   |

## no-client-reflect

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-client-reflect;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements. Route reflection provides a way to decrease BGP control traffic and minimizing the number of update messages sent within the AS.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BGP Route Reflectors on page 3547</a></li> <li>• <a href="#">cluster on page 3650</a></li> </ul>   |

## out-delay

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>out-delay seconds;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>         | <p>Control how often BGP and the routing table exchange route information by specifying how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates and to avoid sending updates too often.</p> <p>Alternatively or in addition, external BGP (EBGP) sessions can also use the route-flap damping mechanism upon the reception of BGP messages coming from an external neighbor.</p> <p>BGP stores the route information it receives from update messages in the routing table, and the routing table exports active routes from the routing table into BGP. BGP then advertises the exported routes to its peers. The <b>out-delay</b> statement enables a form of rate limiting. The delay is added to each update for each prefix individually. When a routing device changes its best path to a destination prefix, the device does not inform its peer about the change unless the route has been present in its routing table for the specified <b>out-delay</b>. If you use <b>out-delay</b> to perform rate-limiting, you can expect a less bursty pattern of updates. You will see a pattern in which updates arrive in a steady flow, and two updates for the same prefix are always spaced by at least the <b>out-delay</b> timer value (for example, 30 seconds). Thus, the <b>out-delay</b> setting is useful for limiting oscillation (sometimes called <i>churn</i>) in a network. Keep in mind that, regardless of the <b>out-delay</b> setting, BGP peers exchange routes immediately after neighbor establishment. The <b>out-delay</b> setting is only designed to delay the exchange of routes between BGP and the local routing table.</p> |

Caution is warranted because an **out-delay** can delay convergence. If your network is configured in a way that avoids oscillation, setting an **out-delay** is not necessary.

When configured, the **out-delay** value displays as **Outbound Timer** when using **show bgp group** or **show bgp group neighbor** commands.


**Default** By default, the exchange of route information between BGP and the routing table occurs immediately after the routes are received. This immediate exchange of route information might cause instabilities in the network reachability information. If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.

**Options** *seconds*—Output delay time.  
**Range:** 0 through 65,535 seconds  
**Default:** 0 seconds

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • [keep on page 3677](#)

## outbound-route-filter

|  |   |
|--|---|
| <b>Syntax</b>  | <pre> outbound-route-filter {     <b>bgp-orf-cisco-mode</b>;     prefix-based {         accept {             (inet   inet6);         }     } } </pre>   |
| <b>Hierarchy Level</b>   | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre> |
| <b>Release Information</b>   | <p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>   | Configure a BGP peer to accept outbound route filters from a remote peer.   |
| <b>Options</b>   | <p><b>accept</b>—Specify that outbound route filters from a BGP peer be accepted.</p> <p><b>inet</b>—Specify that IPv4 prefix-based outbound route filters be accepted.</p> <p><b>inet6</b>—Specify that IPv6 prefix-based outbound route filters be accepted.</p>  |
| <div>  <p><b>NOTE:</b> You can specify that both IPv4 and IPv6 outbound route filters be accepted.</p> </div> |   |
| <p><b>prefix-based</b>—Specify that prefix-based filters be accepted.</p> <p>The <b>bgp-orf-cisco-mode</b> statement is explained separately.</p>  |   |
| <b>Required Privilege Level</b>  | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |

- Related Documentation**
- [Example: Configuring BGP Prefix-Based Outbound Route Filtering on page 3429](#)

## passive (Protocols BGP)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | passive;  |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> <i>neighbor</i> <i>address</i>]</pre> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | Configure the routing device so that active open messages are not sent to the peer. Once you configure the routing device to be passive, the routing device will wait for the peer to issue an open request before a message is sent.   |
| <b>Default</b>                  | If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Preventing BGP Session Flaps When VPN Families Are Configured on page 3591</a></li> </ul>   |

## path-selection

---

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>path-selection {<br/>  (always-compare-med   cisco-non-deterministic   external-router-id);<br/>  as-path-ignore;<br/>  l2vpn-use-bgp-rules;<br/>  med-plus-igp {<br/>    igp-multiplier <i>number</i>;<br/>    med-multiplier <i>number</i>;<br/>  }<br/>}</pre>  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols bgp],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br>[edit protocols bgp],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp]   |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br><b>med-plus-igp</b> option introduced in Junos OS Release 8.1.<br><b>as-path-ignore</b> and <b>l2vpn-use-bgp-rules</b> options introduced in Junos OS Release 10.2. |
| <b>Description</b>         | Configure BGP path selection.   |
| <b>Default</b>             | If the <b>path-selection</b> statement is not included in the configuration, only the multiple exit discriminators (MEDs) of routes that have the same peer ASs are compared.   |
| <b>Options</b>             | <b>always-compare-med</b> —Always compare MEDs whether or not the peer ASs of the compared routes are the same.   |



**NOTE:** We recommend that you configure the **always-compare-med** option.

---

**as-path-ignore**—In the best-path algorithm, skip the step that compares the autonomous system (AS) path lengths. By default, the best-path algorithm evaluates the length of the AS paths and prefers the route with the shortest AS path length.

---



**NOTE:** The **as-path-ignore** statement is not supported with routing instances.

---

**cisco-non-deterministic**—Emulate the Cisco IOS default behavior. This mode evaluates routes in the order that they are received and does not group them according to their neighboring AS. With **cisco-non-deterministic** mode, the active path is always first. All inactive, but eligible, paths follow the active path and are maintained in the order



in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.

As an example, suppose you have three path advertisements for the 192.168.1.0 /24 route:

- Path 1—learned through EBGp; AS Path of 65010; MED of 200
- Path 2—learned through IBGP; AS Path of 65020; MED of 150; IGP cost of 5
- Path 3—learned through IBGP; AS Path of 65010; MED of 100; IGP cost of 10

These advertisements are received in quick succession, within a second, in the order listed. Path 3 is received most recently, so the routing device compares it against path 2, the next most recent advertisement. The cost to the IBGP peer is better for path 2, so the routing device eliminates path 3 from contention. When comparing paths 1 and 2, the routing device prefers path 1 because it is received from an EBGp peer. This allows the routing device to install path 1 as the active path for the route.



**NOTE:** We do not recommend using this configuration option in your network. It is provided solely for interoperability to allow all routing devices in the network to make consistent route selections.

**external-router-id**—Compare the router ID between external BGP paths to determine the active path.

**igp-multiplier *number***—The multiplier value for the IGP cost to a next-hop address. This option is useful for making the MED and IGP cost comparable.

**Range:** 1 through 1000

**Default:** 1

**med-multiplier *number***—The multiplier value for the MED calculation. This option is useful for making the MED and IGP cost comparable.

**Range:** 1 through 1000

**Default:** 1

**med-plus-igp**—Add the IGP cost to the indirect next-hop destination to the MED before comparing MED values for path selection. This statement only affects best-path selection. It does not affect the advertised MED.

The other option is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding BGP Path Selection](#)
- [Example: Ignoring the AS Path Attribute When Selecting the Best Path on page 3449](#)

## peer-as (Protocols BGP)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>peer-as <i>autonomous-system</i>;</code>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>         | <p>Specify the neighbor (peer) autonomous system (AS) number.</p> <p>For EBGP, the peer is in another AS, so the AS number you specify in the <b>peer-as</b> statement must be different from the local router's AS number, which you specify in the <b>autonomous-system</b> statement. For IBGP, the peer is in the same AS, so the two AS numbers that you specify in the <b>autonomous-system</b> and <b>peer-as</b> statements must be the same.</p> <p>The AS numeric range in plain-number format has been extended in Junos OS Release 9.1 to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, <i>BGP Support for Four-octet AS Number Space</i>. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of the Junos OS support 2-byte AS numbers.</p> <p>In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <i>&lt;16-bit high-order value in decimal&gt;.&lt;16-bit low-order value in decimal&gt;</i>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.</p> <p>With the introduction of 4-byte AS numbers, you might have a combination of routers that support 4-byte AS numbers and 2-byte AS numbers. For more information about what happens when establishing BGP peer relationships between 4-byte and 2-byte capable routers, see the following topics:</p> |

- *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview.*

**Options** *autonomous-system*—AS number.  
**Range:** 1 through 4,294,967,295 ( $2^{32} - 1$ ) in plain-number format for 4-byte AS numbers  
**Range:** 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)  
**Range:** 0.0 through 65535.65535 in AS-dot notation format for 4-byte AS numbers

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

## post-policy

**Syntax** `post-policy {  
 exclude-non-eligible;  
}`

**Hierarchy Level** [edit protocols bgp bmp [route-monitoring](#)],  
 [edit protocols bgp group *group-name* bmp route-monitoring],  
 [edit protocols bgp group neighbor *group-name* neighbor *address* bmp route-monitoring],  
 [edit routing-options bmp route-monitoring],  
 [edit routing-options bmp station *station-name* route-monitoring]

**Release Information** Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.  
 Statement introduced in Junos OS Release 13.3.

**Description** For BMP route monitoring, allows you to excludes routes that are non-eligible for the decision process (for example, protocol nexthop not resolved). This represents the view of the BGP routes after running the import policy. If the import policy has rejected the BGP route, the route does not exist in the post policy view.

**Options** *exclude-non-eligible*—Exclude routes that are non-eligible for the decision process.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.


**Related Documentation** • [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

## pre-policy

---


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>pre-policy {<br/>    exclude-non-feasible;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit protocols bgp bmp <a href="#">route-monitoring</a> ],<br>[edit protocols bgp group <i>group-name</i> bmp <a href="#">route-monitoring</a> ],<br>[edit protocols bgp group neighbor <i>group-name</i> neighbor <i>address</i> bmp <a href="#">route-monitoring</a> ],<br>[edit routing-options bmp <a href="#">route-monitoring</a> ],<br>[edit routing-options bmp station <i>station-name</i> <a href="#">route-monitoring</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.  |
| <b>Description</b>              | Excludes routes that are non-feasible from the BMP route monitoring decision process (for example, a route loop). This represents the view of the BGP routes before running the import policy.  |
| <b>Options</b>                  | <b>exclude-non-feasible</b> —Exclude routes that are non-feasible for the decision process.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>  |

## preference (Protocols BGP)


|  |  |
|--|--|
| <b>Syntax</b>  | <code>preference <i>preference</i>;</code>   |
| <b>Hierarchy Level</b>   | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit protocols bgp],<br/>         [edit protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>   | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>   | <p>Specify the preference for routes learned from BGP.</p> <p>At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.</p> <p>At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.</p>   |
| <div>  <b>NOTE:</b> Do not set preference2 for BGP route-policy.         </div> |  |
| <b>Options</b>   | <p><b>preference</b>—Preference to assign to routes learned from BGP or from the group or peer.<br/> <b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)<br/> <b>Default:</b> 170 for the primary preference</p>   |
| <b>Required Privilege Level</b>  | <p>routing—To view this statement in the configuration.<br/>         routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"> <li>• <a href="#">local-preference on page 3687</a></li> <li>• <a href="#">Example: Configuring the Preference Value for BGP Routes on page 3443</a></li> </ul>   |

## priority (Protocols BMP)

---

|   |   |
|---|---|
| <b>Syntax</b>   | priority (high   medium   low);   |
| <b>Hierarchy Level</b>  | [edit logical-systems <i>logical-system-name</i> routing-options <b>bmp</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <b>station station-name</b> ],<br>[edit routing-options <b>bmp</b> ],<br>[edit routing-options bmp <b>station station-name</b> ]   |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.  |
| <b>Description</b>  | Specifies the dispatch priority for BMP. The dispatch priority controls the frequency with which the device is able to forward BMP messages to BMP stations.  |
| <b>Options</b>  | <b>high</b> —Specifies that the routing protocol process handle BMP requests with high urgency.<br><br><b>medium</b> —Specifies that the routing protocol process handle BMP requests with medium urgency.<br><br><b>low</b> —Specifies that the routing protocol process handle BMP requests with low urgency.<br><b>Default:</b> The default dispatch priority is <b>low</b> to minimize interference with other routing protocol process priorities and to match the behavior of previous versions of BMP. |
| <div> <b>NOTE:</b> Setting high or medium priority may reduce the performance of the routing protocol process in its handling route convergence or other work.</div> |   |
| <b>Required Privilege Level</b>   | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>  |

## remove-private

|  |  |
|--|--|
| <b>Syntax</b>  | remove-private all replace nearest;  |
| <b>Hierarchy Level</b>   | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>group</b> <i>group-name</i> neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> <b>neighbor</b> <i>address</i>]</p> |
| <b>Release Information</b>   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>   | <p>When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. If the AS path contains the AS number of the external BGP (EBGP) neighbor, BGP does not remove the private AS number.</p>  |
| <div>  <p><b>NOTE:</b> As of Junos OS 10.0R2 and higher, if there is a need to send prefixes to an EBGP peer that has an AS number that matches an AS number in the AS path, consider using the <code>as-override</code> statement instead of the <code>remove-private</code> statement.</p> </div>                 |  |
| <p>The operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.</p> <p>The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.</p> <p>The set of reserved AS numbers is in the range from 64,512 through 65,535.</p> |  |
| <b>Options</b>   | <p><b>all</b>—Remove all private AS numbers from the original path. Do not stop the process of removing private AS numbers, even if a public AS number is encountered.</p>   |

**nearest**—When you use the **all** and **replace** options, choose the last (right-most) public AS number encountered in the original AS path for the replacement value, as the AS path is processed from left to right. If no public AS number is encountered, the default replacement value is used. (See the **replace** option for information about the default replacement value.)

**replace**—When you use the **all** option, instead of removing private AS numbers, perform a replace operation. The default replacement value for the private AS number is the local AS number at the BGP group level for the BGP peer. If you are unsure about the replacement value, check the local AS value displayed in the output of the **show bgp group group-name** command.

|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Removing Private AS Numbers from AS Paths on page 3457</a></li></ul> |
|------------------------------|---|



## restart-time (BGP Graceful Restart)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>restart-time seconds;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit protocols (bgp   rip   ripng) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (bgp   rip   ripng) <a href="#">graceful-restart (Enabling Globally)</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Configure the duration of the BGP, RIP, or next-generation RIP (RIPng) graceful restart period.  |
| <b>Options</b>                  | <p><b>seconds</b>—Length of time for the graceful restart period.</p> <p><b>Range:</b> 1 through 600 seconds</p> <p><b>Default:</b> Varies by protocol:</p> <ul style="list-style-type: none"> <li>• BGP—120 seconds</li> <li>• RIP and RIPng—60 seconds</li> </ul>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li> <li>• <a href="#">Configuring Graceful Restart Options for RIP and RIPng on page 2266</a></li> <li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li> <li>• <a href="#">stale-routes-time on page 2308</a></li> </ul>   |

## route-monitoring

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>route-monitoring {<br/>    none;<br/>    post-policy {<br/>        exclude-non-eligible;<br/>    }<br/>    pre-policy {<br/>        exclude-non-feasible;<br/>    }<br/>}</pre>   |
| Hierarchy Level          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp <b>bmp</b>],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bmp],<br/>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bmp],<br/>[edit logical-systems <i>logical-system-name</i> routing-options <b>bmp</b>],<br/>[edit logical-systems <i>logical-system-name</i> routing-options bmp <b>station</b> <i>station-name</i>],<br/>[edit protocols bgp <b>bmp</b>],<br/>[edit protocols bgp group <i>group-name</i> bmp],<br/>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bmp],<br/>[edit routing-options <b>bmp</b>],<br/>[edit routing-options bmp <b>station</b> <i>station-name</i>]</p> |
| Release Information      | <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.3.</p>  |
| Description              | <p>Specify whether BMP should send pre-policy route monitoring messages, post-policy route monitoring messages, both types of messages, or none at all. The pre-policy can be configured to exclude routes that are non-feasible for the decision process (for example, a route loop). The post-policy can be configured to exclude routes that are not eligible for the decision process (for example, protocol nexthop not resolved).</p> <p>You can also selectively enable or disable BMP route monitoring at various hierarchy levels (for example, [edit protocols bgp group <i>group-name</i>] or [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]).</p>  |
| Options                  | <p><b>none</b>—Explicitly disables BMP route monitoring.</p> <p><b>Default:</b> If you configure the <b>route-monitoring</b> statement at the [edit routing-options <b>bmp</b>] hierarchy level, the default option is <b>pre-policy</b>. If you configure the <b>route-monitoring</b> statement at any of the [edit protocols bgp] hierarchy levels, the default option is to inherit the configuration from the <b>route-monitoring</b> statement configured at the [edit routing-options <b>bmp</b>] hierarchy level.</p> <p>The other statements are explained separately.</p>   |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>   |

## session-mode

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>session-mode (automatic   multihop   single-hop);</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure BFD session mode to be single-hop or multihop. By default, BGP uses single-hop BFD sessions if the peer is directly connected to the router's interface. BGP uses multihop BFD sessions if the peer is not directly connected to the router's interface. If the peer session's <b>local-address</b> option is configured, the directly connected check is based partly on the source address that would be used for BGP and BFD.</p> <p>For backward compatibility, you can override the default behavior by configuring the <b>single-hop</b> or <b>multihop</b> option. Before Junos OS Release 11.1, the behavior was to assume that IBGP peer sessions were multihop.</p>   |
| <b>Options</b>                  | <p><b>automatic</b>—Configure BGP to use single-hop BFD sessions if the peer is directly connected to the router's interface, and multihop BFD sessions if the peer is not directly connected to the router's interface</p> <p><b>multihop</b>—Configure BGP to use multihop BFD sessions.</p> <p><b>single-hop</b>—Configure BGP to use single-hop BFD sessions.</p> <p><b>Default:</b> automatic</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring BFD Authentication for BGP on page 3473</a></li> <li>• <a href="#">Example: Configuring BFD on Internal BGP Peer Sessions on page 3463</a></li> </ul>  |

- [Example: Configuring BFD Authentication for BGP on page 3473](#)
- [Understanding BFD Authentication for BGP on page 3471](#)

---

## stale-routes-time

---

|                          |  |
|--------------------------|--|
| Syntax                   | stale-routes-time <i>seconds</i> ;   |
| Hierarchy Level          | [edit logical-systems <i>logical-routing-name</i> protocols bgp <a href="#">graceful-restart</a> ],<br>[edit logical-systems <i>logical-routing-name</i> routing-instances <i>routing-instance-name</i> protocols<br>bgp <a href="#">graceful-restart</a> ],<br>[edit protocols bgp <a href="#">graceful-restart</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols bgp <a href="#">graceful-restart</a> ] |
| Release Information      | Statement introduced in Junos OS Release 8.3.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| Description              | Specify the maximum time that stale routes are kept during a restart. The <b>stale-routes-time</b> statement allows you to set the length of time the routing device waits to receive messages from restarting neighbors before declaring them down.   |
| Options                  | <b>seconds</b> —Time the router device waits to receive messages from restarting neighbors before declaring them down.<br><b>Range:</b> 1 through 600 seconds<br><b>Default:</b> 300 seconds   |
| Required Privilege Level | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Graceful Restart Options for BGP on page 2262</a></li><li>• <a href="#">Configuring Graceful Restart for QFabric Systems</a></li><li>• <a href="#">restart-time (BGP Graceful Restart) on page 2307</a></li></ul>  |

## station

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> station <i>station-name</i> {     authentication-algorithm (aes-128-cmac-96   hmac-sha-1-96   md5);     authentication-key <i>key</i>;     authentication-key-chain <i>authentication-key-chain</i>;     connection-mode (active   passive);     hold-down {         seconds;         flaps <i>flaps</i>;         period <i>seconds</i>;     }     initiation-message <i>text</i>;     local-address <i>address</i>;     local-port <i>port</i>;     monitor (disable   enable);     priority (high   low   medium);     route-monitoring {         none;         post-policy {             exclude-non-eligible;         }         pre-policy {             exclude-non-feasible;         }     }     station-address (<i>ip-address</i>   <i>name</i>);     station-port <i>port-number</i>;     statistics-timeout <i>seconds</i>;     traceoptions {         file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;         flag <i>flag</i> &lt;flag-modifier&gt;;     } } </pre> |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options bmp],<br>[edit routing-options bmp]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.  |
| <b>Description</b>              | Specify and configure a BMP monitoring station. Be aware that each BMP monitoring station can use a significant amount of a device's resources. You can configure up to 3 BMP monitoring stations.  |
| <b>Options</b>                  | <p><b><i>station-name</i></b>—Specify a name for the BMP station.</p> <p>The other statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |

## station-address

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>station-address (address   station-name);</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station</a> <i>station-name</i> ],<br>[edit routing-options <a href="#">bmp</a> ],<br>[edit routing-options bmp <a href="#">station</a> <i>station-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.  |
| <b>Description</b>              | Specify the name or address for the BMP monitoring station. You can specify one or the other but not both.  |
| <b>Options</b>                  | <b><i>station-address</i></b> —Specify the address for the BMP station. The address should be a valid IPv4 or IPv6 address.<br><br><b><i>station-name</i></b> —Specify the name for the BMP station.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>  |

## station-port

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>station-port <i>port</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station</a> <i>station-name</i> ],<br>[edit routing-options <a href="#">bmp</a> ],<br>[edit routing-options bmp <a href="#">station</a> <i>station-name</i> ]       |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.  |
| <b>Description</b>              | Specify the port number for the BMP monitoring station.   |
| <b>Options</b>                  | <b>port</b> —Specify the port number for the BMP monitoring station. If the <a href="#">connection-mode</a> statement is configured as <b>active</b> a station port number is required. If the <b>connection-mode</b> statement is configured as <b>passive</b> , you must not configure a station port number.<br><b>Range:</b> 1 though 65535 |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li> <li>• <a href="#">connection-mode on page 3651</a></li> </ul>  |

## statistics-timeout

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>statistics-timeout <i>seconds</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> routing-options <a href="#">bmp</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-options bmp <a href="#">station</a> <i>station-name</i> ],<br>[edit routing-options <a href="#">bmp</a> ],<br>[edit routing-options bmp <a href="#">station</a> <i>station-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br>Statement introduced in Junos OS Release 13.3.  |
| <b>Description</b>              | Specify how often statistics messages are sent to the BMP monitoring station. If you configure a value of 0, no statistics messages are sent.   |
| <b>Options</b>                  | <b><i>seconds</i></b> —Specify the number for the BMP monitoring station.<br><b>Default:</b> 3600 seconds<br><b>Range:</b> 15 though 65535 seconds  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BGP Monitoring Protocol Version 3 on page 3307</a></li></ul>  |



## tcp-mss (Protocols BGP)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>tcp-mss <i>segment-size</i>;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols bgp],</p> <p>[edit protocol bgp group <i>group-name</i>],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors.</p> <p>The MSS is only valid in increments of 2 KB. The value used is based on the value set, but is rounded down to the nearest multiple of 2048.</p>  |
| <b>Options</b>                  | <p><b><i>segment-size</i></b>—MSS for the TCP connection.</p> <p><b>Range:</b> 1 through 4096</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Limiting TCP Segment Size for BGP on page 3586</a></li> </ul>   |

## threshold (detection-time)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | threshold <i>milliseconds</i> ;  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection detection-time],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection detection-time]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPNs and VPLS.</p>  |
| <b>Description</b>         | Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.  |



**NOTE:** The threshold value must be equal to or greater than the transmit interval.

The threshold time must be equal to or greater than the value specified in the `minimum-interval` or the `minimum-receive-interval` statement.

|                                 |  |
|---------------------------------|--|
| <b>Options</b>                  | <b><i>milliseconds</i></b> —Value for the detection time adaptation threshold.<br><b>Range:</b> 1 through 255,000  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for Layer 2 VPN and VPLS</a></li><li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li></ul> |

## threshold (transmit-interval)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <code>threshold <i>milliseconds</i>;</code>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection transmit-interval],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection transmit-interval]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>   |
| <b>Description</b>         | Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.  |

**Options** *milliseconds*—Value for the transmit interval adaptation threshold.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )




**NOTE:** The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring BFD for Layer 2 VPN and VPLS](#)
- [Example: Configuring BFD for Static Routes on page 2918](#)
- [bfd-liveness-detection on page 3643](#)

## traceoptions (Protocols BGP)

|  |   |
|--|---|
| <b>Syntax</b>  | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>   |
| <b>Hierarchy Level</b>   | <pre> [edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp <i>group</i> <i>group-name</i>], [edit protocols bgp <i>group</i> <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp <i>group</i> <i>group-name</i> neighbor   <i>address</i>] </pre> |
| <b>Release Information</b>   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>4byte-as</b> statement introduced in Junos OS Release 9.2.</p> <p><b>4byte-as</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p>  |
| <b>Description</b>   | Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.   |
| <div>  <b>NOTE:</b> The <b>traceoptions</b> statement is not supported on QFabric systems. </div> |   |
| <b>Default</b>   | The default BGP protocol-level tracing options are inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level. The default group-level trace options are inherited from the BGP protocol-level <b>traceoptions</b> statement. The default peer-level trace options are inherited from the group-level <b>traceoptions</b> statement.  |
| <b>Options</b>   | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place BGP tracing output in the file <b>bgp-log</b>.</p>   |

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file.0*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### BGP Tracing Flags

- **4byte-as**—4-byte AS events.
- **bfd**—BFD protocol events.
- **damping**—Damping operations.
- **graceful-restart**—Graceful restart events.
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets.
- **refresh**—BGP refresh packets.
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **filter**—Provide filter trace information. Applies only to **route**, **damping**, and **update** tracing flags.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.                                      |
|                                 | routing-control and trace-control—To add this statement to the configuration.                       |
| <b>Related Documentation</b>    | • <a href="#">log-updown on page 3688</a> statement   |
|                                 | • <a href="#">Tracing Nonstop Active Routing Synchronization Events on page 2278</a>                |
|                                 | • <a href="#">Understanding Trace Operations for BGP Protocol Traffic on page 3619</a>              |
|                                 | • <a href="#">Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 4102</a> |



## traceoptions (Protocols BMP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>file-name</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;; }</pre>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> routing-options <b>bmp</b>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-options bmp <b>station</b> <i>station-name</i>],<br/>         [edit routing-options <b>bmp</b>],<br/>         [edit routing-options bmp <b>station</b> <i>station-name</i>]</p>  |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.<br/>         Statement introduced in Junos OS Release 13.3.</p>   |
| <b>Description</b>         | <p>Configure tracing options for BMP monitoring. To specify more than one tracing operation, include multiple flag statements.</p>  |
| <b>Options</b>             | <p><b>file</b> <i>file-name</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place BMP tracing output in the file <b>bmp-log</b>.</p> <p><b>files</b> <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file.0</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files<br/> <b>Default:</b> 10 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"> <li>• <b>all</b>—Trace all BMP monitoring operations.</li> <li>• <b>down</b>—Down messages.</li> <li>• <b>error</b>—Error conditions.</li> <li>• <b>event</b>—Major events, station establishment, errors, and events.</li> <li>• <b>general</b>—General events.</li> <li>• <b>normal</b>—Normal events.</li> <li>• <b>packets</b>—All messages.</li> <li>• <b>policy</b>—Policy processing.</li> <li>• <b>route</b>—Routing information.</li> <li>• <b>route-monitoring</b>—Route monitoring messages.</li> <li>• <b>state</b>—State transitions.</li> </ul> |

- **statistics**—Statistics messages.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.
- **up**—Up messages.
- **write**—Writing of messages.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable the tracing flag.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | routing and trace—To view this statement in the configuration.<br>routing-control and trace-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Tracing BMP Operations on page 3625</a></li><li>• <a href="#">Understanding Trace Operations for BGP Protocol Traffic on page 3619</a></li><li>• <a href="#">Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 4102</a></li></ul> |

## transmit-interval (BFD Liveness Detection)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>transmit-interval {     minimum-interval milliseconds;     threshold milliseconds; }</pre>   |
| <b>Hierarchy Level</b>     | <pre>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>   bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>   bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   l2vpn oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   vpls oam bfd-liveness-detection], [edit protocols bgp bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> bfd-liveness-detection], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>   bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor   <i>address</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam   bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>   neighbor <i>neighbor-id</i> oam bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</pre> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.2.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for BFD authentication introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>  |
| <b>Description</b>         | <p>Specify the transmit interval for the <b>bfd-liveness-detection</b> statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum time that it requires between packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its</p>  |

peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |  |
|------------------------------|--|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Configuring BFD for Layer 2 VPN and VPLS</i></li><li>• <a href="#">Example: Configuring BFD for Static Routes on page 2918</a></li><li>• <a href="#">bfd-liveness-detection on page 3643</a></li><li>• <a href="#">threshold on page 3738</a></li><li>• <a href="#">minimum-interval on page 3697</a></li><li>• <a href="#">minimum-receive-interval on page 3699</a></li></ul> |
|------------------------------|--|

## version (BFD Liveness Detection)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | version (0   1   automatic);   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit logical-system <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection],</p> <p>[edit protocols bgp bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i> neighbor <i>neighbor-id</i> oam bfd-liveness-detection],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls oam bfd-liveness-detection]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2 for Layer 2 VPN and VPLS.</p>  |
| <b>Description</b>              | Specify the BFD version for detection. You can explicitly configure BFD version 0, version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version, which is either 0 or 1.  |
| <b>Options</b>                  | <p>Configure the BFD version to detect: <b>0</b> (BFD version 0), <b>1</b> (BFD version 1), or <b>automatic</b> (autodetect the BFD version)</p> <p><b>Default:</b> automatic</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |

**Related  
Documentation**

- *Configuring BFD for Layer 2 VPN and VPLS*
- [Example: Configuring BFD Authentication for BGP on page 3473](#)
- [Example: Configuring BFD on Internal BGP Peer Sessions on page 3463](#)
- [Example: Configuring BFD Authentication for BGP on page 3473](#)
- [Understanding BFD Authentication for BGP on page 3471](#)

## CHAPTER 43

# Administration

- [Routine Monitoring on page 3749](#)
- [Operational Commands on page 3749](#)

## Routine Monitoring

---

- [Monitoring BGP Routing Information on page 3749](#)

### Monitoring BGP Routing Information

**Purpose** Use the monitoring functionality to monitor BGP routing information on the routing device.

**Action** To view BGP routing information in the CLI, enter the following commands:

- **show bgp summary**
- **show bgp neighbor**

**Related Documentation**

- [show bgp neighbor on page 2344](#)
- [show bgp summary on page 3778](#)

## Operational Commands

---

- **clear bgp damping**
- **clear bgp neighbor**
- **clear bgp table**
- **show bgp bmp**
- **show bgp group**
- **show bgp neighbor**
- **show bgp summary**
- **show policy damping**
- **show route damping**
- **show route detail**

## clear bgp damping

---

|  |  |
|--|--|
| List of Syntax                           | <a href="#">Syntax on page 3750</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3750</a>   |
| Syntax                                   | <code>clear bgp damping</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code><br><code>&lt;prefix&gt;</code>  |
| Syntax (EX Series Switch and QFX Series) | <code>clear bgp damping</code><br><code>&lt;prefix&gt;</code>  |
| Release Information                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.   |
| Description                              | Clear BGP route flap damping information.  |
| Options                                  | <b>none</b> —Clear all BGP route flap damping information.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>prefix</b> —(Optional) Clear route flap damping information for only the specified destination prefix. |
| Required Privilege Level                 | clear  |
| Related Documentation                    | <ul style="list-style-type: none"><li>• <a href="#">show policy damping on page 3784</a></li><li>• <a href="#">show route damping on page 3098</a></li></ul>   |
| List of Sample Output                    | <a href="#">clear bgp damping on page 3750</a>   |
| Output Fields                            | When you enter this command, you are provided feedback on the status of your request.  |

## Sample Output

### clear bgp damping

```
user@host> clear bgp damping
```



## clear bgp neighbor

|   |  |
|---|--|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3751</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3751</a>   |
| <b>Syntax</b>                                   | <pre>clear bgp neighbor &lt;as <i>as-number</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;malformed-route&gt; &lt;neighbor&gt; &lt;soft   soft-inbound&gt; &lt;soft-minimum-igp&gt;</pre>  |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>clear bgp neighbor &lt;as <i>as-number</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;malformed-route&gt; &lt;neighbor&gt; &lt;soft   soft-inbound&gt; &lt;soft-minimum-igp&gt;</pre>  |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>malformed-route</b> option introduced in Junos OS Release 13.2.</p>  |
| <b>Description</b>                              | <p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> <li>• Change the state of one or more BGP neighbors to <b>IDLE</b>. For neighbors in the <b>ESTABLISHED</b> state, this command drops the TCP connection to the neighbors and then reestablishes the connection.</li> <li>• (<b>soft</b> keyword only) Reapply export policies or import policies, respectively, to one or more BGP neighbors without changing their state.</li> <li>• (<b>soft-inbound</b> keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.</li> </ul>   |
| <b>Options</b>                                  | <p><b>none</b>—Change the state of all BGP neighbors to <b>IDLE</b>.</p> <p><b>as <i>as-number</i></b>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p><b>instance <i>instance-name</i></b>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>malformed-route</b>—(Optional) Remove malformed routes. If a specific neighbor is provided, Junos OS removes malformed routes for that particular neighbor. Otherwise, Junos OS removes malformed routes for all BGP neighbors. To find routes that have</p> |

malformed attributes, run the **show route hidden** command, and look for routes marked with **MalformedAttr** in the AS path field.

**neighbor**—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

**soft**—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

**soft-inbound**—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.

**soft-minimum-igp**—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

**Required Privilege Level**

clear

**Related Documentation**

- [show bgp neighbor on page 2344](#)

**List of Sample Output**

[clear bgp neighbor on page 3752](#)

**Output Fields**

When you enter this command, you are provided feedback on the status of your request.

## Sample Output

**clear bgp neighbor**

```
user@host> clear bgp neighbor
```

## clear bgp table

|   |  |
|---|--|
| <b>Syntax</b>                                   | <code>clear bgp table <i>table-name</i></code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <code>clear bgp table <i>table-name</i></code>   |
| <b>Release Information</b>                      | Command introduced in Junos OS Release 9.0.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>                              | Request that BGP refresh routes in a specified routing table.  |
| <b>Options</b>                                  | <b><code>logical-system (all   <i>logical-system-name</i>)</code></b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b><code>table-name</code></b> —Request that BGP refresh routes in the specified table.  |
| <b>Additional Information</b>                   | In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the <b>clear bgp table</b> command to request that BGP refresh routes in a VPN instance table. |
| <b>Required Privilege Level</b>                 | clear  |
| <b>List of Sample Output</b>                    | <a href="#">clear bgp table private.inet.0 on page 3753</a><br><a href="#">clear bgp table inet.6 logical-system all on page 3753</a><br><a href="#">clear bgp table private.inet.6 logical-system ls1 on page 3753</a><br><a href="#">clear bgp table logical-system all inet.0 on page 3753</a><br><a href="#">clear bgp table logical-system ls2 private.inet.0 on page 3754</a>                      |
| <b>Output Fields</b>                            | This command produces no output.   |

## Sample Output

`clear bgp table private.inet.0`

```
user@host> clear bgp table private.inet.0
```

`clear bgp table inet.6 logical-system all`

```
user@host> clear bgp table inet.6 logical-system all
```

`clear bgp table private.inet.6 logical-system ls1`

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

`clear bgp table logical-system all inet.0`

```
user@host> clear bgp table logical-system all inet.0
```

### clear bgp table logical-system ls2 private.inet.0

```
user@host> clear bgp table logical-system ls2 private.inet.0
```

## show bgp bmp

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show bgp bmp</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series. |
| <b>Description</b>              | Display information about the BGP Monitoring Protocol (BMP).  |
| <b>Options</b>                  | This command has no options.  |
| <b>Required Privilege Level</b> | view  |
| <b>List of Sample Output</b>    | <a href="#">show bgp bmp on page 3755</a>   |
| <b>Output Fields</b>            | <a href="#">Table 298 on page 3755</a> lists the output fields for the <b>show bgp bmp</b> command. Output fields are listed in the approximate order in which they appear.                 |

**Table 298: show bgp bmp Output Fields**

| Field Name                          | Field Description  |
|-------------------------------------|--|
| <b>BMP station address/port</b>     | IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.   |
| <b>BMP session state</b>            | Status of the BMP session: <b>UP</b> or <b>DOWN</b> .  |
| <b>Memory consumed by BMP</b>       | Memory used by the active BMP session.   |
| <b>Statistics timeout</b>           | Amount of time, in seconds, between transmissions of BMP data to the monitoring station.   |
| <b>Memory limit</b>                 | Threshold, in bytes, at which the routing device stops collecting BMP data.  |
| <b>Memory-connect retry timeout</b> | Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded. |

## Sample Output

### show bgp bmp

```

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```



## show bgp group

|   |  |
|---|--|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3757</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3757</a>   |
| <b>Syntax</b>                                   | <pre>show bgp group &lt;brief   detail   summary&gt; &lt;group-name&gt; &lt;exact-instance instance-name&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;rtf&gt;</pre>   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show bgp group &lt;brief   detail   summary&gt; &lt;group-name&gt; &lt;exact-instance instance-name&gt; &lt;instance instance-name&gt;</pre>  |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>exact-instance</b> option introduced in Junos OS Release 11.4.</p>   |
| <b>Description</b>                              | Display information about the configured BGP groups.   |
| <b>Options</b>                                  | <p><b>none</b>—Display group information about all BGP groups.</p> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group information for the specified group.</p> <p><b>exact-instance instance-name</b>—(Optional) Display information for the specified instance only.</p> <p><b>instance instance-name</b>—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show bgp group instance cust1</b> command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>rtf</b>—(Optional) Display BGP group route targeting information.</p> |
| <b>Required Privilege Level</b>                 | view   |
| <b>List of Sample Output</b>                    | <a href="#">show bgp group on page 3761</a><br><a href="#">show bgp group brief on page 3761</a><br><a href="#">show bgp group detail on page 3762</a>   |

[show bgp group rtf detail on page 3763](#)  
[show bgp group summary on page 3763](#)

**Output Fields** [Table 299 on page 3758](#) describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

**Table 299: show bgp group Output Fields**

| Field Name                                | Field Description  | Level of Output             |
|---|--|-----------------------------|
| <b>Group Type or Group</b>                | Type of BGP group: <b>Internal</b> or <b>External</b> .  | All levels                  |
| <b>group-index</b>                        | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.  | <b>rtf detail</b>           |
| <b>AS</b>                                 | AS number of the peer. For internal BGP (IBGP), this number is the same as <b>Local AS</b> .   | <b>brief detail</b><br>none |
| <b>Local AS</b>                           | AS number of the local routing device.   | <b>brief detail</b><br>none |
| <b>Name</b>                               | Name of a specific BGP group.  | <b>brief detail</b><br>none |
| <b>Index</b>                              | Unique index number of a BGP group.  | <b>brief detail</b><br>none |
| <b>Flags</b>                              | Flags associated with the BGP group. This field is used by Juniper Networks customer support.  | <b>brief detail</b><br>none |
| <b>Remove-private options</b>             | Options associated with the <a href="#">remove-private</a> statement.  | <b>brief detail</b><br>none |
| <b>Holdtime</b>                           | Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable. | <b>brief detail</b><br>none |
| <b>Export</b>                             | Export policies configured for the BGP group with the <b>export</b> statement.   | <b>brief detail</b><br>none |
| <b>MED tracks IGP metric update delay</b> | Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire   | All levels                  |
| <b>Traffic Statistics Interval</b>        | Time between sample periods for labeled-unicast traffic statistics, in seconds.  | <b>brief detail</b><br>none |
| <b>Total peers</b>                        | Total number of peers in the group.  | <b>brief detail</b><br>none |



Table 299: show bgp group Output Fields (*continued*)

| Field Name                             | Field Description  | Level of Output |
|--|--|-----------------|
| <b>Established</b>                     | Number of peers in the group that are in the established state.  | All levels      |
| <b>Active/Received/Accepted/Damped</b> | <p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> <li>If a peer is not established, the field shows the state of the peer session: <b>Active</b>, <b>Connect</b>, or <b>Idle</b>.</li> <li>If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the <b>inet.0</b> (main) and <b>inet.2</b> (multicast) routing tables. For example, <b>8/10/10/2</b> and <b>2/4/4/0</b> indicate the following: <ul style="list-style-type: none"> <li>8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the <b>inet.0</b> routing table.</li> <li>2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the <b>inet.2</b> routing table.</li> </ul> </li> </ul> | <b>summary</b>  |
| <b>ip-addresses</b>                    | List of peers who are members of the group. The address is followed by the peer's port number.   | All levels      |
| <b>Route Queue Timer</b>               | Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.  | <b>detail</b>   |
| <b>Route Queue</b>                     | Number of prefixes that are queued up for sending to the peers in the group.   | <b>detail</b>   |
| <b>inet.number</b>                     | <p>Number of active, received, accepted, and damped routes in the routing table. For example, <b>inet.0: 7/10/9/0</b> indicates the following:</p> <ul style="list-style-type: none"> <li>7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the <b>inet.0</b> routing table.</li> </ul>   | none            |

Table 299: show bgp group Output Fields (*continued*)

| Field Name               | Field Description   | Level of Output     |
|--------------------------|---|---------------------|
| <b>Table inet.number</b> | Information about the routing table. <ul style="list-style-type: none"> <li>• <b>Received prefixes</b>—Total number of prefixes from the peer, both active and inactive, that are in the routing table.</li> <li>• <b>Active prefixes</b>—Number of prefixes received from the peer that are active in the routing table.</li> <li>• <b>Suppressed due to damping</b>—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.</li> <li>• <b>Advertised prefixes</b>—Number of prefixes advertised to a peer.</li> <li>• <b>Received external prefixes</b>—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table.</li> <li>• <b>Active external prefixes</b>—Number of prefixes received from the EBGP peers that are active in the routing table.</li> <li>• <b>Externals suppressed</b>—Number of routes received from EBGP peers currently inactive because of damping or other reasons.</li> <li>• <b>Received internal prefixes</b>—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table.</li> <li>• <b>Active internal prefixes</b>—Number of prefixes received from the IBGP peers that are active in the routing table.</li> <li>• <b>Internals suppressed</b>—Number of routes received from IBGP peers currently inactive because of damping or other reasons.</li> <li>• <b>RIB State</b>—Status of the graceful restart process for this routing table: <b>BGP restart is complete</b>, <b>BGP restart in progress</b>, <b>VPN restart in progress</b>, or <b>VPN restart is complete</b>.</li> </ul> | <b>detail</b>       |
| <b>Groups</b>            | Total number of groups.   | All levels          |
| <b>Peers</b>             | Total number of peers.  | All levels          |
| <b>External</b>          | Total number of external peers.   | All levels          |
| <b>Internal</b>          | Total number of internal peers.   | All levels          |
| <b>Down peers</b>        | Total number of unavailable peers.  | All levels          |
| <b>Flaps</b>             | Total number of flaps that occurred.  | All levels          |
| <b>Table</b>             | Name of a routing table.  | <b>brief</b> , none |
| <b>Tot Paths</b>         | Total number of routes.   | <b>brief</b> , none |
| <b>Act Paths</b>         | Number of active routes.  | <b>brief</b> , none |
| <b>Suppressed</b>        | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.  | <b>brief</b> , none |

Table 299: show bgp group Output Fields (*continued*)

| Field Name   | Field Description   | Level of Output |
|--------------|---|-----------------|
| History      | Number of withdrawn routes stored locally to keep track of damping history.   | brief, none     |
| Damp State   | Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.   | brief, none     |
| Pending      | Routes being processed by the BGP import policy.  | brief, none     |
| Group        | Group the peer belongs to in the BGP configuration.   | detail          |
| Receive mask | Mask of the received target included in the advertised route.   | detail          |
| Entries      | Number of route entries received.   | detail          |
| Target       | Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer. | detail          |
| Mask         | Mask which specifies that the peer receive routes with the given route target.  | detail          |

## Sample Output

### show bgp group

```

user@host> show bgp group
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0           0         0         0

bgp.13vpn.0
          0         0         0           0         0         0

bgp.rtarget.0
          2         0         0           0         0         0

```

### show bgp group brief

```

user@host> show bgp group brief
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0           0         0         0

bgp.13vpn.0
          0         0         0           0         0         0

bgp.rtarget.0
          2         0         0           0         0         0

```

## show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal   AS: 1                               Local AS: 1
Name: ibgp             Index: 0                             Flags: <Export Eval>
Holdtime: 0
Total peers: 3         Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1 Peers: 3   External: 0   Internal: 3   Down peers: 3   Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0

```

```

Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

#### show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target
    100:100/64
    200:201/64
    Mask
    00000002
    (Group)
    Entries: 2
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target
    200:201/64
    Mask
    (Group)
    Entries: 1

```

#### show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers      Established      Active/Received/Accepted/Damped
ibgp       Internal  3          0
Groups: 1  Peers: 3      External: 0      Internal: 3      Down peers: 3      Flaps: 3
  bgp.l3vpn.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  bgp.mdt.0        : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.inet.0     : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
  VPN-A.mdt.0      : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

## show bgp neighbor

---

|   |  |
|---|--|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3764</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3764</a>   |
| <b>Syntax</b>                                   | <pre>show bgp neighbor &lt;exact-instance <i>instance-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>neighbor-address</i>&gt; &lt;orf (detail   <i>neighbor-address</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show bgp neighbor &lt;instance <i>instance-name</i>&gt; &lt;exact-instance <i>instance-name</i>&gt; &lt;<i>neighbor-address</i>&gt; &lt;orf (<i>neighbor-address</i>   detail)&gt;</pre>  |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>orf</b> option introduced in Junos OS Release 9.2.</p> <p><b>exact-instance</b> option introduced in Junos OS Release 11.4.</p>  |
| <b>Description</b>                              | Display information about BGP peers.   |
| <b>Options</b>                                  | <p><b>none</b>—Display information about all BGP peers.</p> <p><b>exact-instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, <b>cust1</b>, <b>cust11</b>, and <b>cust111</b> are all displayed when you run the <b>show bgp neighbor instance cust1</b> command).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>neighbor-address</i></b>—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p><b>orf (detail   <i>neighbor-address</i>)</b>—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the <b>detail</b> option to display detailed output.</p> |
| <b>Additional Information</b>                   | For information about the <b>local-address</b> , <b>nlri</b> , <b>hold-time</b> , and <b>preference</b> statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .  |
| <b>Required Privilege Level</b>                 | view   |

**Related Documentation** • [clear bgp neighbor on page 3751](#)

**List of Sample Output** [show bgp neighbor on page 3771](#)  
[show bgp neighbor \(CLNS\) on page 3772](#)  
[show bgp neighbor \(Layer 2 VPN\) on page 3773](#)  
[show bgp neighbor \(Layer 3 VPN\) on page 3775](#)  
[show bgp neighbor neighbor-address on page 3775](#)  
[show bgp neighbor neighbor-address on page 3776](#)  
[show bgp neighbor orf neighbor-address detail on page 3777](#)

**Output Fields** [Table 210 on page 2345](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

**Table 300: show bgp neighbor Output Fields**

| Field Name   | Field Description   |
|--------------|---|
| <b>Peer</b>  | Address of the BGP neighbor. The address is followed by the neighbor port number.   |
| <b>AS</b>    | AS number of the peer.  |
| <b>Local</b> | Address of the local routing device. The address is followed by the peer port number.   |
| <b>Type</b>  | Type of peer: <b>Internal</b> or <b>External</b> .  |
| <b>State</b> | Current state of the BGP session: <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul> |

Table 300: show bgp neighbor Output Fields (*continued*)

| Field Name        | Field Description   |
|-------------------|---|
| <b>Flags</b>      | <p>Internal BGP flags:</p> <ul style="list-style-type: none"> <li>• <b>Aggregate Label</b>—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label.</li> <li>• <b>CleanUp</b>—The peer session is being shut down.</li> <li>• <b>Delete</b>—This peer has been deleted.</li> <li>• <b>Idled</b>—This peer has been permanently idled.</li> <li>• <b>ImportEval</b>—At the last commit operation, this peer was identified as needing to reevaluate all received routes.</li> <li>• <b>Initializing</b>—The peer session is initializing.</li> <li>• <b>SendRtn</b>—Messages are being sent to the peer.</li> <li>• <b>Sync</b>—This peer is synchronized with the rest of the peer group.</li> <li>• <b>TryConnect</b>—Another attempt is being made to connect to the peer.</li> <li>• <b>Unconfigured</b>—This peer is not configured.</li> <li>• <b>WriteFailed</b>—An attempt to write to this peer failed.</li> </ul>   |
| <b>Last state</b> | <p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message.</li> <li>• <b>Connect</b>—BGP is waiting for the transport protocol connection to be completed.</li> <li>• <b>Established</b>—The BGP session has been established, and the peers are exchanging update messages.</li> <li>• <b>Idle</b>—This is the first stage of a connection. BGP is waiting for a Start event.</li> <li>• <b>OpenConfirm</b>—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.</li> <li>• <b>OpenSent</b>—BGP has sent an open message and is waiting to receive an open message from the peer.</li> </ul>   |
| <b>Last event</b> | <p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Closed</b>—The BGP session closed.</li> <li>• <b>ConnectRetry</b>—The transport protocol connection failed, and BGP is trying again to connect.</li> <li>• <b>HoldTime</b>—The session ended because the hold timer expired.</li> <li>• <b>KeepAlive</b>—The local routing device sent a BGP keepalive message to the peer.</li> <li>• <b>Open</b>—The local routing device sent a BGP open message to the peer.</li> <li>• <b>OpenFail</b>—The local routing device did not receive an acknowledgment of a BGP open message from the peer.</li> <li>• <b>RecvKeepAlive</b>—The local routing device received a BGP keepalive message from the peer.</li> <li>• <b>RecvNotify</b>—The local routing device received a BGP notification message from the peer.</li> <li>• <b>RecvOpen</b>—The local routing device received a BGP open message from the peer.</li> <li>• <b>RecvUpdate</b>—The local routing device received a BGP update message from the peer.</li> <li>• <b>Start</b>—The peering session started.</li> <li>• <b>Stop</b>—The peering session stopped.</li> <li>• <b>TransportError</b>—A TCP error occurred.</li> </ul> |



Table 300: show bgp neighbor Output Fields (*continued*)

| Field Name                  | Field Description   |
|-----------------------------|---|
| Last error                  | <p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> <li>• <b>Cease</b>—An error occurred, such as a version mismatch, that caused the session to close.</li> <li>• <b>Finite State Machine Error</b>—In setting up the session, BGP received a message that it did not understand.</li> <li>• <b>Hold Time Expired</b>—The session's hold time expired.</li> <li>• <b>Message Header Error</b>—The header of a BGP message was malformed.</li> <li>• <b>Open Message Error</b>—A BGP open message contained an error.</li> <li>• <b>None</b>—No errors occurred in the BGP session.</li> <li>• <b>Update Message Error</b>—A BGP update message contained an error.</li> </ul>   |
| Export                      | Name of the export policy that is configured on the peer.   |
| Import                      | Name of the import policy that is configured on the peer.   |
| Options                     | <p>Configured BGP options:</p> <ul style="list-style-type: none"> <li>• <b>AddressFamily</b>—Configured address family: <b>inet</b> or <b>inet-vpn</b>.</li> <li>• <b>AuthKeyChain</b>—Authentication key change is enabled.</li> <li>• <b>DropPathAttributes</b>—Certain path attributes are configured to be dropped from neighbor updates during inbound processing.</li> <li>• <b>GracefulRestart</b>—Graceful restart is configured.</li> <li>• <b>HoldTime</b>—Hold time configured with the <b>hold-time</b> statement. The hold time is three times the interval at which keepalive messages are sent.</li> <li>• <b>IgnorePathAttributes</b>—Certain path attributes are configured to be ignored in neighbor updates during inbound processing.</li> <li>• <b>Local Address</b>—Address configured with the <b>local-address</b> statement.</li> <li>• <b>Multihop</b>—Allow BGP connections to external peers that are not on a directly connected network.</li> <li>• <b>NLRI</b>—Configured MBGP state for the BGP group: <b>multicast</b>, <b>unicast</b>, or both if you have configured <b>nlri any</b>.</li> <li>• <b>Peer AS</b>—Configured peer autonomous system (AS).</li> <li>• <b>Preference</b>—Preference value configured with the <b>preference</b> statement.</li> <li>• <b>Refresh</b>—Configured to refresh automatically when the policy changes.</li> <li>• <b>Rib-group</b>—Configured routing table group.</li> </ul> |
| Path-attributes dropped     | Path attribute codes that are dropped from neighbor updates.  |
| Path-attributes ignored     | Path attribute codes that are ignored during neighbor updates.  |
| Authentication key change   | (appears only if the <b>authentication-keychain</b> statement has been configured) Name of the authentication keychain enabled.   |
| Authentication algorithm    | (appears only if the <b>authentication-algorithm</b> statement has been configured) Type of authentication algorithm enabled: <b>hmac</b> or <b>md5</b> .   |
| Address families configured | Names of configured address families for the VPN.   |

Table 300: show bgp neighbor Output Fields (*continued*)

| Field Name                             | Field Description   |
|--|---|
| Local Address                          | Address of the local routing device.  |
| Remove-private options                 | Options associated with the <code>remove-private</code> statement.  |
| Holdtime                               | Hold time configured with the <code>hold-time</code> statement. The hold time is three times the interval at which keepalive messages are sent.   |
| Flags for NLRI<br>inet-label-unicast   | Flags related to labeled-unicast: <ul style="list-style-type: none"> <li>• <b>TrafficStatistics</b>—Collection of statistics for labeled-unicast traffic is enabled.</li> </ul>   |
| Traffic statistics                     | Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> <li>• <b>Options</b>—Options configured for collecting statistics about labeled-unicast traffic.</li> <li>• <b>File</b>—Name and location of statistics log files.</li> <li>• <b>size</b>—Size of all the log files, in bytes.</li> <li>• <b>files</b>—Number of log files.</li> </ul> |
| Traffic Statistics<br>Interval         | Time between sample periods for labeled-unicast traffic statistics, in seconds.   |
| Preference                             | Preference value configured with the <code>preference</code> statement.   |
| Outbound Timer                         | Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the <code>out-delay</code> parameter is configured to a non-zero value.   |
| Number of flaps                        | Number of times the BGP session has gone down and then come back up.  |
| Peer ID                                | Router identifier of the peer.  |
| Group index                            | Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.   |
| Peer index                             | Index that is unique within the BGP group to which the peer belongs.  |
| Local ID                               | Router identifier of the local routing device.  |
| Local Interface                        | Name of the interface on the local routing device.  |
| Active holdtime                        | Hold time that the local routing device negotiated with the peer.   |
| Keepalive Interval                     | Keepalive interval, in seconds.   |
| BFD                                    | Status of BFD failure detection.  |
| Local Address                          | Name of directly connected interface over which direct EBGP peering is established.   |
| NLRI for restart<br>configured on peer | Names of address families configured for restart.   |

Table 300: show bgp neighbor Output Fields (*continued*)

| Field Name                                      | Field Description   |
|---|---|
| NLRI advertised by peer                         | Address families supported by the peer: <b>unicast</b> or <b>multicast</b> .  |
| NLRI for this session                           | Address families being used for this session.   |
| Peer supports Refresh capability                | Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .  |
| Restart time configured on peer                 | Configured time allowed for restart on the neighbor.  |
| Stale routes from peer are kept for             | When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.  |
| Peer does not support Restarter functionality   | Graceful restart restarter-mode is disabled on the peer.  |
| Peer does not support Receiver functionality    | Graceful restart helper-mode is disabled on the peer.   |
| Restart time requested by this peer             | Restart time requested by this neighbor during capability negotiation.  |
| Restart flag received from the peer             | When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the <b>end-of-rib</b> marker from the speaker before advertising routing information to the speaker.                               |
| NLRI that peer supports restart for             | Neighbor supports graceful restart for this address family.   |
| NLRI peer can save forwarding state             | Neighbor supporting this address family saves all forwarding states.  |
| NLRI that peer saved forwarding for             | Neighbor saves all forwarding states for this address family.   |
| NLRI that restart is negotiated for             | Router supports graceful restart for this address family.   |
| NLRI of received end-of-rib markers             | Address families for which end-of-routing-table markers are received from the neighbor.   |
| NLRI of all end-of-rib markers sent             | Address families for which end-of-routing-table markers are sent to the neighbor.   |
| Peer supports 4 byte AS extension (peer-as1)    | Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.  |
| NLRIs for which peer can receive multiple paths | Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route.<br><br>Possible value is <b>inet-unicast</b> . |

Table 300: show bgp neighbor Output Fields (*continued*)

| Field Name   | Field Description   |
|--|---|
| NLRIs for which peer can send multiple paths: inet-unicast | Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route.<br><br>Possible value is <b>inet-unicast</b> .  |
| Table inet.number  | Information about the routing table: <ul style="list-style-type: none"> <li>• <b>RIB State</b>—BGP is in the graceful restart process for this routing table: <b>restart is complete</b> or <b>restart in progress</b>.</li> <li>• <b>Bit</b>—Number that represents the entry in the routing table for this peer.</li> <li>• <b>Send state</b>—State of the BGP group: <b>in sync</b>, <b>not in sync</b>, or <b>not advertising</b>.</li> <li>• <b>Active prefixes</b>—Number of prefixes received from the peer that are active in the routing table.</li> <li>• <b>Received prefixes</b>—Total number of prefixes from the peer, both active and inactive, that are in the routing table.</li> <li>• <b>Accepted prefixes</b>—Total number of prefixes from the peer that have been accepted by a routing policy.</li> <li>• <b>Suppressed due to damping</b>—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.</li> </ul> |
| Last traffic (seconds)                                     | Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.   |
| Input messages   | Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.   |
| Output messages  | Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.   |
| Input dropped path attributes                              | Information about dropped path attributes: <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>   |
| Input ignored path attributes                              | Information about ignored path attributes: <ul style="list-style-type: none"> <li>• <b>Code</b>—Path attribute code.</li> <li>• <b>Count</b>—Path attribute count.</li> </ul>   |
| Output queue   | Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.<br><br>It also specifies the routing table name and the NLRI they represent in the format ( <b>routing table name, NLRI</b> ).<br><br><b>NOTE:</b> The output queues of routing tables that are not advertised, will only show up at <b>extensive</b> output level.   |
| Trace options  | Configured tracing of BGP protocol packets and operations.  |

Table 300: show bgp neighbor Output Fields (*continued*)

| Field Name              | Field Description   |
|-------------------------|---|
| Trace file              | Name of the file to receive the output of the tracing operation.  |
| Filter Updates rcv      | (orf option only) Number of outbound-route filters received for each configured address family.<br><br><b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.   |
| Immediate               | (orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes.<br><br><b>NOTE:</b> The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list. |
| Filter                  | (orf option only) Type of prefix filter received: <b>prefix-based</b> or <b>extended-community</b> .  |
| Received filter entries | (orf option only) List of received filters displayed.   |
| seq                     | (orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.  |
| prefix                  | (orf option only) Address for the prefix entry that matches the filter.   |
| minlength               | (orf option only) Minimum prefix length, in bits, required to match this prefix.  |
| maxlength               | (orf option only) Maximum prefix length, in bits, required to match this prefix.  |
| match                   | (orf option only) For this prefix match, whether to <b>permit</b> or <b>deny</b> route updates.   |

## Sample Output

### show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast

```

```

NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 10)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages: Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214 Local ID: 10.255.167.205 Active Holdtime: 90
  Keepalive Interval: 30 Group index: 0 Peer index: 1

```

### show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1 Local ID: 10.245.245.3 Active Holdtime: 90
  Keepalive Interval: 30 Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          3
    Received prefixes:        3
    Suppressed due to damping: 0
    Advertised prefixes:      3
  Table aaaa.iso.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:          3
    Received prefixes:        3
    Suppressed due to damping: 0
  Last traffic (seconds): Received 6    Sent 5    Checked 5
  Input messages: Total 1736    Updates 4    Refreshes 0    Octets 33385
  Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305

```

```
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)
```

### show bgp neighbor (Layer 2 VPN)

```
user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established  Flags: <ImportEval>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
```

```

Received prefixes:          1
Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            1
Received prefixes:          1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            2
Received prefixes:          2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            1
Received prefixes:          1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:            1
Received prefixes:          1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)

```



```
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)
```

### show bgp neighbor (Layer 3 VPN)

```
user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
    Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10

  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 15      Sent 20      Checked 20
  Input messages: Total 40      Updates 2      Refreshes 0      Octets 856
  Output messages: Total 44      Updates 2      Refreshes 0      Octets 1066
  Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10
```

### show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.1.111
```

```

Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
  Refresh>
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
  Keepalive Interval: 30
BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3 Sent 3 Checked 3
  Input messages: Total 9 Updates 6 Refreshes 0 Octets 403
  Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
  Output Queue[0]: 0 (inet.0, inet-unicast)
  Output Queue[1]: 0 (inet6.0, inet6-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpr size 131072 files 10

```

### show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External State: Established Flags: <Sync>
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6 Local ID: 10.255.245.5 Active Holdtime: 60000
  Keepalive Interval: 20000 Peer index: 0
  BFD: disabled, down

```

```

Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           8
  Received prefixes:        10
  Accepted prefixes:        10
  Suppressed due to damping: 0
  Advertised prefixes:       3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:           0
  Received prefixes:         0
  Accepted prefixes:         0
  Suppressed due to damping: 0
  Advertised prefixes:       0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

#### show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:           1 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:           0 Immediate:           1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

## show bgp summary

---

|  |  |
|--|--|
| List of Syntax                           | <a href="#">Syntax on page 3778</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3778</a>   |
| Syntax                                   | <pre>show bgp summary &lt;exact-instance <i>instance-name</i>&gt; &lt;group <i>group-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| Syntax (EX Series Switch and QFX Series) | <pre>show bgp summary &lt;exact-instance <i>instance-name</i>&gt; &lt;instance <i>instance-name</i>&gt;</pre>  |
| Release Information                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.<br><b>exact-instance</b> option introduced in Junos OS Release 11.4.<br><b>group</b> option introduced in Junos OS Release 13.3.  |
| Description                              | Display BGP summary information.   |
| Options                                  | <b>none</b> —Display BGP summary information for all routing instances.<br><br><b>exact-instance <i>instance-name</i></b> —(Optional) Display information for the specified instance only.<br><br><b>group</b> —Display overview of bgp information for a particular group<br><br><b>instance <i>instance-name</i></b> —(Optional) Display information for all routing instances whose name begins with this string (for example, <b>cust1</b> , <b>cust11</b> , and <b>cust111</b> are all displayed when you run the <b>show bgp summary instance cust1</b> command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level                 | view   |
| List of Sample Output                    | <a href="#">show bgp summary (When a Peer Is Not Established) on page 3781</a><br><a href="#">show bgp summary (When a Peer Is Established) on page 3781</a><br><a href="#">show bgp summary (CLNS) on page 3781</a><br><a href="#">show bgp summary (Layer 2 VPN) on page 3782</a><br><a href="#">show bgp summary (Layer 3 VPN) on page 3782</a><br><a href="#">show bgp summary group on page 3782</a>  |
| Output Fields                            | <a href="#">Table 301 on page 3779</a> describes the output fields for the <b>show bgp summary</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 301: show bgp summary Output Fields

| Field Name          | Field Description  |
|---------------------|--|
| <b>Groups</b>       | Number of BGP groups.  |
| <b>Peers</b>        | Number of BGP peers.   |
| <b>Down peers</b>   | Number of down BGP peers.  |
| <b>Table</b>        | Name of routing table.   |
| <b>Tot Paths</b>    | Total number of paths.   |
| <b>Act Paths</b>    | Number of active routes.   |
| <b>Suppressed</b>   | Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. |
| <b>History</b>      | Number of withdrawn routes stored locally to keep track of damping history.  |
| <b>Damp State</b>   | Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.               |
| <b>Pending</b>      | Routes in process by BGP import policy.  |
| <b>Peer</b>         | Address of each BGP peer. Each peer has one line of output.  |
| <b>AS</b>           | Peer's AS number.  |
| <b>InPkt</b>        | Number of packets received from the peer.  |
| <b>OutPkt</b>       | Number of packets sent to the peer.  |
| <b>OutQ</b>         | Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.                       |
| <b>Flaps</b>        | Number of times the BGP session has gone down and then come back up.   |
| <b>Last Up/Down</b> | Last time since the neighbor transitioned to or from the established state.  |

Table 301: show bgp summary Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| <b>State #Active<br/>/Received/Accepted<br/>/Damped</b> | <p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> <li>If a peer is not established, the field shows the state of the peer session: <b>Active</b>, <b>Connect</b>, or <b>Idle</b>. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: <b>Idle (Removal in progress)</b> or <b>Idle (LicenseFailure)</b>.</li> <li>If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the <b>inet.0</b> (main) and <b>inet.2</b> (multicast) routing tables. For example, <b>8/10/10/2</b> and <b>2/4/4/0</b> indicate the following: <ul style="list-style-type: none"> <li>8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the <b>inet.0</b> routing table.</li> <li>2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the <b>inet.2</b> routing table.</li> </ul> </li> <li>If a BGP session is established in a routing instance, the field indicates the established (<b>Establ</b>) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, <b>Establ VPN-AB.inet.0: 2/4/0</b> indicates the following: <ul style="list-style-type: none"> <li>The BGP session is established.</li> <li>Routes are received in the <b>VPN-AB.inet.0</b> routing table.</li> <li>The local routing device has two active routes, four received routes, and no damped routes from a BGP peer.</li> </ul> </li> </ul> <p>When a BGP session is established, the peers are exchanging update messages.</p> |

## Sample Output

### show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3        65002        86        90        0        2        42:54 0/0/0

0/0/0
10.0.0.4        65002        90        91        0        1        42:54 0/2/0

0/0/0
10.0.0.6        65002        87        90        0        3          3 Active
10.1.12.1       65001        89        89        0        1        42:54 4/4/0

0/0/0

```

### show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0          6          4          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2        65002    88675    88652        0        2        42:38 2/4/0

0/0/0
10.0.0.3        65002    54528    54532        0        1       2w4d22h 0/0/0

0/0/0
10.0.0.4        65002    51597    51584        0        0       2w3d22h 2/2/0

0/0/0

user@host> show bgp summary logical-system R3
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
bgp.13vpn.0      2          2          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2          2        204      206        0        0        1:30:59
Establ
  bgp.13vpn.0: 2/2/2/0
  red.inet.0: 2/2/2/0
10.1.1.10        3        206      207        0        0        1:31:36
Establ
  red.inet.0: 2/2/2/0

```

### show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1    200     1735     1737        0        0       14:26:12 Establ

```

```

bgp.isovpn.0: 3/3/0
aaaa.iso.0: 3/3/0

```

### show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l2vpn.0 1 1 0 0 0 0
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last
Up/Dwn State|#Active/Received/Damped...
10.255.245.35 65299 72 74 0 1 19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36 65299 2164 2423 0 4 19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37 65299 36 37 0 4 17:07 Establ
  inet.0: 0/0/0
10.255.245.39 65299 138 168 0 6 53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69 65299 134 140 0 6 53:42 Establ
  inet.0: 0/0/0

```

### show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5 2 21 22 0 0 6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15 1 19 21 0 0 6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0

```

### show bgp summary group

```

user@host> show bgp summary group Group2
Groups: 3 Peers: 3 Down peers: 3
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1 56 0 0 0 0 51
Idle

user@host> show bgp summary logical-system R3 group toR4
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.10 3 207 207 0 0 1:31:40

```



```
Estab1  
red.inet.0: 2/2/2/0
```

## Sample Output

## show policy damping

---

|   |  |
|---|--|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3784</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3784</a>   |
| <b>Syntax</b>                                   | <code>show policy damping</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <code>show policy damping</code>   |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>                              | Display information about BGP route flap damping parameters.   |
| <b>Options</b>                                  | <b>none</b> —Display information about BGP route flap damping parameters.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.   |
| <b>Additional Information</b>                   | In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes. |
| <b>Required Privilege Level</b>                 | view   |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"><li>• “Configuring BGP Flap Damping Parameters” in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li><li>• <a href="#">clear bgp damping on page 3750</a></li><li>• <a href="#">show route damping on page 3098</a></li></ul>  |
| <b>List of Sample Output</b>                    | <a href="#">show policy damping on page 3785</a>   |
| <b>Output Fields</b>                            | <a href="#">Table 302 on page 3785</a> describes the output fields for the <b>show policy damping</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 302: show policy damping Output Fields

| Field Name                   | Field Description  |
|------------------------------|--|
| <b>Halflife</b>              | Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes. |
| <b>Reuse merit</b>           | Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.  |
| <b>Suppress/cutoff merit</b> | Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.   |
| <b>Maximum suppress time</b> | Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.   |
| <b>Computed values</b>       | <ul style="list-style-type: none"> <li>• <b>Merit ceiling</b>—Maximum merit that a flapping route can collect.</li> <li>• <b>Maximum decay</b>—Maximum decay half-life, in minutes.</li> </ul>   |

## Sample Output

### show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```

## show route damping

---

|   |   |
|---|---|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 3786</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 3786</a>  |
| <b>Syntax</b>                                   | <code>show route damping (decayed   history   suppressed)<br/>&lt;brief   detail   extensive   terse&gt;<br/>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>   |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <code>show route damping (decayed   history   suppressed)<br/>&lt;brief   detail   extensive   terse&gt;</code>   |
| <b>Release Information</b>                      | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 11.3 for the QFX Series.  |
| <b>Description</b>                              | Display the BGP routes for which updates might have been reduced because of route flap damping.   |
| <b>Options</b>                                  | <b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.<br><br><b>decayed</b> —Display route damping entries that might no longer be valid, but are not suppressed.<br><br><b>history</b> —Display entries that have already been withdrawn, but have been logged.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.<br><br><b>suppressed</b> —Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols. |
| <b>Required Privilege Level</b>                 | view  |
| <b>Related Documentation</b>                    | <ul style="list-style-type: none"><li>• <a href="#">clear bgp damping on page 3750</a></li><li>• <a href="#">show policy damping on page 3784</a></li></ul>   |
| <b>List of Sample Output</b>                    | <a href="#">show route damping decayed detail on page 3789</a><br><a href="#">show route damping history on page 3790</a><br><a href="#">show route damping history detail on page 3790</a>   |
| <b>Output Fields</b>                            | <a href="#">Table 280 on page 3099</a> lists the output fields for the <b>show route damping</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 303: show route damping Output Fields

| Field Name                                   | Field Description   | Level of Output         |
|--|---|-------------------------|
| <i>routing-table-name</i>                    | Name of the routing table—for example, <b>inet.0</b> .  | All levels              |
| <b>destinations</b>                          | Number of destinations for which there are routes in the routing table.   | All levels              |
| <b>number routes</b>                         | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li>• <b>active</b></li> <li>• <b>holdddown</b> (routes that are in a pending state before being declared inactive)</li> <li>• <b>hidden</b> (the routes are not used because of a routing policy)</li> </ul>  | All levels              |
| <b>destination-prefix (entry, announced)</b> | Destination prefix. The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination.   | <b>detail extensive</b> |
| <b>[protocol, preference]</b>                | Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p> | All levels              |
| <b>Next-hop reference count</b>              | Number of references made to the next hop.  | <b>detail extensive</b> |
| <b>Source</b>                                | IP address of the route source.   | <b>detail extensive</b> |
| <b>Next hop</b>                              | Network layer address of the directly reachable neighboring system.   | <b>detail extensive</b> |
| <b>via</b>                                   | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b> .  | <b>detail extensive</b> |
| <b>Protocol next hop</b>                     | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.  | <b>detail extensive</b> |
| <b>Indirect next hop</b>                     | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.   | <b>detail extensive</b> |
| <b>State</b>                                 | Flags for this route. For a description of possible values for this field, see the output field table for the <a href="#">show route detail</a> command.  | <b>detail extensive</b> |

Table 303: show route damping Output Fields (*continued*)

| Field Name        | Field Description   | Level of Output  |
|-------------------|---|------------------|
| Local AS          | AS number of the local routing device.  | detail extensive |
| Peer AS           | AS number of the peer routing device.   | detail extensive |
| Age               | How long the route has been known.  | detail extensive |
| Metric            | Metric for the route.   | detail extensive |
| Task              | Name of the protocol that has added the route.  | detail extensive |
| Announcement bits | List of protocols that announce this route. <i>n-Resolve inet</i> indicates that the route is used for route resolution for next hops found in the routing table. <i>n</i> is an index used by Juniper Networks customer support only.  | detail extensive |
| AS path           | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• I—IGP.</li> <li>• E—EGP.</li> <li>• ?—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• [ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured.</li> <li>• { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• ( )—Parentheses enclose a confederation.</li> <li>• ( [ ] )—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> | All levels       |
| to                | Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.   | brief none       |
| via               | Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word <b>Selected</b> .  | brief none       |
| Communities       | Community path attribute for the route. See the output field table for the <a href="#">show route detail</a> command.   | detail extensive |
| Localpref         | Local preference value included in the route.   | All levels       |
| Router ID         | BGP router ID as advertised by the neighbor in the open message.  | detail extensive |

Table 303: show route damping Output Fields (*continued*)

| Field Name                     | Field Description  | Level of Output         |
|--------------------------------|--|-------------------------|
| <b>Merit (last update/now)</b> | Last updated and current figure-of-merit value.  | <b>detail extensive</b> |
| <b>damping-parameters</b>      | Name that identifies the damping parameters used, which is defined in the damping statement at the <b>[edit policy-options]</b> hierarchy level.                         | <b>detail extensive</b> |
| <b>Last update</b>             | Time of most recent change in path attributes.   | <b>detail extensive</b> |
| <b>First update</b>            | Time of first change in path attributes, which started the route damping process.  | <b>detail extensive</b> |
| <b>Flaps</b>                   | Number of times the route has gone up or down or its path attributes have changed.   | <b>detail extensive</b> |
| <b>Suppressed</b>              | ( <b>suppressed</b> keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it. | All levels              |
| <b>Reusable in</b>             | ( <b>suppressed</b> keyword only) Time when a suppressed route will again be available.  | All levels              |
| <b>Preference will be</b>      | ( <b>suppressed</b> keyword only) Preference value that will be applied to the route when it is again active.  | All levels              |

## Sample Output

### show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

  6-Resolve tree 2 7-Resolve tree 3
    AS path: 65490 65520 65525 65525 65525 65525 I ()
    Communities: 65501:390 65501:2000 65501:3000 65504:701
    Localpref: 100
    Router ID: 172.23.2.129
    Merit (last update/now): 1934/1790
    damping-parameters: damping-high

```

Last update: 00:03:28 First update: 00:06:40  
Flaps: 2

### show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0
```

### show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update: 00:01:05 First update: 00:01:05
        Flaps: 1
```



## show route detail

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 3791</a><br><a href="#">Syntax (EX Series Switches) on page 3791</a>   |
| <b>Syntax</b>                      | show route detail<br><destination-prefix><br><logical-system (all   logical-system-name)>   |
| <b>Syntax (EX Series Switches)</b> | show route detail<br><destination-prefix>   |
| <b>Release Information</b>         | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.   |
| <b>Description</b>                 | Display detailed information about the active entries in the routing tables.  |
| <b>Options</b>                     | <b>none</b> —Display all active entries in the routing table on all systems.<br><br><b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.<br><br><b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show route detail on page 3800</a><br><a href="#">show route detail (with BGP Multipath) on page 3806</a><br><a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 3806</a><br><a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 3807</a>  |
| <b>Output Fields</b>               | Table 281 on page 3103 describes the output fields for the <b>show route detail</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 304: show route detail Output Fields**

| Field Name                 | Field Description   |
|----------------------------|---|
| <i>routing-table-name</i>  | Name of the routing table (for example, inet.0).  |
| <i>number destinations</i> | Number of destinations for which there are routes in the routing table.   |
| <i>number routes</i>       | Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> <li><b>active</b> (routes that are active)</li> <li><b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li><b>hidden</b> (routes that are not used because of a routing policy)</li> </ul> |

Table 304: show route detail Output Fields (*continued*)

| Field Name                                     | Field Description   |
|--|---|
| <i>route-destination</i><br>(entry, announced) | <p>Route destination (for example:10.0.0.1/24). The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul> |
| <b>label stacking</b>                          | <p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>  |
| <b>[protocol, preference]</b>                  | <p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+—</b>A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>—</b>A hyphen indicates the last active route.</li> <li>• <b>*—</b>An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>   |
| <b>Level</b>                                   | <p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>  |
| <b>Route Distinguisher</b>                     | IP subnet augmented with a 64-bit prefix.   |
| <b>PMSI</b>                                    | Provider multicast service interface (MVPN routing table).  |
| <b>Next-hop type</b>                           | Type of next hop. For a description of possible values for this field, see <a href="#">Table 282 on page 3108</a> .   |

Table 304: show route detail Output Fields (*continued*)

| Field Name   | Field Description   |
|--|---|
| <b>Next-hop reference count</b>                      | Number of references made to the next hop.  |
| <b>Flood nexthop branches exceed maximum message</b> | Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.   |
| <b>Source</b>  | IP address of the route source.   |
| <b>Next hop</b>                                      | Network layer address of the directly reachable neighboring system.   |
| <b>via</b>   | <p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul> |
| <b>Label-switched-path<br/>lsp-path-name</b>         | Name of the LSP used to reach the next hop.   |
| <b>Label operation</b>                               | MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).  |
| <b>Interface</b>                                     | (Local only) Local interface name.  |
| <b>Protocol next hop</b>                             | Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.  |
| <b>Indirect next hop</b>                             | Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.   |
| <b>State</b>   | State of the route (a route can be in more than one state). See <a href="#">Table 283 on page 3109</a> .  |
| <b>Local AS</b>                                      | AS number of the local routing device.  |
| <b>Age</b>   | How long the route has been known.  |
| <b>AIGP</b>  | Accumulated interior gateway protocol (AIGP) BGP attribute.   |
| <b>Metricn</b>                                       | Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.   |

Table 304: show route detail Output Fields (*continued*)

| Field Name                 | Field Description   |
|----------------------------|---|
| <b>MED-plus-IGP</b>        | Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.   |
| <b>TTL-Action</b>          | <p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <a href="#">show route table</a>.</p>   |
| <b>Task</b>                | Name of the protocol that has added the route.  |
| <b>Announcement bits</b>   | List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.  |
| <b>AS path</b>             | <p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li><b>I</b>—IGP.</li> <li><b>E</b>—EGP.</li> <li><b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li><b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li><b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893.</li> <li><b>[ ]</b>—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li><b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li><b>( )</b>—Parentheses enclose a confederation.</li> <li><b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p> |
| <b>validation-state</b>    | <p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li><b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li><b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li><b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li><b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>   |
| <b>FECs bound to route</b> | Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.   |

Table 304: show route detail Output Fields (*continued*)

| Field Name                | Field Description   |
|---------------------------|---|
| Primary Upstream          | When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path. |
| RPF Nexthops              | When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.   |
| Label                     | Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.                         |
| weight                    | Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.  |
| VC Label                  | MPLS label assigned to the Layer 2 circuit virtual connection.  |
| MTU                       | Maximum transmission unit (MTU) of the Layer 2 circuit.   |
| VLAN ID                   | VLAN identifier of the Layer 2 circuit.   |
| Prefixes bound to route   | Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.  |
| Communities               | Community path attribute for the route. See <a href="#">Table 284 on page 3111</a> for all possible values for this field.  |
| Layer2-info: encaps       | Layer 2 encapsulation (for example, VPLS).  |
| control flags             | Control flags: <b>none</b> or <b>Site Down</b> .  |
| mtu                       | Maximum transmission unit (MTU) information.  |
| Label-Base, range         | First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.  |
| status vector             | Layer 2 VPN and VPLS network layer reachability information (NLRI).   |
| Accepted Multipath        | Current active path when BGP multipath is configured.   |
| Accepted MultipathContrib | Path currently contributing to BGP multipath.   |
| Localpref                 | Local preference value included in the route.   |
| Router ID                 | BGP router ID as advertised by the neighbor in the open message.  |
| Primary Routing Table     | In a routing table group, the name of the primary routing table in which the route resides.   |
| Secondary Tables          | In a routing table group, the name of one or more secondary tables in which the route resides.  |

Table 282 on page 3108 describes all possible values for the Next-hop Types output field.

**Table 305: Next-hop Types Output Field Values**

| Next-Hop Type                   | Description  |
|---------------------------------|--|
| <b>Broadcast (bcast)</b>        | Broadcast next hop.  |
| <b>Deny</b>                     | Deny next hop.   |
| <b>Discard</b>                  | Discard next hop.  |
| <b>Flood</b>                    | Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast. |
| <b>Hold</b>                     | Next hop is waiting to be resolved into a unicast or multicast type.   |
| <b>Indexed (idxd)</b>           | Indexed next hop.  |
| <b>Indirect (indr)</b>          | Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.   |
| <b>Interface</b>                | Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.   |
| <b>Local (locl)</b>             | Local address on an interface. This next-hop type causes packets with this destination address to be received locally.   |
| <b>Multicast (mcst)</b>         | Wire multicast next hop (limited to the LAN).  |
| <b>Multicast discard (mdsc)</b> | Multicast discard.   |
| <b>Multicast group (mgrp)</b>   | Multicast group member.  |
| <b>Receive (recv)</b>           | Receive.   |
| <b>Reject (rjct)</b>            | Discard. An ICMP unreachable message was sent.   |
| <b>Resolve (rslv)</b>           | Resolving next hop.  |
| <b>Routed multicast (mcrt)</b>  | Regular multicast next hop.  |

Table 305: Next-hop Types Output Field Values (*continued*)

| Next-Hop Type         | Description  |
|-----------------------|--|
| <b>Router</b>         | <p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> <li>• Must not be a direct or local subnet for the routing device.</li> <li>• Must have a next hop that is directly connected to the routing device.</li> </ul> |
| <b>Table</b>          | Routing table next hop.  |
| <b>Unicast (ucst)</b> | Unicast.   |
| <b>Unilist (ulst)</b> | List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.  |

Table 283 on page 3109 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 306: State Output Field Values

| Value  | Description  |
|--|--|
| <b>Accounting</b>                            | Route needs accounting.  |
| <b>Active</b>                                | Route is active.   |
| <b>Always Compare MED</b>                    | Path with a lower multiple exit discriminator (MED) is available.                |
| <b>AS path</b>                               | Shorter AS path is available.  |
| <b>Cisco Non-deterministic MED selection</b> | Cisco nondeterministic MED is enabled, and a path with a lower MED is available. |
| <b>Clone</b>                                 | Route is a clone.  |
| <b>Cluster list length</b>                   | Length of cluster list sent by the route reflector.                              |
| <b>Delete</b>                                | Route has been deleted.  |
| <b>Ex</b>                                    | Exterior route.  |
| <b>Ext</b>                                   | BGP route received from an external BGP neighbor.                                |

Table 306: State Output Field Values (*continued*)

| Value  | Description  |
|--|--|
| <b>FlashAll</b>  | Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes. |
| <b>Hidden</b>  | Route not used because of routing policy.  |
| <b>IfCheck</b>   | Route needs forwarding RPF check.  |
| <b>IGP metric</b>  | Path through next hop with lower IGP metric is available.  |
| <b>Inactive reason</b>                                   | Flags for this route, which was not selected as best for a particular destination.   |
| <b>Initial</b>   | Route being added.   |
| <b>Int</b>   | Interior route.  |
| <b>Int Ext</b>   | BGP route received from an internal BGP peer or a BGP confederation peer.  |
| <b>Interior &gt; Exterior &gt; Exterior via Interior</b> | Direct, static, IGP, or EBGP path is available.  |
| <b>Local Preference</b>                                  | Path with a higher local preference value is available.  |
| <b>Martian</b>   | Route is a martian (ignored because it is obviously invalid).  |
| <b>MartianOK</b>   | Route exempt from martian filtering.   |
| <b>Next hop address</b>                                  | Path with lower metric next hop is available.  |
| <b>No difference</b>                                     | Path from neighbor with lower IP address is available.   |
| <b>NoReadvrt</b>   | Route not to be advertised.  |
| <b>NotBest</b>   | Route not chosen because it does not have the lowest MED.  |
| <b>Not Best in its group</b>                             | Incoming BGP AS is not the best of a group (only one AS can be the best).  |
| <b>NotInstall</b>  | Route not to be installed in the forwarding table.   |
| <b>Number of gateways</b>                                | Path with a greater number of next hops is available.  |
| <b>Origin</b>  | Path with a lower origin code is available.  |
| <b>Pending</b>   | Route pending because of a hold-down configured on another route.  |



Table 306: State Output Field Values (*continued*)

| Value                                 | Description   |
|---------------------------------------|---|
| <b>Release</b>                        | Route scheduled for release.  |
| <b>RIB preference</b>                 | Route from a higher-numbered routing table is available.  |
| <b>Route Distinguisher</b>            | 64-bit prefix added to IP subnets to make them unique.  |
| <b>Route Metric or MED comparison</b> | Route with a lower metric or MED is available.  |
| <b>Route Preference</b>               | Route with lower preference value is available  |
| <b>Router ID</b>                      | Path through a neighbor with lower ID is available.   |
| <b>Secondary</b>                      | Route not a primary route.  |
| <b>Unusable path</b>                  | Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> <li>• The route is damped.</li> <li>• The route is rejected by an import policy.</li> <li>• The route is unresolved.</li> </ul> |
| <b>Update source</b>                  | Last tiebreaker is the lowest IP address value.   |

Table 284 on page 3111 describes the possible values for the Communities output field.

Table 307: Communities Output Field Values

| Value   | Description   |
|---|---|
| <i>area-number</i>                                      | 4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.  |
| <b>bandwidth: local AS number:link-bandwidth-number</b> | Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute. |
| <b>domain-id</b>  | Unique configurable number that identifies the OSPF domain.   |
| <b>domain-id-vendor</b>                                 | Unique configurable number that further identifies the OSPF domain.   |
| <i>link-bandwidth-number</i>                            | Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).   |
| <i>local AS number</i>                                  | Local AS number: from 1 through 65,535.   |
| <i>options</i>  | 1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.   |

Table 307: Communities Output Field Values (*continued*)

| Value                                | Description   |
|--------------------------------------|---|
| <b>origin</b>                        | (Used with VPNs) Identifies where the route came from.  |
| <b>ospf-route-type</b>               | 1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses. |
| <b>route-type-vendor</b>             | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x8000</b> . The format is <b>area-number:ospf-route-type:options</b> .  |
| <b>rte-type</b>                      | Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute <b>0x0306</b> . The format is <b>area-number:ospf-route-type:options</b> .  |
| <b>target</b>                        | Defines which VPN the route participates in; <b>target</b> has the format <b>32-bit IP address:16-bit number</b> . For example, 10.19.0.0:100.  |
| <b>unknown IANA</b>                  | Incoming IANA codes with a value between <b>0x1</b> and <b>0x7fff</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.  |
| <b>unknown OSPF vendor community</b> | Incoming IANA codes with a value above <b>0x8000</b> . This code of the BGP extended community attribute is accepted, but it is not recognized.   |

## Sample Output

### show route detail

```

user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10

```

```

Next-hop reference count: 1
Next hop: via so-0/3/0.0, selected
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:30:17 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: IGMP
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

```

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
          Label-switched-path green-r1-r3
          Label operation: Push 100096
          State: <Active Int>
          Local AS: 69
          Age: 1:25:49   Metric: 2
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS: 69
          Age: 1:25:49   Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
          State: <Active Int>
          Local AS: 69
          Age: 1:31:44
          Task: IF
          AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
  *MPLS   Preference: 0
          Next hop type: Receive
          Next-hop reference count: 6
          State: <Active Int>
          Local AS: 69
          Age: 1:31:45   Metric: 1
          Task: MPLS
          Announcement bits (1): 0-KRT
          AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

299840 (1 entry, 1 announced)
```

```

TSI:
KRT in-kerne 299840 /52 -> {indirect(1048575)}
    *RSVP Preference: 7/2
        Next hop type: Flood
        Address: 0x9174a30
        Next-hop reference count: 4
        Next hop type: Router, Next hop index: 798
        Address: 0x9174c28
        Next-hop reference count: 2
        Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
        Label-switched-path R2-to-R4-2p2mp
        Label operation: Pop
        Next hop type: Router, Next hop index: 1048574
        Address: 0x92544f0
        Next-hop reference count: 2
        Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
        Label-switched-path R2-to-R200-p2mp
        Label operation: Pop
        Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
        Label operation: Pop
        State: <Active Int>
        Age: 1:29 Metric: 1
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I...

800010 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:29:30
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:29:30 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected

```

```
State: <Active Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

ff02::2/128 (1 entry, 1 announced)
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::d/128 (1 entry, 1 announced)
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::16/128 (1 entry, 1 announced)
*MLD Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:43
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:31:44
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS: 69 Peer AS: 69
        Age: 1:25:49 Metric2: 1
        AIGP 210
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-green-l2vpn
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Label-base: 800008, range: 8
        Localpref: 100
        Router ID: 10.255.70.103
        Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
        mtu: 0
        Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
        Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp
        Label operation: Push 100000[0]
        Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
        State: <Active Int>
        Local AS: 99
        Age: 10:21
        Task: l2 circuit

```

```
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512
```

### show route detail (with BGP Multipath)

```
user@host> show route detail
```

```
10.1.1.8/30 (2 entries, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 262142
        Address: 0x901a010
        Next-hop reference count: 2
        Source: 10.1.1.2
        Next hop: 10.1.1.2 via ge-0/3/0.1, selected
        Next hop: 10.1.1.6 via ge-0/3/0.5
        State: <Active Ext>
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.2+59955
        Announcement bits (1): 0-KRT
        AS path: 2 I
        Accepted Multipath
        Localpref: 100
        Router ID: 1.1.1.2
  BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 678
        Address: 0x8f97520
        Next-hop reference count: 9
        Source: 10.1.1.6
        Next hop: 10.1.1.6 via ge-0/3/0.5, selected
        State: <NotBest Ext>
        Inactive reason: Not Best in its group - Active preferred
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.6+58198
        AS path: 2 I
        Accepted MultipathContrib
        Localpref: 100
        Router ID: 1.1.1.3
```

### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Next-hop reference count: 3
        Address: 0x9097d90
        Next hop: via vt-0/1/0.1
        Next-hop index: 661
        Label operation: Pop
        Address: 0x9172130
        Next hop: via so-0/0/3.0
        Next-hop index: 654
        Label operation: Swap 299872
        State: **Active Int>
        Local AS: 1001
```



```

Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
      RPF Nexthops :
        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
      RPF Nexthops :
        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```



## PART 12

# Intermediate System to Intermediate System

- [Overview on page 3811](#)
- [Configuration on page 3823](#)
- [Administration on page 3977](#)



## CHAPTER 44

# Overview

- [IS-IS Overview on page 3811](#)

### IS-IS Overview

---

- [IS-IS Overview on page 3812](#)
- [Understanding BFD Authentication for IS-IS on page 3816](#)
- [Understanding Hitless Authentication Key Rollover for IS-IS on page 3818](#)
- [Understanding Loop-Free Alternate Routes for IS-IS on page 3819](#)

## IS-IS Overview

The IS-IS protocol is an interior gateway protocol (IGP) that uses link-state information to make routing decisions.

IS-IS is a link-state IGP that uses the shortest-path-first (SPF) algorithm to determine routes. IS-IS evaluates the topology changes and determines whether to perform a full SPF recalculation or a partial route calculation (PRC). This protocol originally was developed for routing International Organization for Standardization (ISO) Connectionless Network Protocol (CLNP) packets.

Like OSPF routing, IS-IS uses hello packets that allow network convergence to occur quickly when network changes are detected. IS-IS uses the SPF algorithm to determine routes. Using SPF, IS-IS evaluates network topology changes and determines if a full or partial route calculation is required.



**NOTE:** Because IS-IS uses ISO addresses, the configuration of IP version 6 (IPv6) and IP version 4 (IPv4) implementations of IS-IS is identical.



**NOTE:** See *Platforms/FPCs That Cannot Forward TCC Encapsulated ISO Traffic* to find a list of those devices and FPC configurations that cannot pass ISO traffic when encapsulated in TCC format.

This section discusses the following topics:

- [IS-IS Terminology on page 3812](#)
- [ISO Network Addresses on page 3813](#)
- [IS-IS Packets on page 3814](#)
- [Persistent Route Reachability on page 3815](#)
- [IS-IS Support for Multipoint Network Clouds on page 3815](#)
- [Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels on page 3816](#)

---

### IS-IS Terminology

An IS-IS network is a single autonomous system (AS), also called a *routing domain*, that consists of *end systems* and *intermediate systems*. End systems are network entities that send and receive packets. Intermediate systems send and receive packets and relay (forward) packets. (Intermediate system is the Open System Interconnection [OSI] term for a router.) ISO packets are called network PDUs.

In IS-IS, a single AS can be divided into smaller groups called *areas*. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring *Level 1* and *Level 2* intermediate systems. Level 1 systems route within an area; when the destination is

outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

### ISO Network Addresses

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, and is called a *network service access point (NSAP)*.

IS-IS supports multiple NSAP addresses on the loopback lo0 interface.

An end system can have multiple NSAP addresses, in which case the addresses differ only by the last byte (called the *n-selector*). Each NSAP represents a service that is available at that node. In addition to having multiple services, a single node can belong to multiple areas.

Each network entity also has a special network address called a *network entity title (NET)*. Structurally, an NET is identical to an NSAP address but has an n-selector of 00. Most end systems and intermediate systems have one NET. Intermediate systems that participate in multiple areas can have multiple NETs.

The following ISO addresses illustrate the IS-IS address format:

```
49.0001.00a0.c96b.c490.00
49.0001.2081.9716.9018.00
```

NETs take several forms, depending on your network requirements. NET addresses are hexadecimal and range from 8 octets to 20 octets in length. Generally, the format consists of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier, and a selector. The simplest format omits the domain ID and is 10 octets long. For example, the NET address 49.0001.1921.6800.1001.00 consists of the following parts:

- 49—AFI
- 0001—Area ID
- 1921.6800.1001—System identifier
- 00—Selector

The system identifier must be unique within the network. For an IP-only network, we recommend using the IP address of an interface on the router. Configuring a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

The first portion of the address is the area number, which is a variable number from 1 through 13 bytes. The first byte of the area number (49) is the authority and format indicator (AFI). The next bytes are the assigned domain (area) identifier, which can be from 0 through 12 bytes. In the examples above, the area identifier is 0001.

The next six bytes form the system identifier. The system identifier can be any six bytes that are unique throughout the entire domain. The system identifier commonly is the media access control (MAC) address (as in the first example, 00a0.c96b.c490) or the IP address expressed in binary-coded decimal (BCD) (as in the second example, 2081.9716.9018, which corresponds to IP address 208.197.169.18). The last byte (00) is the n-selector.



**NOTE:** The system identifier cannot be 0000.0000.0000. All 0s is an illegal setting, and the adjacency is not formed with this setting.

To provide help with IS-IS debugging, the Junos® operating system (Junos OS) supports dynamic mapping of ISO system identifiers to the hostname. Each system can be configured with a hostname, which allows the system identifier-to-hostname mapping to be carried in a dynamic hostname type, length, and value (TLV) tuple in IS-IS link-state PDUs. This enables intermediate systems in the routing domain to learn about the ISO system identifier of a particular intermediate system.

---

### IS-IS Packets

---

Each IS-IS PDU shares a common header. IS-IS uses the following PDUs to exchange protocol information:

- IS-IS hello (IIH) PDUs—Broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.

IS-IS hello PDUs establish adjacencies with other routers and have three different formats: one for point-to-point hello packets, one for Level 1 broadcast links, and one for Level 2 broadcast links. Level 1 routers must share the same area address to form an adjacency, while Level 2 routers do not have this limitation. The request for adjacency is encoded in the Circuit type field of the PDU.

Hello PDUs have a preset length assigned to them. The IS-IS router does not resize any PDU to match the maximum transmission unit (MTU) on a router interface. Each interface supports the maximum IS-IS PDU of 1492 bytes, and hello PDUs are padded to meet the maximum value. When the hello is sent to a neighboring router, the connecting interface supports the maximum PDU size.

- Link-state PDUs—Contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.

Also included is metric and IS-IS neighbor information. Each link-state PDU must be refreshed periodically on the network and is acknowledged by information within a sequence number PDU.

On point-to-point links, each link-state PDU is acknowledged by a partial sequence number PDU (PSNP), but on broadcast links, a complete sequence number PDU



(CSNP) is sent out over the network. Any router that finds newer link-state PDU information in the CSNP then purges the out-of-date entry and updates the link-state database.

Link-state PDUs support variable-length subnet mask addressing.

- Complete sequence number PDUs (CSNPs)—Contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU.

Contained within the CSNP is a link-state PDU identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database. Also, the advertisement of CSNPs occurs when an adjacency is formed with another router. Like IS-IS hello PDUs, CSNPs come in two types: Level 1 and Level 2.

When a device receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the device requests specific link-state PDU details using a partial sequence number PDU (PSNP).

- Partial sequence number PDUs (PSNPs)—Sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device.

A PSNP is used by an IS-IS router to request link-state PDU information from a neighboring router. A PSNP can also explicitly acknowledge the receipt of a link-state PDU on a point-to-point link. On a broadcast link, a CSNP is used as implicit knowledge. Like hello PDUs and CSNPs, the PSNP also has two types: Level 1 and Level 2.

When a device compares a CSNP to its local database and determines that a link-state PDU is missing, the router issues a PSNP for the missing link-state PDU, which is returned in a link-state PDU from the router sending the CSNP. The received link-state PDU is then stored in the local database, and an acknowledgment is sent back to the originating router.

### Persistent Route Reachability

IPv4 and IPv6 route reachability information in IS-IS link-state PDUs is preserved when you commit a configuration. IP prefixes are preserved with their original packet fragment upon link-state PDU regeneration.


### IS-IS Support for Multipoint Network Clouds

IS-IS does not support multipoint configurations. Therefore, when configuring Frame Relay or Asynchronous Transfer Mode (ATM) networks, you must configure them as collections of point-to-point links, not as multipoint clouds.

### Installing a Default Route to the Nearest Routing Device That Operates at Both IS-IS Levels

When a routing device that operates as both a Level 1 and Level 2 router (Router B) determines that it can reach at least one area other than its own (for example, in Area Y), it sets the ATTACHED bit in its Level 1 link-state PDU. Thereafter, the Level 1 router (Router A) introduces a default route pointing to the nearest attached routing device that operates as both a Level 1 and Level 2 router (Router B). See [Figure 98 on page 3816](#).

**Figure 98: Install Default Route to Nearest Routing Device That Operates at Both Level 1 and Level 2**



**Related Documentation**

- [IS-IS Feature Guide for Routing Devices](#)

### Understanding BFD Authentication for IS-IS

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IS-IS. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 3817](#)
- [Security Authentication Keychains on page 3817](#)
- [Strict Versus Loose Authentication on page 3818](#)

## BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords might be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

## Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session,

and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### Strict Versus Loose Authentication

---

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

**Related Documentation** • [Example: Configuring BFD Authentication for IS-IS on page 3855](#)

## Understanding Hitless Authentication Key Rollover for IS-IS

IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in routing. By default, authentication is disabled. The authentication algorithm creates an encoded checksum that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet's checksum.

If you configure authentication for all peers, each peer in that group inherits the group's authentication.

You can update authentication keys without resetting any IS-IS neighbor sessions. This is referred to as *hitless authentication key rollover*.

Hitless authentication key rollover uses authentication keychains, which consist of the authentication keys that are being updated. The keychain includes multiple keys. Each key in the keychain has a unique start time. At the next key's start time, a rollover occurs from the current key to the next key, and the next key becomes the current key.

You can choose the algorithm through which authentication is established. You can configure MD5 or SHA-1 authentication. You associate a keychain and the authentication algorithm with an IS-IS neighboring session. Each key contains an identifier and a secret password.

The sending peer chooses the active key based on the system time and the start times of the keys in the keychain. The receiving peer determines the key with which it authenticates based on the incoming key identifier.

You can configure either RFC 5304-based encoding or RFC 5310-based encoding for the IS-IS protocol transmission encoding format.

**Related Documentation** • [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837](#)

## Understanding Loop-Free Alternate Routes for IS-IS

In Junos OS Release 9.5 and later, support for IS-IS loop-free alternate routes enables IP fast-reroute capability for IS-IS. Junos OS precomputes loop-free backup routes for all IS-IS routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives recomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. In contrast, global repair can take up to 800 milliseconds to compute a new route. Local repair and global repair are thus complementary. Local repair enables traffic to continue to be routed using a backup path until global repair is able to calculate a new route.

A loop-free path is one that does not forward traffic back through the routing device to reach a given destination. That is, a neighbor whose shortest path to the destination traverses the routing device is not used as a backup route to that destination. To determine loop-free alternate paths for IS-IS routes, Junos OS runs shortest-path-first (SPF) calculations on each one-hop neighbor. You can enable support for alternate loop-free routes on any IS-IS interface. Because it is common practice to enable LDP on an interface for which IS-IS is already enabled, this feature also provides support for LDP label-switched paths (LSPs).



**NOTE:** If you enable support for alternate loop-free routes on an interface configured for both LDP and IS-IS, you can use the `traceroute` command to trace the active path to the primary next hop.

The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.


Junos OS provides two mechanisms for route redundancy for IS-IS through alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS interface, Junos OS creates a single alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection. Use link protection when you assume that only a single link might become unavailable but that the neighboring node on the primary path would still be available through another interface.

Node-link protection establishes an alternate path through a different routing device altogether. Use node-link protection when you assume that access to a node is lost when a link is no longer available. As a result, Junos OS calculates a backup path that avoids the primary next-hop routing device. In Junos OS Release 9.4 and earlier, only the RSVP protocol supports Packet Forwarding Engine local repair and fast reroute as well as link protection and node protection.

In [Figure 99 on page 3820](#), Case 2 shows how link protection allows source Router A to switch to Link B when the primary next hop Link A to destination Router C fails. However,

if Router B fails, Link B also fails, and the protected Link A is lost. If node-link protection is enabled, Router A is able to switch to Link D on Router D and bypass the failed Router B altogether. As shown in Case 1, with node-link protection enabled, Router A has a node-link protection alternate path available through Router D to destination Router C. That means that if Router B fails, Router A can still reach Router C because the path from Router A to Link D remains available as an alternate backup path.

**Figure 99: Link Protection and Node-Link Protection Comparison for IS-IS Routes**



The Junos OS implementation of support for loop-free alternate paths for IS-IS routes is based on the following standards:

- RFC 5286, *Basic Specification for IP Fast-Reroute: Loop-free Alternates*
- RFC 5714, *IP Fast Reroute Framework*

### Configuring Link Protection for IS-IS

You can configure link protection on any interface for which IS-IS is enabled. When you enable link protection, Junos OS creates one alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection assumes that only a single link becomes unavailable but that the neighboring node would still be available through another interface.



**NOTE:** You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable link protection, include the **link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
      link-protection;
    }
  }
}
```

### Configuring Node-Link Protection for IS-IS

You can configure node-link protection on any interface for which IS-IS is enabled. Node-link protection establishes an alternate path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.



**NOTE:** You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable node-link protection, include the **node-link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
      node-link-protection;
    }
  }
}
```

### Excluding an IS-IS Interface as a Backup for Protected Interfaces

By default, all IS-IS interfaces that belong to the master instance or a specific routing instance are eligible as backup interfaces for protected interfaces. You can specify that any IS-IS interface be excluded from functioning as a backup interface to protected interfaces. To exclude an IS-IS interface as a backup interface, include the **no-eligible-backup** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
      no-eligible-backup;
    }
  }
}
```

### Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS

---

Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP label-switched paths (LSPs) by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup path, include the **backup** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      backup;
      to ip-address;
    }
  }
}
```

When configuring an LSP, you must specify the IP address of the egress routing device with the **to** statement. For detailed information about configuring LSPs and RSVP, see the *RSVP Feature Guide for Routing Devices*.

### Using Operational Mode Commands to Monitor Protected IS-IS Routes

---

You can issue operational mode commands that provide more details about your link-protected and node-link-protected IS-IS routes. The following guidelines explain the type of information available from the output of each command:

- **show isis backup label-switched-path**—Displays which MPLS LSPs have been designated as backup paths and the current status of those LSPs.
- **show isis backup spf results**—Displays SPF calculations for each neighbor for a given destination. Indicates whether a specific interface or node has been designated as a backup path and why. Use the **no-coverage** option to display only those nodes that do not have backup coverage.
- **show isis backup coverage**—Displays the percentage of nodes and prefixes for each type of address family that is protected.
- **show isis interface detail**—Displays the type of protection (link or node-link) applied to each protected interface.

#### Related Documentation

- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 3887](#)



# Configuration

- [Configuration Guidelines on page 3823](#)
- [Configuration Examples on page 3828](#)
- [Configuration Tasks on page 3908](#)
- [Configuration Statements on page 3910](#)

## Configuration Guidelines

---

- [Example: Configuring IS-IS on page 3823](#)

### Example: Configuring IS-IS

This example shows how to configure IS-IS.

- [Requirements on page 3823](#)
- [Overview on page 3823](#)
- [Configuration on page 3824](#)
- [Verification on page 3826](#)

#### Requirements

---

No special configuration beyond device initialization is required before configuring this example.

#### Overview


---

In this example, you configure the two IS-IS routing devices in a single area. The devices have NET addresses 49.0002.0192.0168.0001.00 and 49.0002.0192.0168.0002.00 on the lo0 interfaces. Additionally, you configure the ISO family on the IS-IS interfaces.

For Junos OS security devices only, you configure the **mode packet-based** statement at the **[edit security forwarding-options family iso]** hierarchy level.

[Figure 100 on page 3824](#) shows the topology used in this example.

Figure 100: Simple IS-IS Topology



“CLI Quick Configuration” on page 3824 shows the configuration for both of the devices in Figure 100 on page 3824. The section “Step-by-Step Procedure” on page 3824 describes the steps on Device R1.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set security forwarding-options family iso mode packet-based
set interfaces ge-1/2/0 unit 0 description to-R2
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set protocols isis interface ge-1/2/0.0
set protocols isis interface lo0.0

```

**Device R2**

```

set security forwarding-options family iso mode packet-based
set interfaces ge-1/2/0 unit 0 description to-R1
set interfaces ge-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set protocols isis interface ge-1/2/0.0
set protocols isis interface lo0.0

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS:

1. Enable IS-IS if your router is in secure context.
 

```

[edit security forwarding-options family iso]
user@R1# set mode packet-based
      
```
2. Create the interface that connects to Device R2, and configure the ISO family on the interface.
 

```

[edit interfaces ge-1/2/0 unit 0]
user@R1# set description to-R2
user@R1# set family inet address 10.0.0.1/30
user@R1# set family iso
      
```

3. Create the loopback interface, set the IP address, and set the NET address.

```
[edit interfaces lo0 unit 0]
user@R1# set family inet address 192.168.0.1/32
user@R1# set family iso address 49.0002.0192.0168.0001.00
```

4. Enable IS-IS on the interfaces.

```
[edit protocols isis]
user@R1# set interface ge-1/2/0.0
user@R1# set interface lo0.0
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show security
forwarding-options {
  family iso {
    mode packet-based;
  }
}

user@R1# show interfaces
ge-1/2/0 {
  unit 0 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
    }
  }
}

user@R1# show protocols
isis {
  interface ge-1/2/0.0;
  interface lo0.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying IS-IS Interface Configuration on page 3826](#)
- [Verifying IS-IS Interface Configuration in Detail on page 3826](#)
- [Verifying IS-IS Adjacencies on page 3827](#)
- [Verifying IS-IS Adjacencies in Detail on page 3827](#)

### Verifying IS-IS Interface Configuration

**Purpose** Verify the status of the IS-IS-enabled interfaces.

**Action** From operational mode, enter the **show isis interface brief** command.

```
user@R1> show isis interface brief
IS-IS interface database:
Interface          L CirID Level 1 DR          Level 2 DR          L1/L2 Metric
lo0.0              3  0x1 Passive                Passive              0/0
ge-1/2/0.0         3  0x1 R2.02                  R2.02                10/10
```

**Meaning** Verify that the output shows the intended configuration of the interfaces on which IS-IS is enabled.

### Verifying IS-IS Interface Configuration in Detail

**Purpose** Verify the details of IS-IS-enabled interfaces.

**Action** From operational mode, enter the **show isis interface detail** command.

```
user@R1> show isis interface detail
IS-IS interface database:
lo0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1              0      64      0 Passive
    2              0      64      0 Passive
ge-1/2/0.0
  Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 10 s
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1              1      64     10    9.000    27 R2.02 (not us)
    2              1      64     10    9.000    27 R2.02 (not us)
```

**Meaning** Check the following output fields and verify that the output shows the intended configuration of IS-IS-enabled interfaces:

- Interface—Interface configured for IS-IS.
- State—Internal implementation information.
- Circuit id—Circuit identifier.

- Circuit type—Configured level of IS-IS:
  - 1—Level 1 only
  - 2—Level 2 only
  - 3—Level 1 and Level 2
- link-state PDU interval—Time between IS-IS information messages.
- L or Level—Type of adjacency:
  - 1—Level 1 only
  - 2—Level 2 only
  - 3—Level 1 and Level 2
- Adjacencies—Adjacencies established on the interface.
- Priority—Priority value established on the interface.
- Metric—Metric value for the interface.
- Hello(s)—Intervals between hello PDUs.
- Hold(s)—Hold time on the interface.

### *Verifying IS-IS Adjacencies*

**Purpose** Display brief information about IS-IS neighbors.

**Action** From operational mode, enter the **show isis adjacency brief** command.

```
user@R1> show isis adjacency brief
Interface      System      L State      Hold (secs) SNPA
ge-1/2/0.0     R2          1 Up          6  0:5:85:8f:c8:bd
ge-1/2/0.0     R2          2 Up          6  0:5:85:8f:c8:bd
```

**Meaning** Verify the adjacent routers in the IS-IS database.

### *Verifying IS-IS Adjacencies in Detail*

**Purpose** Display extensive information about IS-IS neighbors.

**Action** From operational mode, enter the **show isis adjacency extensive** command.

```
user@R1> show isis adjacency extensive
R2
Interface: ge-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When              State      Event              Down reason
Thu May 31 11:18:48 Up          Seenself
```

R2

```
Interface: ge-1/2/0.0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 00:40:28 ago
Circuit type: 3, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.0.2
Transition log:
When                State      Event      Down reason
Thu May 31 11:18:48  Up        SeenseIf
```

**Meaning** Check the following fields and verify the adjacency information about IS-IS neighbors:

- Interface—Interface through which the neighbor is reachable.
- L or Level—Configured level of IS-IS:
  - 1—Level 1 only
  - 2—Level 2 only
  - 3—Level 1 and Level 2

An exclamation point before the level number indicates that the adjacency is missing an IP address.

- State—Status of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**.
- Event—Message that identifies the cause of a state.
- Down reason—Reason the adjacency is down.
- Restart capable—A neighbor is configured for graceful restart.
- Transition log—List of transitions including **When**, **State**, and **Reason**.

**Related  
Documentation**

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS for GRES with Graceful Restart*
- [Example: Configuring Designated Router Election Priority for IS-IS on page 3906](#)
- *Junos OS Feature Support Reference for SRX Series and J Series Devices*

---

## Configuration Examples

- [Example: Configuring Multi-Level IS-IS on page 3829](#)
- [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837](#)
- [Example: Redistributing OSPF Routes into IS-IS on page 3842](#)
- [Example: Configuring BFD for IS-IS on page 3850](#)
- [Example: Configuring BFD Authentication for IS-IS on page 3855](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859](#)
- [Example: Configuring IS-IS Multicast Topology on page 3867](#)
- [Example: Configuring Link and Node Protection for IS-IS Routes on page 3883](#)
- [Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 3897](#)

- [Example: Configuring IS-IS for CLNS on page 3903](#)
- [Example: Configuring IS-IS Designated Routers on page 3905](#)
- [Example: Enabling Packet Checksums on IS-IS Interfaces on page 3906](#)

## Example: Configuring Multi-Level IS-IS

This example shows how to configure a multi-level IS-IS topology.

- [Requirements on page 3829](#)
- [Overview on page 3829](#)
- [Configuration on page 3830](#)
- [Verification on page 3834](#)

---

### Requirements

No special configuration beyond device initialization is required before configuring this example.

---

### Overview


Like OSPF, the IS-IS protocol supports the partitioning of a routing domain into multiple areas with levels that control interarea flooding. The use of multiple levels improves protocol scalability, as Level 2 (backbone) link-state PDUs are normally not flooded into a Level 1 area.

An IS-IS Level 2 area is analogous to the OSPF backbone area (0), while a Level 1 area operates much like an OSPF totally stubby area, in that a default route is normally used to reach both inter-level and AS external routes.

Unlike OSPF, IS-IS area boundaries occur between routers, such that a given routing device is always wholly contained within a particular area. Level 1 adjacencies can be formed between routers that share a common area number, while a Level 2 adjacency can be formed between routers that might or might not share an area number.

[Figure 101 on page 3830](#) shows the topology used in this example.

### Figure 101: IS-IS Multi-Level Topology



“CLI Quick Configuration” on page 3830 shows the configuration for all of the devices in Figure 101 on page 3830. The section “Step-by-Step Procedure” on page 3832 describes the steps on Device R5.

This example has the following characteristics:

- Device R5 functions as a Level 1/Level 2 router to interconnect the Level 2 backbone area 49.001 and the Level 1 area 49.002 containing Device R6 and Device R7.
- The system ID is based on the devices' IPv4 lo0 addresses.
- Loss of any individual interface does not totally disrupt the IS-IS operation.
- The IPv4 lo0 addresses of all routers are reachable through IS-IS.
- The link between Device R3 and Device S1 appears in area 49.001 as an intra-area route. No IS-IS adjacencies can be established on this interface. This is accomplished by configuring the **passive** statement on Device R3's interface to Device S1.
- The loopback addresses of Level 2 devices do not appear in a Level 1 area.
- There is only one adjacency for each device pairing.

## Configuration

## CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

## Device R3

```
set interfaces fe-1/2/0 unit 0 description to-R4
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.41/30
```



```

set interfaces fe-1/2/2 unit 0 description to-S1
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0003.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable
set protocols isis interface fe-1/2/2.0 passive

```

**Device R4**

```

set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.001.0192.0168.0004.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface lo0.0 level 1 disable

```

**Device R5**

```

set interfaces fe-1/2/0 unit 0 description to-R3
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R4
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 description to-R6
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/3 unit 0 description to-R7
set interfaces fe-1/2/3 unit 0 family inet address 10.0.0.38/30
set interfaces fe-1/2/3 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0005.00
set protocols isis interface fe-1/2/0.0 level 1 disable
set protocols isis interface fe-1/2/1.0 level 1 disable
set protocols isis interface fe-1/2/2.0 level 2 disable
set protocols isis interface fe-1/2/3.0 level 2 disable
set protocols isis interface lo0.0 level 1 disable

```

**Device R6**

```

set interfaces fe-1/2/0 unit 0 description to-R5
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/1 unit 0 description to-R7
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.33/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0006.00
set protocols isis interface fe-1/2/0.0 level 2 disable
set protocols isis interface fe-1/2/1.0 level 2 disable
set protocols isis interface lo0.0 level 2 disable

```

**Device R7**

```

set interfaces fe-1/2/0 unit 0 description to-R6
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.34/30
set interfaces fe-1/2/0 unit 0 family iso

```

```
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.7/32
set interfaces lo0 unit 0 family iso address 49.002.0192.0168.0007.00
set protocols isis interface fe-1/2/0.0 level 2 disable
set protocols isis interface fe-1/2/1.0 level 2 disable
set protocols isis interface lo0.0 level 2 disable
```

**Device S1**      `set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.42/30`  
`set interfaces fe-1/2/0 unit 0 description to-R3`

**Step-by-Step Procedure**      The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multi-level IS-IS:

1.      Configure the network interfaces.

Enable IS-IS on the interfaces by including the ISO address family on each interface.

```
[edit interfaces]
user@R5# set fe-1/2/0 unit 0 description to-R3
user@R5# set fe-1/2/0 unit 0 family inet address 10.0.0.22/30
user@R5# set fe-1/2/0 unit 0 family iso
user@R5# set fe-1/2/1 unit 0 description to-R4
user@R5# set fe-1/2/1 unit 0 family inet address 10.0.0.26/30
user@R5# set fe-1/2/1 unit 0 family iso
user@R5# set fe-1/2/2 unit 0 description to-R6
user@R5# set fe-1/2/2 unit 0 family inet address 10.0.0.29/30
user@R5# set fe-1/2/2 unit 0 family iso
user@R5# set fe-1/2/3 unit 0 description to-R7
user@R5# set fe-1/2/3 unit 0 family inet address 10.0.0.38/30
user@R5# set fe-1/2/3 unit 0 family iso
```

2.      Configure two loopback interface addresses.

One address is for IPv4.

The other is for the IS-IS area 49.002 so that Device R5 can form adjacencies with the other Level 1 devices in area 49.002. Even though Device R5's NET identifies itself as belonging to the Level 1 area 49.002, its loopback interface is not configured as a Level 1 interface. Doing so would cause the route to Device R5's loopback to be injected into the Level 1 area.

```
[edit interfaces lo0 unit 0]
user@R5# set family inet address 192.168.0.5/32
user@R5# set family iso address 49.002.0192.0168.0005.00
```

3.      Specify the IS-IS level on a per-interface basis.

Device R5 becomes adjacent to the other routing devices on the same level on each link.

By default, IS-IS is enabled for IS-IS areas on all interfaces on which the ISO protocol family is enabled (at the `[edit interfaces interface-name unit logical-unit-number]`

hierarchy level). To disable IS-IS at any particular level on an interface, include the **disable** statement.

Device R5's loopback interface is configured to run Level 2 only. If Level 1 operation were enabled on lo0.0, Device R5 would include its loopback address in its Level 1 link-state PDU, which is incorrect for this example in which the loopback addresses of Level 2 devices must not appear in a Level 1 area.

Unlike OSPF, you must explicitly list the router's lo0 interface at the **[edit protocols isis]** hierarchy level, because this interface is the source of the router's NET, and therefore must be configured as an IS-IS interface. In IS-IS, the lo0 interface operates in the passive mode by default, which is ideal because adjacency formation can never occur on a virtual interface.

```
[edit protocols isis]
user@R5# set interface fe-1/2/0.0 level 1 disable
user@R5# set interface fe-1/2/1.0 level 1 disable
user@R5# set interface fe-1/2/0.0 level 2 disable
user@R5# set interface fe-1/2/3.0 level 2 disable
user@R5# set interface lo0.0 level 1 disable
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R5# show interfaces
fe-1/2/0 {
  unit 0 {
    description to-R3;
    family inet {
      address 10.0.0.22/30;
    }
    family iso;
  }
}
fe-1/2/1 {
  unit 0 {
    description to-R4;
    family inet {
      address 10.0.0.26/30;
    }
    family iso;
  }
}
fe-1/2/2 {
  unit 0 {
    description to-R6;
    family inet {
      address 10.0.0.29/30;
    }
    family iso;
  }
}
fe-1/2/3 {
  unit 0 {
```

```
description to-R7;
family inet {
    address 10.0.0.38/30;
}
family iso;
}
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.5/32;
        }
        family iso {
            address 49.002.0192.0168.0005.00;
        }
    }
}
```

user@R5# show protocols

```
isis {
    interface fe-1/2/0.0 {
        level 1 disable;
    }
    interface fe-1/2/1.0 {
        level 1 disable;
    }
    interface fe-1/2/0.0 {
        level 2 disable;
    }
    interface fe-1/2/3.0 {
        level 2 disable;
    }
    interface lo0.0 {
        level 1 disable;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

---

## Verification

Confirm that the configuration is working properly.

- [Checking Interface-to-Area Associations on page 3834](#)
- [Verifying IS-IS Adjacencies on page 3835](#)
- [Examining the IS-IS Database on page 3836](#)

### *Checking Interface-to-Area Associations*

**Purpose** Make sure that the interface-to-area associations are configured as expected.

**Action** From operational mode, enter the **show isis interface** command.

```
user@R5> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
```

|            |   |              |          |       |
|------------|---|--------------|----------|-------|
| lo0.0      | 3 | 0x1 Disabled | Passive  | 0/0   |
| fe-1/2/0.0 | 2 | 0x3 Disabled | R5.03    | 10/10 |
| fe-1/2/1.0 | 2 | 0x2 Disabled | R5.02    | 10/10 |
| fe-1/2/0.0 | 1 | 0x1 R6.02    | Disabled | 10/10 |
| fe-1/2/3.0 | 1 | 0x4 R5.04    | Disabled | 10/10 |

**Meaning** The output shows that Device R5's interfaces have been correctly configured with the ISO family, and that the interfaces have been placed into the correct levels.

You can also see that Device R5 has elected itself as the designated intermediate system (DIS) on its broadcast-capable IS-IS interfaces.

### *Verifying IS-IS Adjacencies*

**Purpose** Verify that the expected adjacencies have formed between Device R5 and its IS-IS neighbors.

**Action** From operational mode, enter the **show isis adjacency detail** command.

```
user@R5> show isis adjacency detail
```

R3

```
Interface: fe-1/2/0.0, Level: 2, State: Up, Expires in 25 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:31 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.03, IP addresses: 10.0.0.21
```

R4

```
Interface: fe-1/2/1.0, Level: 2, State: Up, Expires in 24 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:36 ago
Circuit type: 2, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.02, IP addresses: 10.0.0.25
```

R6

```
Interface: fe-1/2/0.0, Level: 1, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:20:24 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bd
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R6.02, IP addresses: 10.0.0.30
```

R7

```
Interface: fe-1/2/3.0, Level: 1, State: Up, Expires in 21 secs
Priority: 64, Up/Down transitions: 1, Last transition: 03:19:29 ago
Circuit type: 1, Speaks: IP, IPv6, MAC address: 0:5:85:8f:c8:bc
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R5.04, IP addresses: 10.0.0.37
```

**Meaning** These results confirm that Device R5 has two Level 2 adjacencies and two Level 1 adjacencies.

**Examining the IS-IS Database**

**Purpose** Because Device R5 is a L1/L2 attached router, examine the Level 1 link-state database associated with area 49.002 to confirm that loopback addresses from backbone routers are not being advertised into the Level 1 area.

**Action** From operational mode, enter the **show isis database detail** command.

```
user@R5> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```
R5.00-00 Sequence: 0x19, Checksum: 0x7488, Lifetime: 727 secs
  IS neighbor: R5.04                      Metric:      10
  IS neighbor: R6.02                      Metric:      10
  IP prefix: 10.0.0.28/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.36/30                  Metric:      10 Internal Up
```

```
R5.04-00 Sequence: 0x14, Checksum: 0x2668, Lifetime: 821 secs
  IS neighbor: R5.00                      Metric:       0
  IS neighbor: R7.00                      Metric:       0
```

```
R6.00-00 Sequence: 0x17, Checksum: 0xa65, Lifetime: 774 secs
  IS neighbor: R6.02                      Metric:      10
  IS neighbor: R7.02                      Metric:      10
  IP prefix: 10.0.0.28/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.32/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.6/32                Metric:       0 Internal Up
```

```
R6.02-00 Sequence: 0x13, Checksum: 0xd1c0, Lifetime: 908 secs
  IS neighbor: R5.00                      Metric:       0
  IS neighbor: R6.00                      Metric:       0
```

```
R7.00-00 Sequence: 0x17, Checksum: 0xe39, Lifetime: 775 secs
  IS neighbor: R5.04                      Metric:      10
  IS neighbor: R7.02                      Metric:      10
  IP prefix: 10.0.0.32/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.36/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.7/32                Metric:       0 Internal Up
```

```
R7.02-00 Sequence: 0x13, Checksum: 0x404d, Lifetime: 966 secs
  IS neighbor: R6.00                      Metric:       0
  IS neighbor: R7.00                      Metric:       0
```

```
IS-IS level 2 link-state database:
```

```
R3.00-00 Sequence: 0x17, Checksum: 0x5f84, Lifetime: 1085 secs
  IS neighbor: R4.02                      Metric:      10
  IS neighbor: R5.03                      Metric:      10
  IP prefix: 10.0.0.16/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.20/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.40/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.3/32                Metric:       0 Internal Up
```

```
R4.00-00 Sequence: 0x17, Checksum: 0xab3a, Lifetime: 949 secs
  IS neighbor: R4.02                      Metric:      10
  IS neighbor: R5.02                      Metric:      10
  IP prefix: 10.0.0.16/30                  Metric:      10 Internal Up
  IP prefix: 10.0.0.24/30                  Metric:      10 Internal Up
  IP prefix: 192.168.0.4/32                Metric:       0 Internal Up
```

```

R4.02-00 Sequence: 0x14, Checksum: 0xf2a8, Lifetime: 1022 secs
  IS neighbor: R3.00                      Metric:      0
  IS neighbor: R4.00                      Metric:      0

R5.00-00 Sequence: 0x1f, Checksum: 0x20d7, Lifetime: 821 secs
  IS neighbor: R5.02                      Metric:     10
  IS neighbor: R5.03                      Metric:     10
  IP prefix: 10.0.0.20/30                 Metric:     10 Internal Up
  IP prefix: 10.0.0.24/30                 Metric:     10 Internal Up
  IP prefix: 10.0.0.28/30                 Metric:     10 Internal Up
  IP prefix: 10.0.0.32/30                 Metric:     20 Internal Up
  IP prefix: 10.0.0.36/30                 Metric:     10 Internal Up
  IP prefix: 192.168.0.5/32               Metric:      0 Internal Up
  IP prefix: 192.168.0.6/32               Metric:     10 Internal Up
  IP prefix: 192.168.0.7/32               Metric:     10 Internal Up

R5.02-00 Sequence: 0x14, Checksum: 0x6135, Lifetime: 977 secs
  IS neighbor: R4.00                      Metric:      0
  IS neighbor: R5.00                      Metric:      0

R5.03-00 Sequence: 0x14, Checksum: 0x1483, Lifetime: 1091 secs
  IS neighbor: R3.00                      Metric:      0
  IS neighbor: R5.00                      Metric:      0

```

**Meaning** This display indicates that Device R5's loopback interface is correctly configured to run Level 2 only. Had Level 1 operation been enabled on lo0.0, Device R5 would have then included its loopback address in its Level 1 link-state PDU.

You can also see that Device R5 has Level 2 link-state PDUs, received from its adjacent neighbors.

Like an OSPF totally stubby area, no backbone (Level 2) or external prefixes are leaked into a Level 1 area, by default. Level 1 prefixes are leaked up into the IS-IS backbone, however, as can be seen in Device R5's Level 2 link-state PDU.

**Related Documentation**

- [Understanding IS-IS Areas](#)

## Example: Configuring Hitless Authentication Key Rollover for IS-IS

This example shows how to configure hitless authentication key rollover for IS-IS.

- [Requirements on page 3837](#)
- [Overview on page 3838](#)
- [Configuration on page 3839](#)
- [Verification on page 3842](#)

### Requirements

No special configuration beyond device initialization is required before configuring hitless authentication key rollover for IS-IS.

## Overview

Authentication guarantees that only trusted routers participate in routing updates. This keychain authentication method is referred to as hitless because the keys roll over from one to the next without resetting any peering sessions or interrupting the routing protocol. Junos OS supports both RFC 5304, *IS-IS Cryptographic Authentication* and RFC 5310, *IS-IS Generic Cryptographic Authentication*.

This example includes the following statements for configuring the keychain:

- **algorithm**—For each key in the keychain, you can specify an encryption algorithm. The algorithm can be SHA-1 or MD-5.
- **key**—A keychain can have multiple keys. Each key within a keychain must be identified by a unique integer value. The range of valid identifier values is from 0 through 63.
- **key-chain**—For each keychain, you must specify a name. This example defines two keychains: **base-key-global** and **base-key-inter**.
- **options**—For each key in the keychain, you can specify the encoding for the message authentication code: **isis-enhanced** or **basic**. The basic (RFC 5304) operation is enabled by default.

When you configure the **isis-enhanced** option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.

When you configure **basic** (or do not include the **options** statement in the key configuration), Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.

Because this setting is for IS-IS only, the TCP and the BFD protocols ignore the encoding option configured in the key.

- **secret**—For each key in the keychain, you must set a secret password. This password can be entered in either encrypted or plain text format in the **secret** statement. It is always displayed in encrypted format.
- **start-time**—Each key must specify a start time in UTC format. Control gets passed from one key to the next. When a configured start time arrives (based on the routing device's clock), the key with that start time becomes active. Start times are specified in the local time zone for a routing device and must be unique within the key chain.

You can apply a keychain globally to all interfaces or more granularly to specific interfaces.

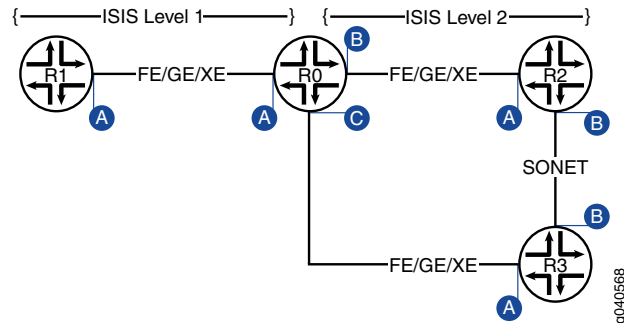
This example includes the following statements for applying the keychain to all interfaces or to particular interfaces:

- **authentication-key-chain**—Enables you to apply a keychain at the global IS-IS level for all Level 1 or all Level 2 interfaces.
- **hello-authentication-key-chain**—Enables you to apply a keychain at the individual IS-IS interface level. The interface configuration overrides the global configuration.



Figure 102 on page 3839 shows the topology used in the example.

Figure 102: Hitless Authentication Key Rollover for IS-IS



This example shows the configuration for Router R0.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces ge-0/0/0 unit 0 description "interface A"
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address fe80::200:f8ff:fe21:67cf/128
set interfaces ge-0/0/1 unit 0 description "interface B"
set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6 address 10FB::C:ABC:1FOC:44DA/128
set interfaces ge-0/0/2 unit 0 description "interface C"
set interfaces ge-0/0/2 unit 0 family inet address 10.0.0.9/30
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
set security authentication-key-chains key-chain base-key-global key 63 secret
"$9$jfkqfTQnCpBDiCt"
set security authentication-key-chains key-chain base-key-global key 63 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-global key 63 algorithm
hmac-sha-1
set security authentication-key-chains key-chain base-key-global key 63 options
isis-enhanced
set security authentication-key-chains key-chain base-key-inter key 0 secret
"$9$8sgx7Vws4ZDkWLGD"
set security authentication-key-chains key-chain base-key-inter key 0 start-time
"2011-8-6.06:54:00-0700"
set security authentication-key-chains key-chain base-key-inter key 0 algorithm md5
set security authentication-key-chains key-chain base-key-inter key 0 options basic
set protocols isis level 1 authentication-key-chain base-key-global
set protocols isis interface ge-0/0/0.0 level 1 hello-authentication-key-chain
base-key-inter
```

### Step-by-Step Procedure

To configure hitless authentication key rollover for IS-IS:

1. Configure the Router R0 interfaces.

```
[edit interfaces ge-0/0/0 unit 0]
user@R0# set description "interface A"
user@R0# set family inet address 10.0.0.1/30
user@R0# set family iso
user@R0# set family inet6 address fe80::200:f8ff:fe21:67cf/128
[edit interfaces ge-0/0/1 unit 0]
user@R0# set interfaces ge-0/0/1 unit 0 description "interface B"
user@R0# set interfaces ge-0/0/1 unit 0 family inet address 10.0.0.5/30
user@R0# set interfaces ge-0/0/1 unit 0 family iso
user@R0# set interfaces ge-0/0/1 unit 0 family inet6 address
10fb::c:abc:1f0c:44da/128
[edit interfaces ge-0/0/2 unit 0]
user@R0# set description "interface C"
user@R0# set family inet address 10.0.0.9/30
user@R0# set interfaces ge-0/0/2 unit 0 family iso
user@R0# set interfaces ge-0/0/2 unit 0 family inet6 address ff06::c3/128
```

2. Configure one or more authentication keys.

```
[edit security authentication-key-chains key-chain base-key-global]
user@R0# set key 63 secret "$9$jkqfTQnCpBDiCt"
user@R0# set key 63 start-time "2011-8-6.06:54:00-0700"
user@R0# set key 63 algorithm hmac-sha-1
user@R0# set key 63 options isis-enhanced
[edit security authentication-key-chains key-chain base-key-inter]
user@R0# set key 0 secret "$9$8sgx7Vws4ZDkWLGD"
user@R0# set key 0 start-time "2011-8-6.06:54:00-0700"
user@R0# set key 0 algorithm md5
user@R0# set key 0 options basic
```

3. Apply the base-key-global keychain to all Level 1 IS-IS interfaces on Router R0.

```
[edit protocols isis level 1]
user@R0# set authentication-key-chain base-key-global
```

4. Apply the base-key-inter keychain to the ge-0/0/0.0 interface on Router R0.

```
[edit protocols isis interface ge-0/0/0.0 level 1]
user@R0# set hello-authentication-key-chain base-key-inter
```

5. If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    description "interface A";
```

```

        family inet {
            address 10.0.0.1/30;
        }
        family iso;
        family inet6 {
            address fe80::200:f8ff:fe21:67cf/128;
        }
    }
}
ge-0/0/1 {
    unit 0 {
        description "interface B";
        family inet {
            address 10.0.0.5/30;
        }
        family iso;
        family inet6 {
            address 10FB::C:ABC:1F0C:44DA/128;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        description "interface C";
        family inet {
            address 10.0.0.9/30;
        }
        family iso;
        family inet6 {
            address ff06::c3/128;
        }
    }
}
}

user@R0# show protocols
isis {
    level 1 authentication-key-chain base-key-global;
    interface ge-0/0/0.0 {
        level 1 hello-authentication-key-chain base-key-inter;
    }
}

user@R0# show security
authentication-key-chains {
    key-chain base-key-global {
        key 63 {
            secret "$9$jfkqfTQnCpBDiCt"; ## SECRET-DATA
            start-time "2011-8-6.06:54:00-0700";
            algorithm hmac-sha-1;
            options isis-enhanced;
        }
    }
    key-chain base-key-inter {
        key 0 {
            secret "$9$8sgx7Vws4ZDkWLGD"; ## SECRET-DATA
            start-time "2011-8-6.06:54:00-0700";
            algorithm md5;
        }
    }
}

```

```
        options basic;  
    }  
}  
}
```

---

### Verification

To verify the configuration, run the following commands:

- [show isis authentication](#)
- [show security keychain](#)

### Related Documentation

- [Understanding Hitless Authentication Key Rollover for IS-IS on page 3818](#)

## Example: Redistributing OSPF Routes into IS-IS

This example shows how to redistribute OSPF routes into an IS-IS network.

- [Requirements on page 3842](#)
- [Overview on page 3842](#)
- [Configuration on page 3843](#)
- [Verification on page 3848](#)

---

### Requirements

No special configuration beyond device initialization is required before configuring this example.

---

### Overview

Export policy can be applied to IS-IS to facilitate route redistribution.


Junos OS does not support the application of import policy for link-state routing protocols like IS-IS because such policies can lead to inconsistent link-state database (LSDB) entries, which in turn can result in routing inconsistencies.

In this example, OSPF routes 192.168.0/24 through 192.168.3/24 are redistributed into IS-IS area 49.0002 from Device R2.

In addition, policies are configured to ensure that Device R1 can reach destinations on the 10.0.0.44/30 network, and that Device R3 can reach destinations on the 10.0.0.36/30 network. This enables end-to-end reachability.

[Figure 103 on page 3843](#) shows the topology used in this example.

Figure 103: IS-IS Route Redistribution Topology



“CLI Quick Configuration” on page 3843 shows the configuration for all of the devices in Figure 103 on page 3843. The section “Step-by-Step Procedure” on page 3844 describes the steps on Device R2. “Step-by-Step Procedure” on page 3845 describes the steps on Device R3.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```
set interfaces fe-1/2/0 unit 0 description to-R7
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.38/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 172.16.3.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0305.00
set protocols isis interface fe-1/2/0.38
set protocols isis interface lo0.0
```

**Device R2**

```
set interfaces fe-1/2/1 unit 0 description to-R5
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.37/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/0 unit 0 description to-OSPF-network
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.45/30
set interfaces lo0 unit 0 family inet address 172.16.9.7/32
set interfaces lo0 unit 0 family iso address 49.0002.0172.0016.0907.00
set protocols isis export ospf-isis
set protocols isis export send-direct-to-isis-neighbors
set protocols isis interface fe-1/2/1.0
set protocols isis interface lo0.0
```

```
set protocols ospf export send-direct-to-ospf-neighbors
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf-isis term 1 from protocol ospf
set policy-options policy-statement ospf-isis term 1 from route-filter 192.168.0.0/22
  longer
set policy-options policy-statement ospf-isis term 1 then accept
set policy-options policy-statement send-direct-to-isis-neighbors from protocol direct
set policy-options policy-statement send-direct-to-isis-neighbors from route-filter
  10.0.0.44/30 exact
set policy-options policy-statement send-direct-to-isis-neighbors then accept
set policy-options policy-statement send-direct-to-ospf-neighbors from protocol direct
set policy-options policy-statement send-direct-to-ospf-neighbors from route-filter
  10.0.0.36/30 exact
set policy-options policy-statement send-direct-to-ospf-neighbors then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.46/30
set interfaces lo0 unit 0 family inet address 192.168.1.1/32
set interfaces lo0 unit 0 family inet address 192.168.2.1/32
set interfaces lo0 unit 0 family inet address 192.168.3.1/32
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols ospf export ospf
set protocols ospf area 0.0.0.1 interface fe-1/2/0.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options policy-statement ospf term 1 from protocol static
set policy-options policy-statement ospf term 1 then accept
set routing-options static route 192.168.0.0/24 discard
set routing-options static route 192.168.1.0/24 discard
set routing-options static route 192.168.3.0/24 discard
set routing-options static route 192.168.2.0/24 discard
```

#### Step-by-Step Procedure

To configure Device R2:

1. Configure the network interfaces.  

```
[edit interfaces]
user@R2# set fe-1/2/1 unit 0 description to-R5
user@R2# set fe-1/2/1 unit 0 family inet address 10.0.0.37/30
user@R2# set fe-1/2/1 unit 0 family iso
user@R2# set fe-1/2/0 unit 0 description to-OSPF-network
user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.45/30
user@R2# set lo0 unit 0 family inet address 172.16.9.7/32
user@R2# set lo0 unit 0 family iso address 49.0002.0172.0016.0907.00
```
2. Configure IS-IS on the interface facing Device R1 and the loopback interface.  

```
[edit protocols isis]
user@R2# set interface fe-1/2/1.0
user@R2# set interface lo0.0
```
3. Configure the policy that enables Device R1 to reach the 10.0.0.44/30 network.  

```
[edit policy-options policy-statement send-direct-to-isis-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.44/30 exact
user@R2# set then accept
```
4. Apply the policy that enables Device R1 to reach the 10.0.0.44/30 network.

```
[edit protocols isis]
user@R2# set export send-direct-to-isis-neighbors
```

5. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@R2# set area 0.0.0.1 interface fe-1/2/0.0
user@R2# set area 0.0.0.1 interface lo0.0 passive
```

6. Configure the OSPF route redistribution policy.

```
[edit policy-options policy-statement ospf-isis term 1]
user@R2# set from protocol ospf
user@R2# set from route-filter 192.168.0.0/22 longer
user@R2# set then accept
```

7. Apply the OSPF route redistribution policy to the IS-IS instance.

```
[edit protocols isis]
user@R2# set export ospf-isis
```

8. Configure the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit policy-options policy-statement send-direct-to-ospf-neighbors]
user@R2# set from protocol direct
user@R2# set from route-filter 10.0.0.36/30 exact
user@R2# set then accept
```

9. Apply the policy that enables Device R3 to reach the 10.0.0.36/30 network.

```
[edit protocols ospf]
user@R2# set export send-direct-to-ospf-neighbors
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multi-level IS-IS:

1. Configure the network interfaces.

Multiple addresses are configured on the loopback interface to simulate multiple route destinations.

```
[edit interfaces]
user@R3# set fe-1/2/0 unit 0 family inet address 10.0.0.46/30
user@R3# set lo0 unit 0 family inet address 192.168.1.1/32
user@R3# set lo0 unit 0 family inet address 192.168.2.1/32
user@R3# set lo0 unit 0 family inet address 192.168.3.1/32
user@R3# set lo0 unit 0 family inet address 192.168.0.1/32
```

2. Configure static routes to the loopback interface addresses.

These are the routes that are redistributed into IS-IS.

```
[edit routing-options static]
user@R3# set route 192.168.0.0/24 discard
user@R3# set route 192.168.1.0/24 discard
user@R3# set route 192.168.3.0/24 discard
user@R3# set route 192.168.2.0/24 discard
```

3. Configure OSPF on the interfaces.  

```
[edit protocols ospf area 0.0.0.1]
user@R3# set interface fe-1/2/0.0
user@R3# set interface lo0.0 passive
```
4. Configure the OSPF policy to export the static routes.  

```
[edit policy-options policy-statement ospf term 1]
user@R3# set from protocol static
user@R3# set then accept
```
5. Apply the OSPF export policy.  

```
[edit protocols ospf]
user@R3# set export ospf
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

**Device R2**

```
user@R2# show interfaces
fe-1/2/1 {
  unit 0 {
    description to-R5;
    family inet {
      address 10.0.0.37/30;
    }
    family iso;
  }
}
fe-1/2/0 {
  unit 0 {
    description to-OSPF-network;
    family inet {
      address 10.0.0.45/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 172.16.9.7/32;
    }
    family iso {
      address 49.0002.0172.0016.0907.00;
    }
  }
}

user@R2# show protocols
isis {
  export [ ospf-isis send-direct-to-isis-neighbors ];
  interface fe-1/2/1.0;
  interface lo0.0;
}
```



```

ospf {
  export send-direct-to-ospf-neighbors;
  area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
      passive;
    }
  }
}

user@R2# show policy-options
policy-statement ospf-isis {
  term 1 {
    from {
      protocol ospf;
      route-filter 192.168.0.0/22 longer;
    }
    then accept;
  }
}
policy-statement send-direct-to-isis-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.44/30 exact;
  }
  then accept;
}
policy-statement send-direct-to-ospf-neighbors {
  from {
    protocol direct;
    route-filter 10.0.0.36/30 exact;
  }
  then accept;
}

```

**Device R3**

```

user@R3# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.46/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.1.1/32;
      address 192.168.2.1/32;
      address 192.168.3.1/32;
      address 192.168.0.1/32;
    }
  }
}

user@R3# show protocols
ospf {

```

```
export ospf;
area 0.0.0.1 {
    interface fe-1/2/0.0;
    interface lo0.0 {
        passive;
    }
}

user@R3# show policy-options
policy-statement ospf {
    term 1 {
        from protocol static;
        then accept;
    }
}

user@R3# show routing-options
static {
    route 192.168.0.0/24 discard;
    route 192.168.1.0/24 discard;
    route 192.168.3.0/24 discard;
    route 192.168.2.0/24 discard;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

---

## Verification

Confirm that the configuration is working properly.

- [Verifying OSPF Route Advertisement on page 3848](#)
- [Verifying Route Redistribution on page 3849](#)
- [Verifying Connectivity on page 3849](#)

### *Verifying OSPF Route Advertisement*

**Purpose** Make sure that the expected routes are advertised by OSPF.

**Action** From operational mode on Device R2, enter the **show route protocol ospf** command.

```
user@R2> show route protocol ospf
```

```
inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.168.0.0/24      *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.0.1/32     *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.1.0/24     *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.1.1/32     *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.2.0/24     *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
```

```

192.168.2.1/32      *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.3.0/24     *[OSPF/150] 03:54:21, metric 0, tag 0
                   > to 10.0.0.46 via fe-1/2/0.0
192.168.3.1/32     *[OSPF/10] 03:54:21, metric 1
                   > to 10.0.0.46 via fe-1/2/0.0
224.0.0.5/32       *[OSPF/10] 03:56:03, metric 1
                   MultiRecv

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

**Meaning** The 192.168/16 routes are advertised by OSPF.

### *Verifying Route Redistribution*

**Purpose** Make sure that the expected routes are redistributed from OSPF into IS-IS.

**Action** From operational mode on Device R1, enter the **show route protocol isis** command.

```
user@R1> show route protocol isis
```

```

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.44/30       *[IS-IS/160] 03:45:24, metric 20
                   > to 10.0.0.37 via fe-1/2/0.0
172.16.9.7/32      *[IS-IS/15] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.0.0/24     *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.0.1/32     *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.1.0/24     *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.1.1/32     *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.2.0/24     *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.2.1/32     *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.3.0/24     *[IS-IS/160] 03:49:46, metric 10
                   > to 10.0.0.37 via fe-1/2/0.0
192.168.3.1/32     *[IS-IS/160] 03:49:46, metric 11, tag2 1
                   > to 10.0.0.37 via fe-1/2/0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

**Meaning** The 192.168/16 routes are redistributed into IS-IS.

### *Verifying Connectivity*

**Purpose** Check that Device R1 can reach the destinations on Device R3.

**Action** From operational mode, enter the **ping** command.

```
user@R1> ping 192.168.1.1
```

```

PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=63 time=2.089 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=63 time=1.270 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=63 time=2.135 ms

```

**Meaning** These results confirm that Device R1 can reach the destinations in the OSPF network.

**Related Documentation**

- [Understanding Routing Policies](#)

## Example: Configuring BFD for IS-IS

This example describes how to configure the Bidirectional Forwarding Detection (BFD) protocol to detect failures in an IS-IS network.

- [Requirements on page 3850](#)
- [Overview on page 3850](#)
- [Configuration on page 3850](#)
- [Verification on page 3853](#)

### Requirements

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 3823](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

### Overview

This example shows two routers connected to each other. A loopback interface is configured on each router. IS-IS and BFD protocols are configured on both routers.

[Figure 104 on page 3850](#) shows the sample network.

**Figure 104: Configuring BFD for IS-IS**



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Router R1**

```

set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 5
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 2
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 3
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic

```

**Router R2**

```

set protocols isis interface so-0/0/0 bfd-liveness-detection detection-time threshold 6
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-interval 3
set protocols isis interface so-0/0/0 bfd-liveness-detection minimum-receive-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection no-adaptation
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval threshold 4
set protocols isis interface so-0/0/0 bfd-liveness-detection transmit-interval
  minimum-interval 1
set protocols isis interface so-0/0/0 bfd-liveness-detection multiplier 2
set protocols isis interface so-0/0/0 bfd-liveness-detection version automatic

```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.



**NOTE:** To simply configure BFD for IS-IS, only the `minimum-interval` statement is required. The BFD protocol selects default parameters for all the other configuration statements when you use the `bfd-liveness-detection` statement without specifying any parameters.



**NOTE:** You can change parameters at any time without stopping or restarting the existing session. BFD automatically adjusts to the new parameter value. However, no changes to BFD parameters take place until the values resynchronize with each BFD peer.

To configure BFD for IS-IS on Routers R1 and R2:

1. Enable BFD failure detection for IS-IS.
 

```

[edit protocols isis]
user@R1# set interface so-0/0/0 bfd-liveness-detection

[edit protocols isis]
user@R2# set interface so-0/0/0 bfd-liveness-detection

```
2. Configure the threshold for the adaptation of the detection time, which must be greater than the multiplier number multiplied by the minimum interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set detection-time threshold 5
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set detection-time threshold 6
```

3. Configure the minimum transmit and receive intervals for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-interval 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-interval 3
```

4. Configure only the minimum receive interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set minimum-receive-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set minimum-receive-interval 1
```

5. Disable BFD adaptation.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set no-adaptation
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set no-adaptation
```

6. Configure the threshold for the transmit interval, which must be greater than the minimum transmit interval.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval threshold 3
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval threshold 4
```

7. Configure the minimum transmit interval for failure detection.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set transmit-interval minimum-interval 1
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set transmit-interval minimum-interval 1
```

8. Configure the multiplier number, which is the number of hello packets not received by the neighbor that causes the originating interface to be declared down.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set multiplier 2
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set multiplier 2
```

9. Configure the BFD version used for detection.

The default is to have the version detected automatically.

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R1# set version automatic
```

```
[edit protocols isis interface so-0/0/0 bfd-liveness-detection]
user@R2# set version automatic
```

## Results

From configuration mode, confirm your configuration by issuing the **show protocols isis interface** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

user@R1# **show protocols isis interface so-0/0/0**

```

bfd-liveness-detection {
  version automatic;
  minimum-interval 2;
  minimum-receive-interval 1;
  multiplier 2;
  no-adaptation;
  transmit-interval {
    minimum-interval 1;
    threshold 3;
  }
  detection-time {
    threshold 5;
  }
}
...

```

user@R2# **show protocols isis interface so-0/0/0**

```

bfd-liveness-detection {
  version automatic;
  minimum-interval 3;
  minimum-receive-interval 1;
  multiplier 2;
  no-adaptation;
  transmit-interval {
    minimum-interval 1;
    threshold 4;
  }
  detection-time {
    threshold 6;
  }
}
...

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1 and R2 on page 3853](#)
- [Verifying That IS-IS Is Configured on page 3854](#)
- [Verifying That BFD Is configured on page 3855](#)

### *Verifying the Connection Between Routers R1 and R2*

**Purpose** Make sure that Routers R1 and R2 are connected to each other.

**Action** Ping the other router to check the connectivity between the two routers as per the network topology.

```
user@R1> ping 10.0.0.2
```

```
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=64 time=1.367 ms
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=1.291 ms
^C
--- 10.0.0.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.291/1.440/1.662/0.160 ms
```

```
user@R2> ping 10.0.0.1
```

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=1.287 ms
64 bytes from 10.0.0.1: icmp_seq=1 ttl=64 time=1.310 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=1.289 ms
^C
--- 10.0.0.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.287/1.295/1.310/0.010 ms
```

**Meaning** Routers R1 and R2 are connected to each other.

### *Verifying That IS-IS Is Configured*

**Purpose** Make sure that the IS-IS instance is running on both routers.

**Action** Use the **show isis database** statement to check if the IS-IS instance is running on both routers, R1 and R2.

```
user@R1> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4a571  0x30c5    1195 L1 L2
R2.00-00    0x4a586  0x4b7e    1195 L1 L2
R2.02-00    0x330ca1 0x3492    1196 L1 L2
  3 LSPs
```

```
IS-IS level 2 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4a856  0x5db0    1194 L1 L2
R2.00-00    0x4a89d  0x149b    1194 L1 L2
R2.02-00    0x1fb2ff 0xd302    1194 L1 L2
  3 LSPs
```

```
user@R2> show isis database
```

```
IS-IS level 1 link-state database:
LSP ID      Sequence Checksum Lifetime Attributes
R1.00-00    0x4b707  0xcc80    1195 L1 L2
R2.00-00    0x4b71b  0xeb37    1198 L1 L2
R2.02-00    0x33c2ce 0xb52d    1198 L1 L2
  3 LSPs
```



IS-IS level 2 link-state database:

| LSP ID   | Sequence | Checksum | Lifetime | Attributes |
|----------|----------|----------|----------|------------|
| R1.00-00 | 0x4b9f2  | 0xee70   | 1192     | L1 L2      |
| R2.00-00 | 0x4ba41  | 0x9862   | 1197     | L1 L2      |
| R2.02-00 | 0x3      | 0x6242   | 1198     | L1 L2      |

3 LSPs

**Meaning** IS-IS is configured on both routers, R1 and R2.

### *Verifying That BFD Is configured*

**Purpose** Make sure that the BFD instance is running on both routers, R1 and R2.

**Action** Use the **show bfd session detail** statement to check if BFD instance is running on the routers.

user@R1> show bfd session detail

| Address  | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------|-------|-----------|-------------|-------------------|------------|
| 10.0.0.2 | Up    | so-0/0/0  | 2.000       | 1.000             | 2          |

Client ISIS R2, TX interval 0.001, RX interval 0.001  
 Client ISIS R1, TX interval 0.001, RX interval 0.001  
 Session down time 00:00:00, previous up time 00:00:15  
 Local diagnostic NbrSignal, remote diagnostic NbrSignal  
 Remote state AdminDown, version 1  
 Router 3, routing table index 17

1 sessions, 2 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

user@R2> show bfd session detail

| Address  | State | Interface | Detect Time | Transmit Interval | Multiplier |
|----------|-------|-----------|-------------|-------------------|------------|
| 10.0.0.1 | Up    | so-0/0/0  | 2.000       | 1.000             | 2          |

Client ISIS R2, TX interval 0.001, RX interval 0.001  
 Session down time 00:00:00, previous up time 00:00:05  
 Local diagnostic NbrSignal, remote diagnostic NbrSignal  
 Remote state AdminDown, version 1  
 Router 2, routing table index 15

1 sessions, 1 clients

Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

**Meaning** BFD is configured on Routers R1 and R2 for detecting failures in the IS-IS network.

**Related Documentation**

- [Understanding BFD for IS-IS](#)

## Example: Configuring BFD Authentication for IS-IS

This example shows how to configure BFD authentication for IS-IS.

- [Requirements on page 3856](#)
- [Overview on page 3856](#)

- [Configuration on page 3856](#)
- [Verification on page 3858](#)

## Requirements


Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 3823](#) for information about the required IS-IS configuration.

## Overview

In this example, a BFD authentication keychain is configured with meticulous keyed MD5 authentication.

[Figure 105 on page 3856](#) shows the topology used in this example.

**Figure 105: IS-IS BFD Authentication Topology**



[“CLI Quick Configuration” on page 3856](#) shows the configuration for both of the devices in [Figure 105 on page 3856](#). The section [“Step-by-Step Procedure” on page 3857](#) describes the steps on Device R1.

## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

|                  |  |
|------------------|--|
| <b>Device R1</b> | <pre> set security authentication-key-chains key-chain secret123 description for-isis-bfd set security authentication-key-chains key-chain secret123 key 1 secret "\$9\$cW-yrv" set security authentication-key-chains key-chain secret123 key 1 start-time   "2012-5-31.13:00:00 -0700" set security authentication-key-chains key-chain secret123 key 2 secret "\$9\$m5T3" set security authentication-key-chains key-chain secret123 key 2 start-time   "2013-5-31.13:00:00 -0700" set security authentication-key-chains key-chain secret123 key 3 secret "\$9\$mTQn" set security authentication-key-chains key-chain secret123 key 3 start-time   "2014-5-31.13:00:00 -0700" set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100 set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain   secret123 set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm   meticulous-keyed-md5 </pre> |
| <b>Device R2</b> | <pre> set security authentication-key-chains key-chain secret123 description for-isis-bfd set security authentication-key-chains key-chain secret123 key 1 secret "\$9\$cW-yrv" set security authentication-key-chains key-chain secret123 key 1 start-time   "2012-5-31.13:00:00 -0700" set security authentication-key-chains key-chain secret123 key 2 secret "\$9\$m5T3" </pre>  |

```

set security authentication-key-chains key-chain secret123 key 2 start-time
"2013-5-31.13:00:00 -0700"
set security authentication-key-chains key-chain secret123 key 3 secret "$9$mTQn"
set security authentication-key-chains key-chain secret123 key 3 start-time
"2014-5-31.13:00:00 -0700"
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection minimum-interval 100
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication key-chain
secret123
set protocols isis interface ge-1/2/0.0 bfd-liveness-detection authentication algorithm
meticulous-keyed-md5

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS BFD authentication:

1. Configure the authentication keychain.
 

```

[edit security authentication-key-chains key-chain secret123]
user@R1# set description for-isis-bfd
user@R1# set key 1 secret "$9$cW-yrv"
user@R1# set key 1 start-time "2012-5-31.13:00:00 -0700"
user@R1# set key 2 secret "$9$m5T3"
user@R1# set key 2 start-time "2013-5-31.13:00:00 -0700"
user@R1# set key 3 secret "$9$mTQn"
user@R1# set key 3 start-time "2014-5-31.13:00:00 -0700"

```
2. Enable BFD.
 

```

[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set minimum-interval 100

```
3. Apply the authentication keychain.
 

```

[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set authentication key-chain secret123

```
4. Set the authentication type.
 

```

[edit protocols isis interface ge-1/2/0.0 bfd-liveness-detection]
user@R1# set authentication algorithm meticulous-keyed-md5

```

**Results** From configuration mode, confirm your configuration by entering the **show protocols** and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@R1# show protocols
isis {
  interface ge-1/2/0.0 {
    bfd-liveness-detection {
      minimum-interval 100;
      authentication {
        key-chain secret123;
        algorithm meticulous-keyed-md5;
      }
    }
  }
}

```

```

    }
  }
  user@R1# show security
  authentication-key-chains {
    key-chain secret123 {
      description for-isis-bfd;
      key 1 {
        secret "$9$cW-yrv"; ## SECRET-DATA
        start-time "2012-5-31.13:00:00 -0700";
      }
      key 2 {
        secret "$9$m5T3"; ## SECRET-DATA
        start-time "2013-5-31.13:00:00 -0700";
      }
      key 3 {
        secret "$9$mTQn"; ## SECRET-DATA
        start-time "2014-5-31.13:00:00 -0700";
      }
    }
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

#### Verifying IS-IS BFD Authentication

**Purpose** Verify the status of IS-IS BFD authentication.

**Action** From operational mode, enter the **show bfd session extensive** command.

```

user@R1> show bfd session extensive

```

| Address  | State | Interface  | Detect Time | Transmit Interval | Multiplier |
|----------|-------|------------|-------------|-------------------|------------|
| 10.0.0.2 | Down  | ge-1/2/0.0 | 0.300       | 1.000             | 3          |

```

Client ISIS L1, TX interval 0.100, RX interval 0.100, Authenticate
  keychain secret123, algo meticulous-keyed-md5, mode strict
Client ISIS L2, TX interval 0.100, RX interval 0.100, Authenticate
  keychain secret123, algo meticulous-keyed-md5, mode strict
Session down time 00:35:13, previous up time 00:12:17
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Logical system 2, routing table index 85
Min async interval 0.100, min slow interval 1.000
Adaptive async TX interval 0.100, RX interval 0.100
Local min TX interval 1.000, minimum RX interval 0.100, multiplier 3
Remote min TX interval 0.100, min RX interval 0.100, multiplier 3
Local discriminator 2, remote discriminator 1
Echo mode disabled/inactive, no-absorb, no-refresh
Authentication enabled/active, keychain secret123, algo meticulous-keyed-md5,
mode strict
Session ID: 0x100101

1 sessions, 2 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 10.0 pps

```

**Meaning** The output shows that BFD authentication is enabled on IS-IS Level 1 and Level 2.

**Related Documentation**

- [Configuring BFD Authentication for IS-IS](#)
- [Example: Configuring BFD for IS-IS](#)

## Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

- [Understanding IS-IS IPv4 and IPv6 Unicast Topologies on page 3859](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859](#)

### Understanding IS-IS IPv4 and IPv6 Unicast Topologies

You can configure IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This enables you to exercise control over the paths that unicast data takes through a network.

A topology is the set of joined nodes. IS-IS evaluates all the paths in a single topology for each IS-IS level and uses the shortest-path-first (SPF) algorithm to determine the best path among all the feasible paths. Topology discovery and SPF calculation is performed in a protocol-neutral fashion because it is done at Layer 2 of the OSI model. If you load the topology with reachability information for a certain protocol (for example, IP), the assumption is that the circuits that are supposed to provide reachability between routing devices can carry the protocol. The SPF algorithm has a per-link orientation, not a per-address family or per-protocol orientation.

Multitopology routing enables you to override this default behavior by enabling a per-address family, per-protocol SPF calculation.

The additional CPU load associated with multiple runs of the SPF algorithm is generally not an issue with the processing power available on today's routing device control planes.

The multitopology extensions alter existing type, length, and value (TLV) tuples by adding a topology ID. Each routing device in a given topology maintains its adjacencies and runs a per-topology SPF calculation.

### Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies

This example shows how to configure IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology.

- [Requirements on page 3860](#)
- [Overview on page 3860](#)
- [Configuration on page 3861](#)
- [Verification on page 3865](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview

This example focuses on IPv4 and IPv6 unicast topologies. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This enables you to exercise control over the paths that unicast data takes through a network.

To enable an IPv6 unicast topology for IS-IS, include the **ipv6-unicast** statement:

```
isis {  
  topologies {  
    ipv6-unicast;  
  }  
}
```

To configure a metric for the IPv6 unicast topology, include the **ipv6-unicast-metric** statement:

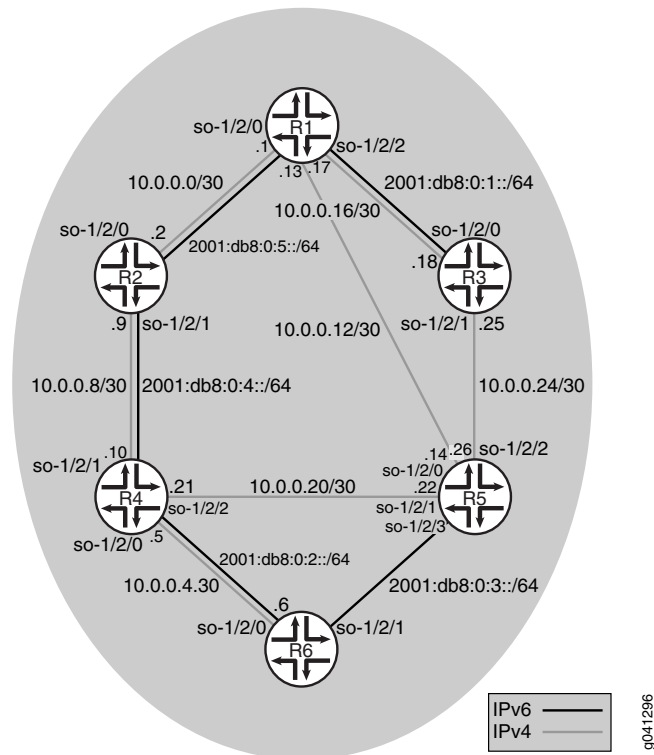
```
isis {  
  interface interface-name {  
    level level-number {  
      ipv6-unicast-metric number;  
    }  
  }  
}
```

To exclude an interface from the IPv6 unicast topologies for IS-IS, include the **no-ipv6-unicast** statement:

```
isis {  
  interface interface-name {  
    no-ipv6-unicast;  
  }  
}
```

[Figure 106 on page 3861](#) shows the topology used in this example. The black lines indicate link membership in the IPv6 topology. The gray lines indicate membership to the IPv4 topology. Using regular TLVs, it would not be possible to build multiple topologies and run an SPF calculation based on them. The multitopology extensions describe an extension to carry the set of supported protocols in the hello packet. After activating multitopology routing support on a link, the link carries all the topologies that the underlying circuit is able to relay.

Figure 106: IS-IS IPv4 and IPv6 Unicast Topologies



“CLI Quick Configuration” on page 3861 shows the configuration for all of the devices in Figure 106 on page 3861. The section “Step-by-Step Procedure” on page 3863 describes the steps on Device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device R1

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.17/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::1/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface so-1/2/2.0
set protocols isis interface lo0.0
```

**Device R2**

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::2/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.2
set protocols isis interface so-1/2/1.0
set protocols isis interface lo0.0
```

**Device R3**

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.18/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.25/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.3/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0003.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::3/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface lo0.0
```

**Device R4**

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.5/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.10/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.21/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces lo0 unit 0 family inet address 192.168.0.4/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0004.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::4/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0
set protocols isis interface so-1/2/2.0 no-ipv6-unicast
set protocols isis interface lo0.0
```

**Device R5**

```
set interfaces so-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet address 10.0.0.22/30
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/2 unit 0 family inet address 10.0.0.26/30
set interfaces so-1/2/2 unit 0 family iso
set interfaces so-1/2/3 unit 0 family iso
set interfaces so-1/2/3 unit 0 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.5/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0005.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::5/128
```



```

set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0 no-ipv6-unicast
set protocols isis interface so-1/2/1.0 no-ipv6-unicast
set protocols isis interface so-1/2/2.0 no-ipv6-unicast
set protocols isis interface so-1/2/3.0
set protocols isis interface lo0.0

```

**Device R6**

```

set interfaces so-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces so-1/2/0 unit 0 family iso
set interfaces so-1/2/0 unit 0 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces so-1/2/1 unit 0 family iso
set interfaces so-1/2/1 unit 0 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 0 family inet address 192.168.0.6/32
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0006.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::6/128
set protocols isis topologies ipv6-unicast
set protocols isis interface so-1/2/0.0
set protocols isis interface so-1/2/1.0
set protocols isis interface lo0.0

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an alternate IPv6 unicast topology:

1. Configure the interfaces.

```

[edit interfaces]
user@R1# set so-1/2/0 unit 0 family inet address 10.0.0.1/30
user@R1# set so-1/2/0 unit 0 family iso
user@R1# set so-1/2/0 unit 0 family inet6 address 2001:db8:0:5::/64 eui-64
user@R1# set so-1/2/1 unit 0 family inet address 10.0.0.13/30
user@R1# set so-1/2/1 unit 0 family iso
user@R1# set so-1/2/2 unit 0 family inet address 10.0.0.17/30
user@R1# set so-1/2/2 unit 0 family iso
user@R1# set so-1/2/2 unit 0 family inet6 address 2001:db8:0:1::/64 eui-64
user@R1# set lo0 unit 0 family inet address 192.168.0.1/32
user@R1# set lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
user@R1# set lo0 unit 0 family inet6 address 2001:db8::1/128

```

2. Enable IS-IS on the interfaces.

```

[edit protocols isis]
user@R1# set interface so-1/2/0.0
user@R1# set interface so-1/2/1.0
user@R1# set interface so-1/2/2.0
user@R1# set interface lo0.0

```

3. Enable multitopology routing on the IS-IS interfaces.

The **ipv6-unicast** statement enables multitopology IS-IS routing on all interfaces that have **family iso** and **family inet6** configured and are listed at the **[edit protocols isis interface]** hierarchy level.

```

[edit protocols isis]
user@R1# set topologies ipv6-unicast

```

4. Disable IPv6 unicast support on a given interface.

If you do not want to run multiprotocol IS-IS routing for IPv6 on a given interface, you can disable multiprotocol routing by including the **no-ipv6-unicast** statement in the IS-IS interface configuration.

```
[edit protocols isis]
user@R1# set interface so-1/2/1.0 no-ipv6-unicast
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
so-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:5::/64 {
        eui-64;
      }
    }
  }
}
so-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.13/30;
    }
    family iso;
  }
}
so-1/2/2 {
  unit 0 {
    family inet {
      address 10.0.0.17/30;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.1/32;
    }
    family iso {
      address 49.0002.0192.0168.0001.00;
```

```

    }
    family inet6 {
        address 2001:db8::1/128;
    }
}

user@R1# show protocols
isis {
    topologies ipv6-unicast;
    interface so-1/2/0.0;
    interface so-1/2/1.0 {
        no-ipv6-unicast;
    }
    interface so-1/2/2.0;
    interface lo0.0;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the Topologies on Neighbors on page 3865](#)
- [Checking the IS-IS SPF Calculations on page 3866](#)
- [Checking the Tcpdump Output on page 3867](#)

### Checking the Topologies on Neighbors

**Purpose** Determine what topologies are supported on neighboring IS-IS devices.

**Action** From operational mode, enter the **show isis adjacency detail** command.

```
user@R1> show isis adjacency detail
```

R2

```

Interface: so-1/2/0.0, Level: 3, State: Up, Expires in 24 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:28:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast, IPV6-Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.2
IPv6 addresses: fe80::2a0:a514:0:24c

```

R5

```

Interface: so-1/2/1.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:27:47 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.14

```

R3

```

Interface: so-1/2/2.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:27:25 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast, IPV6-Unicast

```

```
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 10.0.0.18
IPv6 addresses: fe80::2a0:a514:0:124c
```

**Meaning** As expected, the adjacency with Device R5 only supports the IPv4 unicast topology, while the adjacencies with Device R2 and Device R3 support both the IPv4 and IPv6 topologies.

### *Checking the IS-IS SPF Calculations*

**Purpose** Verify that separate SPF calculations are being run for IPv4 and IPv6.

**Action** From operational mode, enter the **show isis spf brief** command.

```
user@R1> show isis spf brief
```

#### **IPv4 Unicast IS-IS level 1 SPF results:**

| Node  | Metric | Interface  | NH      | Via | SNPA |
|-------|--------|------------|---------|-----|------|
| R6.00 | 20     | so-1/2/1.0 | IPV4 R5 |     |      |
| R4.00 | 20     | so-1/2/0.0 | IPV4 R2 |     |      |
| R5.00 | 10     | so-1/2/1.0 | IPV4 R5 |     |      |
| R3.00 | 10     | so-1/2/2.0 | IPV4 R3 |     |      |
| R2.00 | 10     | so-1/2/0.0 | IPV4 R2 |     |      |
| R1.00 | 0      |            |         |     |      |

6 nodes

#### **IPv4 Unicast IS-IS level 2 SPF results:**

| Node  | Metric | Interface  | NH      | Via | SNPA |
|-------|--------|------------|---------|-----|------|
| R6.00 | 20     | so-1/2/1.0 | IPV4 R5 |     |      |
| R4.00 | 20     | so-1/2/0.0 | IPV4 R2 |     |      |
| R5.00 | 10     | so-1/2/1.0 | IPV4 R5 |     |      |
| R3.00 | 10     | so-1/2/2.0 | IPV4 R3 |     |      |
| R2.00 | 10     | so-1/2/0.0 | IPV4 R2 |     |      |
| R1.00 | 0      |            |         |     |      |

6 nodes

#### **IPv6 Unicast IS-IS level 1 SPF results:**

| Node  | Metric | Interface  | NH      | Via | SNPA |
|-------|--------|------------|---------|-----|------|
| R5.00 | 40     | so-1/2/0.0 | IPV6 R2 |     |      |
| R6.00 | 30     | so-1/2/0.0 | IPV6 R2 |     |      |
| R4.00 | 20     | so-1/2/0.0 | IPV6 R2 |     |      |
| R3.00 | 10     | so-1/2/2.0 | IPV6 R3 |     |      |
| R2.00 | 10     | so-1/2/0.0 | IPV6 R2 |     |      |
| R1.00 | 0      |            |         |     |      |

6 nodes

#### **IPv6 Unicast IS-IS level 2 SPF results:**

| Node  | Metric | Interface  | NH      | Via | SNPA |
|-------|--------|------------|---------|-----|------|
| R5.00 | 40     | so-1/2/0.0 | IPV6 R2 |     |      |
| R6.00 | 30     | so-1/2/0.0 | IPV6 R2 |     |      |
| R4.00 | 20     | so-1/2/0.0 | IPV6 R2 |     |      |
| R3.00 | 10     | so-1/2/2.0 | IPV6 R3 |     |      |
| R2.00 | 10     | so-1/2/0.0 | IPV6 R2 |     |      |
| R1.00 | 0      |            |         |     |      |

6 nodes

**Meaning** As expected, SPF calculations are being performed for IPv4 and IPv6 topologies.

**Checking the Tcpdump Output**

**Purpose** Verify that the link can be a member of both the IPv4 unicast topology and the IPv6 unicast topology.

**Action** user@R1> **monitor traffic** detail interface so-1/2/0.0  
[...]

```
15:52:35.719540 In IS-IS, length 82
p2p IIH, hlen: 20, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
source-id: 0192.0168.0002, holding time: 27s, Flags: [Level 1, Level
2]
circuit-id: 0x01, PDU length: 82
Point-to-point Adjacency State TLV #240, length: 15
Adjacency State: Up (0)
Extended Local circuit-ID: 0x00000054
Neighbor System-ID: 0192.0168.0001
Neighbor Extended Local circuit-ID: 0x00000043
Protocols supported TLV #129, length: 2
NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
IPv4 Interface address(es) TLV #132, length: 4
IPv4 interface address: 10.0.0.2
IPv6 Interface address(es) TLV #232, length: 16
IPv6 interface address: fe80::2a0:a514:0:24c
Area address(es) TLV #1, length: 4
Area address (length: 3): 49.0002
Restart Signaling TLV #211, length: 3
Flags [none], Remaining holding time 0s
Multi Topology TLV #229, length: 4
  IPv4 unicast Topology (0x000), Flags: [none]
  IPv6 unicast Topology (0x002), Flags: [none]
```

**Meaning** The IS-IS hello (IIH) packet shows that IPv4 and IPv6 are supported. The hello packet lists valid IPv4 and IPv6 addresses, and therefore the routing device can create valid next-hop entries. The supported protocols are listed in the multitopology TLV #229.

**Related Documentation**

- [Example: Configuring IS-IS Dual Stacking of IPv4 and IPv6 Unicast Addresses](#)

**Example: Configuring IS-IS Multicast Topology**

- [IS-IS Multicast Topologies Overview on page 3868](#)
- [Example: Configuring IS-IS Multicast Topology on page 3869](#)

## IS-IS Multicast Topologies Overview

Most multicast routing protocols perform a reverse-path forwarding (RPF) check on the source of multicast data packets. If a packet comes in on the interface that is used to send data to the source, the packet is accepted and forwarded to one or more downstream interfaces. Otherwise, the packet is discarded and a notification is sent to the multicast routing protocol running on the interface.

In certain instances, the unicast routing table used for the RPF check is also the table used for forwarding unicast data packets. Thus, unicast and multicast routing are congruent. In other cases, where it is preferred that multicast routing be independent of unicast routing, the multicast routing protocols are configured to perform the RPF check using an alternate unicast routing table `inet.2`.

You can configure IS-IS to calculate an alternate IPv4 multicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to `inet.2`. The IS-IS interface metrics for the multicast topology can be configured independently of the unicast metrics. You can also selectively disable interfaces from participating in the multicast topology while continuing to participate in the regular unicast topology. This enables you to exercise control over the paths that multicast data takes through a network so that it is independent of unicast data paths. You can also configure IS-IS to calculate an alternate IPv6 multicast topology, in addition to the normal IPv6 unicast topology.



**NOTE:** IS-IS only starts advertising the routes when the interface routes are in `inet.2`.

Table 308 on page 3868 lists the various IPv4 statements you can use to configure IS-IS topologies.

**Table 308: IPv4 Statements**

| Statement  | Description   |
|--|---|
| <code>ipv4-multicast</code>                      | Enables an alternate IPv4 multicast topology.                             |
| <code>ipv4-multicast-metric</code> <i>number</i> | Configures the multicast metric for an alternate IPv4 multicast topology. |
| <code>no-ipv4-multicast</code>                   | Excludes an interface from the IPv4 multicast topology.                   |
| <code>no-unicast-topology</code>                 | Excludes an interface from the IPv4 unicast topologies.                   |

Table 309 on page 3868 lists the various IPv6 statements you can use to configure IS-IS topologies.

**Table 309: IPv6 Statements**

| Statement                   | Description                                   |
|-----------------------------|---|
| <code>ipv6-multicast</code> | Enables an alternate IPv6 multicast topology. |

Table 309: IPv6 Statements (*continued*)

| Statement  | Description   |
|--|---|
| <code>ipv6-multicast-metric <i>number</i></code> | Configures the multicast metric for an alternate IPv6 multicast topology. |
| <code>ipv6-unicast-metric <i>number</i></code>   | Configures the unicast metric for an alternate IPv6 multicast topology.   |
| <code>no-ipv6-multicast</code>                   | Excludes an interface from the IPv6 multicast topology.                   |
| <code>no-ipv6-unicast</code>                     | Excludes an interface from the IPv6 unicast topologies.                   |

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

### Example: Configuring IS-IS Multicast Topology

This example shows how to configure a multicast topology for an IS-IS network.

- [Requirements on page 3869](#)
- [Overview on page 3869](#)
- [Configuration on page 3870](#)
- [Verification on page 3874](#)

#### Requirements

Before you begin, configure IS-IS on all routers. See [“Example: Configuring IS-IS” on page 3823](#) for information about the required IS-IS configuration.

This example uses the following hardware and software components:


- Junos OS Release 7.3 or later
- M Series, MX Series, and T Series routers

#### Overview

This example shows an IS-IS multicast topology configuration. Three routers are connected to each other. A loopback interface is configured on each router.

[Figure 107 on page 3870](#) shows the sample network.

Figure 107: Configuring IS-IS Multicast Topology



g040922

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Router R1

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 15
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-0/0/0 level 2 metric 20
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 14
set protocols isis interface so-1/0/0 level 1 metric 13
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-1/0/0 level 2 metric 29
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface fxp0.0 disable
```

#### Router R2

```
set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 13
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 12
set protocols isis interface so-0/0/0 level 2 metric 29
```



```

set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 23
set protocols isis interface so-1/0/0 level 1 metric 14
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 18
set protocols isis interface so-1/0/0 level 2 metric 32
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 26
set protocols isis interface fxp0.0 disable

```

### Router R3

```

set protocols isis traceoptions file isis size 5m world-readable
set protocols isis traceoptions flag error
set protocols isis topologies ipv4-multicast
set protocols isis interface so-0/0/0 level 1 metric 19
set protocols isis interface so-0/0/0 level 1 ipv4-multicast-metric 11
set protocols isis interface so-0/0/0 level 2 metric 27
set protocols isis interface so-0/0/0 level 2 ipv4-multicast-metric 21
set protocols isis interface so-1/0/0 level 1 metric 16
set protocols isis interface so-1/0/0 level 1 ipv4-multicast-metric 26
set protocols isis interface so-1/0/0 level 2 metric 30
set protocols isis interface so-1/0/0 level 2 ipv4-multicast-metric 20
set protocols isis interface fxp0.0 disable

```

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS multicast topologies:

1. Enable the multicast topology for IS-IS by using the **ipv4-multicast** statement.

### Routers R1, R2, and R3

```

[edit protocols isis]
user@host# set traceoptions file isis size 5m world-readable
user@host# set traceoptions flag error
user@host# set topologies ipv4-multicast

```

2. Enable multicast metrics on the first SONET/SDH Interface by using the **ipv4-multicast-metric** statement.

### Router R1

```

[edit protocols isis interface so-0/0/0 ]
user@R1# set level 1 metric 15
user@R1# set level 1 ipv4-multicast-metric 18
user@R1# set level 2 metric 20
user@R1# set level 2 ipv4-multicast-metric 14

```

### Router R2

```

[edit protocols isis interface so-0/0/0]
user@R2# set level 1 metric 13
user@R2# set level 1 ipv4-multicast-metric 12
user@R2# set level 2 metric 29
user@R2# set level 2 ipv4-multicast-metric 23

```

### Router R3

```

[edit protocols isis interface so-0/0/0]

```

```
user@R3# set level 1 metric 19
user@R3# set level 1 ipv4-multicast-metric 11
user@R3# set level 2 metric 27
user@R3# set level 2 ipv4-multicast-metric 21
```

3. Enable multicast metrics on a second sonet Interface by using the **ipv4-multicast-metric** statement.

#### Router R1

```
[edit protocols isis interface so-1/0/0]
user@R1# set level 1 metric 13
user@R1# set level 1 ipv4-multicast-metric 12
user@R1# set level 2 metric 29
user@R1# set level 2 ipv4-multicast-metric 23
```

#### Router R2

```
[edit protocols isis interface so-1/0/0]
user@R2# set level 1 metric 14
user@R2# set level 1 ipv4-multicast-metric 18
user@R2# set level 2 metric 32
user@R2# set level 2 ipv4-multicast-metric 26
```

#### Router R3

```
[edit protocols isis interface so-1/0/0]
user@R3# set level 1 metric 16
user@R3# set level 1 ipv4-multicast-metric 26
user@R3# set level 2 metric 30
user@R3# set level 2 ipv4-multicast-metric 20
```

4. Disable the out-of-band management port, fxp0.

#### Routers R1, R2, and R3

```
[edit protocols isis]
user@host# set interface fxp0.0 disable
```

5. If you are done configuring the routers, commit the configuration.

#### Routers R1, R2, and R3

```
[edit]
user@host# commit
```

**Results** From configuration mode, confirm your configuration by using the **show protocols isis** statement. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

#### Router R1

```
user@R1# show protocols isis

traceoptions {
  file isis size 5m world-readable;
  flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
  level 1 {
```

```

        metric 15;
        ipv4-multicast-metric 18;
    }
    level 2 {
        metric 20;
        ipv4-multicast-metric 14;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 13;
        ipv4-multicast-metric 12;
    }
    level 2 {
        metric 29;
        ipv4-multicast-metric 23;
    }
}
interface fxp0.0 {
    disable;
}

```

### Router R2

user@R2# show protocols isis

```

traceoptions {
    file isis size 5m world-readable;
    flag error;
}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 13;
        ipv4-multicast-metric 12;
    }
    level 2 {
        metric 29;
        ipv4-multicast-metric 23;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 14;
        ipv4-multicast-metric 18;
    }
    level 2 {
        metric 32;
        ipv4-multicast-metric 26;
    }
}
interface fxp0.0 {
    disable;
}

```

### Router R3

user@R3# show protocols isis

```

traceoptions {
    file isis size 5m world-readable;
    flag error;
}

```

```
}
topologies ipv4-multicast;
interface so-0/0/0 {
    level 1 {
        metric 19;
        ipv4-multicast-metric 11;
    }
    level 2 {
        metric 27;
        ipv4-multicast-metric 21;
    }
}
interface so-1/0/0 {
    level 1 {
        metric 16;
        ipv4-multicast-metric 26;
    }
    level 2 {
        metric 30;
        ipv4-multicast-metric 20;
    }
}
interface fxp0.0 {
    disable;
}
```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the Connection Between Routers R1, R2, and R3 on page 3874](#)
- [Verifying That IS-IS Is Configured on page 3876](#)
- [Verifying the Configured Multicast Metric Values on page 3878](#)
- [Verifying the Configuration of the Multicast Topology on page 3879](#)

### **Verifying the Connection Between Routers R1, R2, and R3**

**Purpose** Make sure that Routers R1, R2, and R3 are connected to each other.

**Action** Ping the other two routers from any router, to check the connectivity between the three routers as per the network topology.

```
user@R1> ping 10.0.3.9
```

```
PING 10.0.3.9 (10.0.3.9): 56 data bytes
64 bytes from 10.0.3.9: icmp_seq=0 ttl=64 time=1.299 ms
64 bytes from 10.0.3.9: icmp_seq=1 ttl=64 time=52.304 ms
64 bytes from 10.0.3.9: icmp_seq=2 ttl=64 time=1.271 ms
64 bytes from 10.0.3.9: icmp_seq=3 ttl=64 time=1.343 ms
64 bytes from 10.0.3.9: icmp_seq=4 ttl=64 time=1.434 ms
64 bytes from 10.0.3.9: icmp_seq=5 ttl=64 time=1.306 ms
^C
--- 10.0.3.9 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.271/9.826/52.304/18.997 ms
```

```
user@R1> ping 10.0.3.10
```

```

PING 10.0.3.10 (10.0.3.10): 56 data bytes
64 bytes from 10.0.3.10: icmp_seq=0 ttl=64 time=1.431 ms
64 bytes from 10.0.3.10: icmp_seq=1 ttl=64 time=1.296 ms
64 bytes from 10.0.3.10: icmp_seq=2 ttl=64 time=1.887 ms
^C
--- 10.0.3.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.296/1.538/1.887/0.253 ms

```

```
user@R2> ping 10.0.2.9
```

```

PING 10.0.2.9 (10.0.2.9): 56 data bytes
64 bytes from 10.0.2.9: icmp_seq=0 ttl=64 time=1.365 ms
64 bytes from 10.0.2.9: icmp_seq=1 ttl=64 time=1.813 ms
64 bytes from 10.0.2.9: icmp_seq=2 ttl=64 time=1.290 ms
^C
--- 10.0.2.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.290/1.489/1.813/0.231 ms

```

```
user@R2> ping 10.0.2.10
```

```

PING 10.0.2.10 (10.0.2.10): 56 data bytes
64 bytes from 10.0.2.10: icmp_seq=0 ttl=63 time=1.318 ms
64 bytes from 10.0.2.10: icmp_seq=1 ttl=63 time=1.394 ms
64 bytes from 10.0.2.10: icmp_seq=2 ttl=63 time=1.366 ms
64 bytes from 10.0.2.10: icmp_seq=3 ttl=63 time=1.305 ms
^C
--- 10.0.2.10 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.305/1.346/1.394/0.036 ms

```

```
user@R3> ping 10.0.1.10
```

```

PING 10.0.1.10 (10.0.1.10): 56 data bytes
64 bytes from 10.0.1.10: icmp_seq=0 ttl=63 time=1.316 ms
64 bytes from 10.0.1.10: icmp_seq=1 ttl=63 time=1.418 ms
64 bytes from 10.0.1.10: icmp_seq=2 ttl=63 time=1.277 ms
^C
--- 10.0.1.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.277/1.337/1.418/0.059 ms

```

```
user@R3> ping 10.0.1.9
```

```

PING 10.0.1.9 (10.0.1.9): 56 data bytes
64 bytes from 10.0.1.9: icmp_seq=0 ttl=64 time=1.381 ms
64 bytes from 10.0.1.9: icmp_seq=1 ttl=64 time=1.499 ms
64 bytes from 10.0.1.9: icmp_seq=2 ttl=64 time=1.300 ms
64 bytes from 10.0.1.9: icmp_seq=3 ttl=64 time=1.397 ms
^C
--- 10.0.1.9 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.300/1.394/1.499/0.071 ms

```

**Meaning** Routers R1, R2, and R3 have a peer relationship with each other.

### *Verifying That IS-IS Is Configured*

**Purpose** Make sure that the IS-IS instance is running on Routers R1, R2, and R3, and that they are adjacent to each other.

**Action** Use the **show isis adjacency detail** command to check the adjacency between the routers.

#### **Router R1**

```
user@R1> show isis adjacency detail
```

R2

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:59 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R2

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 8 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:23:58 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.10
```

R3

```
Interface: so-1/0/0, Level: 1, State: Up, Expires in 7 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

R3

```
Interface: so-1/0/0, Level: 2, State: Up, Expires in 6 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:24:20 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R3.02, IP addresses: 10.0.2.10
```

#### **Router R2**

```
user@R2> show isis adjacency detail
```

R1

```
Interface: so-0/0/0, Level: 1, State: Up, Expires in 20 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
Restart capable: Yes, Adjacency advertisement: Advertise
LAN id: R2.02, IP addresses: 10.0.1.9
```

R1

```
Interface: so-0/0/0, Level: 2, State: Up, Expires in 26 secs
Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:50 ago
Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc
Topologies: IPV4-Multicast
```

Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R2.02, IP addresses: 10.0.1.9

R3

Interface: so-1/0/0, Level: 1, State: Up, Expires in 8 secs  
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago  
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd  
 Topologies: IPV4-Multicast  
 Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R3.03, IP addresses: 10.0.3.10

R3

Interface: so-1/0/0, Level: 2, State: Up, Expires in 8 secs  
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:27:22 ago  
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bd  
 Topologies: IPV4-Multicast  
 Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R3.03, IP addresses: 10.0.3.10

### Router R3

user@R3> show isis adjacency detail

R2

Interface: so-0/0/0, Level: 1, State: Up, Expires in 18 secs  
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago  
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc  
 Topologies: IPV4-Multicast  
 Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R3.03, IP addresses: 10.0.3.9

R2

Interface: so-0/0/0, Level: 2, State: Up, Expires in 22 secs  
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:09 ago  
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc  
 Topologies: IPV4-Multicast  
 Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R3.03, IP addresses: 10.0.3.9

R1

Interface: so-1/0/0, Level: 1, State: Up, Expires in 21 secs  
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago  
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc  
 Topologies: IPV4-Multicast  
 Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R3.02, IP addresses: 10.0.2.9

R1

Interface: so-1/0/0, Level: 2, State: Up, Expires in 19 secs  
 Priority: 64, Up/Down transitions: 1, Last transition: 2d 19:33:59 ago  
 Circuit type: 3, Speaks: IP, MAC address: 0:1b:c0:86:54:bc  
 Topologies: IPV4-Multicast  
 Restart capable: Yes, Adjacency advertisement: Advertise  
 LAN id: R3.02, IP addresses: 10.0.2.9

**Meaning** IS-IS is configured on Routers R1, R2, and R3, and they are adjacent to each other.

**Verifying the Configured Multicast Metric Values**

**Purpose** Make sure that the SPF calculations are accurate as per the configured multicast metric values on Routers R1, R2, and R3.

**Action** Use the **show isis spf results** command to check the SPF calculations for the network.

**Router R1**

```
user@R1> show isis spf results
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.03  28         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R2.00  18         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R3.00  17         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  0
      4 nodes

IPv4 Multicast IS-IS level 2 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.03  40         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R3.00  22         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R2.00  14         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bd
R1.00  0
      4 nodes
```

**Router R2**

```
user@R2> show isis spf results
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  29         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R3.00  18         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  12         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R2.02  12
R2.00  0
      5 nodes

IPv4 Multicast IS-IS level 2 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  45         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R3.00  26         so-1/0/0   IPV4 R3  0:1b:c0:86:54:bd
R1.00  23         so-0/0/0   IPV4 R1  0:1b:c0:86:54:bc
R2.02  23
R2.00  0
      5 nodes
```

**Router R3**

```
user@R3> show isis spf results
...
IPv4 Multicast IS-IS level 1 SPF results:
Node  Metric  Interface  NH  Via  SNPA
R3.02  26
R1.00  23         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R2.02  23         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R2.00  11         so-0/0/0   IPV4 R2  0:1b:c0:86:54:bc
R3.03  11
```



```

R3.00 0
      6 nodes

IPv4 Multicast IS-IS level 2 SPF results:
Node Metric Interface NH Via SNPA
R2.02 34 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R2.00 21 so-0/0/0 IPv4 R2 0:1b:c0:86:54:bc
R3.03 21
R1.00 20 so-1/0/0 IPv4 R1 0:1b:c0:86:54:bc
R3.02 20
R3.00 0
      6 nodes

```

**Meaning** The configured multicast metric values are used in SPF calculations for the IS-IS network.

### *Verifying the Configuration of the Multicast Topology*

**Purpose** Make sure that the multicast topology is configured on Routers R1, R2, and R3.

**Action** Use the **show isis database detail** command to verify the multicast topology configuration on the routers.

#### **Router R1**

```
user@R1> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```

R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 663 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up

```

```

R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 883 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13
  IPv4 Unicast IS neighbor: R3.03 Metric: 14
  IPv4 Multicast IS neighbor: R2.02 Metric: 12
  IPv4 Multicast IS neighbor: R3.03 Metric: 18
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 13 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 14 Internal Up

```

```

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 913 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R2.00 Metric: 0

```

```

R3.00-00 Sequence: 0x13c, Checksum: 0xc8de, Lifetime: 488 secs
  IPv4 Unicast IS neighbor: R3.02 Metric: 16
  IPv4 Unicast IS neighbor: R3.03 Metric: 19
  IPv4 Multicast IS neighbor: R3.02 Metric: 26
  IPv4 Multicast IS neighbor: R3.03 Metric: 11
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 16 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 19 Internal Up

```

```

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 625 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0

```

```
R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 714 secs
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0
```

IS-IS level 2 link-state database:

```
R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 816 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 20
  IPv4 Unicast IS neighbor: R3.02 Metric: 31
  IPv4 Multicast IS neighbor: R2.02 Metric: 14
  IPv4 Multicast IS neighbor: R3.02 Metric: 22
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 20 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 31 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 29 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 966 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 29
  IPv4 Unicast IS neighbor: R3.03 Metric: 32
  IPv4 Multicast IS neighbor: R2.02 Metric: 23
  IPv4 Multicast IS neighbor: R3.03 Metric: 26
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 29 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 28 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 32 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 966 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
```

```
R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 805 secs
  IPv4 Unicast IS neighbor: R3.02 Metric: 30
  IPv4 Unicast IS neighbor: R3.03 Metric: 27
  IPv4 Multicast IS neighbor: R3.02 Metric: 20
  IPv4 Multicast IS neighbor: R3.03 Metric: 21
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 31 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 30 Internal Up
  IP IPv4 Unicast prefix: 10.0.3.8/30 Metric: 27 Internal Up
```

```
R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 844 secs
  IPv4 Unicast IS neighbor: R1.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0
```

```
R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 844 secs
  IPv4 Unicast IS neighbor: R2.00 Metric: 0
  IPv4 Unicast IS neighbor: R3.00 Metric: 0
```

## Router R2

```
user@R2> show isis database detail
```

IS-IS level 1 link-state database:

```
R1.00-00 Sequence: 0x142, Checksum: 0xd07, Lifetime: 524 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 15
  IPv4 Unicast IS neighbor: R3.02 Metric: 15
  IPv4 Multicast IS neighbor: R2.02 Metric: 18
  IPv4 Multicast IS neighbor: R3.02 Metric: 17
  IP IPv4 Unicast prefix: 10.0.1.8/30 Metric: 15 Internal Up
  IP IPv4 Unicast prefix: 10.0.2.8/30 Metric: 15 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 748 secs
  IPv4 Unicast IS neighbor: R2.02 Metric: 13
```

```

IPV4 Unicast IS neighbor: R3.03      Metric:      14
IPV4 Multicast IS neighbor: R2.02     Metric:      12
IPV4 Multicast IS neighbor: R3.03     Metric:      18
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      13 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      14 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 777 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1102 secs
IPV4 Unicast IS neighbor: R3.02      Metric:      16
IPV4 Unicast IS neighbor: R3.03      Metric:      19
IPV4 Multicast IS neighbor: R3.02     Metric:      26
IPV4 Multicast IS neighbor: R3.03     Metric:      11
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      16 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      19 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 488 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 577 secs
IPV4 Unicast IS neighbor: R2.00      Metric:      0
IPV4 Unicast IS neighbor: R3.00      Metric:      0

IS-IS level 2 link-state database:

R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 676 secs
IPV4 Unicast IS neighbor: R2.02      Metric:      20
IPV4 Unicast IS neighbor: R3.02      Metric:      31
IPV4 Multicast IS neighbor: R2.02     Metric:      14
IPV4 Multicast IS neighbor: R3.02     Metric:      22
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      20 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      31 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      29 Internal Up

R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 831 secs
IPV4 Unicast IS neighbor: R2.02      Metric:      29
IPV4 Unicast IS neighbor: R3.03      Metric:      32
IPV4 Multicast IS neighbor: R2.02     Metric:      23
IPV4 Multicast IS neighbor: R3.03     Metric:      26
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      29 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      28 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 831 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0
IPV4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 667 secs
IPV4 Unicast IS neighbor: R3.02      Metric:      30
IPV4 Unicast IS neighbor: R3.03      Metric:      27
IPV4 Multicast IS neighbor: R3.02     Metric:      20
IPV4 Multicast IS neighbor: R3.03     Metric:      21
IP IPV4 Unicast prefix: 10.0.1.8/30   Metric:      31 Internal Up
IP IPV4 Unicast prefix: 10.0.2.8/30   Metric:      30 Internal Up
IP IPV4 Unicast prefix: 10.0.3.8/30   Metric:      27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 707 secs
IPV4 Unicast IS neighbor: R1.00      Metric:      0

```

```
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

```
R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 707 secs
```

```
IPv4 Unicast IS neighbor: R2.00    Metric:      0
```

```
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

### Router R3

```
user@R3> show isis database detail
```

```
IS-IS level 1 link-state database:
```

```
R1.00-00 Sequence: 0x143, Checksum: 0xb08, Lifetime: 1155 secs
```

```
IPv4 Unicast IS neighbor: R2.02    Metric:      15
```

```
IPv4 Unicast IS neighbor: R3.02    Metric:      15
```

```
IPv4 Multicast IS neighbor: R2.02   Metric:      18
```

```
IPv4 Multicast IS neighbor: R3.02   Metric:      17
```

```
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      15 Internal Up
```

```
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      15 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0xf02b, Lifetime: 687 secs
```

```
IPv4 Unicast IS neighbor: R2.02    Metric:      13
```

```
IPv4 Unicast IS neighbor: R3.03    Metric:      14
```

```
IPv4 Multicast IS neighbor: R2.02   Metric:      12
```

```
IPv4 Multicast IS neighbor: R3.03   Metric:      18
```

```
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      13 Internal Up
```

```
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      14 Internal Up
```

```
R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 716 secs
```

```
IPv4 Unicast IS neighbor: R1.00    Metric:      0
```

```
IPv4 Unicast IS neighbor: R2.00    Metric:      0
```

```
R3.00-00 Sequence: 0x13d, Checksum: 0xc6df, Lifetime: 1044 secs
```

```
IPv4 Unicast IS neighbor: R3.02    Metric:      16
```

```
IPv4 Unicast IS neighbor: R3.03    Metric:      19
```

```
IPv4 Multicast IS neighbor: R3.02   Metric:      26
```

```
IPv4 Multicast IS neighbor: R3.03   Metric:      11
```

```
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      16 Internal Up
```

```
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      19 Internal Up
```

```
R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 430 secs
```

```
IPv4 Unicast IS neighbor: R1.00    Metric:      0
```

```
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

```
R3.03-00 Sequence: 0x138, Checksum: 0xad56, Lifetime: 519 secs
```

```
IPv4 Unicast IS neighbor: R2.00    Metric:      0
```

```
IPv4 Unicast IS neighbor: R3.00    Metric:      0
```

```
IS-IS level 2 link-state database:
```

```
R1.00-00 Sequence: 0x142, Checksum: 0x2c7c, Lifetime: 617 secs
```

```
IPv4 Unicast IS neighbor: R2.02    Metric:      20
```

```
IPv4 Unicast IS neighbor: R3.02    Metric:      31
```

```
IPv4 Multicast IS neighbor: R2.02   Metric:      14
```

```
IPv4 Multicast IS neighbor: R3.02   Metric:      22
```

```
IP IPv4 Unicast prefix: 10.0.1.8/30 Metric:      20 Internal Up
```

```
IP IPv4 Unicast prefix: 10.0.2.8/30 Metric:      31 Internal Up
```

```
IP IPv4 Unicast prefix: 10.0.3.8/30 Metric:      29 Internal Up
```

```
R2.00-00 Sequence: 0x13f, Checksum: 0x4826, Lifetime: 769 secs
```

```
IPv4 Unicast IS neighbor: R2.02    Metric:      29
```

```
IPv4 Unicast IS neighbor: R3.03    Metric:      32
```

```

IPv4 Multicast IS neighbor: R2.02    Metric:      23
IPv4 Multicast IS neighbor: R3.03    Metric:      26
IP  IPv4 Unicast prefix: 10.0.1.8/30  Metric:     29 Internal Up
IP  IPv4 Unicast prefix: 10.0.2.8/30  Metric:     28 Internal Up
IP  IPv4 Unicast prefix: 10.0.3.8/30  Metric:     32 Internal Up

R2.02-00 Sequence: 0x13c, Checksum: 0x57e2, Lifetime: 769 secs
IPv4 Unicast IS neighbor: R1.00      Metric:      0
IPv4 Unicast IS neighbor: R2.00      Metric:      0

R3.00-00 Sequence: 0x13d, Checksum: 0x1b19, Lifetime: 610 secs
IPv4 Unicast IS neighbor: R3.02      Metric:      30
IPv4 Unicast IS neighbor: R3.03      Metric:      27
IPv4 Multicast IS neighbor: R3.02    Metric:      20
IPv4 Multicast IS neighbor: R3.03    Metric:      21
IP  IPv4 Unicast prefix: 10.0.1.8/30  Metric:     31 Internal Up
IP  IPv4 Unicast prefix: 10.0.2.8/30  Metric:     30 Internal Up
IP  IPv4 Unicast prefix: 10.0.3.8/30  Metric:     27 Internal Up

R3.02-00 Sequence: 0x139, Checksum: 0xfb0e, Lifetime: 649 secs
IPv4 Unicast IS neighbor: R1.00      Metric:      0
IPv4 Unicast IS neighbor: R3.00      Metric:      0

R3.03-00 Sequence: 0x139, Checksum: 0xab57, Lifetime: 649 secs
IPv4 Unicast IS neighbor: R2.00      Metric:      0
IPv4 Unicast IS neighbor: R3.00      Metric:      0

```

**Meaning** Multicast topology is configured on Routers R1, R2, and R3.

**Related Documentation**

- [Example: Configuring Multitopology Routing Based on a Multicast Source](#)
- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859](#)

## Example: Configuring Link and Node Protection for IS-IS Routes

- [Understanding Loop-Free Alternate Routes for IS-IS on page 3883](#)
- [Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN on page 3887](#)

### Understanding Loop-Free Alternate Routes for IS-IS

In Junos OS Release 9.5 and later, support for IS-IS loop-free alternate routes enables IP fast-reroute capability for IS-IS. Junos OS precomputes loop-free backup routes for all IS-IS routes. These backup routes are preinstalled in the Packet Forwarding Engine, which performs a local repair and implements the backup path when the link for a primary next hop for a particular route is no longer available. With local repair, the Packet Forwarding Engine can correct a path failure before it receives recomputed paths from the Routing Engine. Local repair reduces the amount of time needed to reroute traffic to less than 50 milliseconds. In contrast, global repair can take up to 800 milliseconds to compute a new route. Local repair and global repair are thus complementary. Local repair enables traffic to continue to be routed using a backup path until global repair is able to calculate a new route.

A loop-free path is one that does not forward traffic back through the routing device to reach a given destination. That is, a neighbor whose shortest path to the destination traverses the routing device is not used as a backup route to that destination. To determine

loop-free alternate paths for IS-IS routes, Junos OS runs shortest-path-first (SPF) calculations on each one-hop neighbor. You can enable support for alternate loop-free routes on any IS-IS interface. Because it is common practice to enable LDP on an interface for which IS-IS is already enabled, this feature also provides support for LDP label-switched paths (LSPs).



**NOTE:** If you enable support for alternate loop-free routes on an interface configured for both LDP and IS-IS, you can use the `traceroute` command to trace the active path to the primary next hop.


The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.

Junos OS provides two mechanisms for route redundancy for IS-IS through alternate loop-free routes: link protection and node-link protection. When you enable link protection or node-link protection on an IS-IS interface, Junos OS creates a single alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection offers per-link traffic protection. Use link protection when you assume that only a single link might become unavailable but that the neighboring node on the primary path would still be available through another interface.

Node-link protection establishes an alternate path through a different routing device altogether. Use node-link protection when you assume that access to a node is lost when a link is no longer available. As a result, Junos OS calculates a backup path that avoids the primary next-hop routing device. In Junos OS Release 9.4 and earlier, only the RSVP protocol supports Packet Forwarding Engine local repair and fast reroute as well as link protection and node protection.

In [Figure 99 on page 3820](#), Case 2 shows how link protection allows source Router A to switch to Link B when the primary next hop Link A to destination Router C fails. However, if Router B fails, Link B also fails, and the protected Link A is lost. If node-link protection is enabled, Router A is able to switch to Link D on Router D and bypass the failed Router B altogether. As shown in Case 1, with node-link protection enabled, Router A has a node-link protection alternate path available through Router D to destination Router C. That means that if Router B fails, Router A can still reach Router C because the path from Router A to Link D remains available as an alternate backup path.

**Figure 108: Link Protection and Node-Link Protection Comparison for IS-IS Routes**



The Junos OS implementation of support for loop-free alternate paths for IS-IS routes is based on the following standards:

- RFC 5286, *Basic Specification for IP Fast-Reroute: Loop-free Alternates*
- RFC 5714, *IP Fast Reroute Framework*

### Configuring Link Protection for IS-IS

You can configure link protection on any interface for which IS-IS is enabled. When you enable link protection, Junos OS creates one alternate path to the primary next hop for all destination routes that traverse a protected interface. Link protection assumes that only a single link becomes unavailable but that the neighboring node would still be available through another interface.



**NOTE:** You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable link protection, include the **link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```
[edit]
protocols {
  isis {
    interface interface-name {
```

```

    link-protection;
  }
}

```

### Configuring Node-Link Protection for IS-IS

You can configure node-link protection on any interface for which IS-IS is enabled. Node-link protection establishes an alternate path through a different routing device altogether for all destination routes that traverse a protected interface. Node-link protection assumes that the entire routing device, or node, has failed. Junos OS therefore calculates a backup path that avoids the primary next-hop routing device.



**NOTE:** You must also configure a per-packet load-balancing routing policy to ensure that the routing protocol process installs all the next hops for a given route in the routing table.

To enable node-link protection, include the **node-link-protection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```

[edit]
protocols {
  isis {
    interface interface-name {
      node-link-protection;
    }
  }
}

```

### Excluding an IS-IS Interface as a Backup for Protected Interfaces

By default, all IS-IS interfaces that belong to the master instance or a specific routing instance are eligible as backup interfaces for protected interfaces. You can specify that any IS-IS interface be excluded from functioning as a backup interface to protected interfaces. To exclude an IS-IS interface as a backup interface, include the **no-eligible-backup** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level:

```

[edit]
protocols {
  isis {
    interface interface-name {
      no-eligible-backup;
    }
  }
}

```

### Configuring RSVP Label-Switched Paths as Backup Paths for IS-IS

Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP label-switched paths (LSPs) by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup



path, include the **backup** statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      backup;
      to ip-address;
    }
  }
}
```

When configuring an LSP, you must specify the IP address of the egress routing device with the **to** statement. For detailed information about configuring LSPs and RSVP, see the *RSVP Feature Guide for Routing Devices*.

### **Using Operational Mode Commands to Monitor Protected IS-IS Routes**

You can issue operational mode commands that provide more details about your link-protected and node-link-protected IS-IS routes. The following guidelines explain the type of information available from the output of each command:

- **show isis backup label-switched-path**—Displays which MPLS LSPs have been designated as backup paths and the current status of those LSPs.
- **show isis backup spf results**—Displays SPF calculations for each neighbor for a given destination. Indicates whether a specific interface or node has been designated as a backup path and why. Use the **no-coverage** option to display only those nodes that do not have backup coverage.
- **show isis backup coverage**—Displays the percentage of nodes and prefixes for each type of address family that is protected.
- **show isis interface detail**—Displays the type of protection (link or node-link) applied to each protected interface.

### **Example: Configuring Node-Link Protection for IS-IS Routes in a Layer 3 VPN**

Node-link protection establishes an alternate path through a different routing device. Use node-link protection when you assume that access to a node is lost when a link is no longer available. Junos OS calculates a backup path that avoids the primary next-hop routing device.

- [Requirements on page 3887](#)
- [Overview on page 3888](#)
- [Configuration on page 3888](#)
- [Verification on page 3895](#)

### **Requirements**

This example requires Junos OS Release 9.5 or later.

No special configuration beyond device initialization is required before configuring this example.


### Overview

In this example, core-facing interfaces are enabled for IS-IS Level 2, LDP, and RSVP. Node-link protection is enabled on all the core-facing interfaces, which means that if the primary next hop for any destination that traverses the interfaces becomes unavailable, Junos OS uses a backup link that avoids the next-hop router altogether if necessary.

You also need to configure a routing policy that requires all traffic to use per-packet load balancing in order to enable Packet Forwarding Engine local repair. With local repair, the Packet Forwarding Engine can correct a path failure and implement a backup loop-free alternate route before it receives recomputed paths from the Routing Engine.

Figure 109 on page 3888 shows the topology used in this example.

Figure 109: IS-IS Node-Link Protection Topology



On Device PE1, an RSVP LSP is configured as a backup path for IS-IS. Relying on the shortest-path-first (SPF) calculation of backup paths for one-hop neighbors might result in less than 100 percent backup coverage for a specific network topology. You can enhance coverage of IS-IS and LDP LSPs by configuring RSVP LSPs as backup paths. To configure a specific RSVP LSP as a backup path, include the **backup** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

“CLI Quick Configuration” on page 3888 shows the configuration for all of the devices in Figure 109 on page 3888. The section “Step-by-Step Procedure” on page 3892 describes the steps on Device P1.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device CE1**      **set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30**  
                      **set interfaces lo0 unit 0 family inet address 10.255.1.1/32**

**Device PE1**      **set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30**  
                      **set interfaces fe-1/2/0 unit 0 family iso**  
                      **set interfaces fe-1/2/0 unit 0 family mpls**

```

set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.5/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.21/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.2.2/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0202.00
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface fe-1/2/1.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path to-p2 backup
set protocols mpls label-switched-path to-p2 to 10.255.4.4
set protocols mpls label-switched-path to-p2 ldp-tunneling
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.2.2
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.5.5
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface fe-1/2/0.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A routing-options static route 10.255.1.1/32 next-hop 10.0.0.1
set routing-options autonomous-system 65534
set routing-options forwarding-table export ecmp

```

**Device P1**

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.6/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.9/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.25/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.3.3/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0303.00

```

```
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp
```

Device P2

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.10/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.13/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.4.4/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0404.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp
```

Device P3

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.22/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.30/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.7.7/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0707.00
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
```

```

set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-options forwarding-table export ecmp

```

#### Device PE2

```

set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.14/30
set interfaces fe-1/2/0 unit 0 family iso
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 10.0.0.17/30
set interfaces fe-1/2/1 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family inet address 10.0.0.29/30
set interfaces fe-1/2/2 unit 0 family iso
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.5.5/32
set interfaces lo0 unit 0 family iso address 49.0001.0010.0000.0505.00
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols bgp group l3vpn type internal
set protocols bgp group l3vpn local-address 10.255.5.5
set protocols bgp group l3vpn family inet-vpn unicast
set protocols bgp group l3vpn peer-as 65534
set protocols bgp group l3vpn local-as 65534
set protocols bgp group l3vpn neighbor 10.255.2.2
set protocols isis spf-options delay 1000
set protocols isis interface all node-link-protection
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 level 2 metric 0
set protocols ldp deaggregate
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement ecmp term 1 then load-balance per-packet
set routing-instances VPN-A instance-type vrf
set routing-instances VPN-A interface fe-1/2/1.0
set routing-instances VPN-A route-distinguisher 65534:1234
set routing-instances VPN-A vrf-target target:65534:1234
set routing-instances VPN-A routing-options static route 10.255.1.1/32 next-hop 10.0.0.18
set routing-options autonomous-system 65534
set routing-options forwarding-table export ecmp

```

**Device CE2**      **set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.18/30**  
**set interfaces lo0 unit 0 family inet address 10.255.6.6/32**

**Step-by-Step Procedure**    The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure multi-level IS-IS:

1.    Configure the interfaces.

Enable IS-IS and MPLS.

```
[edit interfaces]
user@P1# set fe-1/2/0 unit 0 family inet address 10.0.0.6/30
user@P1# set fe-1/2/0 unit 0 family iso
user@P1# set fe-1/2/0 unit 0 family mpls
user@P1# set fe-1/2/1 unit 0 family inet address 10.0.0.9/30
user@P1# set fe-1/2/1 unit 0 family iso
user@P1# set fe-1/2/1 unit 0 family mpls
user@P1# set fe-1/2/2 unit 0 family inet address 10.0.0.25/30
user@P1# set fe-1/2/2 unit 0 family iso
user@P1# set fe-1/2/2 unit 0 family mpls
user@P1# set lo0 unit 0 family inet address 10.255.3.3/32
user@P1# set lo0 unit 0 family iso address 49.0001.0010.0000.0303.00
```

2.    Configure the IS-IS interfaces for Level 2.

```
[edit protocols]
user@P1# set isis interface all level 2 metric 10
user@P1# set isis interface all level 1 disable
user@P1# set isis interface fxp0.0 disable
user@P1# set isis interface lo0.0 level 2 metric 0
```

3.    Enable IS-IS node-link protection, which also automatically extends backup coverage to all LDP LSPs.

```
[edit protocols]
user@P1# set isis interface all node-link-protection
```

4.    (Optional) Configure a 1000-millisecond time interval between the detection of a topology change and when the SPF algorithm runs.

```
[edit protocols]
user@P1# set isis spf-options delay 1000
```

5.    Configure MPLS to use both RSVP and LDP label-switched paths (LSPs).

```
[edit protocols]
user@P1# set mpls interface all
user@P1# set mpls interface fxp0.0 disable
user@P1# set rsvp interface all
user@P1# set rsvp interface fxp0.0 disable
user@P1# set ldp interface all
user@P1# set ldp interface fxp0.0 disable
```

6.    (Optional) For LDP, enable forwarding equivalence class (FEC) deaggregation, which results in faster global convergence.

```
[edit protocols]
user@P1# set ldp deaggregate
```

7. To enable Packet Forwarding Engine local repair, establish a policy that forces the routing protocol process to install all the next hops for a given route.

This policy ensures that the backup route is installed in the forwarding table used by the Packet Forwarding Engine to forward traffic to a given destination.

```
[edit policy-options policy-statement ecmp term 1]
user@P1# set then load-balance per-packet
```

8. Apply the policy to the forwarding table of the local router with the **export** statement.

```
[edit routing-options forwarding-table]
user@P1# set export ecmp
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P1# show interfaces
fe-1/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.6/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
      address 10.0.0.9/30;
    }
    family iso;
    family mpls;
  }
}
fe-1/2/2 {
  unit 0 {
    family inet {
      address 10.0.0.25/30;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.3.3/32;
    }
    family iso {
```

```
        address 49.0001.0010.0000.0303.00;
    }
}

user@P1# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
isis {
    spf-options delay 1000;
    interface all {
        node-link-protection;
        level 2 metric 10;
        level 1 disable;
    }
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        level 2 metric 0;
    }
}
ldp {
    deaggregate;
    interface all;
    interface fxp0.0 {
        disable;
    }
}

user@P1# show policy-options
policy-statement ecmp {
    term 1 {
        then {
            load-balance per-packet;
        }
    }
}

user@P1# show routing-options
forwarding-table {
    export ecmp;
}
```

If you are done configuring the device, enter **commit** from configuration mode.



**Verification**

Confirm that the configuration is working properly.

- [Checking the MPLS LSP Backup Path on page 3895](#)
- [Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node on page 3895](#)
- [Checking the Backup Coverage on page 3896](#)
- [Checking the Type of Protection Configured on page 3897](#)

**Checking the MPLS LSP Backup Path**

|                |   |
|----------------|---|
| <b>Purpose</b> | Display information about the MPLS label-switched-paths (LSPs) designated as the backup route for the IS-IS routes.   |
| <b>Action</b>  | <p>On Device PE1, from operational mode, enter the <b>show isis backup label-switched-path</b> command.</p> <pre> user@PE1&gt; show isis backup label-switched-path Backup MPLS LSPs: to-p2, Egress: 10.255.4.4, Status: up, Last change: 01:17:45 TE-metric: 19, Metric: 0, Refcount: 1 </pre> |
| <b>Meaning</b> | The output shows that the backup path is up and operational.  |

**Checking Which Next-Hop Neighbors Are Designated as Backup Paths to the Destination Node**

|                |   |
|----------------|---|
| <b>Purpose</b> | Display SPF calculations for each neighbor for a given destination.   |
| <b>Action</b>  | <p>On Device PE1, from operational mode, enter the <b>show isis backup spf results</b> command.</p> <pre> user@PE1&gt; show isis backup spf results  IS-IS level 1 SPF results: 0 nodes  IS-IS level 2 SPF results: PE2.00 Primary next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd Root: P2, Root Metric: 20, Metric: 10, Root Preference: 0x0 track-item: P2.00-00 Eligible, Backup next-hop: fe-1/2/1.0, LSP, to-p2 Root: P3, Root Metric: 10, Metric: 10, Root Preference: 0x0 Not eligible, Reason: Interface is already covered Root: P1, Root Metric: 10, Metric: 20, Root Preference: 0x0 track-item: P3.00-00 Not eligible, Reason: Interface is already covered P2.00 Primary next-hop: fe-1/2/1.0, IPV4, P1, SNPA: 0:5:85:8f:c8:bd Root: P2, Root Metric: 20, Metric: 0, Root Preference: 0x0 track-item: P2.00-00 Not eligible, Reason: Primary next-hop link fate sharing Root: P1, Root Metric: 10, Metric: 10, Root Preference: 0x0 Not eligible, Reason: Primary next-hop link fate sharing Root: P3, Root Metric: 10, Metric: 20, Root Preference: 0x0 </pre> |

```

        track-item: P1.00-00
        Not eligible, Reason: Primary next-hop node fate sharing
P3.00
    Primary next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
    Root: P2, Root Metric: 20, Metric: 20, Root Preference: 0x0
    track-item: P3.00-00
    track-item: P2.00-00
    track-item: P1.00-00
    Eligible, Backup next-hop: fe-1/2/1.0, LSP, to-p2
    Root: P3, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Not eligible, Reason: Interface is already covered
    Root: P1, Root Metric: 10, Metric: 10, Root Preference: 0x0
    track-item: P3.00-00
    Not eligible, Reason: Interface is already covered
P1.00
    Primary next-hop: fe-1/2/1.0, IPV4, P1, SNPA: 0:5:85:8f:c8:bd
    Root: P2, Root Metric: 20, Metric: 10, Root Preference: 0x0
    track-item: P2.00-00
    track-item: P1.00-00
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: P1, Root Metric: 10, Metric: 0, Root Preference: 0x0
    Not eligible, Reason: Primary next-hop link fate sharing
    Root: P3, Root Metric: 10, Metric: 10, Root Preference: 0x0
    track-item: P1.00-00
    Eligible, Backup next-hop: fe-1/2/2.0, IPV4, P3, SNPA: 0:5:85:8f:c8:bd
4 nodes

```

**Meaning** The output indicates whether a specific interface or node has been designated as a backup path and why.

### *Checking the Backup Coverage*

**Purpose** Check the percentage of protected nodes and prefixes.

**Action** From operational mode, enter the **show isis backup coverage** command.

```
user@PE1> show isis backup coverage
```

```
Backup Coverage:
```

| Topology     | Level | Node   | IPv4   | IPv6  | CLNS  |
|--------------|-------|--------|--------|-------|-------|
| IPV4 Unicast | 1     | 0.00%  | 0.00%  | 0.00% | 0.00% |
| IPV4 Unicast | 2     | 75.00% | 87.50% | 0.00% | 0.00% |

```
user@P1> show isis backup coverage
```

```
Backup Coverage:
```

| Topology     | Level | Node   | IPv4   | IPv6  | CLNS  |
|--------------|-------|--------|--------|-------|-------|
| IPV4 Unicast | 1     | 0.00%  | 0.00%  | 0.00% | 0.00% |
| IPV4 Unicast | 2     | 75.00% | 71.43% | 0.00% | 0.00% |

```
user@P2> show isis backup coverage
```

```
Backup Coverage:
```

| Topology     | Level | Node   | IPv4   | IPv6  | CLNS  |
|--------------|-------|--------|--------|-------|-------|
| IPV4 Unicast | 1     | 0.00%  | 0.00%  | 0.00% | 0.00% |
| IPV4 Unicast | 2     | 50.00% | 37.50% | 0.00% | 0.00% |

```
user@P3> show isis backup coverage
```

```
Backup Coverage:
```

| Topology     | Level | Node   | IPv4   | IPv6  | CLNS  |
|--------------|-------|--------|--------|-------|-------|
| IPV4 Unicast | 1     | 0.00%  | 0.00%  | 0.00% | 0.00% |
| IPV4 Unicast | 2     | 75.00% | 71.43% | 0.00% | 0.00% |

```

user@PE2> show isis backup coverage
Backup Coverage:
Topology      Level  Node   IPv4   IPv6   CLNS
IPv4 Unicast  1     0.00%  0.00%  0.00%  0.00%
IPv4 Unicast  2     50.00% 37.50%  0.00%  0.00%

```

**Meaning** The level of backup coverage available through IS-IS routes depends on the actual network topology and is typically less than 100 percent for all destinations on any given routing device. You can extend backup coverage to include RSVP LSPs.

### *Checking the Type of Protection Configured*

**Purpose** On all nodes in the IS-IS domain, check the percentage of protected nodes and prefixes.

**Action** From operational mode, enter the **show isis interface detail** command.

```

user@PE1> show isis interface detail

IS-IS interface database:
lo0.0
  Index: 76, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Adjacency advertisement: Advertise
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1           0       64      0 Passive
    2           0       64      0 Passive
fe-1/2/2.0
  Index: 79, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s
  Adjacency advertisement: Advertise
  Protection Type: Node Link
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    2           1       64      10    9.000    27 P3.03 (not us)
fe-1/2/1.0
  Index: 77, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s
  Adjacency advertisement: Advertise
  Protection Type: Node Link
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    2           1       64      10    9.000    27 P1.02 (not us)

```

**Meaning** The output shows that node-link protection is configured on the interfaces.

**Related Documentation**

- *Example: Configuring BFD for IS-IS*

## Example: Configuring an IS-IS Default Route Policy on Logical Systems

This example shows logical systems configured on a single physical router and explains how to configure a default route on one logical system.

- [Requirements on page 3898](#)
- [Overview on page 3898](#)

- [Configuration on page 3898](#)
- [Verification on page 3901](#)

## Requirements

No special configuration beyond device initialization is required before configuring this example.


## Overview

This example shows a logical system redistributing a default route to other logical systems. All logical systems are running IS-IS. A common reason for a default route is to provide a path for sending traffic destined outside the IS-IS domain.

In this example, the default route is not used for forwarding traffic. The **no-install** statement prevents the route from being installed in the forwarding table of Logical System LS3. If you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols. The **discard** statement silently drops packets without notice.

[Figure 110 on page 3898](#) shows the sample network.

**Figure 110: IS-IS Logical Systems with a Default Route to an ISP**



## Configuration

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set logical-systems LS3 interfaces lt-1/2/0 unit 3 description LS3->LS2
set logical-systems LS3 interfaces lt-1/2/0 unit 3 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 3 peer-unit 4
```

```

set logical-systems LS3 interfaces lt-1/2/0 unit 3 family inet address 10.0.2.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 3 family iso
set logical-systems LS3 interfaces lt-1/2/0 unit 5 description LS3->LS1
set logical-systems LS3 interfaces lt-1/2/0 unit 5 encapsulation ethernet
set logical-systems LS3 interfaces lt-1/2/0 unit 5 peer-unit 0
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family inet address 10.0.1.1/30
set logical-systems LS3 interfaces lt-1/2/0 unit 5 family iso
set logical-systems LS3 interfaces lo0 unit 3 family iso address 49.0001.1234.1600.2231.00
set logical-systems LS3 protocols isis export isis-default
set logical-systems LS3 protocols isis interface lt-1/2/0.3
set logical-systems LS3 protocols isis interface lt-1/2/0.5
set logical-systems LS3 protocols isis interface lo0.3 passive
set logical-systems LS3 routing-options static route 0.0.0.0/0 discard
set logical-systems LS3 routing-options static route 0.0.0.0/0 no-install
set logical-systems LS3 policy-options policy-statement isis-default from protocol static
set logical-systems LS3 policy-options policy-statement isis-default from route-filter
    0.0.0.0/0 exact
set logical-systems LS3 policy-options policy-statement isis-default then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an IS-IS default route policy on logical systems:

1. Configure the logical tunnel interfaces.

```

[edit logical-systems LS3 interfaces lt-1/2/0]
user@R1# set unit 3 description LS3->LS2
user@R1# set unit 3 encapsulation ethernet
user@R1# set unit 3 peer-unit 4
user@R1# set unit 3 family inet address 10.0.2.1/30
user@R1# set unit 3 family iso
user@R1# set unit 5 description LS3->LS1
user@R1# set unit 5 encapsulation ethernet
user@R1# set unit 5 peer-unit 0
user@R1# set unit 5 family inet address 10.0.1.1/30
user@R1# set unit 5 family iso
[edit logical-systems LS3 interfaces lo0 unit 3]
user@R1# set family iso address 49.0001.1234.1600.2231.00

```

2. Enable IS-IS on the interfaces.

```

[edit logical-systems LS3 protocols isis]
user@R1# set interface lt-1/2/0.3
user@R1# set interface lt-1/2/0.5
user@R1# set interface lo0.3 passive

```

3. Configure the default route on Logical System LS3.

```

[edit logical-systems LS3 routing-options]
user@R1# set static route 0.0.0.0/0 discard
user@R1# set static route 0.0.0.0/0 no-install

```

4. Configure the default route policy on Logical System LS3.

```

[edit logical-systems LS3 policy-options]
user@R1# set policy-statement isis-default from protocol static

```

```
user@R1# set policy-statement isis-default from route-filter 0.0.0.0/0 exact
user@R1# set policy-statement isis-default then accept
```

5. Apply the export policy to IS-IS on Logical System LS3.

```
[edit logical-systems LS3 protocols isis]
user@R1# set export isis-default
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@R1# commit
```

### Results

From configuration mode, confirm your configuration by issuing the **show logical-systems LS3** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show logical-systems LS3
interfaces {
  lt-1/2/0 {
    unit 3 {
      description LS3->LS2;
      encapsulation ethernet;
      peer-unit 4;
      family inet {
        address 10.0.2.1/30;
      }
      family iso;
    }
    unit 5 {
      description LS3->LS1;
      encapsulation ethernet;
      peer-unit 0;
      family inet {
        address 10.0.1.1/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 3 {
      family iso {
        address 49.0001.1234.1600.2231.00;
      }
    }
  }
}
protocols {
  isis {
    export isis-default;
    interface lt-1/2/0.3;
    interface lt-1/2/0.5;
    interface lo0.3 {
      passive;
    }
  }
}
```

```
    }  
  }  
}  
policy-options {  
  policy-statement isis-default {  
    from {  
      protocol static;  
      route-filter 0.0.0.0/0 exact;  
    }  
    then accept;  
  }  
}  
routing-options {  
  static {  
    route 0.0.0.0/0 {  
      discard;  
      no-install;  
    }  
  }  
}
```

### Verification

Confirm that the configuration is working properly.

#### *Verifying That the Static Route Is Redistributed*

**Purpose** Make sure that the IS-IS policy is working by checking the routing tables.

```

Action user@R1> show route logical-system LS3
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:00:45
                   Discard
10.0.0.0/30        *[IS-IS/15] 1w0d 10:14:14, metric 20
                   to 10.0.2.2 via lt-1/2/0.3
                   > to 10.0.1.2 via lt-1/2/0.5
10.0.1.0/30        *[Direct/0] 1w0d 10:15:18
                   > via lt-1/2/0.5
10.0.1.1/32        *[Local/0] 1w0d 10:15:18
                   Local via lt-1/2/0.5
10.0.2.0/30        *[Direct/0] 1w0d 10:15:18
                   > via lt-1/2/0.3
10.0.2.1/32        *[Local/0] 1w0d 10:15:18
                   Local via lt-1/2/0.3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1234.1600.2231/72
                   *[Direct/0] 1w0d 10:17:19
                   > via lo0.3

user@R1> show route logical-system LS2
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:01:38, metric 10
                   > to 10.0.2.1 via lt-1/2/0.4
10.0.0.0/30        *[Direct/0] 1w0d 10:16:11
                   > via lt-1/2/0.1
10.0.0.2/32        *[Local/0] 1w0d 10:16:11
                   Local via lt-1/2/0.1
10.0.1.0/30        *[IS-IS/15] 1w0d 10:15:07, metric 20
                   > to 10.0.0.1 via lt-1/2/0.1
                   to 10.0.2.1 via lt-1/2/0.4
10.0.2.0/30        *[Direct/0] 1w0d 10:16:11
                   > via lt-1/2/0.4
10.0.2.2/32        *[Local/0] 1w0d 10:16:11
                   Local via lt-1/2/0.4

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.2002/72
                   *[Direct/0] 1w0d 10:18:12
                   > via lo0.2

user@R1> show route logical-system LS1
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[IS-IS/160] 00:02:01, metric 10
                   > to 10.0.1.1 via lt-1/2/0.0
10.0.0.0/30        *[Direct/0] 1w0d 10:16:34
                   > via lt-1/2/0.2
10.0.0.1/32        *[Local/0] 1w0d 10:16:34
                   Local via lt-1/2/0.2

```



```

10.0.1.0/30      *[Direct/0] 1w0d 10:16:34
                  > via lt-1/2/0.0
10.0.1.2/32      *[Local/0] 1w0d 10:16:34
                  Local via lt-1/2/0.0
10.0.2.0/30      *[IS-IS/15] 1w0d 10:15:55, metric 20
                  to 10.0.1.1 via lt-1/2/0.0
                  > to 10.0.0.2 via lt-1/2/0.2

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

49.0001.1720.1600.1001/72
                  *[Direct/0] 1w0d 10:18:35
                  > via lo0.1

```

**Meaning** The routing table on Logical System LS3 contains the default 0.0.0.0/0 route from protocol **Static**. The routing tables on Logical System LS1 and Logical System LS2 contain the default 0.0.0.0/0 route from protocol **IS-IS**. If Logical System LS1 and Logical System LS2 receive packets destined for networks not specified in their routing tables, those packets will be sent to Logical System LS3 for further processing. This configuration assumes that Logical System LS3 has a connection to an ISP or another external network.

**Related Documentation**

- [Example: Creating an Interface on a Logical System](#)

## Example: Configuring IS-IS for CLNS

- [Understanding IS-IS for CLNS on page 3903](#)
- [Example: Configuring IS-IS for CLNS on page 3903](#)

### Understanding IS-IS for CLNS

IS-IS extensions provide the basic interior gateway protocol (IGP) support for collecting intradomain routing information for Connectionless Network Service (CLNS) destinations within a CLNS network. Routers that learn host addresses through End System-to-Intermediate System (ES-IS) can advertise the addresses to other routers (intermediate systems) by using IS-IS.

For more information about IS-IS, see the ISO 10589 standard.

### Example: Configuring IS-IS for CLNS

This example shows how to create a routing instance and enable the IS-IS protocol on all interfaces.

- [Requirements on page 3903](#)
- [Overview on page 3904](#)
- [Configuration on page 3904](#)
- [Verification on page 3905](#)

#### Requirements

Before you begin, configure the network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.

### Overview

The configuration instructions in this topic describe how to create a routing instance called `aaaa`, enable IS-IS on all interfaces, define the BGP export policy name (`dist-bgp`), family (ISO), and protocol (BGP), and apply the export policy to IS-IS.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set routing-instances aaaa protocols isis clns-routing
set routing-instances aaaa protocols isis interface all
set routing-instances aaaa protocols isis no-ipv4-routing no-ipv6-routing
set policy-options policy-statement dist-bgp from family iso protocol bgp
set policy-options policy-statement dist-bgp then accept
set routing-instances aaaa protocols isis export dist-bgp
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS for CLNS:

1. Enable CLNS routing.  

```
[edit routing-instances aaaa]
user@host# set protocols isis clns-routing
```
2. Enable IS-IS on all interfaces.  

```
[edit routing-instances aaaa]
user@host# set protocols isis interface all
```
3. (Optional) Disable IPv4 and IPv6 routing to configure a pure CLNS network.  

```
[edit routing-instances aaaa]
user@host# set protocols isis no-ipv4-routing no-ipv6-routing
```
4. Define the BGP export policy name, family, and protocol.  

```
[edit policy-options]
user@host# set policy-statement dist-bgp from family iso protocol bgp
```
5. Define the action for the export policy.  

```
[edit policy-options]
user@host# set policy-statement dist-bgp then accept
```
6. Apply the export policy to IS-IS.  

```
[edit routing-instances aaaa]
user@host# set protocols isis export dist-bgp
```

**Results** From configuration mode, confirm your configuration by entering the **show routing-instances** and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-instances
aaaa {
  protocols {
    isis {
      export dist-bgp;
      no-ipv4-routing;
      no-ipv6-routing;
      clns-routing;
      interface all;
    }
  }
}

user@host# show policy-options
policy-statement dist-bgp {
  from {
    family iso;
    protocol bgp;
  }
  then accept;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the ISO Routes on page 3905](#)
- [Checking the SPF Calculations on page 3905](#)

### **Verifying the ISO Routes**

**Purpose** Verify that the expected ISO routes are displayed in the IS-IS routing table.

**Action** From operational mode, enter the **show isis route** command.

### **Checking the SPF Calculations**

**Purpose** Display information about IS-IS shortest-path-first (SPF) calculations.

**Action** From operational mode, enter the **show isis spf** command.

## **Example: Configuring IS-IS Designated Routers**

- [Understanding IS-IS Designated Routers on page 3905](#)
- [Example: Configuring Designated Router Election Priority for IS-IS on page 3906](#)

### **Understanding IS-IS Designated Routers**

A router advertises its priority to become a designated router in its hello packets. On all multiaccess networks (physical networks that support the attachment of more than two

routers, such as Ethernet networks), IS-IS uses the advertised priorities to elect a designated router for the network. This router is responsible for sending network link-state advertisements, which describe all the routers attached to the network. These advertisements are flooded throughout a single area. The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

A router's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127, which you configure on the IS-IS interface. The router with the highest priority becomes the designated router for the area (Level 1, Level 2, or both), also configured on the IS-IS interface. If routers in the network have the same priority, then the router with the highest MAC address is elected as the designated router. By default, routers have a priority value of 64.

---

### Example: Configuring Designated Router Election Priority for IS-IS

---

This example shows how to configure the designated router election priority for IS-IS.

Before you begin:

- Configure network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Enable IS-IS on the interfaces. See [“Example: Configuring IS-IS” on page 3823](#).

In this example, you configure the priority for logical interface ge-0/0/1.0 to be 100 and the level number to be 1. If this interface has the highest priority value, the router becomes the designated router for the Level 1 area.

To configure a designated router election priority for IS-IS:

```
[edit]
user@host# set protocols isis interface ge-0/0/1.0 level 1 priority 100
```

#### Related Documentation

- [Example: Configuring IS-IS](#)

### Example: Enabling Packet Checksums on IS-IS Interfaces

This example shows how to enable packet checksums for IS-IS interfaces.

- [Requirements on page 3906](#)
- [Overview on page 3906](#)
- [Configuration on page 3907](#)
- [Verification on page 3908](#)

---

#### Requirements

---

Before you begin, configure IS-IS on both routers. See [“Example: Configuring IS-IS” on page 3823](#) for information about the sample IS-IS configuration.

---

#### Overview

---


Junos OS supports IS-IS checksums as documented in RFC 3358, *Optional Checksums in Intermediate System to Intermediate System (ISIS)*.

IS-IS protocol data units (PDUs) include link-state PDUs, complete sequence number PDUs (CSNPs), partial sequence number PDUs (PSNPs), and IS-IS hello (IIH) packets. These PDUs can be corrupt due to faulty implementations of Layer 2 hardware or lack of checksums on a specific network technology. Corruption of length or type, length, and value (TLV) fields can lead to the generation of extensive numbers of empty link-state PDUs in the receiving node. Because authentication is not a replacement for a checksum mechanism, you might want to enable the optional checksum TLV on your IS-IS interfaces.

The checksum cannot be enabled with MD5 hello authentication on the same interface.

Figure 111 on page 3907 shows the topology used in this example.

**Figure 111: IS-IS Checksum Topology**



This example describes the steps on Device R1.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set protocols isis traceoptions file isis
set protocols isis traceoptions flag all
set protocols isis interface fe-1/2/0.1 checksum
  
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IS-IS checksums:

1. Enable checksums.
 

```

[edit protocols isis interface fe-1/2/0.1]
user@R1# set checksum
      
```
2. (Optional) Enable tracing for tracking checksum operations.
 

```

[edit protocols isis traceoptions]
user@R1# set file isis
user@R1# set flag all
      
```

**Results** From configuration mode, confirm your configuration by entering the **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
isis {
  traceoptions {
    file isis;
    flag all;
  }
  interface fe-1/2/0.1 {
    checksum;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

---

### Verification

Confirm that the configuration is working properly.

#### Verifying Checksums

**Purpose** Verify that checksums are performed.

**Action** From operational mode, enter the **show log isis | match checksum** command.

```
user@R1> show log isis | match checksum
```

```
May 31 16:47:39.513267      sequence 0x49 checksum 0x8e64
May 31 16:47:39.513394      sequence 0x4e checksum 0x34b3
May 31 16:47:39.513517      sequence 0x50 checksum 0x9dcb
May 31 16:47:46.563781      sequence 0x45 checksum 0x7e1a
May 31 16:47:46.563970      sequence 0x46 checksum 0x226d
May 31 16:47:46.564104      sequence 0x52 checksum 0x99cd
May 31 16:47:46.581087      sequence 0x49 checksum 0x8e64
May 31 16:47:46.581222      sequence 0x4e checksum 0x34b3
May 31 16:47:46.581353      sequence 0x50 checksum 0x9dcb
May 31 16:47:55.799090      sequence 0x45 checksum 0x7e1a
May 31 16:47:55.799223      sequence 0x46 checksum 0x226d
May 31 16:47:55.799347      sequence 0x52 checksum 0x99cd
May 31 16:47:55.818255      sequence 0x49 checksum 0x8e64
May 31 16:47:55.818473      sequence 0x4e checksum 0x34b3
May 31 16:47:55.818606      sequence 0x50 checksum 0x9dcb
May 31 16:48:03.455816      sequence 0x49 checksum 0x8e64
May 31 16:48:03.455973      sequence 0x4e checksum 0x34b3
```

**Meaning** The output shows that checksum information is captured in the IS-IS trace log file.

**Related Documentation**

- [Understanding Checksums on IS-IS Interfaces](#)

---

## Configuration Tasks

- [Configuring IS-IS Authentication on page 3909](#)
- [Configuring Authentication Without Network-Wide Deployment on page 3910](#)

## Configuring IS-IS Authentication

All IS-IS protocol exchanges can be authenticated to guarantee that only trusted routing devices participate in the autonomous system (AS) routing. By default, IS-IS authentication is disabled on the routing device.

To configure IS-IS authentication, you must define an authentication password and specify the authentication type.

You can configure one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving routing device uses an authentication key (password) to verify the packet. Simple authentication is included for compatibility with existing IS-IS implementations. However, we recommend that you do *not* use this authentication method because it is insecure (the text can be “sniffed”).



**CAUTION:** A simple password that exceeds 254 characters is truncated.

- HMAC-MD5 authentication—Uses an iterated cryptographic hash function. The receiving routing device uses an authentication key (password) to verify the packet.

You can also configure more fine-grained interface-level authentication for hello packets.

To enable authentication and specify an authentication method, include the **authentication-type** statement, specifying the **simple** or **md5** authentication type:

**authentication-type** *authentication*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure a password, include the **authentication-key** statement. The authentication password for all routing devices in a domain must be the same.

**authentication-key** *key*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

To configure hitless authentication key rollover, include the **authentication-key-chain (Protocols IS-IS)** statement.

The password can contain up to 255 characters. If you include spaces, enclose all characters in quotation marks (" ").

If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces that are shared with a Junos OS implementation.

Authentication of hello packets, partial sequence number PDU (PSNP), and complete sequence number PDU (CSNP) can be suppressed to enable interoperability with the

routing software of different vendors. Different vendors handle authentication in various ways, and suppressing authentication for different PDU types might be the simplest way to allow compatibility within the same network.

To configure IS-IS to generate authenticated packets, but not to check the authentication on received packets, include the **no-authentication-check** statement:

**no-authentication-check;**

To suppress authentication of IS-IS hello packets, include the **no-hello-authentication** statement:

**no-hello-authentication;**

To suppress authentication of PSNPs, include the **no-psnp-authentication** statement:

**no-psnp-authentication;**

To suppress authentication of CSNPs, include the **no-csnp-authentication** statement:

**no-csnp-authentication;**

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.



**NOTE:** The **authentication** and the **no-authentication** statements must be configured at the same hierarchy level. Configuring authentication at the [edit protocols isis interface *interface-name*] hierarchy level and configuring **no-authentication** at the [edit protocols isis] hierarchy level has no effect.

---

**Related  
Documentation**

- [Configuring Authentication Without Network-Wide Deployment on page 3910](#)

## Configuring Authentication Without Network-Wide Deployment

To allow the use of authentication without requiring network-wide deployment, include the **loose-authentication-check** statement:

**loose-authentication-check;**

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

**Related  
Documentation**

- [Example: Configuring Hitless Authentication Key Rollover for IS-IS](#)

## Configuration Statements

---

- [authentication-key \(Protocols IS-IS\) on page 3913](#)
- [authentication-key-chain \(Protocols IS-IS\) on page 3914](#)
- [authentication-type \(Protocols IS-IS\) on page 3915](#)
- [bfd-liveness-detection \(Protocols IS-IS\) on page 3916](#)



- [checksum \(Protocols IS-IS\) on page 3918](#)
- [csnp-interval on page 3919](#)
- [disable \(Protocols IS-IS\) on page 3920](#)
- [export \(Protocols IS-IS\) on page 3921](#)
- [external-preference \(Protocols IS-IS\) on page 3922](#)
- [family \(Protocols IS-IS\) on page 3923](#)
- [graceful-restart \(Protocols IS-IS\) on page 3924](#)
- [hello-authentication-key on page 3925](#)
- [hello-authentication-key-chain on page 3926](#)
- [hello-authentication-type on page 3927](#)
- [hello-interval \(Protocols IS-IS\) on page 3928](#)
- [hello-padding on page 3929](#)
- [hold-time \(Protocols IS-IS\) on page 3931](#)
- [ignore-attached-bit on page 3932](#)
- [interface \(Protocols IS-IS\) on page 3933](#)
- [ipv4-multicast on page 3935](#)
- [ipv4-multicast-metric on page 3936](#)
- [ipv6-multicast on page 3936](#)
- [ipv6-multicast-metric on page 3937](#)
- [ipv6-unicast on page 3938](#)
- [ipv6-unicast-metric on page 3939](#)
- [isis on page 3940](#)
- [level \(Global IS-IS\) on page 3941](#)
- [link-protection \(Protocols IS-IS\) on page 3942](#)
- [loose-authentication-check on page 3942](#)
- [lsp-interval on page 3943](#)
- [lsp-lifetime on page 3944](#)
- [max-areas on page 3945](#)
- [mesh-group \(Protocols IS-IS\) on page 3946](#)
- [metric \(Protocols IS-IS\) on page 3947](#)
- [no-adjacency-holddown on page 3948](#)
- [no-authentication-check on page 3949](#)
- [no-csnp-authentication on page 3949](#)
- [no-eligible-backup \(Protocols IS-IS\) on page 3950](#)
- [no-hello-authentication on page 3950](#)
- [no-ipv4-multicast on page 3951](#)
- [no-ipv4-routing on page 3952](#)

- [no-ipv6-multicast](#) on page 3953
- [no-ipv6-routing](#) on page 3954
- [no-ipv6-unicast](#) on page 3955
- [no-psnp-authentication](#) on page 3955
- [no-unicast-topology](#) on page 3956
- [node-link-protection](#) (Protocols IS-IS) on page 3956
- [overload](#) (Protocols IS-IS) on page 3957
- [passive](#) (Protocols IS-IS) on page 3960
- [point-to-point](#) on page 3961
- [preference](#) (Protocols IS-IS) on page 3962
- [prefix-export-limit](#) (Protocols IS-IS) on page 3963
- [priority](#) (Protocols IS-IS) on page 3964
- [reference-bandwidth](#) (Protocols IS-IS) on page 3965
- [rib-group](#) (Protocols IS-IS) on page 3966
- [spf-options](#) (Protocols IS-IS) on page 3967
- [topologies](#) (Protocols IS-IS) on page 3968
- [traceoptions](#) (Protocols IS-IS) on page 3969
- [traffic-engineering](#) (Protocols IS-IS) on page 3972
- [wide-metrics-only](#) on page 3975

## authentication-key (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | authentication-key <i>key</i> ;  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis <b>level</b> <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>level</b> <i>level-number</i> ],<br>[edit protocols isis <b>level</b> <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i> ]   |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>         | <p>Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the <b>authentication-type</b> statement.</p> <p>All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Juniper Networks routing device.</p> |
| <b>Default</b>             | If you do not include this statement and the <b>authentication-type</b> statement, IS-IS authentication is disabled.   |
| <b>Options</b>             | <b>key</b> —Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").   |



**CAUTION:** A simple password for authentication is truncated if it exceeds 254 characters.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> </ul> |

## authentication-key-chain (Protocols IS-IS)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | authentication-key-chain <i>key-chain-name</i> ;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>name</i> protocols isis level <i>level-number</i> ],<br>[edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis level <i>level-number</i> ],<br>[edit protocols isis level <i>level-number</i> ],<br>[edit routing-instances <i>instance-name</i> protocols isis level <i>level-number</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Apply and enable an authentication keychain to the routing device.   |
| <b>Options</b>                  | <b>key-chain</b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837</a></li><li>• <a href="#">Example: Configuring Route Authentication for BGP on page 3572</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Understanding Hitless Authentication Key Rollover for IS-IS on page 3818</a></li></ul> |

## authentication-type (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>authentication-type <i>authentication</i>;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit protocols isis <a href="#">level level-number</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the <b>authentication-key</b> statement.  |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, IS-IS authentication is disabled.   |
| <b>Options</b>                  | <p><b><i>authentication</i></b>—Authentication scheme:</p> <ul style="list-style-type: none"> <li>• <b>md5</b>—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</li> <li>• <b>simple</b>—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you <i>not</i> use this authentication method.</li> </ul> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> <li>• <a href="#">authentication-key on page 3913</a></li> <li>• <a href="#">no-authentication-check on page 3949</a></li> </ul>  |

## bfd-liveness-detection (Protocols IS-IS)

---

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>bfd-liveness-detection {   authentication {     algorithm <i>algorithm-name</i>;     key-chain <i>key-chain-name</i>;     loose-check;   }   detection-time {     threshold <i>milliseconds</i>;   }   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   multiplier <i>number</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   }   version (1   automatic); }</pre>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a>],<br/>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/>isis <a href="#">interface interface-name</a>],<br/>[edit protocols isis <a href="#">interface interface-name</a>],<br/>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.<br/>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/><b>detection-time threshold</b> and <b>transmit-interval threshold</b> options added in Junos OS Release 8.2.<br/>Support for logical systems introduced in Junos OS Release 8.3.<br/><b>no-adaptation</b> statement introduced in Junos OS Release 9.0.<br/><b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> options introduced in Junos OS Release 9.6.<br/>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.   |
| <b>Options</b>             | <p><b>authentication algorithm <i>algorithm-name</i></b>—Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain <i>key-chain-name</i></b>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the <b>[edit security]</b> hierarchy level.</p> <p><b>authentication loose-check</b>—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication might not be configured at both ends of the BFD session.</p> |

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

**Range:** 1 through 255,000

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version)

**Default:** automatic

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring BFD for IS-IS</i></li> <li>• <i>Example: Configuring BFD Authentication for IS-IS</i></li> </ul> |
|------------------------------|---|

## checksum (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | checksum;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Enable checksums for packets on this interface.<br><br>Junos OS supports IS-IS checksums as documented in RFC 3358, <i>Optional Checksums in Intermediate System to Intermediate System (ISIS)</i> .<br><br>The checksum cannot be enabled with MD5 hello authentication on the same interface.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i></li></ul>   |




## csnp-interval

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>csnp-interval (seconds   disable);</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>]</p>                             |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure the interval between complete sequence number PDUs (CSNPs) on a LAN interface.</p> <p>If the routing device is the designated router on a LAN, IS-IS sends CSN packets every 10 seconds. If the routing device is on a point-to-point interface, it sends CSN packets every 5 seconds multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.</p> <p>To configure the interface not to send any CSNPs, specify the <b>disable</b> option.</p> |
| <b>Default</b>                  | By default, IS-IS sends CSNPs periodically. If the routing device is the designated router on a LAN, IS-IS sends CSNPs every 10 seconds. If the routing device is on a point-to-point interface, it sends CSNPs every 5 seconds multiplied by the number of IS-IS adjacencies over point-to-point links, which are in UP state.  |
| <b>Options</b>                  | <p><b>disable</b>—Do not send CSNPs on this interface.</p> <p><b>seconds</b>—Number of seconds between the sending of CSNPs.</p> <p><b>Range:</b> 1 through 65,535 seconds</p> <p><b>Default:</b> 10 seconds on LAN broadcast links. 5 seconds on point-to-point links.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces</i></li> </ul>  |

## disable (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | disable;   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <b>isis</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>traffic-engineering</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>traffic-engineering</b>],</p> <p>[edit protocols <b>isis</b>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis <b>traffic-engineering</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>traffic-engineering</b>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Disable IS-IS on the routing device, on an interface, or on a level.</p> <p>At the <b>[edit protocols isis traffic-engineering]</b> hierarchy level, disable IS-IS support for traffic engineering.</p> <p>Enabling IS-IS on an interface (by including the <b>interface</b> statement at the <b>[edit protocols isis]</b> or the <b>[edit routing-instances routing-instance-name protocols isis]</b> hierarchy level), disabling it (by including the <b>disable</b> statement), and not actually having IS-IS run on an interface (by including the <b>passive</b> statement) are mutually exclusive states.</p>   |
| <b>Default</b>                  | <p>IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which <b>family iso</b> is enabled.</p> <p>IS-IS support for traffic engineering is enabled.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multi-Level IS-IS on page 3829</a></li> <li>• <a href="#">IS-IS Overview on page 3812</a></li> </ul>   |

## export (Protocols IS-IS)

|   |   |
|---|---|
| <b>Syntax</b>   | <code>export [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>  | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>  |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>  | <p>Apply one or more policies to routes being exported from the routing table into IS-IS.</p> <p>All routing protocols store the routes that they learn in the routing table. The routing table uses this collected route information to determine the active routes to destinations. The routing table then installs the active routes into its forwarding table and exports them into the routing protocols. It is these exported routes that the protocols advertise.</p> <p>For each protocol, you control which routes the protocol stores in the routing table and which routes the routing table exports into the protocol from the routing table by defining a <i>routing policy</i> for that protocol.</p> |
| <div>  <p><b>NOTE:</b> For IS-IS, you cannot apply routing policies that affect how routes are imported into the routing table; doing so with a link-state protocol can easily lead to an inconsistent topology database.</p> </div> |   |
| <b>Options</b>  | <i>policy-names</i> —Name of one or more policies.  |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">Example: Redistributing OSPF Routes into IS-IS</a></li> <li>• <a href="#">Example: Configuring an IS-IS Default Route Policy on Logical Systems on page 3897</a></li> </ul>  |

## external-preference (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>external-preference <i>preference</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure the preference of external routes.  |
| <b>Options</b>                  | <i>preference</i> —Preference value.<br><b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )<br><b>Default:</b> 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Route Preferences Overview</i></li><li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li><li>• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i></li><li>• <a href="#">preference on page 3962</a></li></ul>  |

## family (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>family inet {   shortcuts {     multicast-rpf-routes;   } } family inet6 {   shortcuts; }</pre>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis traffic-engineering],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering],</p> <p>[edit protocols isis traffic-engineering],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis traffic-engineering]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3.</p> <p>Support for IPv6 for IGP shortcuts introduced in Junos OS Release 9.3.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Configure the address family for traffic engineering IS-IS interior gateway protocol (IGP) shortcuts.  |
| <b>Options</b>                  | <p>inet—IPv4 address family</p> <p>inet6—IPv6 address family</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>•</li> </ul>  |

## graceful-restart (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>graceful-restart {<br/>    disable;<br/>    helper-disable;<br/>    restart-duration <i>seconds</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | <p>Configure graceful restart parameters for IS-IS.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the <b>[edit routing-options]</b> hierarchy level. When graceful restart is enabled, the restarting routing device is not removed from the network topology during the restart period. The adjacencies are reestablished after restart is complete.</p> <p>On LAN interfaces where IS-IS is configured on a transit router that serves as the designated router (DR), a graceful restart causes:</p> <ul style="list-style-type: none"><li>• The ingress router of the label-switched path (LSP), which passes through the DR, to break the LSP.</li><li>• The ingress router to re-signal the LSP.</li></ul> |
| <b>Options</b>                  | <p><b>disable</b>—Disable graceful restart for IS-IS.</p> <p><b>helper-disable</b>—Disable graceful restart helper capability. Helper mode is enabled by default.</p> <p><b>restart-duration <i>seconds</i></b>—Time period for the restart to last, in seconds.<br/><b>Range:</b> 30 through 300 seconds<br/><b>Default:</b> 30 seconds</p>  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Routing Protocols Graceful Restart on page 2261</a></li></ul>   |

## hello-authentication-key

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>hello-authentication-key password;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the <b>hello-authentication-type</b> statement.   |
| <b>Default</b>                  | By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.   |
| <b>Options</b>                  | <b>password</b> —Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">authentication-key on page 3913</a></li> <li>• <a href="#">authentication-type on page 3915</a></li> <li>• <a href="#">hello-authentication-type on page 3927</a></li> </ul>   |


## hello-authentication-key-chain

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | hello-authentication-key-chain <i>key-chain-name</i> ;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.2.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Apply an authentication keychain to the IS-IS interface.  |
| <b>Options</b>                  | <i>key-chain-name</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li></ul>   |



## hello-authentication-type

|   |  |
|---|--|
| <b>Syntax</b>   | hello-authentication-type (md5   simple);  |
| <b>Hierarchy Level</b>  | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>number</i>]</p> |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>  | <p>Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the <b>hello-authentication-key</b> statement.</p> <p>You can configure authentication for a given IS-IS level on an interface. On a point-to-point link, if you enable hello authentication for both IS-IS levels, the password configured for Level 1 is used for both levels.</p>   |
| <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>CAUTION:</b> If no authentication is configured for Level 1 on a point-to-point link with both levels enabled, the hello packets are sent without any password, regardless of the Level 2 authentication configurations.</p> </div> </div> |  |
| <b>Default</b>  | By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.  |
| <b>Options</b>  | <p><b>md5</b>—Specifies Message Digest 5 as the packet verification type.</p> <p><b>simple</b>—Specifies simple authentication as the packet verification type.</p>  |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">authentication-key on page 3913</a></li> <li>• <a href="#">authentication-type on page 3915</a></li> <li>• <a href="#">hello-authentication-key on page 3925</a></li> </ul>   |

## hello-interval (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>hello-interval <i>seconds</i>;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Modify the frequency with which the routing device sends hello packets out of an interface, in seconds.</p> <p>Routing devices send hello packets at a fixed interval on all interfaces to establish and maintain neighbor relationships. This interval is advertised in the hello interval field in the hello packet.</p> <p>You can send out hello packets in subsecond intervals. To send out hello packets every 333 milliseconds, set the <b>hello-interval</b> value to 1.</p>   |
| <b>Options</b>                  | <b><i>seconds</i></b> —Frequency of transmission for hello packets.<br><b>Range:</b> 1 through 20,000 seconds<br><b>Default:</b> 3 seconds (for designated intermediate system [DIS] routers), 9 seconds (for non-DIS routers)  |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><i>hold-time</i></li></ul>  |

## hello-padding

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | hello-padding (adaptive   disable   loose   strict);  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]  |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>         | <p>Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.</p> <p>This helps to prevent a premature adjacency Up state when one routing device's MTU does not meet the requirements to establish the adjacency.</p> <p>As an OSI Layer 2 protocol, IS-IS does not support data fragmentation. Therefore, maximum packet sizes must be established and supported between two routers. During adjacency establishment, the IS-IS protocol makes sure that the link supports a packet size of 1492 bytes by padding outgoing hello packets up to the maximum packet size of 1492 bytes.</p> <p>This is the default behavior of the Junos OS IS-IS implementation. However, Junos OS provides an option to disable hello padding that can override this behavior.</p> <p>There are four types of hello padding:</p> <ul style="list-style-type: none"> <li>Adaptive padding—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state type, length, and value (TLV) tuple. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. Adaptive padding has more overhead than loose padding and is able to detect MTU asymmetry from one side of the connection. This one-sided detection can result in generation of extra link-state PDUs that are flooded throughout the network. Specify the <b>adaptive</b> option to configure enough padding to establish an adjacency to neighbors.</li> <li>Disabled padding—Padding is disabled on all types of interfaces for all adjacency states. Specify the <b>disable</b> option to accommodate interfaces that support less than the default packet size of 1492 bytes.</li> <li>Loose padding (the default)—The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. Loose padding might not be able to detect certain situations such as asymmetrical MTUs between the routing devices. Specify the <b>loose</b> option to configure enough padding to initialize an adjacency to neighbors.</li> </ul> |

- **Strict padding**—Padding is done on all interface types and for all adjacency states, and is continuous. Strict padding has the most overhead. The advantage is that strict padding detects MTU issues on both sides of a link. Specify the **strict** option to configure padding to allow all adjacency states with neighbors.

**Options**    **adaptive**—Configure padding until the neighbor adjacency is established and active.

**disable**—Disable padding on all types of interfaces for all adjacency states.

**loose**—Configure padding until the state of the adjacency is initialized.

**strict**—Configure padding for all adjacency states.

**Required Privilege Level**    routing—To view this statement in the configuration.  
   routing-control—To add this statement to the configuration.

**Related Documentation**    • *Example: Configuring IS-IS*

## hold-time (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>hold-time seconds;</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receive another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.</p> <p>The hold time specifies how long a neighbor should consider this routing device to be operative without receiving another hello packet. If the neighbor does not receive a hello packet from this routing device within the hold time, it marks the routing device as being unavailable.</p> <p>For systems configured with graceful routing switchover (GRES) with Graceful Restart, the hold time for Master and Backup Routing Engines should be set to a value higher than 40 seconds. This ensures that adjacencies between the Routing Engine and the neighboring peer 'helper' routers do not time out, stopping graceful restart, and all traffic.</p> |
| <b>Options</b>                  | <p><b>seconds</b>—Hold-time value, in seconds.</p> <p><b>Range:</b> 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds</p> <p><b>Default:</b> 9 seconds (for designated intermediate system [DIS] routers), 27 seconds (for non-DIS routers; three times the default hello interval)</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Graceful Routing Engine Switchover on page 2268</a></li> <li>• <a href="#">Example: Configuring IS-IS</a></li> <li>• <a href="#">Example: Configuring IS-IS for GRES with Graceful Restart</a></li> <li>• <a href="#">hello-interval on page 3928</a></li> </ul>   |

## ignore-attached-bit

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ignore-attached-bit;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | <p>Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement enables the routing device to ignore the attached bit on incoming Level 1 link-state PDUs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.</p> <p>There might be times, such as during a denial-of-service (DoS) attack, that you do not want a Level 1 router to be able to forward traffic based on a default route.</p> <p>To prevent a routing device from being able to reach interarea destinations, you can prevent the routing device from installing the default route without affecting the status of its IS-IS adjacencies. The <b>ignore-attached-bit</b> statement is used to tell the routing device to ignore the presence of the attached bit in Level 1 link-state PDUs, which blocks the installation of the IS-IS default route.</p> |
| <b>Default</b>                  | The <b>ignore-attached-bit</b> statement is disabled by default.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>•</li></ul>  |

## interface (Protocols IS-IS)

```

Syntax  interface (all | interface-name) {
        disable;
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            transmit-interval {
                threshold milliseconds;
                minimum-interval milliseconds;
            }
            multiplier number;
        }
        checksum;
        csnp-interval (seconds | disable);
        hello-padding (adaptive | loose | strict);
        ldp-synchronization {
            disable;
            hold-time seconds;
        }
        lsp-interval milliseconds;
        mesh-group (value | blocked);
        no-adjacency-holddown;
        no-ipv4-multicast;
        no-ipv6-multicast;
        no-ipv6-unicast;
        no-unicast-topology;
        passive;
        point-to-point;
        level level-number {
            disable;
            hello-authentication-key key;
            hello-authentication-key-chain key-chain-name;
            hello-authentication-type authentication;
            hello-interval seconds;
            hold-time seconds;
            ipv4-multicast-metric metric;
            ipv6-multicast-metric metric;
            ipv6-unicast-metric metric;
            metric metric;
            passive;
            priority number;
            te-metric metric;
        }
    }

```

|                                 |   |
|---------------------------------|---|
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <b>isis</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>isis</b> ],<br>[edit protocols <b>isis</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure interface-specific IS-IS properties. To configure more than one interface, include the <b>interface</b> statement multiple times.</p> <p>Enabling IS-IS on an interface (by including the <b>interface</b> statement at the [edit protocols <b>isis</b>] or the [edit routing-instances <i>routing-instance-name</i> protocols <b>isis</b>] hierarchy level), disabling it (by including the <b>disable</b> statement), and not actually having IS-IS run on an interface (by including the <b>passive</b> statement) are mutually exclusive states.</p> |
| <b>Options</b>                  | <p><b>all</b>—Have Junos OS create IS-IS interfaces automatically. If you include this option, disable IS-IS on the management interface (fxp0).</p> <p><b>interface-name</b>—Name of an interface. Specify the full interface name, including the physical and logical address components.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring IS-IS</i></li><li>• <i>Example: Configuring Multi-Level IS-IS</i></li></ul>   |



## ipv4-multicast

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | ipv4-multicast;  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>         | Configure alternate IPv4 multicast topologies.   |



**NOTE:** The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.

|                                 |  |
|---------------------------------|--|
| <b>Default</b>                  | Multicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li> </ul> |

## ipv4-multicast-metric

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ipv4-multicast-metric <i>metric</i> ;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Specify the multicast topology metric value for the level.  |
| <b>Options</b>                  | <i>metric</i> —Metric value.<br><b>Range:</b> 0 through 16,777,215  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li></ul>  |

## ipv6-multicast

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ipv6-multicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>topologies</b> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>topologies</b> ],<br>[edit protocols isis <b>topologies</b> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>topologies</b> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | Configure alternate IPv6 multicast topologies.  |
| <b>Default</b>                  | Multicast topologies are disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li></ul>  |


## ipv6-multicast-metric

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>ipv6-multicast-metric <i>metric</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>   |
| <b>Description</b>              | Specify the IPv6 alternate multicast topology metric value for the level.  |
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 0 through 16,777,215</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li> </ul>   |

## ipv6-unicast

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | ipv6-unicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ],<br>[edit protocols isis <a href="#">topologies</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">topologies</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | <p>Configure alternate IPv6 unicast topologies.</p> <p>This statement causes IS-IS to calculate an alternate IPv6 unicast topology, in addition to the normal IPv4 unicast topology, and add the corresponding routes to inet6.0.</p>   |
|                                 | <div> <b>NOTE:</b> The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics. You can also selectively disable interfaces from participating in the IPv6 topology while continuing to participate in the IPv4 topology. This lets you exercise control over the paths that unicast data takes through a network.</div> |
| <b>Default</b>                  | IPv6 unicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859</a></li></ul>  |

## ipv6-unicast-metric

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>ipv6-unicast-metric <i>metric</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>   |
| <b>Description</b>              | Specify the IPv6 unicast topology metric value for the level. The IS-IS interface metrics for the IPv4 topology can be configured independently of the IPv6 metrics.   |
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 0 through 16,777,215</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859</a></li> </ul>   |

## isis

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | isis { ... }  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Enable IS-IS routing on the routing device or for a routing instance.<br><br>The <b>isis</b> statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.                                      |
| <b>Default</b>                  | IS-IS is disabled on the routing device.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS on page 3823</a></li><li>• <a href="#">Example: Configuring Multi-Level IS-IS on page 3829</a></li></ul>   |

## level (Global IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> level <i>level-number</i> {     authentication-key <i>key</i>;     authentication-key-chain (Protocols IS-IS) <i>key-chain-name</i>;     authentication-type <i>type</i>;     disable;     external-preference <i>preference</i>;     no-csnp-authentication;     no-hello-authentication;     no-psnp-authentication;     preference <i>preference</i>;     wide-metrics-only; } </pre>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure the global-level properties.</p> <p>You can administratively divide a single AS into smaller groups called areas. You configure each routing device interface to be in an area. Any interface can be in any area. The area address applies to the entire routing device. You cannot specify one interface to be in one area and another interface in a different area. To route between areas, you must have two adjacent Level 2 routers that communicate with each other.</p> <p>Level 1 routers can only route within their IS-IS area. To send traffic outside their area, Level 1 routers must send packets to the nearest intra-area Level 2 router. A routing device can be a Level 1 router, a Level 2 router, or both. You specify the router level on a per-interface basis, and a routing device becomes adjacent to other routing devices on the same level on that link only.</p> <p>You can configure one Level 1 routing process and one Level 2 routing process on each interface, and you can configure the two levels differently.</p> |
| <b>Options</b>                  | <p><b><i>level-number</i></b>—IS-IS level number.</p> <p><b>Values:</b> 1 or 2</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |

- Related Documentation**
- [Example: Configuring IS-IS](#)
  - [Example: Configuring Multi-Level IS-IS](#)

## link-protection (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | link-protection;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Enable link protection on the specified IS-IS interface. Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 3883</a></li><li>• <a href="#">node-link-protection on page 3956</a></li></ul>   |

## loose-authentication-check

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | loose-authentication-check;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Allow the use of MD5 authentication without requiring network-wide deployment.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS</a></li></ul>  |



## lsp-interval

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>lsp-interval <i>milliseconds</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit protocols isis <b>interface</b> <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.<br/>         Statement introduced in Junos OS Release 9.0 for EX Series switches.<br/>         Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure the link-state PDU interval time.</p> <p>By default, the routing device sends one link-state PDU packet out an interface every 100 milliseconds. To disable the transmission of all link-state PDUs, set the interval to 0.</p> <p>Link-state PDU throttling by use of the <b>lsp-interval</b> statement controls the flooding pace to neighboring routing devices in order to not overload them.</p> <p>Also, consider that control traffic (such as link-state PDUs and related packets) might delay user traffic (information packets) because control traffic always has precedence in terms of scheduling on the routing device interface cards. Unfortunately, the control traffic transmission rate is not decreased on low-bandwidth interfaces, such as DS-0 or fractional T1 and E1 interface. Line control traffic stays the same. On a low-bandwidth circuit that is transmitting 30 full-MTU-sized packets, there is not much bandwidth left over for other types of packets.</p> |
| <b>Default</b>                  | By default, the routing device sends one link-state PDU out an interface every 100 milliseconds.  |
| <b>Options</b>                  | <p><b>milliseconds</b>—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.</p> <p><b>Range:</b> 0 through 1000 milliseconds</p> <p><b>Default:</b> 100 milliseconds</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li> </ul>  |

## **lsp-lifetime**

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>lsp-lifetime seconds;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.</p> <p>Because link-state PDUs have a maximum lifetime, they need to be refreshed. Refreshing means that a routing device needs to re-originate its link-state PDUs periodically. The re-origination interval must be less than the link-state PDU's lifetime. For example, if the link-state PDU is valid for 1200 seconds, the routing device needs to refresh the link-state PDU in less than 1200 seconds to avoid removal of the link-state PDU from the link-state database by other routing devices. The recommended maximum link-state PDU origination interval is the lifetime minus 300 seconds. So, in a default environment this would be 900 seconds. In Junos OS, the refresh interval is derived from the lifetime and is equal to the lifetime minus 317 seconds. You can change the lifetime to a higher value to reduce the number of refreshes in the network. (You would rarely want to increase the number of refreshes.) Often these periodic link-state PDU refreshes are referred to as refresh noise, and network administrators want to reduce this noise as much as possible.</p> <p>The <a href="#">show isis overview</a> command displays the link-state PDU lifetime.</p> |
| <b>Default</b>                  | By default, link-state PDUs are maintained in network databases for 1200 seconds (20 minutes) before being considered invalid. This length of time, called the <i>LSP lifetime</i> , normally is sufficient to guarantee that link-state PDUs never expire.   |
| <b>Options</b>                  | <b>seconds</b> —link-state PDU lifetime, in seconds.<br><b>Range:</b> 350 through 65,535 seconds<br><b>Default:</b> 1200 seconds  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li><li>• <a href="http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf">http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf</a></li></ul>  |

## max-areas

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>max-areas <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Modify the maximum number of IS-IS areas advertised.</p> <p>This value is included in the Maximum Address Area field of the IS-IS common PDU header included in all outgoing PDUs.</p> <p>The maximum number of areas you can advertise is restricted to 36 to ensure that the IIH PDUs have enough space to include other type, length, and value (TLV) fields, such as the Authentication and IPv4 and IPv6 Interface Address TLVs.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.</p> <p><b>Range:</b> 3 through 36</p> <p><b>Default:</b> 3</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Example: Configuring Multi-Level IS-IS</i></li> </ul>  |

## mesh-group (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | mesh-group (blocked   <i>value</i> );   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>interface</b> <i>interface-name</i> ],<br>[edit protocols isis <b>interface</b> <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure an interface to be part of a mesh group, which is a set of fully connected nodes.</p> <p>A <i>mesh group</i> is a set of routing devices that are fully connected. That is, they have a fully meshed topology. When link-state PDUs are being flooded throughout an area, each router within a mesh group receives only a single copy of a link-state PDU instead of receiving one copy from each neighbor, thus minimizing the overhead associated with the flooding of link-state PDUs.</p> <p>To create a mesh group and designate that an interface be part of the group, assign a mesh-group number to all the routing device interfaces in the group. To prevent an interface in the mesh group from flooding link-state PDUs, configure blocking on that interface.</p> |
| <b>Options</b>                  | <p><b>blocked</b>—Configure the interface so that it does not flood link-state PDUs.</p> <p><b>value</b>—Number that identifies the mesh group.</p> <p><b>Range:</b> 1 through 4,294,967,295 (<math>2^{32} - 1</math>; 32 bits are allocated to identify a mesh group)</p>  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Mesh Groups of IS-IS Interfaces</i></li></ul>   |

## metric (Protocols IS-IS)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <code>metric <i>metric</i>;</code>  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>         | Specify the metric value for the level.   |

All IS-IS routes have a cost, which is a routing metric that is used in the IS-IS link-state calculation. The cost is an arbitrary, dimensionless integer that can be from 1 through 63, or from 1 through 16,777,215 ( $2^{24} - 1$ ) if you are using wide metrics.

Similar to other routing protocols, IS-IS provides a way of exporting routes from the routing table into the IS-IS network. When a route is exported into the IS-IS network without a specified metric, IS-IS uses default metric values for the route, depending on the protocol that was used to learn the route.

[Table 310 on page 3947](#) depicts IS-IS route export metric default values.

**Table 310: Default Metric Values for Routes Exported into IS-IS**

| Protocol Used for Learning the Route | Default Metric Value  |
|--------------------------------------|---|
| Direct                               | 10  |
| Static                               | Same as reported by the protocol used for exporting the route |
| Aggregate                            | 10  |
| Generate                             | 10  |
| RIP                                  | Same as reported by the protocol used for exporting the route |
| OSPF                                 | Same as reported by the protocol used for exporting the route |
| BGP                                  | 10  |

The default metric values behavior can be customized by using routing policies.

|                |  |
|----------------|--|
| <b>Options</b> | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics) |
|----------------|--|

**Default:** 10 (for all interfaces except lo0), 0 (for the lo0 interface)

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li><li>• <i>te-metric</i></li><li>• <a href="#">wide-metrics-only on page 3975</a></li></ul> |

---

## no-adjacency-holddown

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-adjacency-holddown;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Disable the hold-down timer for IS-IS adjacencies.<br><br>A hold-down timer delays the advertising of adjacencies by waiting until a time period has elapsed before labeling adjacencies in the up state. You can disable this hold-down timer, which labels adjacencies up faster. However, disabling the hold-down timer creates more frequent link-state PDU updates and SPF computation. |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hold-time on page 3931</a></li></ul>   |

## no-authentication-check

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-authentication-check;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">csnp-interval on page 3919</a></li> <li>• <a href="#">hello-authentication-type on page 3927</a></li> </ul>  |

## no-csnp-authentication

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-csnp-authentication;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Suppress authentication check on complete sequence number PDU (CSNP) packets.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">csnp-interval on page 3919</a></li> </ul>  |

## no-eligible-backup (Protocols IS-IS)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-eligible-backup;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 3883</a></li><li>• <a href="#">link-protection on page 3942</a></li><li>• <a href="#">node-link-protection on page 3956</a></li></ul>   |

## no-hello-authentication

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-hello-authentication;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Suppress authentication check on complete sequence number hello packets.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">hello-authentication-type on page 3927</a></li></ul>   |




## no-ipv4-multicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv4-multicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Exclude an interface from IPv4 multicast topologies.   |
| <b>Default</b>                  | Multicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li> </ul>   |

## no-ipv4-routing

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv4-routing;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | <p>Disable IP version 4 (IPv4) routing.</p> <p>Disabling IPv4 routing has the following results:</p> <ul style="list-style-type: none"><li>• The routing device does not advertise the network layer protocol identifier (NLPID) for IPv4 in the Junos OS link-state PDU fragment zero.</li><li>• The routing device does not advertise any IPv4 prefixes in Junos OS link-state PDUs.</li><li>• The routing device does not advertise the NLPID for IPv4 in Junos OS hello packets.</li><li>• The routing device does not advertise any IPv4 addresses in Junos OS hello packets.</li><li>• The routing device does not calculate any IPv4 routes.</li></ul> <div> <b>NOTE:</b> Note: Even when no-ipv4-routing is configured, an IS-IS traceoptions log can list rejected IPv4 addresses. When a configuration is committed, IS-IS schedules a scan of the routing table to determine whether any routes need to be exported into the IS-IS link state database. The implicit default export policy action is to reject everything. IPv4 addresses from the routing table are examined for export, rejected by the default policy, and the rejections are logged.</div> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859</a></li></ul>   |

## no-ipv6-multicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv6-multicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | Exclude an interface from the IPv6 multicast topologies.   |
| <b>Default</b>                  | Multicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li> </ul>   |

## no-ipv6-routing

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-ipv6-routing;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>              | <p>Disable IP version 6 (IPv6) routing.</p> <p>Disabling IPv6 routing has the following results:</p> <ul style="list-style-type: none"><li>• The routing device does not advertise the network layer protocol identifier (NLPID) for IPv6 in the Junos OS link-state PDU fragment zero.</li><li>• The routing device does not advertise any IPv6 prefixes in Junos OS link-state PDUs.</li><li>• The routing device does not advertise the NLPID for IPv6 in Junos OS hello packets.</li><li>• The routing device does not advertise any IPv6 addresses in Junos OS hello packets.</li><li>• The routing device does not calculate any IPv6 routes.</li></ul> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859</a></li></ul>  |

## no-ipv6-unicast

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-ipv6-unicast;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | Exclude an interface from the IPv6 unicast topologies. This enables you to exercise control over the paths that unicast data takes through a network.  |
| <b>Default</b>                  | IPv6 unicast topologies are disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859</a></li> </ul>   |

## no-psnp-authentication

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-psnp-authentication;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Suppress authentication check on partial sequence number PDU (PSNP) packets.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring IS-IS Authentication on page 3909</a></li> </ul>  |

## no-unicast-topology

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-unicast-topology;  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Exclude an interface from the IPv4 unicast topologies.  |
| <b>Default</b>                  | IPv4 unicast topologies are disabled.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Multicast Topology on page 3867</a></li></ul>  |

## node-link-protection (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | node-link-protection;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ],<br>[edit protocols isis interface <i>interface-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.5.<br>Statement introduced in Junos OS Release 9.5 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Enable node-link protection on the specified IS-IS interface. Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 3883</a></li><li>• <a href="#">link-protection on page 3942</a></li></ul>  |

## overload (Protocols IS-IS)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> overload {     advertise-high-metrics;     allow-route-leaking;     timeout <i>seconds</i>; } </pre>  |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols <i>isis</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i>],</p> <p>[edit protocols <i>isis</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i>]</p>  |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>         | <p>Configure the local routing device so that it appears to be overloaded. This statement causes the routing device to continue participating in IS-IS routing, but prevents it from being used for transit traffic. Traffic destined to immediately attached subnets continues to transit the routing device.</p> <p>You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.</p> <p>You configure or disable overload mode in IS-IS with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the IS-IS instance started is less than the specified timeout.</p> <p>A timer is started for the difference between the timeout and the time elapsed since the instance started. If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set. When the timer expires, overload mode is cleared. In overload mode, the routing device IS-IS advertisements are originated with the overload bit set. This causes the transit traffic to take paths around the routing device. However, the overloaded routing device's own links are still accessible.</p> <p>The value of the overload bit depends on these three scenarios:</p> <ol style="list-style-type: none"> <li>1. When the overload bit has already been set to a given value and the routing process is restarted: Link-state PDUs are regenerated with the overload bit cleared.</li> <li>2. When the overload bit is reset to a lesser value while the routing process is running: Link-state PDUs are regenerated with the overload bit cleared.</li> <li>3. When the overload bit is reset to a greater value while the routing process is running: Link-state PDUs are regenerated with the overload bit set to the difference between the old and new value.</li> </ol> <p>In overload mode, the routing device advertisement is originated with all the transit routing device links (except stub) set to a metric of 0xFFFF. The stub routing device links are</p> |

advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and take paths around the routing device.

To understand the reason for setting the overload bit, consider that BGP converges slowly. It is not very good at detecting that a neighbor is down because it has slow-paced keepalive timers. Once the BGP neighbor is determined to be down, it can take up to 2 minutes for a BGP router to declare the neighbor down. IS-IS is much quicker. IS-IS only takes 10-30 seconds to detect absent peers. It is the slowness of BGP, more precisely the slowness of internal BGP (IBGP), that necessitates the use of the overload bit. IS-IS and BGP routing are mutually dependent on each other. If both do not converge at the same time, traffic is dropped without notification (black holed).

You might want to configure the routing device so that it appears to be overloaded when you are restarting routing on the device. Setting the overload bit for a fixed amount of time right after a restart of the routing protocol process (rpd) ensures that the router does not receive transit traffic while the routing protocols (especially IBGP) are still converging.

Setting the overload bit is useful when performing hardware or software maintenance work on a routing device. After the maintenance work, clear the overload bit to carry on forwarding transit traffic. Manual clearing of the overload bit is not always possible. What is needed is an automated way of clearing the overload bit after some amount of time. Most networks use a time value of 300 seconds. This 5-minute value provides a good balance, allowing time to bring up even large internal IBGP meshes, while still relatively quick.

Another appropriate application for setting for the overload bit is on dedicated devices such as BGP route reflectors, which are intentionally not meant to carry any transit traffic. In this case, you would not use the timer.

You can verify that the overload bit is set by running the **show isis database** command.



**Options** **advertise-high-metrics**—Advertise maximum link metrics in NLRI instead of setting the overload bit.

The **advertise-high-metric** setting is only valid while the routing device is in overload mode.

When **advertise-high-metric** is configured, IS-IS does not set the overload bit. Rather, it sets the metric to 63 or 16,777,214, depending whether wide metrics are enabled. This allows the overloaded routing device to be used for transit as a last resort.

An L1-L2 router in overload mode stops leaking route information between L1 and L2 levels and clears its attached bit. This is also true when **advertise-high-metrics** is configured.

**allow-route-leaking**—Enable leaking of route information into the network even if the overload bit is set.



**NOTE:** The **allow-route-leaking** option does not work if the routing device is in dynamic overload mode. Dynamic overload can occur if the device has exceeded its resource limits, such as the prefix limit.

**timeout seconds**—Number of seconds at which the overloading is reset.

**Range:** 60 through 1800 seconds


**Default:** 0 seconds

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring IS-IS*

## passive (Protocols IS-IS)

|   |  |
|---|--|
| <b>Syntax</b>   | <code>passive;</code>  |
| <b>Hierarchy Level</b>  | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i>]</p> |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>  | <p>Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.</p> <p>This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the <b>interface</b> statement at the [edit protocols isis] or the [edit routing-instances <i>routing-instance-name</i> protocols isis] hierarchy level. To disable it, include the <b>disable</b> statement at those hierarchy levels. The three states—enabling, disabling, or not running IS-IS on an interface—are mutually exclusive.</p>  |
| <div>  <p><b>NOTE:</b> Configuring IS-IS on a loopback interface automatically renders it as a passive interface, irrespective of whether the <b>passive</b> statement was used in the configuration of the interface.</p> </div>  |  |
| <p>If neither passive mode nor the <b>family iso</b> option is configured on the IS-IS interface, then the routing device treats the interface as not being operational, and no direct IPv4/IPv6 routes are exported into IS-IS. (You configure the <b>family iso</b> option at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.)</p> |  |
| <b>Default</b>  | By default, IS-IS must be configured on an interface or a level for direct interface addresses to be advertised into that level.   |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Multi-Level IS-IS on page 3829</a></li> <li>• <a href="#">disable</a></li> </ul>   |

## point-to-point

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | point-to-point;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ],<br>[edit protocols isis <a href="#">interface interface-name</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">interface interface-name</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure an IS-IS interface to behave like a point-to-point connection.</p> <p>You can use the <b>point-to-point</b> statement to configure a LAN interface to act like a point-to-point interface for IS-IS. You do not need an unnumbered LAN interface, and it has no effect if configured on an interface that is already point-to-point.</p> <p>The <b>point-to-point</b> statement affects only IS-IS protocol procedures on that interface. All other protocols continue to treat the interface as a LAN interface. Only two IS-IS routing devices can be connected to the LAN interface, and both must be configured as point-to-point.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">IS-IS Overview on page 3812</a></li> <li>• <a href="#">Understanding IS-IS Designated Routers on page 3905</a></li> <li>• <a href="#">Example: Configuring Synchronization Between IS-IS and LDP</a></li> </ul>  |

## preference (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>preference preference;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ],<br>[edit protocols isis <a href="#">level level-number</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a> ]                               |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure the preference of internal routes.</p> <p>Route preferences (also known as administrative distances) are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.</p> <p>To change the preference values, include the <b>preference</b> statement (for internal routes) or the <b>external-preference</b> statement.</p> |
| <b>Options</b>                  | <p><b>preference</b>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Route Preferences Overview</i></li><li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li><li>• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i></li><li>• <a href="#">external-preference on page 3922</a></li></ul>   |

## prefix-export-limit (Protocols IS-IS)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>prefix-export-limit <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>],</p> <p>[edit protocols isis <a href="#">level level-number</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols isis <a href="#">level level-number</a>]</p>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure a limit to the number of prefixes exported into IS-IS.</p> <p>By default, there is no limit to the number of prefixes that can be exported into IS-IS. To configure a limit to the number of prefixes that can be exported into IS-IS, include the <b>prefix-export-limit</b> statement. The <b>prefix-export-limit</b> statement protects the rest of the network from a malicious policy by applying a threshold filter for exported routes.</p> <p>The number of prefixes depends on the size of your network. Good design advice is to set it to double the total number of IS-IS Level 1 and Level 2 routing devices in your network.</p> <p>If the number of prefixes exported into IS-IS exceeds the configured limit, the overload bit is set and the overload state is reached. When other routers detect that this bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes. The overload state can be cleared by using the <a href="#">clear isis overload</a> command.</p> <p>The <a href="#">show isis overview</a> command displays the prefix export limit when it is configured.</p> |
| <b>Options</b>                  | <p><b><i>number</i></b>—Prefix limit.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> None</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Redistributing BGP Routes with a Specific Community Tag into IS-IS</i></li> <li>• <i>Example: Redistributing OSPF Routes into IS-IS</i></li> </ul>  |

## priority (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>priority <i>number</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit protocols isis interface <i>interface-name</i> level <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis interface <i>interface-name</i> level <i>level-number</i> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure the interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.</p> <p>The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.</p> <p>A routing device advertises its priority to become a designated router in its hello packets. On all multiaccess networks, IS-IS uses the advertised priorities to elect a designated router for the network. This routing device is responsible for sending network link-state advertisements, which describe all the routing devices attached to the network. These advertisements are flooded throughout a single area.</p> <p>A routing device's priority for becoming the designated router is indicated by an arbitrary number from 0 through 127. Routing devices with a higher value are more likely to become the designated router.</p> |
| <b>Options</b>                  | <i>number</i> —Priority value.<br><b>Range:</b> 0 through 127<br><b>Default:</b> 64   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IS-IS Designated Routers on page 3905</a></li></ul>  |

## reference-bandwidth (Protocols IS-IS)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>reference-bandwidth <i>reference-bandwidth</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>],</p> <p>[edit protocols <a href="#">isis</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Optimize routing based on bandwidth by setting the reference bandwidth used in calculating the default interface cost.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics. When there are several equal-cost routes to a destination, traffic is distributed equally among them.</p> <p>The cost of a route is described by a single dimensionless metric that is determined using the following formula:</p> $\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$ <p>For example, if you set the reference bandwidth to 1 Gbps (that is, <i>reference-bandwidth</i> is set to 1,000,000,000), a 100-Mbps interface has a routing metric of 10.</p> <p>All IS-IS interfaces have a cost, which is a routing metric that is used in the IS-IS link-state calculation. Routes with lower total path metrics are preferred over those with higher path metrics.</p> |
| <b>Options</b>                  | <p><i>reference-bandwidth</i>—Reference bandwidth value in bits per second.</p> <p><b>Range:</b> 9600 through 1,000,000,000,000 bps</p> <p><b>Default:</b> None</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring IS-IS</i></li> <li>• <a href="http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf">http://www.juniper.net/us/en/training/certification/JNCIP_studyguide.pdf</a></li> </ul>  |

## rib-group (Protocols IS-IS)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>rib-group {<br/>    inet <i>group-name</i>;<br/>    inet6 <i>group-name</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ]             |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | <p>Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.</p> <p>Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.</p> |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of the routing table group.</p> <p><b>inet</b>—Install IPv4 IS-IS routes.</p> <p><b>inet6</b>—Install IPv6 IS-IS routes.</p>   |
| <b>Required Privilege Level</b> | <p><b>routing</b>—To view this statement in the configuration.</p> <p><b>routing-control</b>—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li><li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li><li>• <i>Understanding Multiprotocol BGP</i></li></ul>   |



## spf-options (Protocols IS-IS)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>spf-options {     delay <i>milliseconds</i>;     holddown <i>milliseconds</i>;     rapid-runs <i>number</i>; }</pre>  |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <i>isis</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ],<br>[edit protocols <i>isis</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <i>isis</i> ]  |
| <b>Release Information</b> | Statement introduced in Junos OS Release 8.5.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>         | <p>Configure options for running the shortest-path-first (SPF) algorithm.</p> <p>Running the SPF algorithm is usually the beginning of a series of larger system-wide events. For example, the SPF algorithm can lead to interior gateway protocol (IGP) prefix changes, which then lead to BGP nexthop resolution changes. Consider what happens if there are rapid link changes in the network. The local routing device can become overwhelmed. This is why it sometimes makes sense to throttle the scheduling of the SPF algorithm.</p> <p>You can configure the following SPF options:</p> <ul style="list-style-type: none"> <li>• The delay in the time between the detection of a topology change and when the SPF algorithm actually runs.</li> <li>• The maximum number of times that the SPF algorithm can run in succession before the hold-down timer begins.</li> <li>• The time to hold down, or wait, before running another SPF calculation after the SPF algorithm has run in succession the configured maximum number of times.</li> </ul> <p>If the network stabilizes during the hold-down period and the SPF algorithm does not need to run again, the system reverts to the configured values for the <b>delay</b> and <b>rapid-runs</b> statements.</p> |
| <b>Options</b>             | <p><b>delay <i>milliseconds</i></b>—Time interval between the detection of a topology change and when the SPF algorithm runs.</p> <p><b>Range:</b> 50 through 1000 milliseconds</p> <p><b>Default:</b> 200 milliseconds</p> <p><b>holddown <i>milliseconds</i></b>—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.</p> <p><b>Range:</b> 2000 through 10,000 milliseconds</p> <p><b>Default:</b> 5000 milliseconds</p>   |

**rapid-runs *number***—Maximum number of times the SPF algorithm can run in succession.  
After the maximum is reached, the holddown interval begins.

**Range:** 1 through 5

**Default:** 3

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Link and Node Protection for IS-IS Routes on page 3883](#)

---

## topologies (Protocols IS-IS)

---

**Syntax**

```
topologies {  
    ipv4-multicast;  
    ipv6-multicast;  
    ipv6-unicast;  
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [isis](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [isis](#)],  
[edit protocols [isis](#)],  
[edit routing-instances *routing-instance-name* protocols [isis](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure alternate IS-IS topologies.  
  
The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring IS-IS IPv4 and IPv6 Unicast Topologies on page 3859](#)
- [Example: Configuring IS-IS Multicast Topology on page 3867](#)

## traceoptions (Protocols IS-IS)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>name</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">isis</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ],<br>[edit protocols <a href="#">isis</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">isis</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>         | Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple <b>flag</b> statements.  |



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

|                |   |
|----------------|---|
| <b>Default</b> | The default IS-IS protocol-level tracing options are those inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.  |
| <b>Options</b> | <p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>name</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory <b>/var/log</b>. We recommend that you place IS-IS tracing output in the file <b>isis-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 10 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one flag, include multiple <b>flag</b> statements.</p> |

### IS-IS Protocol-Specific Tracing Flags

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets
- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDUs
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations, including adjacency changes

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

***flag-modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **send**—Trace the packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                              |   |
|------------------------------|---|
| <b>Required Privilege</b>    | routing and trace—To view this statement in the configuration.  |
| <b>Level</b>                 | routing-control and trace-control—To add this statement to the configuration.   |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>Example: Configuring the Transmission Frequency for CSNPs on IS-IS Interfaces</i></li> <li>• <i>Example: Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces</i></li> <li>• <i>Example: Enabling Packet Checksums on IS-IS Interfaces</i></li> </ul> |

## traffic-engineering (Protocols IS-IS)

---

**Syntax**

```
traffic-engineering {
  disable;
  credibility-protocol-preference;
  family inet {
    shortcuts {
      multicast-rpf-routes;
    }
  }
  family inet6 {
    shortcuts;
  }
  multipath {
    lsp-equal-cost;
  }
}
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [isis](#)],  
[edit protocols [isis](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Support for the **family** statement introduced in Junos OS Release 9.3.  
Support for the **credibility-protocol-preference** statement introduced in Junos OS Release 9.4.  
Support for the **multipath** statement introduced in Junos OS Release 9.6.  
Support for the **lsp-equal-cost** statement introduced in Junos OS Release 9.6.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure traffic engineering properties for IS-IS.

IS-IS always performs shortest-path-first (SPF) calculations to determine next hops. For prefixes reachable through a particular next hop, IS-IS places that next hop for that prefix in the inet.0 routing table. In addition, for routers running MPLS, IS-IS installs the prefix for IPv4 routes in the inet.3 routing table as well. The inet.3 table, which is present on the ingress router, contains the host address of each MPLS label-switched path (LSP) egress router. BGP uses this routing table to resolve next-hop addresses.

If you enable IS-IS traffic engineering shortcuts and if there is a label-switched path to a point along the path to that prefix, IS-IS installs the prefix in the inet.3 routing table and uses the LSP as a next hop. The net result is that for BGP egress routers for which there is no LSP, BGP automatically uses an LSP along the path to reach the egress router.

In Junos OS Release 9.3 and later, IS-IS traffic engineering shortcuts support IPv6 routes. LSPs to be used for shortcuts continue to be signaled using IPv4. However, by default, shortcut routes calculated through IPv6 routes are added to the inet6.3 routing table. The default behavior is for only BGP to use LSPs in its calculations. If you configure MPLS so that both BGP and interior gateway protocols use LSPs for forwarding traffic, shortcut routes calculated through IPv6 are added to the inet6.0 routing table. IS-IS ensures that the IPv6 routes running over the IPv4 MPLS LSP are correctly de-encapsulated at the

tunnel egress by pushing an extra IPv6 explicit null label between the IPv6 payload and the IPv4 transport label.

RSVP LSPs with a higher preference than IS-IS routes are not considered during the computation of traffic engineering shortcuts.

To configure IS-IS so that it uses LSPs as shortcuts when installing information in the inet.3 or inet6.3 routing table, include the following statements:

```
family inet {
  shortcuts {
    multicast-rpf-routes;
  }
}
family inet6 {
  shortcuts;
}
```

For IPv4 traffic, include the **inet** statement. For IPv6 traffic, include the **inet6** statement.

To configure load balancing across multiple LSPs, include the **multipath** statement.

When traffic engineering shortcuts are used, RSVP first looks at the **metric2** value, which is derived from the IGP cost. After this, RSVP considers the LSP metric value. So, if a certain path changes for an LSP and the cost changes, not all LSPs are used to load-balance the network.

When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default behavior for load balancing, include the **lsp-equal-cost** statement to retain the equal cost path information in the routing table.

```
multipath {
  lsp-equal-cost;
}
```

Because the inet.3 routing table is present only on ingress routers, you can configure LSP shortcuts only on these routers.

**Default** IS-IS traffic engineering support is enabled.

By default, IS-IS supports traffic engineering by exchanging basic information with the traffic engineering database. To disable this support, and to disable IS-IS shortcuts if they are configured, include the **disable** statement.

**Options**    **credibility-protocol-preference**—Specify that IS-IS should use the configured protocol preference for IGP routes to determine the traffic engineering database credibility value. By default, the traffic engineering database prefers IS-IS routes even when the routes of another IGP are configured with a lower, that is, more preferred value. Use this statement to override this default behavior.

The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure IS-IS to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.



**NOTE:** This feature is also supported for OSPFv2.

---

**lsp-equal-cost**—Configure LSPs to be retained as equal cost paths for load balancing when a better path metric is found during the IS-IS internal routing table calculation. When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. This is undesirable, because certain equal-cost multipath (ECMP) combinations can be lost during route calculation. To override this default IS-IS behavior, include the **lsp-equal-cost** statement for load balancing, so that the equal cost path information is retained in the routing table.

**multipath**—Enable load balancing for multiple LSPs.

The remaining statements are explained separately.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration. |
|---------------------------------|---|

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Enabling OSPF Traffic Engineering Support on page 4153</a></li><li>• <a href="#">Example: Enabling IS-IS Traffic Engineering Support</a></li><li>• <a href="#">traffic-engineering (OSPF) on page 4240</a></li></ul> |
|------------------------------|---|



## wide-metrics-only

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | wide-metrics-only;   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols isis <b>level</b> <i>level-number</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br>isis <b>level</b> <i>level-number</i> ],<br>[edit protocols isis <b>level</b> <i>level-number</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols isis <b>level</b> <i>level-number</i> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis.<br><br>Normally, IS-IS metrics can have values up to 63, and IS-IS generates two type, length, and value (TLV) tuples, one for an IS-IS adjacency and the second for an IP prefix. To allow IS-IS to support traffic engineering, a second pair of TLVs has been added to IS-IS, one for IP prefixes and the second for IS-IS adjacency and traffic engineering information. With these TLVs, IS-IS metrics can have values up to 16,777,215 ( $2^{24} - 1$ ).<br><br>To configure IS-IS to generate only the new pair of TLVs and thus to allow the wider range of metric values, include the <b>wide-metrics-only</b> statement. |
| <b>Default</b>                  | By default, Junos OS supports the sending and receiving of wide metrics. Junos OS allows a maximum metric value of 63 and generates both pairs of TLVs.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li> <li>• <i>te-metric</i></li> </ul>  |



## CHAPTER 46

# Administration

- [Operational Commands on page 3977](#)

### Operational Commands

---

- [clear isis adjacency](#)
- [clear isis database](#)
- [clear isis overload](#)
- [clear isis statistics](#)
- [show isis adjacency](#)
- [show isis authentication](#)
- [show isis backup coverage](#)
- [show isis backup label-switched-path](#)
- [show isis backup spf results](#)
- [show isis database](#)
- [show isis hostname](#)
- [show isis interface](#)
- [show isis overview](#)
- [show isis route](#)
- [show isis spf](#)
- [show isis statistics](#)

## clear isis adjacency

---

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3978</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3978</a>   |
| <b>Syntax</b>                                     | <pre>clear isis adjacency &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis adjacency &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>   |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>                                | Remove entries from the IS-IS adjacency database.  |
| <b>Options</b>                                    | <p><b>none</b>—Remove all entries from the adjacency database.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear all adjacencies for the specified routing instance only.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear all adjacencies for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear adjacencies for the specified neighbor only.</p> |
| <b>Required Privilege Level</b>                   | clear  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis adjacency on page 3986</a></li></ul>   |
| <b>List of Sample Output</b>                      | <a href="#">clear isis adjacency on page 3978</a>  |
| <b>Output Fields</b>                              | See <a href="#">show isis adjacency</a> for an explanation of output fields.   |

## Sample Output

### clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State          Hold (secs) SNPA
```

```
so-1/0/0.0    karaku1      3 Up                26
so-1/1/3.0    1921.6800.5080 3 Up                23
so-5/0/0.0    1921.6800.5080 3 Up                19
```

```
user@host> clear isis adjacency karaku1
```

```
user@host> show isis adjacency
```

```
IS-IS adjacency database:
```

| Interface  | System         | L State        | Hold (secs) | SNPA |
|------------|----------------|----------------|-------------|------|
| so-1/0/0.0 | karaku1        | 3 Initializing | 26          |      |
| so-1/1/3.0 | 1921.6800.5080 | 3 Up           | 24          |      |
| so-5/0/0.0 | 1921.6800.5080 | 3 Up           | 21          |      |

## clear isis database

---

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3980</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3980</a>   |
| <b>Syntax</b>                                     | <pre>clear isis database &lt;entries&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>clear isis database &lt;entries&gt; &lt;instance <i>instance-name</i>&gt;</pre>   |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>                                | Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.   |
| <b>Options</b>                                    | <b>none</b> —Remove all entries from the IS-IS link-state database for all routing instances.<br><br><b>entries</b> —(Optional) Name of the database entry.<br><br><b>instance <i>instance-name</i></b> —(Optional) Clear all entries for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show isis database on page 4000</a></li></ul>  |
| <b>List of Sample Output</b>                      | <a href="#">clear isis database on page 3980</a>   |
| <b>Output Fields</b>                              | See <a href="#">show isis database</a> for an explanation of output fields.  |

## Sample Output

### clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12   0x84dd             1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
```

|                      |      |        |      |
|----------------------|------|--------|------|
| crater.00-00         | 0x19 | 0xe92c | 1134 |
| badlands.00-00       | 0x16 | 0x1454 | 985  |
| carlsbad.00-00       | 0x33 | 0x220b | 1015 |
| ranier.00-00         | 0x2e | 0xfc31 | 1007 |
| 1921.6800.5066.00-00 | 0x11 | 0x7313 | 566  |
| 1921.6800.5067.00-00 | 0x14 | 0xd9d4 | 939  |

6 LSPs

user@host> clear isis database

user@host> show isis database

IS-IS level 1 link-state database:

| LSP ID | Sequence | Checksum | Lifetime (secs) |
|--------|----------|----------|-----------------|
|--------|----------|----------|-----------------|

IS-IS level 2 link-state database:

| LSP ID | Sequence | Checksum | Lifetime (secs) |
|--------|----------|----------|-----------------|
|--------|----------|----------|-----------------|

## clear isis overload

---

|  |   |
|--|---|
| List of Syntax                             | <a href="#">Syntax on page 3982</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3982</a>  |
| Syntax                                     | <code>clear isis overload</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches and QFX Series) | <code>clear isis overload</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>  |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| Description                                | <p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display <b>overload</b> after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p> |
| Options                                    | <p><b>none</b>—Reset the IS-IS dynamic overload bit.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>  |
| Required Privilege Level                   | clear   |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">show isis database on page 4000</a></li></ul>   |
| List of Sample Output                      | <a href="#">clear isis overload on page 3982</a>  |
| Output Fields                              | See <a href="#">show isis database</a> for an explanation of output fields.   |

## Sample Output

### clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                               Sequence Checksum Lifetime Attributes
```



```
pro3-c.00-00          0x4   0x10db    1185 L1 L2 Overload
```

```
1 LSPs
```

```
IS-IS level 2 link-state database:
```

| LSP ID       | Sequence | Checksum | Lifetime | Attributes            |
|--------------|----------|----------|----------|-----------------------|
| pro3-c.00-00 | 0x5      | 0x429f   | 1185     | L1 L2 <b>Overload</b> |

```
pro2-a.00-00          0x91e  0x2589    874 L1 L2
```

```
pro2-a.02-00          0x1    0xcbc    874 L1 L2
```

```
3 LSPs
```

```
user@host> clear isis overload
```

```
user@host> show isis database
```

```
IS-IS level 1 link-state database:
```

| LSP ID       | Sequence | Checksum | Lifetime | Attributes |
|--------------|----------|----------|----------|------------|
| pro3-c.00-00 | 0xa      | 0x429e   | 1183     | L1 L2      |

```
1 LSPs
```

```
IS-IS level 2 link-state database:
```

| LSP ID       | Sequence | Checksum | Lifetime | Attributes |
|--------------|----------|----------|----------|------------|
| pro3-c.00-00 | 0xc      | 0x9c39   | 1183     | L1 L2      |
| pro2-a.00-00 | 0x91e    | 0x2589   | 783      | L1 L2      |
| pro2-a.02-00 | 0x1      | 0xcbc    | 783      | L1 L2      |

```
3 LSPs
```

## clear isis statistics

---

|  |   |
|--|---|
| List of Syntax                             | <a href="#">Syntax on page 3984</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3984</a>  |
| Syntax                                     | <code>clear isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code><br><code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>  |
| Syntax (EX Series Switches and QFX Series) | <code>clear isis statistics</code><br><code>&lt;instance <i>instance-name</i>&gt;</code>  |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| Description                                | Set statistics about IS-IS traffic to zero.   |
| Options                                    | <b>none</b> —Set IS-IS traffic statistics to zero for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| Required Privilege Level                   | view  |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">show isis statistics on page 4030</a></li></ul>   |
| List of Sample Output                      | <a href="#">clear isis statistics on page 3984</a>  |
| Output Fields                              | See <a href="#">show isis statistics</a> for an explanation of output fields.   |

## Sample Output

### clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

| PDU type | Received | Processed | Drops | Sent   | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP      | 12793    | 12793     | 0     | 8666   | 719    |
| IIH      | 116751   | 116751    | 0     | 118834 | 0      |
| CSNP     | 203956   | 203956    | 0     | 204080 | 0      |
| PSNP     | 7356     | 7350      | 6     | 8635   | 0      |
| Unknown  | 0        | 0         | 0     | 0      | 0      |
| Totals   | 340856   | 340850    | 6     | 340215 | 719    |

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0  
LSP queue length: 0 Drops: 0

SPF runs: 1064  
Fragments rebuilt: 1087  
LSP regenerations: 436  
Purges initiated: 0

user@host> clear isis statistics

user@host> show isis statistics  
IS-IS statistics for merino:

| PDU type | Received | Processed | Drops | Sent | Rexmit |
|----------|----------|-----------|-------|------|--------|
| LSP      | 0        | 0         | 0     | 0    | 0      |
| IIH      | 3        | 3         | 0     | 3    | 0      |
| CSNP     | 2        | 2         | 0     | 4    | 0      |
| PSNP     | 0        | 0         | 0     | 0    | 0      |
| Unknown  | 0        | 0         | 0     | 0    | 0      |
| Totals   | 5        | 5         | 0     | 7    | 0      |

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0  
LSP queue length: 0 Drops: 0

SPF runs: 0  
Fragments rebuilt: 0  
LSP regenerations: 0  
Purges initiated: 0

## show isis adjacency

---

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3986</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3986</a>   |
| <b>Syntax</b>                                     | <pre>show isis adjacency &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis adjacency &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt;</pre>  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>                                | Display information about IS-IS neighbors.   |
| <b>Options</b>                                    | <p><b>none</b>—Display standard information about IS-IS neighbors for all routing instances.</p> <p><b><i>system id</i></b>—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p><b>brief   detail   extensive</b>—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis adjacency on page 3978</a></li></ul>  |
| <b>List of Sample Output</b>                      | <a href="#">show isis adjacency on page 3988</a><br><a href="#">show isis adjacency brief on page 3988</a><br><a href="#">show isis adjacency detail on page 3989</a><br><a href="#">show isis adjacency extensive on page 3989</a>  |
| <b>Output Fields</b>                              | <a href="#">Table 311 on page 3987</a> describes the output fields for the <b>show isis adjacency</b> command. Output fields are listed in the approximate order in which they appear.   |

Table 311: show isis adjacency Output Fields

| Field Name                                | Field Description   | Level of Output         |
|---|---|-------------------------|
| <b>Interface</b>                          | Interface through which the neighbor is reachable.  | All levels              |
| <b>System</b>                             | System identifier ( <b>sysid</b> ), displayed as a name, if possible.   | <b>brief</b>            |
| <b>L or Level</b>                         | Level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address. | All levels              |
| <b>State</b>                              | State of the adjacency: <b>Up</b> , <b>Down</b> , <b>New</b> , <b>One-way</b> , <b>Initializing</b> , or <b>Rejected</b> .  | All levels              |
| <b>Hold (secs)</b>                        | Remaining hold time of the adjacency.   | <b>brief</b>            |
| <b>SNPA</b>                               | Subnetwork point of attachment (MAC address of the next hop).   | <b>brief</b>            |
| <b>Expires in</b>                         | How long until the adjacency expires, in seconds.   | <b>detail</b>           |
| <b>Priority</b>                           | Priority to become the designated intermediate system.  | <b>detail extensive</b> |
| <b>Up/Down transitions</b>                | Count of adjacency status changes from <b>Up</b> to <b>Down</b> or from <b>Down</b> to <b>Up</b> .  | <b>detail</b>           |
| <b>Last transition</b>                    | Time of the last <b>Up/Down</b> transition.   | <b>detail</b>           |
| <b>Circuit type</b>                       | Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.  | <b>detail</b>           |
| <b>Speaks</b>                             | Protocols supported by this neighbor.   | <b>detail extensive</b> |
| <b>MAC address</b>                        | MAC address of the interface.   | <b>detail extensive</b> |
| <b>Topologies</b>                         | Supported topologies.   | <b>detail extensive</b> |
| <b>Restart capable</b>                    | Whether a neighbor is capable of graceful restart: <b>Yes</b> or <b>No</b> .  | <b>detail extensive</b> |
| <b>Adjacency advertisement: Advertise</b> | This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.   | <b>detail extensive</b> |
| <b>Adjacency advertisement: Suppress</b>  | This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.   | <b>detail extensive</b> |
| <b>IP addresses</b>                       | IP address of this neighbor.  | <b>detail extensive</b> |

Table 311: show isis adjacency Output Fields (*continued*)

| Field Name     | Field Description  | Level of Output |
|----------------|--|-----------------|
| Transition log | <p>List of recent transitions, including:</p> <ul style="list-style-type: none"> <li>• <b>When</b>—Time at which an IS-IS adjacency transition occurred.</li> <li>• <b>State</b>—Current state of the IS-IS adjacency (<b>up</b>, <b>down</b>, or <b>rejected</b>). <ul style="list-style-type: none"> <li>• <b>Up</b>—Adjacency is up and operational.</li> <li>• <b>Down</b>—Adjacency is down and not available.</li> <li>• <b>Rejected</b>—Adjacency has been rejected.</li> </ul> </li> <li>• <b>Event</b>—Type of transition that occurred. <ul style="list-style-type: none"> <li>• <b>Seenself</b>—Possible routing loop has been detected.</li> <li>• <b>Interface down</b>—IS-IS interface has gone down and is no longer available.</li> <li>• <b>Error</b>—Adjacency error.</li> </ul> </li> <li>• <b>Down reason</b>—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> <li>• <b>3-Way Handshake Failed</b>—Connection establishment failed.</li> <li>• <b>Address Mismatch</b>—Address mismatch caused link failure.</li> <li>• <b>Aged Out</b>—Link expired.</li> <li>• <b>ISO Area Mismatch</b>—IS-IS area mismatch caused link failure.</li> <li>• <b>Bad Hello</b>—Unacceptable hello message caused link failure.</li> <li>• <b>BFD Session Down</b>—Bidirectional failure detection caused link failure.</li> <li>• <b>Interface Disabled</b>—IS-IS interface is disabled.</li> <li>• <b>Interface Down</b>—IS-IS interface is unavailable.</li> <li>• <b>Interface Level Disabled</b>—IS-IS level is disabled.</li> <li>• <b>Level Changed</b>—IS-IS level has changed on the adjacency.</li> <li>• <b>Level Mismatch</b>—Levels on adjacency are not compatible.</li> <li>• <b>MPLS LSP Down</b>—Label-switched path (LSP) is unavailable.</li> <li>• <b>MT Topology Changed</b>—IS-IS topology has changed.</li> <li>• <b>MT Topology Mismatch</b>—IS-IS topology is mismatched.</li> <li>• <b>Remote System ID Changed</b>—Adjacency peer system ID changed.</li> <li>• <b>Protocol Shutdown</b>—IS-IS protocol is disabled.</li> <li>• <b>CLI Command</b>—Adjacency brought down by user.</li> <li>• <b>Unknown</b>—Unknown.</li> </ul> </li> </ul> | extensive       |

## Sample Output

### show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up          23

```

### show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 3988](#).

### show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
```

### show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
Transition log:
When           State      Event      Down reason
Wed Nov  8 21:24:25  Up        Seenself
```

## show isis authentication

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 3990</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 3990</a>   |
| <b>Syntax</b>                                     | <pre>show isis authentication &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis authentication &lt;instance <i>instance-name</i>&gt;</pre>  |
| <b>Release Information</b>                        | <p>Command introduced in Junos OS Release 7.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for hitless authentication key rollover introduced in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>                                | Display information about IS-IS authentication.  |
| <b>Options</b>                                    | <p><b>none</b>—Display information about IS-IS authentication.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view   |
| <b>List of Sample Output</b>                      | <a href="#">show isis authentication on page 3991</a><br><a href="#">show isis authentication (With Hitless Authentication Key Rollover Configured) on page 3991</a>   |
| <b>Output Fields</b>                              | <p><a href="#">Table 312 on page 3990</a> describes the output fields for the <b>show isis authentication</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

**Table 312: show isis authentication Output Fields**

| Field Name       | Field Description  |
|------------------|--|
| <b>Interface</b> | Interface name.  |
| <b>Level</b>     | IS-IS level.   |
| <b>IIH Auth</b>  | <p>IS-IS Hello (IIH) packet authentication type.</p> <p>Displays the name of the active keychain if hitless authentication key rollover is configured.</p> |
| <b>CSN Auth</b>  | Complete sequence number authentication type.  |



Table 312: show isis authentication Output Fields (*continued*)

| Field Name            | Field Description                            |
|-----------------------|--|
| PSN Auth              | Partial sequence number authentication type. |
| L1 LSP Authentication | Layer 1 link-state PDU authentication type.  |
| L2 LSP Authentication | Layer 2 link-state PDU authentication type.  |

## Sample Output

### show isis authentication

```

user@host> show isis authentication
Interface      Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0     1      Simple    Simple    Simple
                2      MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

### show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface      Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0     2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

## show isis backup coverage

|   |   |
|---|---|
| <b>Syntax</b>                                     | <b>show isis backup coverage</b><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <b>show isis backup coverage</b><br><instance <i>instance-name</i> >  |
| <b>Release Information</b>                        | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>                                | Display information about the level of backup coverage available.   |
| <b>Options</b>                                    | <p><b>none</b>—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> |
| <b>Required Privilege Level</b>                   | view  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 3883</a></li> <li>• <a href="#">show isis backup label-switched-path on page 3994</a></li> </ul>  |
| <b>List of Sample Output</b>                      | <a href="#">show isis backup coverage on page 3993</a>  |
| <b>Output Fields</b>                              | Table 313 on page 3992 lists the output fields for the <b>show isis backup coverage</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 313: show isis backup coverage Output Fields**

| Field Name      | Field Description  |
|-----------------|--|
| <b>Topology</b> | Type of topology or address family: <b>IPV4 Unicast</b> or <b>IPV6 Unicast</b> .                             |
| <b>Level</b>    | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1</li> <li>• 2—Level 2</li> </ul>              |
| <b>Node</b>     | By topology, the percentage of all routes configured on the node that are protected through backup coverage. |

Table 313: show isis backup coverage Output Fields (*continued*)

| Field Name | Field Description  |
|------------|--|
| IPv4       | Percentage of IPv4 unicast routes that are protected through backup coverage.                          |
| IPv6       | Percentage of IPv6 unicast routes that are protected through backup coverage.                          |
| CLNS       | Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage. |

## Sample Output

show isis backup coverage

```
user@host> show isis backup coverage
Backup Coverage:
  Topology  Level  Node   IPv4   IPv6   CLNS
  IPV4 Unicast    2  28.57% 22.22% 0.00% 0.00%
  IPV6 Unicast    2   0.00% 0.00% 0.00% 0.00%
```

## show isis backup label-switched-path

|   |   |
|---|---|
| <b>Syntax</b>                                     | <b>show isis backup label-switched-path</b><br><b>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</b>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <b>show isis backup label-switched-path</b>   |
| <b>Release Information</b>                        | Command introduced in Junos OS Release 9.5.<br>Command introduced in Junos OS Release 9.5 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>                                | Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.  |
| <b>Options</b>                                    | <b>none</b> —Display information about MPLS LSPs designated as backup routes for IS-IS routes.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 3883</a></li> <li>• <a href="#">show isis backup coverage on page 3992</a></li> </ul>   |
| <b>List of Sample Output</b>                      | <a href="#">show isis backup label-switched-path on page 3995</a>   |
| <b>Output Fields</b>                              | <a href="#">Table 314 on page 3994</a> lists the output fields for the <b>show isis backup label-switched-path</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 314: show isis backup label-switched-path Output Fields**

| Field Name              | Field Description   |
|-------------------------|---|
| <b>Backup MPLS LSPs</b> | List of MPLS LSPs designated as backup paths for IS-IS routes.  |
| <b>Egress</b>           | IP address of the egress routing device for the LSP.  |
| <b>Status</b>           | State of the LSP: <ul style="list-style-type: none"> <li>• <b>Up</b>—The routing device can detect RSVP hello messages from the neighbor.</li> <li>• <b>Down</b>—The routing device has received one of the following indications:               <ul style="list-style-type: none"> <li>• Communication failure from the neighbor.</li> <li>• Communication from IGP that the neighbor is unavailable.</li> <li>• Change in the sequence numbers in the RSVP hello messages sent by the neighbor.</li> </ul> </li> <li>• <b>Deleted</b>—LSP is no longer available as a backup path.</li> </ul> |

Table 314: show isis backup label-switched-path Output Fields (*continued*)

| Field Name  | Field Description  |
|-------------|--|
| Last change | Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> . |
| TE-metric   | Configured traffic engineering metric.   |
| Metric      | Configured metric.   |

## Sample Output

### show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0
```

## show isis backup spf results

---

|                             |  |
|-----------------------------|--|
| Syntax                      | <pre>show isis backup spf results &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;no-coverage&gt; &lt;topology (ipv4-unicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>  |
| Syntax (EX Series Switches) | <pre>show isis backup spf results &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;no-coverage&gt; &lt;topology (ipv4-unicast   unicast)&gt;</pre>  |
| Release Information         | Command introduced in Junos OS Release 9.5.  |
| Description                 | Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.   |
| Options                     | <p><b>none</b>—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p><b>level (1   2)</b>—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p><b>no-coverage</b>—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display SPF calculations for backup paths for the specified topology only.</p> |
| Required Privilege Level    | view   |
| Related Documentation       | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Link and Node Protection for IS-IS Routes on page 3883</a></li><li>• <a href="#">show isis backup coverage on page 3992</a></li></ul>   |
| List of Sample Output       | <ul style="list-style-type: none"><li>• <a href="#">show isis backup spf results on page 3997</a></li><li>• <a href="#">show isis backup spf results no-coverage on page 3998</a></li></ul>  |
| Output Fields               | <p><a href="#">Table 315 on page 3997</a> lists the output fields for the <b>show isis backup spf results</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

Table 315: show isis backup spf results Output Fields

| Field Name              | Field Description  |
|-------------------------|--|
| <i>node-name</i>        | Name of the destination node.  |
| <b>Address</b>          | Address of the destination node.   |
| <b>Primary next-hop</b> | Interface and name of the node of the primary next hop to reach the destination.                   |
| <b>Root</b>             | Name of the next-hop neighbor.   |
| <b>Metric</b>           | Metric to the node.  |
| <b>Eligible</b>         | Indicates that the next-hop neighbor has been designated as a backup path to the destination node. |
| <b>Backup next-hop</b>  | Name of the interface of the backup next hop.  |
| <b>SNPA</b>             | Subnetwork point of attachment (MAC address of the next hop).                                      |
| <b>LSP</b>              | Name of the MPLS label-switched path (LSP) designated as a backup path.                            |
| <b>Not eligible</b>     | Indicates that the next-hop neighbor cannot function as a backup path to the destination.          |
| <b>Reason</b>           | Describes why the next-hop neighbor is designated as <b>Not eligible</b> as a backup path.         |

## Sample Output

### show isis backup spf results

```

user@host> show isis backup spf results
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0

```

```

        track-item: pro-bng3-j.00-00
        track-item: pro-bng3-i.00-00
        Not eligible, IPV4, Reason: Path loops
        Not eligible, IPV6, Reason: Path loops
    2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
2 nodes

```

### show isis backup spf results no-coverage

```

user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21

```



```
Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes
```

## show isis database

---

|  |   |
|--|---|
| List of Syntax                             | <a href="#">Syntax on page 4000</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4000</a>  |
| Syntax                                     | <pre>show isis database &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>  |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis database &lt;system-id&gt; &lt;brief   detail   extensive&gt; &lt;level (1   2)&gt; &lt;instance <i>instance-name</i>&gt;</pre>  |
| Release Information                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| Description                                | Display the entries in the IS-IS link-state database, which contains data about PDU packets.  |
| Options                                    | <p><b>none</b>—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><b><i>system id</i></b>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p><b><i>brief   detail   extensive</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>instance instance-name</i></b>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p><b><i>level (1   2)</i></b>—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p><b><i>logical-system (all   logical-system-name)</i></b>—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p> |
| Required Privilege Level                   | view  |
| Related Documentation                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis database on page 3980</a></li></ul>  |
| List of Sample Output                      | <a href="#">show isis database on page 4002</a><br><a href="#">show isis database brief on page 4003</a><br><a href="#">show isis database detail on page 4003</a>  |

[show isis database extensive on page 4003](#)

**Output Fields** [Table 316 on page 4001](#) describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

**Table 316: show isis database Output Fields**

| Field Name            | Field Description  | Level of Output         |
|-----------------------|--|-------------------------|
| <b>Interface name</b> | Name of the interface on which the link-state PDU has been received; always <b>IS-IS</b> for this command.   | All levels              |
| level                 | Level of intermediate system: <ul style="list-style-type: none"> <li>• <b>1</b>—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system.</li> <li>• <b>2</b>—Intermediate system routes between areas and toward other ASs.</li> </ul>   | All levels              |
| LSP ID                | Link-state PDU identifier.   | All levels              |
| Sequence              | Sequence number of the link-state PDU.   | All levels              |
| Checksum              | Checksum value of the link-state PDU.  | All levels              |
| Lifetime (secs)       | Remaining lifetime of the link-state PDU, in seconds.  | All levels              |
| Attributes            | Attributes of the specified database: <b>L1</b> , <b>L2</b> , <b>Overload</b> , or <b>Attached</b> (L1 only).  | none <b>brief</b>       |
| # LSPs                | Total number of link-state PDUs in the specified link-state database.  | none <b>brief</b>       |
| IP prefix             | Prefix advertised by this link-state PDU.  | <b>detail extensive</b> |
| IS neighbor           | IS-IS neighbor of the advertising system.  | <b>detail extensive</b> |
| ES neighbor           | (J Series routers only) An ES-IS neighbor of the advertising system.   | <b>detail extensive</b> |
| IP prefix             | IPv4 prefix advertised by this link-state PDU.   | <b>detail extensive</b> |
| V6 prefix             | IPv6 prefix advertised by this link-state PDU.   | <b>detail extensive</b> |
| Metric                | Metric of the prefix or neighbor.  | <b>detail extensive</b> |
| Header                | <ul style="list-style-type: none"> <li>• <b>LSP ID</b>—Link state PDU identifier of the header.</li> <li>• <b>Length</b>—Header length.</li> <li>• <b>Allocated Length</b>—Amount of length available for the header.</li> <li>• <b>Router ID</b>—Address of the local routing device.</li> <li>• <b>Remaining Lifetime</b>—Remaining lifetime of the link-state PDU, in seconds.</li> </ul> | <b>extensive</b>        |

Table 316: show isis database Output Fields (*continued*)

| Field Name | Field Description  | Level of Output |
|------------|--|-----------------|
| Packet     | <ul style="list-style-type: none"> <li>• <b>LSP ID</b>—The identifier for the link-state PDU.</li> <li>• <b>Length</b>—Packet length.</li> <li>• <b>Lifetime</b>—Remaining lifetime, in seconds.</li> <li>• <b>Checksum</b>—The checksum of the link-state PDU.</li> <li>• <b>Sequence</b>—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments.</li> <li>• <b>Attributes</b>—Packet attributes.</li> <li>• <b>NLPID</b>—Network layer protocol identifier.</li> <li>• <b>Fixed length</b>—Specifies the set length for the packet.</li> </ul>  | extensive       |
| TLVs       | <ul style="list-style-type: none"> <li>• <b>Area Address</b>—Area addresses that the routing device can reach.</li> <li>• <b>Speaks</b>—Supported routing protocols.</li> <li>• <b>IP router id</b>—ID of the routing device (usually the IP address).</li> <li>• <b>IP address</b>—IPv4 address.</li> <li>• <b>Hostname</b>—Assigned name of the routing device.</li> <li>• <b>IP prefix</b>—IP prefix of the routing device.</li> <li>• <b>Metric</b>—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device.</li> <li>• <b>IP extended prefix</b>—Extended IP prefix of the routing device.</li> <li>• <b>IS neighbor</b>—Directly attached neighbor's name and metric.</li> <li>• <b>IS extended neighbor</b>—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index.</li> </ul> <p>The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.</p> | extensive       |

## Sample Output

### show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3    0x3167    1057 L1 L2
camaro.00-00          0x5    0x770e    1091 L1 L2
ranier.00-00          0x4    0xaa95    1091 L1 L2
glacier.00-00         0x4    0x206f    1089 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x3    0x87a2    1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6    0x8d6b    1096 L1 L2
camaro.00-00          0x9    0x877b    1101 L1 L2
ranier.00-00          0x8    0x855d    1103 L1 L2
glacier.00-00         0x7    0xf892    1098 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x6    0x562     1105 L1 L2
  6 LSPs

```

### show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 4002](#).

### show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  ES neighbor: 0015.0015.0015           Metric:      10 Down
  ES neighbor: 0025.0025.0025           Metric:      10 Down
  ES neighbor: 0030.0030.0030           Metric:      10 Down
  ES neighbor: 0040.0040.0040           Metric:      10 Down
  ES neighbor: sisira                     Metric:       0
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 50.50.50.50/32              Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up
```

### show isis database extensive

```
user@host> show isis database extensive
```

IS-IS level 1 link-state database:

```
Router-A.00-00 Sequence: 0x1, Checksum: 0xf75c, Lifetime: 1116 secs
```

IP prefix: 192.168.0.1/32                      Metric:              0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 85 bytes  
Allocated length: 1492 bytes, Router ID: 192.168.0.1  
Remaining lifetime: 1116 secs, Level: 1, Interface: 0  
Estimated free bytes: 1353, Actual free bytes: 1407  
Aging timer expires in: 1116 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 85 bytes, Lifetime : 1200 secs  
Checksum: 0xf75c, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 18, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.1  
IP address: 192.168.0.1  
Hostname: Router-A  
IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up  
IP extended prefix: 192.168.0.1/32 metric 0 up  
No queued transmissions

IS-IS level 2 link-state database:

Router-A.00-00 Sequence: 0x5, Checksum: 0x3196, Lifetime: 1158 secs  
IS neighbor: Router-B.02                      Metric:              10  
Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00  
IS neighbor: Router-E.02                      Metric:              10  
Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00  
IP prefix: 10.0.0.0/30                      Metric:              10 Internal Up  
IP prefix: 10.0.0.4/30                      Metric:              10 Internal Up  
IP prefix: 192.168.0.1/32                      Metric:              0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 208 bytes  
Allocated length: 1492 bytes, Router ID: 192.168.0.1  
Remaining lifetime: 1158 secs, Level: 2, Interface: 0  
Estimated free bytes: 1233, Actual free bytes: 1284  
Aging timer expires in: 1158 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 208 bytes, Lifetime : 1198 secs  
Checksum: 0x3196, Sequence: 0x5, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.1  
IP address: 192.168.0.1  
Hostname: Router-A  
IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up  
IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up  
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up  
IP extended prefix: 192.168.0.1/32 metric 0 up

```

IP extended prefix: 10.0.0.4/30 metric 10 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-B.02, Internal, Metric: default 10
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.1
    Local interface index: 101, Remote interface index: 0
IS extended neighbor: Router-B.02, Metric: default 10
  IP address: 10.0.0.5
    Local interface index: 102, Remote interface index: 0
No queued transmissions

Router-B.00-00 Sequence: 0x5, Checksum: 0xf8f, Lifetime: 1183 secs
  IS neighbor: Router-B.02                      Metric:      10
    Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
  IS neighbor: Router-C.02                      Metric:      10
    Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
IP prefix: 10.0.0.4/30                          Metric:      10 Internal Up
IP prefix: 10.0.0.8/30                          Metric:      10 Internal Up
IP prefix: 192.168.0.2/32                      Metric:      0 Internal Up

Header: LSP ID: Router-B.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.2
  Remaining lifetime: 1183 secs, Level: 2, Interface: 102
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1183 secs
  Protocols: IP, IPv6

Packet: LSP ID: Router-B.00-00, Length: 208 bytes, Lifetime : 1196 secs
  Checksum: 0xf8f, Sequence: 0x5, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0002 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 192.168.0.2
  IP address: 192.168.0.2
  Hostname: Router-B
  IP prefix: 192.168.0.2/32, Internal, Metric: default 0, Up
  IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
  IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
  IP extended prefix: 192.168.0.2/32 metric 0 up
  IP extended prefix: 10.0.0.4/30 metric 10 up
  IP extended prefix: 10.0.0.8/30 metric 10 up
  IS neighbor: Router-B.02, Internal, Metric: default 10
  IS neighbor: Router-C.02, Internal, Metric: default 10
  IS extended neighbor: Router-B.02, Metric: default 10
    IP address: 10.0.0.6
      Local interface index: 108, Remote interface index: 0
  IS extended neighbor: Router-C.02, Metric: default 10
    IP address: 10.0.0.9
      Local interface index: 109, Remote interface index: 0
No queued transmissions

Router-B.02-00 Sequence: 0x1, Checksum: 0x3c7c, Lifetime: 1156 secs
  IS neighbor: Router-A.00                      Metric:      0
    Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
  IS neighbor: Router-B.00                      Metric:      0

```

Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00

Header: LSP ID: Router-B.02-00, Length: 76 bytes  
Allocated length: 284 bytes, Router ID: 0.0.0.0  
Remaining lifetime: 1156 secs, Level: 2, Interface: 102  
Estimated free bytes: 208, Actual free bytes: 208  
Aging timer expires in: 1156 secs

Packet: LSP ID: Router-B.02-00, Length: 76 bytes, Lifetime : 1196 secs  
Checksum: 0x3c7c, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
IS neighbor: Router-B.00, Internal, Metric: default 0  
IS neighbor: Router-A.00, Internal, Metric: default 0  
IS extended neighbor: Router-B.00, Metric: default 0  
IS extended neighbor: Router-A.00, Metric: default 0  
No queued transmissions

Router-C.00-00 Sequence: 0x5, Checksum: 0x255b, Lifetime: 1182 secs  
IS neighbor: Router-C.02 Metric: 10  
Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00  
IS neighbor: Router-D.03 Metric: 10  
Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00  
IP prefix: 10.0.0.8/30 Metric: 10 Internal Up  
IP prefix: 10.0.0.12/30 Metric: 10 Internal Up  
IP prefix: 192.168.0.3/32 Metric: 0 Internal Up

Header: LSP ID: Router-C.00-00, Length: 208 bytes  
Allocated length: 284 bytes, Router ID: 192.168.0.3  
Remaining lifetime: 1182 secs, Level: 2, Interface: 102  
Estimated free bytes: 114, Actual free bytes: 76  
Aging timer expires in: 1182 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-C.00-00, Length: 208 bytes, Lifetime : 1196 secs  
Checksum: 0x255b, Sequence: 0x5, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.3  
IP address: 192.168.0.3  
Hostname: Router-C  
IP prefix: 192.168.0.3/32, Internal, Metric: default 0, Up  
IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up  
IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up  
IP extended prefix: 192.168.0.3/32 metric 0 up  
IP extended prefix: 10.0.0.8/30 metric 10 up  
IP extended prefix: 10.0.0.12/30 metric 10 up  
IS neighbor: Router-C.02, Internal, Metric: default 10  
IS neighbor: Router-D.03, Internal, Metric: default 10  
IS extended neighbor: Router-C.02, Metric: default 10  
IP address: 10.0.0.10  
Local interface index: 105, Remote interface index: 0  
IS extended neighbor: Router-D.03, Metric: default 10



IP address: 10.0.0.13  
 Local interface index: 106, Remote interface index: 0  
 No queued transmissions

Router-C.02-00 Sequence: 0x1, Checksum: 0xaa09, Lifetime: 1181 secs  
 IS neighbor: Router-B.00 Metric: 0  
 Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00  
 IS neighbor: Router-C.00 Metric: 0  
 Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00

Header: LSP ID: Router-C.02-00, Length: 76 bytes  
 Allocated length: 284 bytes, Router ID: 0.0.0.0  
 Remaining lifetime: 1181 secs, Level: 2, Interface: 102  
 Estimated free bytes: 208, Actual free bytes: 208  
 Aging timer expires in: 1181 secs

Packet: LSP ID: Router-C.02-00, Length: 76 bytes, Lifetime : 1194 secs  
 Checksum: 0xaa09, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
 IS neighbor: Router-C.00, Internal, Metric: default 0  
 IS neighbor: Router-B.00, Internal, Metric: default 0  
 IS extended neighbor: Router-C.00, Metric: default 0  
 IS extended neighbor: Router-B.00, Metric: default 0  
 No queued transmissions

Router-D.00-00 Sequence: 0x4, Checksum: 0x8ab7, Lifetime: 1180 secs  
 IS neighbor: Router-D.02 Metric: 10  
 Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00  
 IS neighbor: Router-D.03 Metric: 10  
 Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00  
 IP prefix: 10.0.0.12/30 Metric: 10 Internal Up  
 IP prefix: 10.0.0.20/30 Metric: 10 Internal Up  
 IP prefix: 192.168.0.4/32 Metric: 0 Internal Up

Header: LSP ID: Router-D.00-00, Length: 208 bytes  
 Allocated length: 284 bytes, Router ID: 192.168.0.4  
 Remaining lifetime: 1180 secs, Level: 2, Interface: 102  
 Estimated free bytes: 114, Actual free bytes: 76  
 Aging timer expires in: 1180 secs  
 Protocols: IP, IPv6

Packet: LSP ID: Router-D.00-00, Length: 208 bytes, Lifetime : 1192 secs  
 Checksum: 0x8ab7, Sequence: 0x4, Attributes: 0x3 <L1 L2>  
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
 Area address: 49.0002 (3)  
 LSP Buffer Size: 1492  
 Speaks: IP  
 Speaks: IPV6  
 IP router id: 192.168.0.4  
 IP address: 192.168.0.4  
 Hostname: Router-D  
 IP prefix: 192.168.0.4/32, Internal, Metric: default 0, Up  
 IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up  
 IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up  
 IP extended prefix: 192.168.0.4/32 metric 0 up

IP extended prefix: 10.0.0.12/30 metric 10 up  
IP extended prefix: 10.0.0.20/30 metric 10 up  
IS neighbor: Router-D.02, Internal, Metric: default 10  
IS neighbor: Router-D.03, Internal, Metric: default 10  
IS extended neighbor: Router-D.02, Metric: default 10  
IP address: 10.0.0.22  
Local interface index: 115, Remote interface index: 0  
IS extended neighbor: Router-D.03, Metric: default 10  
IP address: 10.0.0.14  
Local interface index: 114, Remote interface index: 0  
No queued transmissions

Router-D.02-00 Sequence: 0x1, Checksum: 0xebbc, Lifetime: 1128 secs  
IS neighbor: Router-D.00 Metric: 0  
Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00  
IS neighbor: Router-F.00 Metric: 0  
Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-D.02-00, Length: 76 bytes  
Allocated length: 284 bytes, Router ID: 0.0.0.0  
Remaining lifetime: 1128 secs, Level: 2, Interface: 101  
Estimated free bytes: 208, Actual free bytes: 208  
Aging timer expires in: 1128 secs

Packet: LSP ID: Router-D.02-00, Length: 76 bytes, Lifetime : 1160 secs  
Checksum: 0xebbc, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
IS neighbor: Router-D.00, Internal, Metric: default 0  
IS neighbor: Router-F.00, Internal, Metric: default 0  
IS extended neighbor: Router-D.00, Metric: default 0  
IS extended neighbor: Router-F.00, Metric: default 0  
No queued transmissions

Router-D.03-00 Sequence: 0x1, Checksum: 0x129b, Lifetime: 1180 secs  
IS neighbor: Router-C.00 Metric: 0  
Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00  
IS neighbor: Router-D.00 Metric: 0  
Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00

Header: LSP ID: Router-D.03-00, Length: 76 bytes  
Allocated length: 284 bytes, Router ID: 0.0.0.0  
Remaining lifetime: 1180 secs, Level: 2, Interface: 101  
Estimated free bytes: 208, Actual free bytes: 208  
Aging timer expires in: 1180 secs

Packet: LSP ID: Router-D.03-00, Length: 76 bytes, Lifetime : 1192 secs  
Checksum: 0x129b, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:  
IS neighbor: Router-D.00, Internal, Metric: default 0  
IS neighbor: Router-C.00, Internal, Metric: default 0  
IS extended neighbor: Router-D.00, Metric: default 0  
IS extended neighbor: Router-C.00, Metric: default 0  
No queued transmissions

Router-E.00-00 Sequence: 0x4, Checksum: 0x9da9, Lifetime: 1155 secs

```

IS neighbor: Router-E.02                      Metric:      10
  Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
IS neighbor: Router-F.02                      Metric:      20
  Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
IP prefix: 10.0.0.0/30                        Metric:      10 Internal Up
IP prefix: 10.0.0.16/30                       Metric:      20 Internal Up
IP prefix: 192.168.0.5/32                     Metric:       0 Internal Up

```

```

Header: LSP ID: Router-E.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.5
  Remaining lifetime: 1155 secs, Level: 2, Interface: 101
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1155 secs
  Protocols: IP, IPv6

```

```

Packet: LSP ID: Router-E.00-00, Length: 208 bytes, Lifetime : 1185 secs
  Checksum: 0x9da9, Sequence: 0x4, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

#### TLVs:

```

Area address: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 192.168.0.5
IP address: 192.168.0.5
Hostname: Router-E
IP prefix: 192.168.0.5/32, Internal, Metric: default 0, Up
IP prefix: 10.0.0.16/30, Internal, Metric: default 20, Up
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
IP extended prefix: 192.168.0.5/32 metric 0 up
IP extended prefix: 10.0.0.16/30 metric 20 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-F.02, Internal, Metric: default 20
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.2
  Local interface index: 112, Remote interface index: 0
IS extended neighbor: Router-F.02, Metric: default 20
  IP address: 10.0.0.17
  Local interface index: 111, Remote interface index: 0
No queued transmissions

```

```

Router-E.02-00 Sequence: 0x1, Checksum: 0xb4fa, Lifetime: 1130 secs
IS neighbor: Router-A.00                      Metric:       0
  Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
IS neighbor: Router-E.00                      Metric:       0
  Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00

```

```

Header: LSP ID: Router-E.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1130 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1130 secs

```

```

Packet: LSP ID: Router-E.02-00, Length: 76 bytes, Lifetime : 1161 secs
  Checksum: 0xb4fa, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

## TLVs:

IS neighbor: Router-E.00, Internal, Metric: default 0  
IS neighbor: Router-A.00, Internal, Metric: default 0  
IS extended neighbor: Router-E.00, Metric: default 0  
IS extended neighbor: Router-A.00, Metric: default 0

No queued transmissions

Router-F.00-00 Sequence: 0x5, Checksum: 0x94bd, Lifetime: 1153 secs  
IS neighbor: Router-D.02 Metric: 10  
Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00  
IS neighbor: Router-F.02 Metric: 10  
Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00  
IP prefix: 10.0.0.16/30 Metric: 10 Internal Up  
IP prefix: 10.0.0.20/30 Metric: 10 Internal Up  
IP prefix: 192.168.0.6/32 Metric: 0 Internal Up

Header: LSP ID: Router-F.00-00, Length: 208 bytes  
Allocated length: 284 bytes, Router ID: 192.168.0.6  
Remaining lifetime: 1153 secs, Level: 2, Interface: 101  
Estimated free bytes: 76, Actual free bytes: 76  
Aging timer expires in: 1153 secs  
Protocols: IP, IPv6

Packet: LSP ID: Router-F.00-00, Length: 208 bytes, Lifetime : 1183 secs  
Checksum: 0x94bd, Sequence: 0x5, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

## TLVs:

Area address: 49.0002 (3)  
LSP Buffer Size: 1492  
Speaks: IP  
Speaks: IPV6  
IP router id: 192.168.0.6  
IP address: 192.168.0.6  
Hostname: Router-F  
IP prefix: 192.168.0.6/32, Internal, Metric: default 0, Up  
IP prefix: 10.0.0.16/30, Internal, Metric: default 10, Up  
IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up  
IP extended prefix: 192.168.0.6/32 metric 0 up  
IP extended prefix: 10.0.0.16/30 metric 10 up  
IP extended prefix: 10.0.0.20/30 metric 10 up  
IS neighbor: Router-D.02, Internal, Metric: default 10  
IS neighbor: Router-F.02, Internal, Metric: default 10  
IS extended neighbor: Router-D.02, Metric: default 10  
IP address: 10.0.0.21  
Local interface index: 94, Remote interface index: 0  
IS extended neighbor: Router-F.02, Metric: default 10  
IP address: 10.0.0.18  
Local interface index: 93, Remote interface index: 0

No queued transmissions

Router-F.02-00 Sequence: 0x1, Checksum: 0xf5ae, Lifetime: 1153 secs  
IS neighbor: Router-E.00 Metric: 0  
Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00  
IS neighbor: Router-F.00 Metric: 0  
Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-F.02-00, Length: 76 bytes  
Allocated length: 284 bytes, Router ID: 0.0.0.0  
Remaining lifetime: 1153 secs, Level: 2, Interface: 101

Estimated free bytes: 208, Actual free bytes: 208  
Aging timer expires in: 1153 secs

Packet: LSP ID: Router-F.02-00, Length: 76 bytes, Lifetime : 1183 secs  
Checksum: 0xf5ae, Sequence: 0x1, Attributes: 0x3 <L1 L2>  
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes  
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

IS neighbor: Router-F.00, Internal, Metric: default 0  
IS neighbor: Router-E.00, Internal, Metric: default 0  
IS extended neighbor: Router-F.00, Metric: default 0  
IS extended neighbor: Router-E.00, Metric: default 0

No queued transmissions

## show isis hostname

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4012</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4012</a>  |
| <b>Syntax</b>                                     | <pre>show isis hostname &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show isis hostname  |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>                                | <p>Display IS-IS hostname database information.</p> <p>This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the <b>set system host-name</b> command, or a static mapping defined in Junos OS with the <b>set system static-host-mapping <i>hostname</i> sysid</b> command (type static). The local router always has its type set to static even if <b>static-host-mapping</b> is not configured.</p> |
| <b>Options</b>                                    | <p><b>none</b>—Display IS-IS hostname database information.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>   |
| <b>Required Privilege Level</b>                   | view  |
| <b>List of Sample Output</b>                      | <a href="#">show isis hostname on page 4013</a>   |
| <b>Output Fields</b>                              | <p><a href="#">Table 317 on page 4012</a> describes the output fields for the <b>show isis hostname</b> command. Output fields are listed in the approximate order in which they appear.</p>  |

**Table 317: show isis hostname Output Fields**

| Field Name       | Field Description   |
|------------------|---|
| <b>System Id</b> | System identifier mapped to the hostname.   |
| <b>Hostname</b>  | Hostname mapped to the system identifier.   |
| <b>Type</b>      | <p>Type of mapping between system identifier and hostname.</p> <ul style="list-style-type: none"> <li><b>Dynamic</b>—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>.</li> <li><b>Static</b>—Hostname mapping configured by user.</li> </ul> |


## Sample Output

show isis hostname

```
user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname
1921.6800.4201 isis1
1921.6800.4202 isis2
1921.6800.4203 isis3
```

|                      | Type    |
|----------------------|---------|
| 1921.6800.4201 isis1 | Dynamic |
| 1921.6800.4202 isis2 | Static  |
| 1921.6800.4203 isis3 | Dynamic |

```
show isis interface
```

|  |   |
|--|---|
| List of Syntax                             | <a href="#">Syntax on page 4014</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4014</a>  |
| Syntax                                     | <pre>show isis interface &lt;brief   detail   extensive&gt; &lt;interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>   |
| Syntax (EX Series Switches and QFX Series) | <pre>show isis interface &lt;brief   detail   extensive&gt; &lt;interface-name&gt;</pre>  |
| Release Information                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| Description                                | <p>Display status information about IS-IS-enabled interfaces.</p>   |
|  | <div>  <p><b>NOTE:</b> If the configured metric for an IS-IS level is above 63, and the <b>wide-metrics-only</b> statement is not configured, the <b>show isis interface detail</b> command and the <b>show isis interface extensive</b> command display 63 as the metric value for that level. Configure the <b>wide-metrics-only</b> statement to generate metric values greater than 63 on a per IS-IS level basis.</p> <p>The <b>show isis interface</b> command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p> </div> |
| Options                                    | <p><b>none</b>—Display standard information about all IS-IS-enabled interfaces.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified interface only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>   |
| Required Privilege Level                   | view  |
| Related Documentation                      | <ul style="list-style-type: none"> <li><i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i></li> </ul>   |
| List of Sample Output                      | <a href="#">show isis interface on page 4016</a><br><a href="#">show isis interface brief on page 4017</a><br><a href="#">show isis interface detail on page 4017</a><br><a href="#">show isis interface extensive on page 4017</a>   |



**Output Fields** Table 318 on page 4015 describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

**Table 318: show isis interface Output Fields**

| Field Name               | Field Description  | Level of Output  |
|--------------------------|--|------------------|
| <i>interface-name</i>    | Name of the interface.   | detail           |
| <b>Designated router</b> | Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.   | detail           |
| <b>Index</b>             | Interface index assigned by the Junos OS kernel.   | detail           |
| <b>State</b>             | Internal implementation information.   | detail           |
| <b>Circuit id</b>        | Circuit identifier.  | detail           |
| <b>Circuit type</b>      | Circuit type: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul>  | detail           |
| <b>LSP interval</b>      | Interval between link-state PDUs sent from the interface.  | detail           |
| <b>CSNP interval</b>     | Interval between complete sequence number PDUs sent from the interface.  | detail extensive |
| <b>Sysid</b>             | System identifier.   | detail           |
| <b>Interface</b>         | Interface through which the adjacency is made.   | none brief       |
| <b>L or Level</b>        | Level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> <p><b>NOTE:</b> The default IS-IS level on loopback interfaces are always same as the IS-IS level configured on other IS-IS interfaces in a router. You can also configure IS-IS level on loopback interfaces per your requirement.</p> | All levels       |
| <b>CirID</b>             | Circuit identifier.  | none brief       |
| <b>Level 1 DR</b>        | Level 1 designated intermediate system.  | none brief       |
| <b>Level 2 DR</b>        | Level 2 designated intermediate system.  | none brief       |
| <b>L1/L2 Metric</b>      | Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.   | none brief       |

Table 318: show isis interface Output Fields (*continued*)

| Field Name                                | Field Description   | Level of Output         |
|---|---|-------------------------|
| <b>Adjacency advertisement: Advertise</b> | This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).   | <b>detail extensive</b> |
| <b>Adjacency advertisement: Suppress</b>  | This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.   | <b>detail extensive</b> |
| <b>Adjacencies</b>                        | Number of adjacencies established on this interface.  | <b>detail</b>           |
| <b>Priority</b>                           | Priority value for this interface.  | <b>detail</b>           |
| <b>Metric</b>                             | Metric value for this interface.  | <b>detail</b>           |
| <b>Hello(s) / Hello Interval</b>          | Interface's hello interval.   | <b>detail extensive</b> |
| <b>Hold(s) / Hold Time</b>                | Interface's hold time.  | <b>detail extensive</b> |
| <b>Designated Router</b>                  | Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.   | <b>detail</b>           |
| <b>Hello padding</b>                      | Type of hello padding: <ul style="list-style-type: none"> <li>• <b>Adaptive</b>—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface.</li> <li>• <b>Loose</b>—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state.</li> <li>• <b>Strict</b>—Padding is performed on all interface types and for all adjacency states, and is continuous.</li> </ul> | <b>extensive</b>        |
| <b>LDP sync state</b>                     | Current LDP synchronization state: <b>in sync</b> , <b>in holddown</b> , or <b>not supported</b> .  | <b>extensive</b>        |
| <b>reason</b>                             | Reason for being in the LDP sync state.   | <b>extensive</b>        |
| <b>config holdtime</b>                    | Configured value of the hold timer.   | <b>extensive</b>        |
| <b>remaining</b>                          | If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.   | <b>extensive</b>        |

## Sample Output

### show isis interface

```
user@host> show isis interface
```

IS-IS interface database:

| Interface  | L | CirID | Level 1 DR     | Level 2 DR     | L1/L2 Metric |
|------------|---|-------|----------------|----------------|--------------|
| at-2/3/0.0 | 3 | 0x1   | Point to Point | Point to Point | 10/10        |
| lo0.0      | 3 | 0x1   | Passive        | Passive        | 0/0          |

### show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 4016](#).

### show isis interface detail

```
user@host> show isis interface detail
```

IS-IS interface database:

at-2/3/0.0

Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3

LSP interval: 100 ms, CSNP interval: 5 s

| Level | Adjacencies | Priority | Metric | Hello (s) | Hold (s) | Designated Router |
|-------|-------------|----------|--------|-----------|----------|-------------------|
|-------|-------------|----------|--------|-----------|----------|-------------------|

|   |   |    |    |       |    |  |
|---|---|----|----|-------|----|--|
| 1 | 1 | 64 | 10 | 9.000 | 27 |  |
|---|---|----|----|-------|----|--|

|   |   |    |    |       |    |  |
|---|---|----|----|-------|----|--|
| 2 | 1 | 64 | 10 | 9.000 | 27 |  |
|---|---|----|----|-------|----|--|

lo0.0

Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0

LSP interval: 100 ms, CSNP interval: disabled

Adjacency advertisement: Advertise

Protection Type: Node Link, No eligible Backup

| Level | Adjacencies | Priority | Metric | Hello (s) | Hold (s) | Designated Router |
|-------|-------------|----------|--------|-----------|----------|-------------------|
|-------|-------------|----------|--------|-----------|----------|-------------------|

|   |   |    |   |         |  |  |
|---|---|----|---|---------|--|--|
| 1 | 0 | 64 | 0 | Passive |  |  |
|---|---|----|---|---------|--|--|

|   |   |    |   |         |  |  |
|---|---|----|---|---------|--|--|
| 2 | 0 | 64 | 0 | Passive |  |  |
|---|---|----|---|---------|--|--|

### show isis interface extensive

```
user@host> show isis interface extensive
```

IS-IS interface database:

xe-6/1/0.0

Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2

LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding

Adjacency advertisement: Advertise

Level 1

Adjacencies: 0, Priority: 64, Metric: 10

Disabled

Level 2

Adjacencies: 1, Priority: 64, Metric: 10

Hello Interval: 20.000 s, Hold Time: 60 s

Designated Router: nemean.03

## show isis overview

|   |   |
|---|---|
| <b>Syntax</b>                                     | <b>show isis overview</b><br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <b>show isis overview</b><br><instance <i>instance-name</i> >   |
| <b>Release Information</b>                        | Command introduced in Junos OS Release 8.5.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>                                | Display IS-IS overview information.   |
| <b>Options</b>                                    | <b>none</b> —Display standard overview information about IS-IS for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display overview information for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view  |
| <b>List of Sample Output</b>                      | <a href="#">show isis overview on page 4020</a>   |
| <b>Output Fields</b>                              | <a href="#">Table 319 on page 4018</a> lists the output fields for the <b>show isis overview</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 319: show isis overview Output Fields**

| Field Name              | Field Description  |
|-------------------------|--|
| Instance                | IS-IS routing instance.  |
| Router ID               | Router ID of the routing device.   |
| Adjacency holddown      | Adjacency holddown capability: <b>enabled</b> or <b>disabled</b> .   |
| Maximum Areas           | Maximum number of IS-IS areas advertised by the routing device.  |
| LSP life time           | Lifetime of the link-state PDU, in seconds.  |
| Attached bit evaluation | Attached bit capability: <b>enabled</b> or <b>disabled</b> .   |
| SPF delay               | Delay before performing consecutive shortest-path-first (SPF) calculations.  |
| SPF holddown            | Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached. |

Table 319: show isis overview Output Fields (*continued*)

| Field Name                     | Field Description   |
|--------------------------------|---|
| SPF rapid runs                 | Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.                |
| Overload bit at startup is set | Overload bit capability is enabled.   |
| Overload high metrics          | Overload high metrics capability: <b>enabled</b> or <b>disabled</b> .   |
| Overload timeout               | Time period after which overload is reset and the time that remains before the timer is set to expire.                  |
| Traffic engineering            | Traffic engineering capability: <b>enabled</b> or <b>disabled</b> .   |
| Restart                        | Graceful restart capability: <b>enabled</b> or <b>disabled</b> .  |
| Restart duration               | Time period for complete reacquisition of IS-IS neighbors.  |
| Helper mode                    | Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .   |
| Level                          | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1 information</li> <li>• 2—Level 2 information</li> </ul> |
| IPv4 is enabled                | IP Protocol version 4 capability is enabled.  |
| IPv6 is enabled                | IP Protocol version 6 capability is enabled.  |
| CLNS is enabled                | (J Series routers only) OSI CLNP capability is enabled.   |
| Internal route preference      | Preference value of internal routes.  |
| External route preference      | Preference value of external routes.  |
| Prefix export limit            | Number of prefixes allowed to be exported, as configured by the <a href="#">prefix-export-limit</a> statement.          |
| Prefix export count            | Number of prefixes exported.  |
| Wide area metrics are enabled  | Wide area metrics capability is enabled.  |
| Narrow metrics are enabled     | Narrow metrics capability is enabled.   |

## Sample Output

### show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 10.255.107.183
  Adjacency holddown: disabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export limit: 5, Prefix export count: 5
  Wide metrics are enabled
```

## show isis route

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4021</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4021</a>  |
| <b>Syntax</b>                                     | <pre>show isis route &lt;destination&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show isis route &lt;destination&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>   |
| <b>Release Information</b>                        | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>                                | Display the routes in the IS-IS routing table.  |
| <b>Options</b>                                    | <p><b>none</b>—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p><b><i>destination</i></b>—(Optional) Destination address for the route.</p> <p><b>inet   inet6</b>—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routes for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p> |
| <b>Required Privilege Level</b>                   | view  |
| <b>List of Sample Output</b>                      | <a href="#">show isis route logical-system on page 4022</a><br><a href="#">show isis route (CLNS) on page 4022</a><br><a href="#">show isis route on page 4023</a>  |
| <b>Output Fields</b>                              | <p><a href="#">Table 320 on page 4022</a> describes the output fields for the <b>show isis route</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

Table 320: show isis route Output Fields

| Field Name             | Field Description  |
|------------------------|--|
| <b>Current version</b> | Number of the current version of the IS-IS routing table.  |
| <b>L1</b>              | Version of Level 1 SPF that was run.   |
| <b>L2</b>              | Version of Level 2 SPF that was run.   |
| <b>Prefix</b>          | Destination of the route.  |
| <b>L</b>               | IS-IS level: <ul style="list-style-type: none"> <li>• 1—Level 1 only</li> <li>• 2—Level 2 only</li> <li>• 3—Level 1 and Level 2</li> </ul> |
| <b>Version</b>         | Version of SPF that generated the route.   |
| <b>Metric</b>          | Metric value associated with the route.  |
| <b>Type</b>            | Metric type: <b>int</b> (internal) or <b>ext</b> (external).   |
| <b>Interface</b>       | Interface to the next hop.   |
| <b>Via</b>             | System identifier of the next hop, displayed as a name if possible.  |
| <b>ISO Routes</b>      | ISO routing table entries.   |
| <b>snpa</b>            | MAC address.   |

## Sample Output

### show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table           Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2    11    20 int  gr-0/2/0.0  h
10.9.201.1/32 2    11    60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table   Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2    11    20 int  gr-0/2/0.0  h

```

### show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table           Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0   1    10    10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa

```



```

0/0
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
  1      10      20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
  1      10      0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
  1      10      10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

### show isis route

```
user@host> show isis route
```

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                 IPV4 camaro
10.255.71.238/32           2    13        20   int  so-6/0/0.0           IPV4 olympic
                           as0.0                 IPV4 glacier
10.255.71.239/32           2    13        20   int  so-6/0/0.0           IPV4 olympic
                           ae0.0                 IPV4 camaro
10.255.71.242/32           2    13        10   int  as0.0                 IPV4 glacier
10.255.71.243/32           2    13        10   int  so-6/0/0.0           IPV4 olympic
12.13.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
12.15.0.0/30                2    13        20   int  so-6/0/0.0           IPV4 olympic
13.15.0.0/30                2    13        30   int  ae0.0                 IPV4 camaro
                           so-6/0/0.0           IPV4 olympic
                           as0.0                 IPV4 glacier
13.16.0.0/30                2    13        25   int  as0.0                 IPV4 glacier
14.15.0.0/30                2    13        20   int  ae0.0                 IPV4 camaro
192.2.1.0/30                2    13        30   int  so-6/0/0.0           IPV4 olympic
                           as0.0                 IPV4 glacier
1eee::/64                   2    13        30   int  so-6/0/0.0           IPV6 olympic
                           as0.0                 IPV6 glacier
abcd::10:255:71:52/128     2    13        10   int  ae0.0                 IPV6 camaro
abcd::10:255:71:238/128   2    13        20   int  so-6/0/0.0           IPV6 olympic

```

|                         |   |    |    |     |            |              |
|-------------------------|---|----|----|-----|------------|--------------|
|                         |   |    |    |     | as0.0      | IPV6 glacier |
| abcd::10:255:71:239/128 | 2 | 13 | 20 | int | so-6/0/0.0 | IPV6 olympic |
|                         |   |    |    |     | ae0.0      | IPV6 camaro  |
| abcd::10:255:71:242/128 | 2 | 13 | 10 | int | as0.0      | IPV6 glacier |
| abcd::10:255:71:243/128 | 2 | 13 | 10 | int | so-6/0/0.0 | IPV6 olympic |

## show isis spf

|                                    |   |
|------------------------------------|---|
| <b>List of Syntax</b>              | <a href="#">Syntax on page 4025</a><br><a href="#">Syntax (EX Series Switches) on page 4025</a>   |
| <b>Syntax</b>                      | <pre>show isis spf (brief   log   results) &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>  |
| <b>Syntax (EX Series Switches)</b> | <pre>show isis spf (brief   log   results) &lt;instance <i>instance-name</i>&gt; &lt;level (1   2)&gt; &lt;topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)&gt;</pre>  |
| <b>Release Information</b>         | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>                 | Display information about IS-IS shortest-path-first (SPF) calculations.   |
| <b>Options</b>                     | <p><b>brief</b>—Display an overview of SPF calculations.</p> <p><b>instance <i>instance instance-name</i></b>—(Optional) Display SPF calculations for the specified routing instance.</p> <p><b>level (1   2)</b>—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p><b>log</b>—Display the log of SPF calculations.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>results</b>—Display the results of SPF calculations.</p> <p><b>topology (ipv4-multicast   ipv6-multicast   ipv6-unicast   unicast)</b>—(Optional) Display SPF calculations for the specified topology only.</p> |
| <b>Required Privilege Level</b>    | view  |
| <b>List of Sample Output</b>       | <a href="#">show isis spf log on page 4026</a><br><a href="#">show isis spf results logical-system on page 4027</a><br><a href="#">show isis spf results (CLNS) on page 4028</a>  |
| <b>Output Fields</b>               | <p><a href="#">Table 321 on page 4025</a> describes the output fields for the <b>show isis spf</b> command. Output fields are listed in the approximate order in which they appear.</p>   |

**Table 321: show isis spf Output Fields**

| Field Name | Field Description    |
|------------|----------------------|
| Node       | System ID of a node. |

Table 321: show isis spf Output Fields (*continued*)

| Field Name     | Field Description   |
|----------------|---|
| Metric         | Metric to the node.   |
| Interface      | Interface of the next hop.  |
| Via            | System ID of the next hop.  |
| SNPA           | Subnetwork point of attachment (MAC address of the next hop).                           |
| Start time     | (log option only) Time that the SPF computation started.                                |
| Elapsed (secs) | (log option only) Length of time, in seconds, required to complete the SPF computation. |
| Count          | (log option only) Number of times the SPF was triggered.                                |
| Reason         | (log option only) Reason that the SPF computation was completed.                        |

## Sample Output

### show isis spf log

```

user@host> show isis spf log logical-system lsl
IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000069    1 Reconfig
Fri Oct 31 12:41:18   0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000050    3 Address change on so-1/2/0.0
Fri Oct 31 12:41:23   0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58   0.000198    1 Periodic SPF
Fri Oct 31 13:10:29   0.000209    1 Periodic SPF
IS-IS level 2 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000035    1 Reconfig
Fri Oct 31 12:41:18   0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23   0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37   0.000195    1 Periodic SPF
Fri Oct 31 12:55:50   0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14   0.000198    1 Periodic SPF
IPv6 Unicast IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000028    1 Reconfig
Fri Oct 31 12:41:18   0.000043    3 Updated LSP fix.00-00

```

```

Fri Oct 31 12:41:18      0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000041    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28      0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59      0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30      0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08      0.000289    1 Periodic SPF
Fri Oct 31 13:11:07      0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

```

```

Start time      Elapsed (secs) Count Reason
Fri Oct 31 12:41:18      0.000027    1 Reconfig
Fri Oct 31 12:41:18      0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18      0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59      0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30      0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50      0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55      0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58      0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46      0.000201    1 Periodic SPF
...

```

#### show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
```

```
IS-IS level 1 SPF results:
```

| Node    | Metric | Interface     | Via  | SNPA             |
|---------|--------|---------------|------|------------------|
| scat.00 | 10     | ge-1/1/0.0    | scat | 0:90:69:a6:48:9d |
|         | 20     | 10.9.1.0/30   |      |                  |
| fix.02  | 10     |               |      |                  |
| fix.00  | 0      |               |      |                  |
|         | 10     | 10.9.1.0/30   |      |                  |
|         | 10     | 10.9.5.0/30   |      |                  |
|         | 10     | 10.9.6.0/30   |      |                  |
|         | 20     | 10.9.7.0/30   |      |                  |
|         | 60     | 10.9.201.1/32 |      |                  |
| 3 nodes |        |               |      |                  |

```
IS-IS level 2 SPF results:
```

| Node    | Metric | Interface     | Via | SNPA |
|---------|--------|---------------|-----|------|
| skag.00 | 20     | gr-0/2/0.0    | h   |      |
|         | 30     | 10.9.7.0/30   |     |      |
| skag.02 | 20     | gr-0/2/0.0    | h   |      |
| h.00    | 10     | gr-0/2/0.0    | h   |      |
|         | 20     | 10.9.6.0/30   |     |      |
|         | 20     | 10.9.7.0/30   |     |      |
|         | 60     | 10.9.201.1/32 |     |      |
| fix.00  | 0      |               |     |      |
|         | 10     | 10.9.1.0/30   |     |      |
|         | 10     | 10.9.5.0/30   |     |      |
|         | 10     | 10.9.6.0/30   |     |      |
| 4 nodes |        |               |     |      |

```
IPV6 Unicast IS-IS level 1 SPF results:
```

| Node    | Metric | Interface            | Via  | SNPA             |
|---------|--------|----------------------|------|------------------|
| scat.00 | 10     | ge-1/1/0.0           | scat | 0:90:69:a6:48:9d |
|         |        | ge-1/1/0.0           | scat | 0:90:69:a6:48:9d |
|         | 20     | 8009:1::a09:1400/126 |      |                  |
| fix.02  | 10     |                      |      |                  |

```

fix.00          0
                10      8009:1::a09:1400/126
                10      8009:2::a09:1e00/126
                20      8009:3::a09:3200/126
                10      8009:4::a09:2800/126
3 nodes

IPv6 Unicast IS-IS level 2 SPF results:
Node      Metric  Interface      Via      SNPA
skag.00    20      gr-0/2/0.0     h
           30      8009:3::a09:3200/126
skag.02    20      gr-0/2/0.0     h
           20      gr-0/2/0.0     h
h.00       10      gr-0/2/0.0     h
           20      8009:3::a09:3200/126
           20      8009:4::a09:2800/126
fix.00     0
           10      8009:1::a09:1400/126
           10      8009:2::a09:1e00/126
           10      8009:4::a09:2800/126
4 nodes

Multicast IS-IS level 1 SPF results:
Node      Metric  Interface      Via      SNPA
scat.00    10      ge-1/1/0.0     scat     0:90:69:a6:48:9d
fix.02     10
fix.00     0
3 nodes

Multicast IS-IS level 2 SPF results:
Node      Metric  Interface      Via      SNPA
skag.00    20      gr-0/2/0.0     h
skag.02    20      gr-0/2/0.0     h
h.00       10      gr-0/2/0.0     h
fix.00     0
4 nodes
...

```

**show isis spf results (CLNS)**

```

user@host> show isis spf results
IS-IS level 1 SPF results:
Node      Metric  Interface      Via      SNPA
skag.00 10      fe-0/0/1.0     toothache 0:12:0:34:0:56
           20      fe-0/0/1.0     toothache 0:12:0:34:0:56
           10      192.168.37.64/29
           20      192.168.37.64/29
           10      192.168.37.64/29
pro1-a.02 10
pro1-a.00 0
           0      10.255.245.1/32
           10      192.168.37.64/29
           0      192.168.37.64/29
3 nodes

IS-IS level 2 SPF results:
Node      Metric  Interface      Via      SNPA
skag.00 10      fe-0/0/1.0     toothache 0:12:0:34:0:56
           0      fe-0/0/1.0     toothache 0:12:0:34:0:56

```

|           |    |                                      |
|-----------|----|--------------------------------------|
|           | 20 | 10.255.245.1/32                      |
|           | 20 | 192.168.37.64/29                     |
|           | 20 | 47.0005.80ff.f800.0000.0109.0010/104 |
| pro1-a.02 | 10 |                                      |
| pro1-a.00 | 0  |                                      |
|           | 0  | 10.255.245.1/32                      |
|           | 10 | 192.168.37.64/29                     |
| 3 nodes   |    |                                      |

## show isis statistics

---

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4030</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4030</a>   |
| <b>Syntax</b>                                     | show isis statistics<br><instance <i>instance-name</i> ><br><logical-system (all   <i>logical-system-name</i> )>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show isis statistics<br><instance <i>instance-name</i> >   |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>                                | Display statistics about IS-IS traffic.  |
| <b>Options</b>                                    | <b>none</b> —Display IS-IS traffic statistics for all routing instances.<br><br><b>instance <i>instance-name</i></b> —(Optional) Display statistics for the specified routing instance.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear isis statistics on page 3984</a></li></ul>   |
| <b>List of Sample Output</b>                      | <a href="#">show isis statistics on page 4032</a>  |
| <b>Output Fields</b>                              | <a href="#">Table 322 on page 4031</a> describes the output fields for the <b>show isis statistics</b> command. Output fields are listed in the approximate order in which they appear.  |



Table 322: show isis statistics Output Fields

| Field Name                  | Field Description   |
|-----------------------------|---|
| PDU type                    | <p>PDU type:</p> <ul style="list-style-type: none"> <li>• <b>CSNP</b>—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU.</li> <li>• <b>IIH</b>—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems.</li> <li>• <b>LSP</b>—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area.</li> <li>• <b>PSNP</b>—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device.</li> <li>• <b>Unknown</b>—The PDU type is unknown.</li> </ul> |
| Received                    | Number of PDUs received since IS-IS started or since the statistics were set to zero.   |
| Processed                   | Number of PDUs received less the number dropped.  |
| Drops                       | Number of PDUs dropped.   |
| Sent                        | Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.  |
| Rexmit                      | Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.  |
| Total packets received/sent | Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.   |
| SNP queue length            | Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.   |
| LSP queue length            | Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.   |
| SPF runs                    | Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.   |
| Fragments rebuilt           | Number of link-state PDU fragments that the local system has computed.  |
| LSP regenerations           | Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.  |
| Purges initiated            | Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.  |

## Sample Output

### show isis statistics

```
user@host> show isis statistics
```

```
IS-IS statistics for merino:
```

| PDU type | Received | Processed | Drops | Sent   | Rexmit |
|----------|----------|-----------|-------|--------|--------|
| LSP      | 12227    | 12227     | 0     | 8184   | 683    |
| IIH      | 113808   | 113808    | 0     | 115817 | 0      |
| CSNP     | 198868   | 198868    | 0     | 198934 | 0      |
| PSNP     | 6985     | 6979      | 6     | 8274   | 0      |
| Unknown  | 0        | 0         | 0     | 0      | 0      |
| Totals   | 331888   | 331882    | 6     | 331209 | 683    |

```
Total packets received: 331888 Sent: 331892
```

```
SNP queue length:      0 Drops:      0  
LSP queue length:      0 Drops:      0
```

```
SPF runs:              1014  
Fragments rebuilt:     1038  
LSP regenerations:     425  
Purges initiated:      0
```

## PART 13

# Open Shortest Path First

- [Overview on page 4035](#)
- [Configuration on page 4047](#)
- [Administration on page 4243](#)



## CHAPTER 47

# Overview

- [OSPF Overview on page 4035](#)

### OSPF Overview

---

- [OSPF Overview on page 4036](#)
- [OSPF Areas and Router Functionality Overview on page 4041](#)
- [Packets Overview on page 4043](#)
- [OSPF External Metrics Overview on page 4046](#)

## OSPF Overview

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

Junos OS supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including virtual links, stub areas, and for OSPFv2, authentication. Junos OS does not support type-of-service (ToS) routing.

OSPF was designed for the Transmission Control Protocol/Internet Protocol (TCP/IP) environment and as a result explicitly supports IP subnetting and the tagging of externally derived routing information. OSPF also provides for the authentication of routing updates.

OSPF routes IP packets based solely on the destination IP address contained in the IP packet header. OSPF quickly detects topological changes, such as when router interfaces become unavailable, and calculates new loop-free routes quickly and with a minimum of routing overhead traffic.



**NOTE:** On SRX Series devices, when only one link-protection is configured under the OSPF interface, the device does not install an alternative route in the forwarding table. When the per-packet load-balancing is enabled as a workaround, the device does not observe both the OSPF metric and sending the traffic through both the interfaces.

An OSPF AS can consist of a single area, or it can be subdivided into multiple areas. In a single-area OSPF network topology, each router maintains a database that describes the topology of the AS. Link-state information for each router is flooded throughout the AS. In a multiarea OSPF topology, each router maintains a database that describes the topology of its area, and link-state information for each router is flooded throughout that area. All routers maintain summarized topologies of other areas within an AS. Within each area, OSPF routers have identical topological databases. When the AS or area topology changes, OSPF ensures that the contents of all routers' topological databases converge quickly.

All OSPFv2 protocol exchanges can be authenticated. OSPFv3 relies on IPsec to provide this functionality. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used. A single authentication scheme is configured for each area, which enables some areas to use stricter authentication than others.

Externally derived routing data (for example, routes learned from BGP) is passed transparently throughout the AS. This externally derived data is kept separate from the OSPF link-state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.



**NOTE:** By default, Junos OS is compatible with RFC 1583, *OSPF Version 2*. In Junos OS Release 8.5 and later, you can disable compatibility with RFC 1583 by including the `no-rfc-1583` statement. For more information, see [“Example: Disabling OSPFv2 Compatibility with RFC 1583” on page 4074](#).

This topic describes the following information:

- [OSPF Default Route Preference Values on page 4038](#)
- [OSPF Routing Algorithm on page 4038](#)
- [OSPF Three-Way Handshake on page 4039](#)
- [OSPF Version 3 on page 4040](#)

### OSPF Default Route Preference Values

The Junos OS routing protocol process assigns a default preference value to each route that the routing table receives. The default value depends on the source of the route. The preference value is from 0 through 4,294,967,295 ( $2^{32} - 1$ ), with a lower value indicating a more preferred route. [Table 323 on page 4038](#) lists the default preference values for OSPF.

**Table 323: Default Route Preference Values for OSPF**

| How Route Is Learned    | Default Preference | Statement to Modify Default Preference |
|-------------------------|--------------------|--|
| OSPF internal route     | 10                 | OSPF <code>preference</code>           |
| OSPF AS external routes | 150                | OSPF <code>external-preference</code>  |

### OSPF Routing Algorithm

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topological databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a routing device starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The routing device then uses the OSPF hello protocol to acquire neighbors, by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routing devices), the OSPF hello protocol elects a designated router for the network. This routing device is responsible for sending *link-state advertisements* (LSAs) that describe the network, which reduces the amount of network traffic and the size of the routing devices' topological databases.

The routing device then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated



router form adjacencies with other routing devices.) Adjacencies determine the distribution of routing protocol packets. Routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A routing device sends LSA packets to advertise its state periodically and when its state changes. These packets include information about the routing device's adjacencies, which allows detection of nonoperational routing devices.

Using a reliable algorithm, the routing device floods LSAs throughout the area, which ensures that all routing devices in an area have exactly the same topological database. Each routing device uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The routing device then uses this tree to route network traffic.


The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area (*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. The area border routers use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

Autonomous system (AS) boundary routers flood information about external autonomous systems throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

### OSPF Three-Way Handshake

OSPF creates a topology map by flooding LSAs across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in [Figure 112 on page 4039](#).

**Figure 112: OSPF Three-Way Handshake**



In [Figure 112 on page 4039](#), Router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that Router B can receive traffic from Router A. Router B generates a response to Router A to acknowledge receipt of the hello packet. When Router A receives the response, it establishes that Router B can receive traffic from Router A. Router A then generates a final response packet to inform Router B that Router A can receive traffic from Router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

### OSPF Version 3

---

OSPFv3 is a modified version of OSPF that supports IP version 6 (IPv6) addressing. OSPFv3 differs from OSPFv2 in the following ways:

- All neighbor ID information is based on a 32-bit router ID.
- The protocol runs per link rather than per subnet.
- Router and network link-state advertisements (LSAs) do not carry prefix information.
- Two new LSA types are included: link-LSA and intra-area-prefix-LSA.
- Flooding scopes are as follows:
  - Link-local
  - Area
  - AS
- Link-local addresses are used for all neighbor exchanges except virtual links.
- Authentication is removed. The IPv6 authentication header relies on the IP layer.
- The packet format has changed as follows:
  - Version number 2 is now version number 3.
  - The **db** option field has been expanded to 24 bits.
  - Authentication information has been removed.
  - Hello messages do not have address information.
  - Two new option bits are included: **R** and **V6**.
- Type 3 summary LSAs have been renamed *inter-area-prefix-LSAs*.
- Type 4 summary LSAs have been renamed *inter-area-router-LSAs*.

#### Related Documentation

- [Understanding OSPF Areas and Backbone Areas on page 4052](#)
- [OSPF Configuration Overview](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4074](#)

## OSPF Areas and Router Functionality Overview

In OSPF, a single autonomous system (AS) can be divided into smaller groups called *areas*. This reduces the number of link-state advertisements (LSAs) and other OSPF overhead traffic sent on the network, and it reduces the size of the topology database that each router must maintain. The routing devices that participate in OSPF routing perform one or more functions based on their location in the network.

This topic describes the following OSPF area types and routing device functions:

- [Areas on page 4041](#)
- [Area Border Routers on page 4041](#)
- [Backbone Areas on page 4041](#)
- [AS Boundary Routers on page 4042](#)
- [Backbone Router on page 4042](#)
- [Internal Router on page 4042](#)
- [Stub Areas on page 4042](#)
- [Not-So-Stubby Areas on page 4042](#)
- [Transit Areas on page 4043](#)

---

### Areas

An *area* is a set of networks and hosts within an AS that have been administratively grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Also, routing within the area is determined only by the area's topology, providing the area with some protection from bad routing data.

All routing devices within an area have identical topology databases.

---

### Area Border Routers

Routing devices that belong to more than one area and connect one or more OSPF areas to the backbone area are called *area border routers* (ABRs). At least one interface is within the backbone while another interface is in another area. ABRs also maintain a separate topological database for each area to which they are connected.

---

### Backbone Areas

An OSPF *backbone area* consists of all networks in area ID 0.0.0.0, their attached routing devices, and all ABRs. The backbone itself does not have any ABRs. The backbone distributes routing information between areas. The backbone is simply another area, so the terminology and rules of areas apply: a routing device that is directly connected to the backbone is an internal router on the backbone, and the backbone's topology is hidden from the other areas in the AS.

The routing devices that make up the backbone must be physically contiguous. If they are not, you must configure *virtual links* to create the appearance of backbone connectivity. You can create virtual links between any two ABRs that have an interface to a common nonbackbone area. OSPF treats two routing devices joined by a virtual link as if they were connected to an unnumbered point-to-point network.

---

### AS Boundary Routers

Routing devices that exchange routing information with routing devices in non-OSPF networks are called *AS boundary routers*. They advertise externally learned routes throughout the OSPF AS. Depending on the location of the AS boundary router in the network, it can be an ABR, a backbone router, or an internal router (with the exception of stub areas). Internal routers within a stub area cannot be an AS boundary router because stub areas cannot contain any Type 5 LSAs.

Routing devices within the area where the AS boundary router resides know the path to that AS boundary router. Any routing device outside the area only knows the path to the nearest ABR that is in the same area where the AS boundary router resides.

---

### Backbone Router

*Backbone routers* are routing devices that have one or more interfaces connected to the OSPF backbone area (area ID 0.0.0.0).

---

### Internal Router

Routing devices that connect to only one OSPF area are called *internal routers*. All interfaces on internal routers are directly connected to networks within a single area.

---

### Stub Areas

*Stub areas* are areas through which or into which AS external advertisements are not flooded. You might want to create stub areas when much of the topological database consists of AS external advertisements. Doing so reduces the size of the topological databases and therefore the amount of memory required on the internal routers in the stub area.

Routing devices within a stub area rely on the default routes originated by the area's ABR to reach external AS destinations. You must configure the **default-metric** option on the ABR before it advertises a default route. Once configured, the ABR advertises a default route in place of the external routes that are not being advertised within the stub area, so that routing devices in the stub area can reach destinations outside the area.

The following restrictions apply to stub areas: you cannot create a virtual link through a stub area, a stub area cannot contain an AS boundary router, the backbone cannot be a stub area, and you cannot configure an area as both a stub area and a not-so-stubby area.

---

### Not-So-Stubby Areas

An OSPF stub area has no external routes in it, so you cannot redistribute from another protocol into a stub area. A *not-so-stubby area* (NSSA) allows external routes to be

flooded within the area. These routes are then leaked into other areas. However, external routes from other areas still do not enter the NSSA.

The following restriction applies to NSSAs: you cannot configure an area as both a stub area and an NSSA.

### Transit Areas

*Transit areas* are used to pass traffic from one adjacent area to the backbone (or to another area if the backbone is more than two hops away from an area). The traffic does not originate in, nor is it destined for, the transit area.

#### Related Documentation

- [OSPF Overview on page 4036](#)
- [Packets Overview on page 4043](#)
- [OSPF Configuration Overview](#)
- [Understanding OSPF Areas and Backbone Areas on page 4052](#)
- [Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas on page 4059](#)

## Packets Overview

There are several types of link-state advertisement (LSA) packets.

This topic describes the following information:

- [OSPF Packet Header on page 4043](#)
- [Hello Packets on page 4044](#)
- [Database Description Packets on page 4044](#)
- [Link-State Request Packets on page 4044](#)
- [Link-State Update Packets on page 4044](#)
- [Link-State Acknowledgment Packets on page 4045](#)
- [Link-State Advertisement Packet Types on page 4045](#)

### OSPF Packet Header

All OSPFv2 packets have a common 24-byte header, and OSPFv3 packets have a common 16-byte header, that contains all information necessary to determine whether OSPF should accept the packet. The header consists of the following fields:

- Version number—The current OSPF version number. This can be either **2** or **3**.
- Type—Type of OSPF packet.
- Packet length—Length of the packet, in bytes, including the header.
- Router ID—IP address of the router from which the packet originated.
- Area ID—Identifier of the area in which the packet is traveling. Each OSPF packet is associated with a single area. Packets traveling over a virtual link are labeled with the backbone area ID, 0.0.0.0.

- Checksum—Fletcher checksum.
- Authentication—(OSPFv2 only) Authentication scheme and authentication information.
- Instance ID—(OSPFv3 only) Identifier used when there are multiple OSPFv3 realms configured on a link.

### Hello Packets

---

Routers periodically send hello packets on all interfaces, including virtual links, to establish and maintain neighbor relationships. Hello packets are multicast on physical networks that have a multicast or broadcast capability, which enables dynamic discovery of neighboring routers. (On nonbroadcast networks, dynamic neighbor discovery is not possible, so you must configure all neighbors statically as described in [“Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network” on page 4079.](#))

Hello packets consist of the OSPF header plus the following fields:

- Network mask—(OSPFv2 only) Network mask associated with the interface.
- Hello interval—How often the router sends hello packets. All routers on a shared network must use the same hello interval.
- Options—Optional capabilities of the router.
- Router priority—The router’s priority to become the designated router.
- Router dead interval—How long the router waits without receiving any OSPF packets from a router before declaring that router to be down. All routers on a shared network must use the same router dead interval.
- Designated router—IP address of the designated router.
- Backup designated router—IP address of the backup designated router.
- Neighbor—IP addresses of the routers from which valid hello packets have been received within the time specified by the router dead interval.

### Database Description Packets

---

When initializing an adjacency, OSPF exchanges database description packets, which describe the contents of the topological database. These packets consist of the OSPF header, packet sequence number, and the link-state advertisement’s header.

### Link-State Request Packets

---

When a router detects that portions of its topological database are out of date, it sends a link-state request packet to a neighbor requesting a precise instance of the database. These packets consist of the OSPF header plus fields that uniquely identify the database information that the router is seeking.

### Link-State Update Packets

---

Link-state update packets carry one or more link-state advertisements one hop farther from their origin. The router multicasts (floods) these packets on physical networks that support multicast or broadcast mode. The router acknowledges all link-state update

packets and, if retransmission is necessary, sends the retransmitted advertisements unicast.

Link-state update packets consist of the OSPF header plus the following fields:

- Number of advertisements—Number of link-state advertisements included in this packet.
- Link-state advertisements—The link-state advertisements themselves.

### Link-State Acknowledgment Packets

The router sends link-state acknowledgment packets in response to link-state update packets to verify that the update packets have been received successfully. A single acknowledgment packet can include responses to multiple update packets.

Link-state acknowledgment packets consist of the OSPF header plus the link-state advertisement header.

### Link-State Advertisement Packet Types

Link-state request, link-state update, and link-state acknowledgment packets are used to reliably flood link-state advertisement packets. OSPF sends the following types of link-state advertisements:

- Router link advertisements—Are sent by all routers to describe the state and cost of the router's links to the area. These link-state advertisements are flooded throughout a single area only.
- Network link advertisements—Are sent by designated routers to describe all the routers attached to the network. These link-state advertisements are flooded throughout a single area only.
- Summary link advertisements—Are sent by area border routers to describe the routes that they know about in other areas. There are two types of summary link advertisements: those used when the destination is an IP network, and those used when the destination is an AS boundary router. Summary link advertisements describe interarea routes, that is, routes to destinations outside the area but within the AS. These link-state advertisements are flooded throughout the advertisement's associated areas.
- AS external link advertisement—Are sent by AS boundary routers to describe external routes that they know about. These link-state advertisements are flooded throughout the AS (except for stub areas).

Each link-state advertisement type describes a portion of the OSPF routing domain. All link-state advertisements are flooded throughout the AS.

Each link-state advertisement packet begins with a common 20-byte header.

#### Related Documentation

- [OSPF Overview on page 4036](#)
- [OSPF Areas and Router Functionality Overview on page 4041](#)
- [OSPF Configuration Overview](#)

- [OSPF Designated Router Overview on page 4047](#)
- [Understanding OSPFv2 Authentication](#)
- [OSPF Timers Overview on page 4117](#)

## OSPF External Metrics Overview

When OSPF exports route information from external autonomous systems (ASs), it includes a cost, or *external metric*, in the route. OSPF supports two types of external metrics: Type 1 and Type 2. The difference between the two metrics is how OSPF calculates the cost of the route. Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router. Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router. By default, OSPF uses the Type 2 external metric.



## CHAPTER 48

# Configuration

- [Basic OSPF Area Configuration on page 4047](#)
- [Advanced OSPF Area Configuration on page 4058](#)
- [OSPF Interface Configuration on page 4075](#)
- [OSPF Route Control Configuration on page 4094](#)
- [OSPF Fault Detection Configuration on page 4117](#)
- [OSPF Redundancy Features Configuration on page 4134](#)
- [OSPF Traffic Engineering Configuration on page 4150](#)
- [OSPF Database Protection Configuration on page 4162](#)
- [OSPF Policy Configuration on page 4164](#)
- [OSPF Monitoring Configuration on page 4198](#)
- [Configuration Statements on page 4205](#)

### Basic OSPF Area Configuration

---

- [Examples: Configuring OSPF Designated Routers on page 4047](#)
- [Examples: Configuring OSPF Areas on page 4052](#)

### Examples: Configuring OSPF Designated Routers

- [OSPF Designated Router Overview on page 4047](#)
- [Example: Configuring an OSPF Router Identifier on page 4048](#)
- [Example: Controlling OSPF Designated Router Election on page 4050](#)

#### OSPF Designated Router Overview

---

Large LANs that have many routing devices and therefore many OSPF adjacencies can produce heavy control-packet traffic as link-state advertisements (LSAs) are flooded across the network. To alleviate the potential traffic problem, OSPF uses designated routers on all multiaccess networks (broadcast and nonbroadcast multiaccess [NBMA] networks types). Rather than broadcasting LSAs to all their OSPF neighbors, the routing devices send their LSAs to the designated router. Each multiaccess network has a designated router, which performs two main functions:

- Originate network link advertisements on behalf of the network.

- Establish adjacencies with all routing devices on the network, thus participating in the synchronizing of the link-state databases.

In LANs, the election of the designated router takes place when the OSPF network is initially established. When the first OSPF links are active, the routing device with the highest router identifier (defined by the **router-id** configuration value, which is typically the IP address of the routing device, or the loopback address) is elected the designated router. The routing device with the second highest router identifier is elected the backup designated router. If the designated router fails or loses connectivity, the backup designated router assumes its role and a new backup designated router election takes place between all the routers in the OSPF network.

OSPF uses the router identifier for two main purposes: to elect a designated router, unless you manually specify a priority value, and to identify the routing device from which a packet is originated. At designated router election, the router priorities are evaluated first, and the routing device with the highest priority is elected designated router. If router priorities tie, the routing device with the highest router identifier, which is typically the routing device's IP address, is chosen as the designated router. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

At least one routing device on each logical IP network or subnet must be eligible to be the designated router for OSPFv2. At least one routing device on each logical link must be eligible to be the designated router for OSPFv3.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router. A priority of 255 means the routing device is always the designated router.

---

### Example: Configuring an OSPF Router Identifier

---

This example shows how to configure an OSPF router identifier.

- [Requirements on page 4048](#)
- [Overview on page 4049](#)
- [Configuration on page 4049](#)
- [Verification on page 4050](#)

#### **Requirements**

Before you begin:

- Identify the interfaces on the routing device that will participate in OSPF. You must enable OSPF on all interfaces within the network on which OSPF traffic is to travel.
- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.

### Overview

The router identifier is used by OSPF to identify the routing device from which a packet originated. Junos OS selects a router identifier according to the following set of rules:

1. By default, Junos OS selects the lowest configured physical IP address of an interface as the router identifier.
2. If a loopback interface is configured, the IP address of the loopback interface becomes the router identifier.
3. If multiple loopback interfaces are configured, the lowest loopback address becomes the router identifier.
4. If a router identifier is explicitly configured using the **router-id address** statement under the **[edit routing-options]** hierarchy level, the above three rules are ignored.



**NOTE:** If the router identifier is modified in a network, the link-state advertisements (LSAs) advertised by the previous router identifier are retained in the OSPF database until the LSA retransmit interval has timed out.

If the router identifier is not configured explicitly and an interface IP address is used as the router identifier, the established OSPF adjacency flaps when the interface goes down, or when it is brought back into the network. When the interface is brought back into the network, or a new interface is introduced into the network, the router identifier is selected again based on the rules stated above. Hence, it is strongly recommended that you explicitly configure the router identifier under the **[edit routing-options]** hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.



**NOTE:** The router identifier behavior described here holds good even when configured under **[edit routing-instances routing-instance-name routing-options]** and **[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options]** hierarchy levels.

In this example, you configure the OSPF router identifier by setting its router ID value to the IP address of the device, which is 177.162.4.24.

### Configuration

#### CLI Quick Configuration

To quickly configure an OSPF router identifier, copy the following command and paste it into the CLI.

```
[edit]
set routing-options router-id 177.162.4.24
```

#### Step-by-Step Procedure

To configure an OSPF router identifier:

1. Configure the OSPF router identifier by entering the **[router-id]** configuration value.
- ```
[edit]
```

```
user@host# set routing-options router-id 177.162.4.24
```

2. If you are done configuring the device, commit the configuration.

```
[edit]  
user@host# commit
```

### Results

Confirm your configuration by entering the **show routing-options router-id** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options router-id  
router-id 177.162.4.24;
```

### Verification

After you configure the router ID and activate OSPF on the routing device, the router ID is referenced by multiple OSPF operational mode commands that you can use to monitor and troubleshoot the OSPF protocol. The router ID fields are clearly marked in the output.

### Example: Controlling OSPF Designated Router Election

---

This example shows how to control OSPF designated router election.

- [Requirements on page 4050](#)
- [Overview on page 4050](#)
- [Configuration on page 4050](#)
- [Verification on page 4051](#)

### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.

### Overview

This example shows how to control OSPF designated router election. Within the example, you set the OSPF interface to **ge-/0/0/1** and the device priority to 200. The higher the priority value, the greater likelihood the routing device will become the designated router.

By default, routing devices have a priority of 128. A priority of 0 marks the routing device as ineligible to become the designated router. A priority of 1 means the routing device has the least chance of becoming a designated router.

### Configuration

#### CLI Quick Configuration

To quickly configure an OSPF designated router election, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

**Step-by-Step Procedure** To control OSPF designated router election:

1. Configure an OSPF interface and specify the device priority.



**NOTE:** To specify an OSPFv3 interface, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.3 interface ge-0/0/1 priority 200
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

### Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.3 {
  interface ge-0/0/1.0 {
    priority 200;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

### Verification

Confirm that the configuration is working properly.

- [Verifying the Designated Router Election on page 4051](#)

### Verifying the Designated Router Election

**Purpose** Based on the priority you configured for a specific OSPF interface, you can confirm the address of the area's designated router. The DR ID, DR, or DR-ID field displays the address of the area's designated router. The BDR ID, BDR, or BDR-ID field displays the address of the backup designated router.

**Action** From operational mode, enter the `show ospf interface` and the `show ospf neighbor` commands for OSPFv2, and enter the `show ospf3 interface` and the `show ospf3 neighbor` commands for OSPFv3.

**Related Documentation** • [OSPF Areas and Router Functionality Overview on page 4041](#)

- [OSPF Configuration Overview](#)

## Examples: Configuring OSPF Areas


- [Understanding OSPF Areas and Backbone Areas on page 4052](#)
- [Example: Configuring a Single-Area OSPF Network on page 4053](#)
- [Example: Configuring a Multiarea OSPF Network on page 4055](#)

### Understanding OSPF Areas and Backbone Areas

OSPF networks in an autonomous system (AS) are administratively grouped into *areas*. Each area within an AS operates like an independent network and has a unique 32-bit area ID, which functions similar to a network address. Within an area, the topology database contains only information about the area, link-state advertisements (LSAs) are flooded only to nodes within the area, and routes are computed only within the area. The topology of an area is hidden from the rest of the AS, thus significantly reducing routing traffic in the AS. Subnetworks are divided into other areas, which are connected to form the whole of the main network. Routing devices that are wholly within an area are called *internal routers*. All interfaces on internal routers are directly connected to networks within the area.

The central area of an AS, called the *backbone area*, has a special function and is always assigned the area ID 0.0.0.0. (Within a simple, single-area network, this is also the ID of the area.) Area IDs are unique numeric identifiers, in dotted decimal notation, but they are not IP addresses. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by a routing device that has interfaces in more than one area. These connecting routing devices are called *area border routers* (ABRs). [Figure 113 on page 4052](#) shows an OSPF topology of three areas connected by two ABRs.

Figure 113: Multiarea OSPF Topology



Because all areas are adjacent to the backbone area, OSPF routers send all traffic not destined for their own area through the backbone area. The ABRs in the backbone area are then responsible for transmitting the traffic through the appropriate ABR to the destination area. The ABRs summarize the link-state records of each area and advertise destination address summaries to neighboring areas. The advertisements contain the ID of the area in which each destination lies, so that packets are routed to the appropriate ABR. For example, in the OSPF areas shown in [Figure 113 on page 4052](#), packets sent from Router A to Router C are automatically routed through ABR B.

Junos OS supports active backbone detection. Active backbone detection is implemented to verify that ABRs are connected to the backbone. If the connection to the backbone area is lost, then the routing device's default metric is not advertised, effectively rerouting traffic through another ABR with a valid connection to the backbone. Active backbone detection enables transit through an ABR with no active backbone connection. An ABR advertises to other routing devices that it is an ABR even if the connection to the backbone is down, so that the neighbors can consider it for interarea routes.

An OSPF restriction requires all areas to be directly connected to the backbone area so that packets can be properly routed. All packets are routed first to the backbone area by default. Packets that are destined for an area other than the backbone area are then routed to the appropriate ABR and on to the remote host within the destination area.

### Example: Configuring a Single-Area OSPF Network

This example shows how to configure a single-area OSPF network.

- [Requirements on page 4053](#)
- [Overview on page 4053](#)
- [Configuration on page 4054](#)
- [Verification on page 4055](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 4048.

#### Overview


To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

In an autonomous system (AS), the backbone area is always assigned area ID 0.0.0.0 (within a simple, single-area network, this is also the ID of the area). Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an

AS. All other networks or areas in the AS must be directly connected to the backbone area by area border routers that have interfaces in more than one area. You must also create a backbone area if your network consists of multiple areas. In this example, you create the backbone area and add interfaces, such as **ge-0/0/0**, as needed to the OSPF area.

To use OSPF on the device, you must configure at least one OSPF area, such as the one shown in [Figure 114 on page 4054](#).

**Figure 114: Typical Single-Area OSPF Network Topology**



#### Configuration

##### CLI Quick Configuration

To quickly configure a single-area OSPF network, copy the following command and paste it into the CLI. You repeat this configuration for all interfaces that are part of the OSPF area.

```
[edit]
set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

##### Step-by-Step Procedure

To configure a single-area OSPF network:

1. Configure the single-area OSPF network by specifying the area ID and associated interface.



**NOTE:** For a single-area OSPFv3 network, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```



### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Interfaces in the Area

- |                |                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured. |
| <b>Action</b>  | From operational mode, enter the <b>show ospf interface</b> command for OSPFv2, and enter the <b>show ospf3 interface</b> command for OSPFv3.                  |

### Example: Configuring a Multiarea OSPF Network

This example shows how to configure a multiarea OSPF network. To reduce traffic and topology maintenance for the devices in an OSPF autonomous system (AS), you can group the OSPF-enabled routing devices into multiple areas.

- [Requirements on page 4055](#)
- [Overview on page 4056](#)
- [Configuration on page 4056](#)
- [Verification on page 4058](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.

### Overview


To activate OSPF on a network, you must enable the OSPF protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF, you must configure one or more interfaces on the device within an OSPF area. Once the interfaces are configured, OSPF LSAs are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

Each OSPF area consists of routing devices configured with the same area number. The backbone area is always assigned area ID 0.0.0.0. (All area identifiers (IDs) must be unique within an AS.) All other networks or areas in the AS must be directly connected to the backbone area by a router that has interfaces in more than one area. In

[Figure 115 on page 4056](#), Devices A and C are in the backbone area (area 0), and Devices D and E are in area 2. Device B has a special role. This is the area border router that connects area 0 and area 2. The area border router maintains a separate topological database for each area to which it is connected.

To reduce traffic and topology maintenance for the devices in an OSPF AS, you can group them into multiple areas as shown in [Figure 115 on page 4056](#). In this example, you create the backbone area, create an additional area (area 2) and assign it unique area ID 0.0.0.2, and you configure Device B as the area border router, where interface **ge-0/0/0** participates in OSPF area 0 and interface **ge-0/0/2** participates in OSPF area 2.

**Figure 115: Typical Multiarea OSPF Network Topology**



### Configuration

**CLI Quick Configuration** To quickly configure a multiarea OSPF network, copy the following commands and paste them into the CLI. You repeat this configuration for all interfaces that are part of the OSPF area.

**Device A** [edit]  
 set protocols ospf area 0.0.0.0 interface ge-0/0/0  
 set protocols ospf area 0.0.0.0 interface ge-0/0/1

**Device C** [edit]  
 set protocols ospf area 0.0.0.0 interface ge-0/0/0

**Device B** [edit]

```
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

**Device D**      [edit]  
 set protocols ospf area 0.0.0.2 interface ge-0/0/0  
 set protocols ospf area 0.0.0.2 interface ge-0/0/2

**Device E**      [edit]  
 set protocols ospf area 0.0.0.2 interface ge-0/0/2

**Step-by-Step Procedure**      To configure a multiarea OSPF network:

1.      Configure the backbone area.



**NOTE:** For an OSPFv3 network, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/0
user@A# set protocols ospf area 0.0.0.0 interface ge-0/0/1
```

```
[edit]
user@C# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

```
[edit]
user@B# set protocols ospf area 0.0.0.0 interface ge-0/0/0
```

2.      Configure an additional area for your OSPF network.

```
[edit]
user@B# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

```
[edit]
user@D# set protocols ospf area 0.0.0.2 interface ge-0/0/0
user@D# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

```
[edit]
user@E# set protocols ospf area 0.0.0.2 interface ge-0/0/2
```

3.      If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

### Results

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@A# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
  interface ge-0/0/1.0;
}
```

```
user@C# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
}

user@B# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
}
area 0.0.0.2 {
  interface ge-0/0/2.0;
}

user@D# show protocols ospf
area 0.0.0.2 {
  interface ge-0/0/0.0;
  interface ge-0/0/2.0;
}

user@E# show protocols ospf
area 0.0.0.2 {
  interface ge-0/0/2.0;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4058](#)

### **Verifying the Interfaces in the Area**

**Purpose** Verify that the interface for OSPF or OSPFv3 has been configured for the appropriate area. Confirm that the Area field displays the value that you configured.

**Action** From operational mode, enter the **show ospf interface** command for OSPFv2, and enter the **show ospf3 interface** command for OSPFv3.

**Related Documentation**

- [OSPF Areas and Router Functionality Overview on page 4041](#)
- [OSPF Configuration Overview](#)

---

## **Advanced OSPF Area Configuration**

- [Examples: Configuring OSPF Stub and Not-So-Stubby Areas on page 4059](#)
- [Example: Configuring OSPF Multiarea Adjacency on page 4069](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4073](#)


## Examples: Configuring OSPF Stub and Not-So-Stubby Areas

- [Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas on page 4059](#)
- [Example: Configuring OSPF Stub and Totally Stubby Areas on page 4060](#)
- [Example: Configuring OSPF Not-So-Stubby Areas on page 4064](#)

### Understanding OSPF Stub Areas, Totally Stubby Areas, and Not-So-Stubby Areas

Figure 116 on page 4059 shows an autonomous system (AS) across which many external routes are advertised. If external routes make up a significant portion of a topology database, you can suppress the advertisements in areas that do not have links outside the network. By doing so, you can reduce the amount of memory the nodes use to maintain the topology database and free it for other uses.

Figure 116: OSPF AS Network with Stub Areas and NSSAs



To control the advertisement of external routes into an area, OSPF uses stub areas. By designating an area border router (ABR) interface to the area as a stub interface, you suppress external route advertisements through the ABR. Instead, the ABR advertises a default route (through itself) in place of the external routes and generates network summary (Type 3) link-state advertisements (LSAs). Packets destined for external routes are automatically sent to the ABR, which acts as a gateway for outbound traffic and routes the traffic appropriately.



**NOTE:** You must explicitly configure the ABR to generate a default route when attached to a stub or not-so-stubby-area (NSSA). To inject a default route with a specified metric value into the area, you must configure the `default-metric` option and specify a metric value.

For example, area 0.0.0.3 in [Figure 116 on page 4059](#) is not directly connected to the outside network. All outbound traffic is routed through the ABR to the backbone and then to the

destination addresses. By designating area 0.0.0.3 as a stub area, you reduce the size of the topology database for that area by limiting the route entries to only those routes internal to the area.

A stub area that only allows routes internal to the area and restricts Type 3 LSAs from entering the stub area is often called a *totally stubby area*. You can convert area 0.0.0.3 to a totally stubby area by configuring the ABR to only advertise and allow the default route to enter into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area.



**NOTE:** If you incorrectly configure a totally stubby area, you might encounter network connectivity issues. You should have advanced knowledge of OSPF and understand your network environment before configuring totally stubby areas.

Similar to area 0.0.0.3 in [Figure 116 on page 4059](#), area 0.0.0.4 has no external connections. However, area 0.0.0.4 has static customer routes that are not internal OSPF routes. You can limit the external route advertisements to the area and advertise the static customer routes by designating the area an NSSA. In an NSSA, the AS boundary router generates NSSA external (Type 7) LSAs and floods them into the NSSA, where they are contained. Type 7 LSAs allow an NSSA to support the presence of AS boundary routers and their corresponding external routing information. The ABR converts Type 7 LSAs into AS external (Type 5) LSAs and leaks them to the other areas, but external routes from other areas are not advertised within the NSSA.

---

### Example: Configuring OSPF Stub and Totally Stubby Areas

---

This example shows how to configure an OSPF stub area and a totally stubby area to control the advertisement of external routes into an area.

- [Requirements on page 4060](#)
- [Overview on page 4061](#)
- [Configuration on page 4062](#)
- [Verification on page 4063](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

### Overview

The backbone area, which is 0 in [Figure 117 on page 4062](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an autonomous system (AS). All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by area border routers (ABRs) that have interfaces in more than one area.

Stub areas are areas through which or into which OSPF does not flood AS external link-state advertisements (Type 5 LSAs). You might create stub areas when much of the topology database consists of AS external advertisements and you want to minimize the size of the topology databases on the internal routers in the stub area.

The following restrictions apply to stub areas:

- You cannot create a virtual link through a stub area.
- A stub area cannot contain an AS boundary router.
- You cannot configure the backbone as a stub area.
- You cannot configure an area as both a stub area and a not-so-stubby area (NSSA).

In this example, you configure each routing device in area 7 (area ID 0.0.0.7) as a stub router and some additional settings on the ABR:

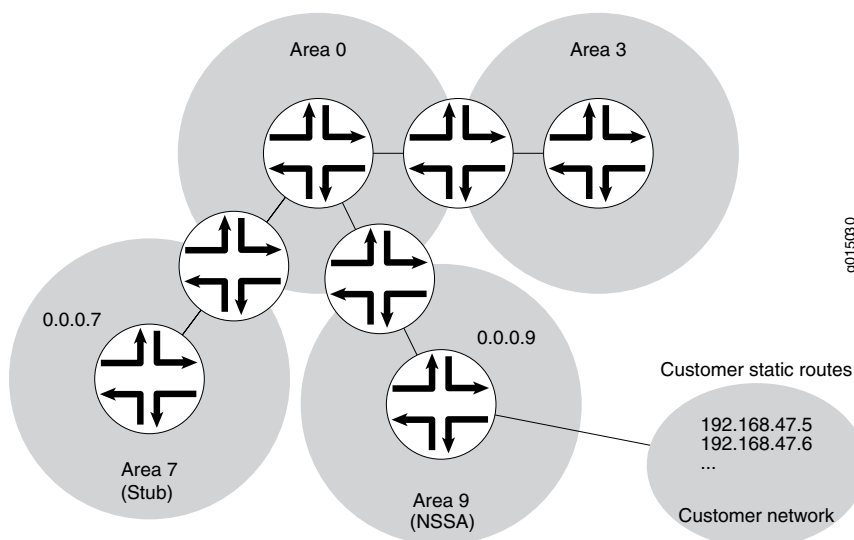
- **stub**—Specifies that this area become a stub area and not be flooded with Type 5 LSAs. You must include the **stub** statement on all routing devices that are in area 7 because this area has no external connections.
- **default-metric**—Configures the ABR to generate a default route with a specified metric into the stub area. This default route enables packet forwarding from the stub area to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to a stub. You must explicitly configure this option to generate a default route.
- **no-summaries**—(Optional) Prevents the ABR from advertising summary routes into the stub area by converting the stub area into a totally stubby area. If configured in combination with the **default-metric** statement, a totally stubby area only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into a totally stubby area. Only the ABR requires this additional configuration because it is the only routing device within the totally stubby area that creates Type 3 LSAs used to receive and send traffic from outside of the area.

**NOTE:**

In Junos OS Release 8.5 and later, the following applies:

- A router-identifier interface that is not configured to run OSPF is no longer advertised as a stub network in OSPF LSAs.
- OSPF advertises a local route with a prefix length of 32 as a stub link if the loopback interface is configured with a prefix length other than 32. OSPF also advertises the direct route with the configured mask length, as in earlier releases.

**Figure 117: OSPF Network Topology with Stub Areas and NSSAs**



### Configuration

#### CLI Quick Configuration

- To quickly configure an OSPF stub area, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the stub area.

[edit]

```
set protocols ospf area 0.0.0.7 stub
```

- To quickly configure the ABR to inject a default route into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

[edit]

```
set protocols ospf area 0.0.0.7 stub default-metric 10
```

- (Optional) To quickly configure the ABR to restrict all summary advertisements and allow only internal routes and default route advertisements into the area, copy the following command and paste it into the CLI. You apply this configuration only on the ABR.

[edit]

```
set protocols ospf area 0.0.0.7 stub no-summaries
```



**Step-by-Step  
Procedure**

To configure OSPF stub areas:

1. On all routing devices in the area, configure an OSPF stub area.



**NOTE:** To specify an OSPFv3 stub area, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub
```

2. On the ABR, inject a default route into the area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub default-metric 10
```

3. (Optional) On the ABR, restrict summary LSAs from entering the area. This step converts the stub area into a totally stubby area.

```
[edit]
user@host# set protocols ospf area 0.0.0.7 stub no-summaries
```

4. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

**Results**

Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices:

```
user@host# show protocols ospf
area 0.0.0.7 {
  stub;
}
```

Configuration on the ABR (the output also includes the optional setting):

```
user@host# show protocols ospf
area 0.0.0.7 {
  stub default-metric 10 no-summaries;
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4064](#)
- [Verifying the Type of OSPF Area on page 4064](#)

### ***Verifying the Interfaces in the Area***

**Purpose** Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub as the type of OSPF area.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

### ***Verifying the Type of OSPF Area***

**Purpose** Verify that the OSPF area is a stub area. Confirm that the output displays Normal Stub as the Stub type.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

---

### **Example: Configuring OSPF Not-So-Stubby Areas**

This example shows how to configure an OSPF not-so-stubby area (NSSA) to control the advertisement of external routes into an area.

- [Requirements on page 4064](#)
- [Overview on page 4064](#)
- [Configuration on page 4066](#)
- [Verification on page 4068](#)

#### ***Requirements***

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### ***Overview***

The backbone area, which is 0 in [Figure 118 on page 4066](#), has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation. Area IDs need only be unique within an AS. All other networks or areas (such as 3, 7, and 9) in the AS must be directly connected to the backbone area by ABRs that have interfaces in more than one area.

An OSPF stub area has no external routes, so you cannot redistribute routes from another protocol into a stub area. OSPF NSSAs allow external routes to be flooded within the area.

In addition, you might have a situation when exporting Type 7 LSAs into the NSSA is unnecessary. When an AS boundary router is also an ABR with an NSSA attached, Type 7 LSAs are exported into the NSSA by default. If the ABR is attached to multiple NSSAs, a separate Type 7 LSA is exported into each NSSA by default. During route redistribution, this routing device generates both Type 5 LSAs and Type 7 LSAs. You can disable exporting Type 7 LSAs into the NSSA.



**NOTE:** The following restriction applies to NSSAs: You cannot configure an area as both a stub area and an NSSA.

You configure each routing device in area 9 (area ID 0.0.0.9) with the following setting:


- **nssa**—Specifies an OSPF NSSA. You must include the **nssa** statement on all routing devices in area 9 because this area only has external connections to static routes.

You also configure the ABR in area 9 with the following additional settings:

- **no-summaries**—Prevents the ABR from advertising summary routes into the NSSA. If configured in combination with the **default-metric** statement, the NSSA only allows routes internal to the area and advertises the default route into the area. External routes and destinations to other areas are no longer summarized or allowed into the NSSA. Only the ABR requires this additional configuration because it is the only routing device within the NSSA that creates Type 3 LSAs used to receive and send traffic from outside the area.
- **default-lsa**—Configures the ABR to generate a default route into the NSSA. In this example, you configure the following:
  - **default-metric**—Specifies that the ABR generate a default route with a specified metric into the NSSA. This default route enables packet forwarding from the NSSA to external destinations. You configure this option only on the ABR. The ABR does not automatically generate a default route when attached to an NSSA. You must explicitly configure this option for the ABR to generate a default route.
  - **metric-type**—(Optional) Specifies the external metric type for the default LSA, which can be either Type 1 or Type 2. When OSPF exports route information from external ASs, it includes a cost, or external metric, in the route. The difference between the two metrics is how OSPF calculates the cost of the route. Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. Type 2 external metrics use only the external cost assigned by the AS boundary router. By default, OSPF uses the Type 2 external metric.
  - **type-7**—(Optional) Floods Type 7 default LSAs into the NSSA if the **no-summaries** statement is configured. By default, when the **no-summaries** statement is configured, a Type 3 LSA is injected into NSSAs for Junos OS release 5.0 and later. To support backward compatibility with earlier Junos OS releases, include the **type-7** statement.

The second example also shows the optional configuration required to disable exporting Type 7 LSAs into the NSSA by including the **no-nssa-abr** statement on the routing device that performs the functions of both an ABR and an AS boundary router.

Figure 118: OSPF Network Topology with Stub Areas and NSSAs



#### Configuration

- [Configuring Routing Devices to Participate in a Not-So-Stubby-Area on page 4066](#)
- [Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas on page 4068](#)

#### Configuring Routing Devices to Participate in a Not-So-Stubby-Area

**CLI Quick Configuration** To quickly configure an OSPF NSSA, copy the following command and paste it into the CLI. You must configure all routing devices that are part of the NSSA.

```
[edit]
set protocols ospf area 0.0.0.9 nssa
```

To quickly configure an ABR that participates in an OSPF NSSA, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.9 nssa default-lsa default-metric 10
set protocols ospf area 0.0.0.9 nssa default-lsa metric-type 1
set protocols ospf area 0.0.0.9 nssa default-lsa type-7
set protocols ospf area 0.0.0.9 nssa no-summaries
```

**Step-by-Step Procedure** To configure OSPF NSSAs:

1. On all routing devices in the area, configure an OSPF NSSA.



**NOTE:** To specify an OSPFv3 NSSA area, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.9 nssa
```

2. On the ABR, enter OSPF configuration mode and specify the NSSA area 0.0.0.9 that you already created.

```
[edit ]
user@host# edit protocols ospf area 0.0.0.9 nssa
```

3. On the ABR, inject a default route into the area.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa default-metric 10
```

4. (Optional) On the ABR, specify the external metric type for the default route.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa metric-type 1
```

5. (Optional) On the ABR, specify the flooding of Type 7 LSAs.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set default-lsa type-7
```

6. On the ABR, restrict summary LSAs from entering the area.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# set no-summaries
```

7. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.9 nssa]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on all routing devices in the area:

```
user@host# show protocols ospf
area 0.0.0.9 {
  nssa;
}
```

Configuration on the ABR. The output also includes the optional **metric-type** and **type-7** statements.

```
user@host# show protocols ospf
area 0.0.0.9 {
  nssa {
    default-lsa {
      default-metric 10;
      metric-type 1;
      type-7;
    }
    no-summaries;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

#### *Disabling the Export of Type 7 Link State Advertisements into Not-So-Stubby Areas*

**CLI Quick Configuration** To quickly disable exporting Type 7 LSAs into the NSSA, copy the following command and paste it into the CLI. You configure this setting on an AS boundary router that is also an ABR with an NSSA area attached.

```
[edit]
set protocols ospf no-nssa-abr
```

**Step-by-Step Procedure** You can configure this setting if you have an AS boundary router that is also an ABR with an NSSA area attached.

1. Disable exporting Type 7 LSAs into the NSSA.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf no-nssa-abr
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
no-nssa-abr;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

#### *Verification*

Confirm that the configuration is working properly.

- [Verifying the Interfaces in the Area on page 4068](#)
- [Verifying the Type of OSPF Area on page 4069](#)
- [Verifying the Type of LSAs on page 4069](#)

#### *Verifying the Interfaces in the Area*

**Purpose** Verify that the interface for OSPF has been configured for the appropriate area. Confirm that the output includes Stub NSSA as the type of OSPF area.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Verifying the Type of OSPF Area**

**Purpose** Verify that the OSPF area is a stub area. Confirm that the output displays Not so Stubby Stub as the Stub type.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

**Verifying the Type of LSAs**

**Purpose** Verify the type of LSAs that are in the area. If you disabled exporting Type 7 LSAs into an NSSA, confirm that the Type field does not include NSSA as a type of LSA.

**Action** From operational mode, enter the **show ospf database** command for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

- Related Documentation**
- *Example: Configuring OSPFv3 Stub and Totally Stubby Areas*
  - [OSPF Areas and Router Functionality Overview on page 4041](#)
  - *OSPF Configuration Overview*

**Example: Configuring OSPF Multiarea Adjacency**

- [Multiarea Adjacency for OSPF on page 4069](#)
- [Example: Configuring Multiarea Adjacency for OSPF on page 4070](#)

**Multiarea Adjacency for OSPF**

An area is a set of networks and hosts within an autonomous system (AS) that have been administratively grouped together. By default, a single interface can belong to only one OSPF area. However, in some situations, you might want to configure an interface to belong to more than one area. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. For example, you can configure an interface to belong to multiple areas with a high-speed backbone link between two area border routers (ABRs) so you can create multiarea adjacencies that belong to different areas.

In Junos OS Release 9.2 and later, you can configure a logical interface to belong to more than one OSPFv2 area. Support for OSPFv3 was introduced in Junos OS Release 9.4. As defined in RFC 5185, *OSPF Multi-Area Adjacency*, the ABRs establish multiple adjacencies belonging to different areas over the same logical interface. Each multiarea adjacency is announced as a point-to-point unnumbered link in the configured area by the routers connected to the link. For each area, one of the logical interfaces is treated as primary, and the remaining interfaces that are configured for the area are designated as secondary.

Any logical interface not configured as a secondary interface for an area is treated as the primary interface for that area. A logical interface can be configured as primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.

### Example: Configuring Multiarea Adjacency for OSPF

---

This example shows how to configure multiarea adjacency for OSPF.

- [Requirements on page 4070](#)
- [Overview on page 4070](#)
- [Configuration on page 4071](#)
- [Verification on page 4073](#)

#### Requirements

Before you begin, plan your multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

#### Overview

By default, a single interface can belong to only one OSPF area. You can configure a single interface to belong in multiple OSPF areas. Doing so allows the corresponding link to be considered an intra-area link in multiple areas and to be preferred over other higher-cost intra-area paths. When configuring a secondary interface, consider the following:

- For OSPFv2, you cannot configure point-to-multipoint and nonbroadcast multiaccess (NBMA) network interfaces as a secondary interface because secondary interfaces are treated as a point-to-point unnumbered link.
- Secondary interfaces are supported for LAN interfaces (the primary interface can be a LAN interface, but any secondary interfaces are treated as point-to-point unnumbered links over the LAN). In this scenario, you must ensure that there are only two routing devices on the LAN or that there are only two routing devices on the LAN that have secondary interfaces configured for a specific OSPF area.
- Since the purpose of a secondary interface is to advertise a topological path through an OSPF area, you cannot configure a secondary interface or a primary interface with one or more secondary interfaces to be passive. Passive interfaces advertise their address, but do not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).
- Any logical interface not configured as a secondary interface for an area is treated as a primary interface for that area. A logical interface can be configured as the primary interface only for one area. For any other area for which you configure the interface, you must configure it as a secondary interface.
- You cannot configure the **secondary** statement with the **interface all** statement.
- You cannot configure a secondary interface by its IP address.

In this example, you configure an interface to be in two areas, creating a multiarea adjacency with a link between two ABRs: ABR R1 and ABR R2. On each ABR, area 0.0.0.1 contains the primary interface and is the primary link between the ABRs, and area 0.0.0.2 contains the secondary logical interface, which you configure by including the **secondary** statement. You configure interface **so-0/0/0** on ABR R1 and interface **so-1/0/0** on ABR R2.



**Configuration**

**CLI Quick Configuration** To quickly configure a secondary logical interface for an OSPF area, copy the following commands and paste them into the CLI.

Configuration on ABR R1:

```
[edit]
set interfaces so-0/0/0 unit 0 family inet address 192.168.8.45/30
set routing-options router-id 10.255.0.1
set protocols ospf area 0.0.0.1 interface so-0/0/0
set protocols ospf area 0.0.0.2 interface so-0/0/0 secondary
```

Configuration on ABR R2:

```
[edit]
set interfaces so-1/0/0 unit 0 family inet address 192.168.8.37/30
set routing-options router-id 10.255.0.2
set protocols ospf area 0.0.0.1 interface so-1/0/0
set protocols ospf area 0.0.0.2 interface so-1/0/0 secondary
```

**Step-by-Step Procedure** To configure a secondary logical interface:

1. Configure the device interfaces.



**NOTE:** For OSPFv3, on each interface specify the inet6 address family and include the IPv6 address.

```
[edit]
user@R1# set interfaces so-0/0/0 unit 0 family inet address 192.168.8.45/30
```

```
[edit]
user@R2# set interfaces so-1/0/0 unit 0 family inet address 192.168.8.37/30
```

2. Configure the router identifier.

```
[edit]
user@R1# set routing-options router-id 10.255.0.1
```

```
[edit]
user@R2# set routing-options router-id 10.255.0.2
```

3. On each ABR, configure the primary interface for the OSPF area.



**NOTE:** For OSPFv3, include the ospf3 statement at the [edit protocols] hierarchy level.

```
[edit]
user@R1# set protocols ospf 0.0.0.1 interface so-0/0/0
```

```
[edit ]
user@R2# set protocols ospf 0.0.0.2 interface so-1/0/0
```

4. On each ABR, configure the secondary interface for the OSPF area.

```
[edit ]
user@R1# set protocols ospf area 0.0.0.1 so-0/0/0 secondary
```

```
[edit ]
user@R2# set protocols ospf area 0.0.0.2 so-1/0/0 secondary
```

5. If you are done configuring the devices, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

### Results

Confirm your configuration by entering the **show interfaces**, **show routing-options**, and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R1:

```
user@R1# show interfaces
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.8.45/30;
    }
  }
}

user@R1# show routing-options
router-id 10.255.0.1;

user@R1# show protocols ospf
area 0.0.0.1 {
  interface so-0/0/0.0;
}
area 0.0.0.2 {
  interface so-0/0/0.0 {
    secondary;
  }
}
```

Configuration on ABR R2:

```
user@R2# show interfaces
so-0/0/0 {
  unit 0 {
    family inet {
      address 192.168.8.37/30;
    }
  }
}

user@R2# show routing-options
router-id 10.255.0.2;

user@R2# show protocols ospf
area 0.0.0.1 {
```

```

    interface so-1/0/0.0;
  }
  area 0.0.0.2 {
    interface so-1/0/0.0 {
      secondary;
    }
  }
}

```

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the Secondary Interface on page 4073](#)
- [Verifying the Interfaces in the Area on page 4073](#)
- [Verifying Neighbor Adjacencies on page 4073](#)

### **Verifying the Secondary Interface**

**Purpose** Verify that the secondary interface appears for the configured area. The Secondary field displays if the interface is configured as a secondary interface. The output might also show the same interface listed in multiple areas.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

### **Verifying the Interfaces in the Area**

**Purpose** Verify the interfaces configured for the specified area.

**Action** From operational mode, enter the **show ospf interface area *area-id*** command for OSPFv2, and enter the **show ospf3 interface area *area-id*** command for OSPFv3..

### **Verifying Neighbor Adjacencies**

**Purpose** Verify the primary and secondary neighbor adjacencies. The Secondary field displays if the neighbor is on a secondary interface.

**Action** From operational mode, enter the **show ospf neighbor detail** command for OSPFv2, and enter the **show ospf3 neighbor detail** command for OSPFv3.

**Related Documentation**

- [OSPF Areas and Router Functionality Overview on page 4041](#)
- [Understanding OSPF Areas and Backbone Areas on page 4052](#)
- [OSPF Configuration Overview](#)

## **Example: Disabling OSPFv2 Compatibility with RFC 1583**

- [OSPFv2 Compatibility with RFC 1583 Overview on page 4074](#)
- [Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4074](#)

## OSPFv2 Compatibility with RFC 1583 Overview

---

In the first implementation of OSPF (RFC1583, *OSPF Version 2*), the summary route assumes the cost of the granular route with the lowest cost. OSPF RFC 2328, *OSPF Version 2* changes the behavior so that the summary route assumes the cost of the granular route with the highest cost. OSPF readvertises the summary route whenever the cost of the summary changes. When using the default RFC 1583 behavior, this happens when the granular route with the lowest metric is changed or lost. When RFC 2328 is used, this happens when the granular route with the highest cost is changed or lost.

By default, the Junos OS implementation of OSPF is compatible with RFC 1583. This means that Junos OS maintains a single best route to an autonomous system (AS) boundary router in the OSPF routing table, rather than multiple intra-AS paths, if they are available. You can disable compatibility with RFC 1583. It is preferable to do so when the same external destination is advertised by AS boundary routers that belong to different OSPF areas. When you disable compatibility with RFC 1583, the OSPF routing table maintains the multiple intra-AS paths that are available, which the router uses to calculate AS external routes as defined in RFC 2328. Being able to use multiple available paths to calculate an AS external route can prevent routing loops.

## Example: Disabling OSPFv2 Compatibility with RFC 1583

---

This example shows how to disable OSPFv2 compatibility with RFC 1583 on the routing device.

- [Requirements on page 4074](#)
- [Overview on page 4074](#)
- [Configuration on page 4074](#)
- [Verification on page 4075](#)

### Requirements

No special configuration beyond device initialization is required before disabling OSPFv2 compatibility with RFC 1583.

### Overview

The introduction of RFC 2328 changed the method used to calculate the routes in an OSPF network. By default, the Junos OS implementation of OSPFv2 is compatible with RFC 1583, so OSPF uses the minimum cost to determine the route to any of the networks within the specified range. When you disable RFC 1583 compatibility, OSPF uses the maximum cost to determine the route to any of the networks within the specified range. To minimize the potential for routing loops, configure the same RFC compatibility on all OSPF devices in an OSPF domain.

### Configuration

#### CLI Quick Configuration

To quickly disable OSPFv2 compatibility with RFC 1583, copy the following command and paste it into the CLI. You configure this setting on all devices that are part of the OSPF domain.

[edit]

```
set protocols ospf no-rfc-1583
```

**Step-by-Step Procedure** To disable OSPFv2 compatibility with RFC 1583:

1. Disable RFC 1583.  

```
[edit]  
user@host# set protocols ospf no-rfc-1583
```
2. If you are done configuring the device, commit the configuration.  

```
[edit]  
user@host# commit
```



**NOTE:** Repeat this configuration on each routing device that participates in an OSPF routing domain.

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf  
no-rfc-1583;
```

### Verification

Confirm that the configuration is working properly.

#### Verifying the OSPF Routes

**Purpose** Verify that the OSPF routing table maintains the intra-AS paths with the largest metric, which the router uses to calculate AS external routes.

**Action** From operational mode, enter the **show ospf route detail** command.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)

## OSPF Interface Configuration

- [Examples: Configuring OSPF Interfaces on page 4075](#)
- [Example: Configuring Multiple Address Families for OSPFv3 on page 4090](#)

### Examples: Configuring OSPF Interfaces

- [About OSPF Interfaces on page 4076](#)
- [Example: Configuring an Interface on a Broadcast or Point-to-Point Network on page 4077](#)

- [Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network on page 4079](#)
- [Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network on page 4082](#)
- [Example: Configuring OSPF Demand Circuits on page 4084](#)
- [Example: Configuring a Passive OSPF Interface on page 4086](#)
- [Example: Configuring OSPFv2 Peer interfaces on page 4088](#)

---

### About OSPF Interfaces

To activate OSPF on a network, you must enable the OSPF protocol on one or more interfaces on each device within the network on which traffic is to travel. How you configure the interface depends on whether the interface is connected to a broadcast or point-to-point network, a point-to-multipoint network, a nonbroadcast multiaccess (NBMA) network, or across a demand circuit.

- A broadcast interface behaves as if the routing device is connected to a LAN.
- A point-to-point interface provides a connection between a single source and a single destination (there is only one OSPF adjacency).
- A point-to-multipoint interface provides a connection between a single source and multiple destinations.
- An NBMA interface behaves in a similar fashion to a point-to-multipoint interface, but you might configure an NBMA interface to interoperate with other equipment.
- A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You can also configure an OSPF interface to be passive, to operate in passive traffic engineering mode, or to be a peer interface.

- A passive interface advertises its address, but does not run the OSPF protocol (adjacencies are not formed and hello packets are not generated).
- An interface operating in OSPF passive traffic engineering mode floods link address information within the autonomous system (AS) and makes it available for traffic engineering calculations.
- A peer interface can be configured for OSPFv2 routing devices. A peer interface is required for Generalized MPLS (GMPLS) to transport traffic engineering information through a link separate from the control channel. You establish this separate link by configuring a peer interface. The peer interface name must match the Link Management Protocol (LMP) peer name. A peer interface is optional for a hierarchy of RSVP label-switched paths (LSPs). After you configure the forwarding adjacency, you can configure OSPFv2 to advertise the traffic engineering properties of a forwarding adjacency to a specific peer.

Point-to-point interfaces differ from multipoint in that only one OSPF adjacency is possible. (A LAN, for instance, can have multiple addresses and can run OSPF on each subnet simultaneously.) As such, when you configure a numbered point-to-point interface

to OSPF by name, multiple OSPF interfaces are created. One, which is unnumbered, is the interface on which the protocol is run. An additional OSPF interface is created for each address configured on the interface, if any, which is automatically marked as passive.

For OSPFv3, one OSPF-specific interface must be created per interface name configured under OSPFv3. OSPFv3 does not allow interfaces to be configured by IP address.

Enabling OSPF on an interface (by including the **interface** statement), disabling it (by including the **disable** statement), and not actually having OSPF run on an interface (by including the **passive** statement) are mutually exclusive states.



**NOTE:** When you configure OSPFv2 on an interface, you must also include the **family inet** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. When you configure OSPFv3 on an interface, you must also include the **family inet6** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support address families other than unicast IPv6.

### Example: Configuring an Interface on a Broadcast or Point-to-Point Network

This example shows how to configure an OSPF interface on a broadcast or point-to-point network.

- [Requirements on page 4077](#)
- [Overview on page 4077](#)
- [Configuration on page 4078](#)
- [Verification on page 4079](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### Overview

If the interface on which you are configuring OSPF supports broadcast mode (such as a LAN), or if the interface supports point-to-point mode (such as a PPP interface or a point-to-point logical interface on Frame Relay), you specify the interface by including the IP address or the interface name for OSPFv2, or only the interface name for OSPFv3. In Junos OS Release 9.3 and later, an OSPF point-to-point interface can be an Ethernet

interface without a subnet. If you configure an interface on a broadcast network, designated router and backup designated router election is performed.



**NOTE:**

- Using both the interface name and the IP address of the same interface produces an invalid configuration.
- Including the IP address of loopback0 interface unit may implicitly enable OSPF on unnumbered interfaces with “unnumbered-address lo0.0” configured.

In this example, you configure interface **ge-0/2/0** as an OSPFv2 interface in OSPF area 0.0.0.1.

**Configuration**

**CLI Quick Configuration**

To quickly configure an OSPF interface on a broadcast or point-to-point network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
set protocols ospf area 0.0.0.1 interface ge-0/2/0
```

**Step-by-Step Procedure**

To configure an OSPF interface on a broadcast or point-to-point network:

1. Configure the interface.



**NOTE:** For an OSPFv3 interface, specify an IPv6 address.

```
[edit]
user@host# set interfaces ge-0/2/0 unit 0 family inet address 10.0.0.1
```

2. Create an OSPF area.



**NOTE:** For an OSPFv3 interface, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface ge-0/2/0
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```



### Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
ge-0/2/0 {
  unit 0 {
    family inet {
      address 10.0.0.1/32;
    }
  }
}

user@host# show protocols ospf
area 0.0.0.1 {
  interface ge-0/2/0.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

### Verification

Confirm that the configuration is working properly.

#### Verifying the OSPF Interface

- |                |                                                                                                                                                             |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify the interface configuration. Depending on your deployment, the Type field might display LAN or P2P.                                                  |
| <b>Action</b>  | From operational mode, enter the <b>show ospf interface detail</b> command for OSPFv2, and enter the <b>show ospf3 interface detail</b> command for OSPFv3. |

### Example: Configuring an OSPFv2 Interface on a Nonbroadcast Multiaccess Network

This example shows how to configure an OSPFv2 interface on a nonbroadcast multiaccess (NBMA) network.

- [Requirements on page 4079](#)
- [Overview on page 4080](#)
- [Configuration on page 4081](#)
- [Verification on page 4082](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).

- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

### Overview

When you configure OSPFv2 on an NBMA network, you can use nonbroadcast mode rather than point-to-multipoint mode. Using this mode offers no advantages over point-to-multipoint mode, but it has more disadvantages than point-to-multipoint mode. Nevertheless, you might occasionally find it necessary to configure nonbroadcast mode to interoperate with other equipment. Because there is no autodiscovery mechanism, you must configure each neighbor.

Nonbroadcast mode treats the NBMA network as a partially connected LAN, electing designated and backup designated routers. All routing devices must have a direct connection to both the designated and backup designated routers, or unpredictable results occur.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration. For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as the interface name.

In this example, you configure the Asynchronous Transfer Mode (ATM) interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify the following settings:

- **interface-type nbma**—Sets the interface to run in NBMA mode. You must explicitly configure the interface to run in NBMA mode.
- **neighbor address <eligible>**—Specifies the IP address of the neighboring device. OSPF routing devices normally discover their neighbors dynamically by listening to the broadcast or multicast hello packets on the network. Because an NBMA network does not support broadcast (or multicast), the device cannot discover its neighbors dynamically, so you must configure all the neighbors statically. To configure multiple neighbors, include multiple **neighbor** statements. If you want the neighbor to be a designated router, include the **eligible** keyword.
- **poll-interval**—Specifies the length of time, in seconds, before the routing device sends hello packets out of the interface before it establishes adjacency with a neighbor. Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The range is from 1 through 255 seconds. By default, the device sends hello packets out the interface every 120 seconds before it establishes adjacency with a neighbor.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the **poll-interval** statement to the time specified in the **hello-interval** statement.

### Configuration

**CLI Quick Configuration** To quickly configure an OSPFv2 interface on an NBMA network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 interface-type nbma
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 neighbor 192.0.2.2 eligible
set protocols ospf area 0.0.0.1 interface at-0/1/0.0 poll-interval 130
```

**Step-by-Step Procedure** To configure an OSPFv2 interface on an NBMA network:

1. Configure the interface.  

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.1
```
2. Create an OSPF area.  

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```
3. Assign the interface to the area.  
 In this example, include the **eligible** keyword to allow the neighbor to be a designated router.  

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface at-0/1/0 interface-type nbma neighbor 192.0.2.2 eligible
```
4. Configure the poll interval.  

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface at-0/1/0 poll-interval 130
```
5. If you are done configuring the device, commit the configuration.  

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

### Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/32;
    }
  }
}

user@host# show protocols ospf
area 0.0.0.1 {
  interface at-0/1/0.0 {
```

```
    interface-type nbma;  
    neighbor 192.0.2.2 eligible;  
    poll-interval 130;  
  }  
}
```

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the OSPF Interface**

**Purpose** Verify the interface configuration. Confirm that the Type field displays NBMA.

**Action** From operational mode, enter the **show ospf interface detail** command.

---

### **Example: Configuring an OSPFv2 Interface on a Point-to-Multipoint Network**

This example shows how to configure an OSPFv2 interface on a point-to-multipoint network.

- [Requirements on page 4082](#)
- [Overview on page 4082](#)
- [Configuration on page 4083](#)
- [Verification on page 4083](#)

### **Requirements**

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#)
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

### **Overview**

When you configure OSPFv2 on a nonbroadcast multiaccess (NBMA) network, such as a multipoint Asynchronous Transfer Mode (ATM) or Frame Relay, OSPFv2 operates by default in point-to-multipoint mode. In this mode, OSPFv2 treats the network as a set of point-to-point links. Because there is no autodiscovery mechanism, you must configure each neighbor.

When you configure the interface, specify either the IP address or the interface name. Using both the IP address and the interface name produces an invalid configuration.

In this example, you configure ATM interface **at-0/1/0** as an OSPFv2 interface in OSPF area 0.0.0.1, and you specify 192.0.2.1 as the neighbor's IP address.

**Configuration**

**CLI Quick Configuration** To quickly configure an OSPFv2 interface on a point-to-multipoint network, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
set protocols ospf area 0.0.0.1 interface at-0/1/0 neighbor 192.0.2.1
```

**Step-by-Step Procedure** To configure an OSPFv2 interface on a point-to-multipoint network:

1. Configure the interface.

```
[edit]
user@host# set interfaces at-0/1/0 unit 0 family inet address 192.0.2.2
```

2. Create an OSPF area.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

3. Assign the interface to the area and specify the neighbor.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface at-0/1/0 neighbor 192.0.2.1
```

To configure multiple neighbors, include a **neighbor** statement for each neighbor.

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

**Results**

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
at-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.2/32;
    }
  }
}

user@host# show protocols ospf
area 0.0.0.1 {
  interface at-0/1/0.0 {
    neighbor 192.0.2.1;
  }
}
```

**Verification**

Confirm that the configuration is working properly.

### ***Verifying the OSPF Interface***

**Purpose** Verify the interface configuration. Confirm that the Type field displays P2MP.

**Action** From operational mode, enter the **show ospf interface detail** command.

---

### **Example: Configuring OSPF Demand Circuits**

This example shows how to configure an OSPF demand circuit interface.

- [Requirements on page 4084](#)
- [Overview on page 4084](#)
- [Configuration on page 4085](#)
- [Verification on page 4086](#)

#### ***Requirements***

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.



**NOTE:** If you are using OSPF demand circuits over an ISDN link, you must configure an ISDN interface and enable dial-on-demand routing. See the *Junos OS Network Interfaces Library for Routing Devices*.

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### ***Overview***

OSPF sends periodic hello packets to establish and maintain neighbor adjacencies and uses link-state advertisements (LSAs) to make routing calculations and decisions. OSPF support for demand circuits is defined in RFC 1793, *Extending OSPF to Support Demand Circuits*, and suppresses the periodic hello packets and LSAs. A demand circuit is a connection on which you can limit traffic based on user agreements. The demand circuit can limit bandwidth or access time based on agreements between the provider and user.

You configure demand circuits on an OSPF interface. When the interface becomes a demand circuit, all hello packets and LSAs are suppressed as soon as OSPF synchronization is achieved. LSAs have a DoNotAge bit that stops the LSA from aging and prevents periodic updates from being sent. Hello packets and LSAs are sent and received on a demand-circuit interface only when there is a change in the network topology. This reduces the amount of traffic through the OSPF interface.

Consider the following when configuring OSPF demand circuits:

- Periodic hellos are only suppressed on point-to-point and point-to-multipoint interfaces. If you configure demand circuits on an OSPF broadcast network or on an OSPF nonbroadcast multiaccess (NBMA) network, periodic hello packets are still sent.
- Demand circuit support on an OSPF point-to-multipoint interface resembles that for point-to-point interfaces. If you configure a point-to-multipoint interface as a demand circuit, the device negotiates hello suppression separately on each interface that is part of the point-to-multipoint network.

This example assumes that you have a point-to-point connection between two devices using SONET/SDH interfaces. A demand-circuit interface automatically negotiates the demand-circuit connection with its OSPF neighbor. If the neighbor does not support demand circuits, then no demand circuit connection is established.

In this example, you configure OSPF interface **so-0/1/0** in OSPF area 0.0.0.1 as a demand circuit.

### Configuration

#### CLI Quick Configuration

To quickly configure an OSPF demand circuit interface, copy the following command and paste it into the CLI. You must configure both neighboring interfaces for OSPF demand circuits for the connection to be established.

[edit]

```
set protocols ospf area 0.0.0.1 interface so-0/1/0 demand-circuit
```

#### Step-by-Step Procedure

To configure an OSPF demand circuit interface on one neighboring interface:

1. Create an OSPF area.



**NOTE:** For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

[edit]

```
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the neighboring interface as a demand circuit.

[edit protocols ospf area 0.0.0.1]

```
user@host# set interface so-0/1/0 demand-circuit
```

3. If you are done configuring the device, commit the configuration.

[edit protocols ospf area 0.0.0.1]

```
user@host# commit
```



**NOTE:** Repeat this entire configuration on the other neighboring interface.

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
ospf {
  area 0.0.0.1 {
    interface so-0/1/0.0 {
      demand-circuit;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Status of Neighboring Interfaces

- |                |                                                                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify information about the neighboring interface. When the neighbor is configured for demand circuits, a DC flag displays.                              |
| <b>Action</b>  | From operational mode, enter the <b>show ospf neighbor detail</b> command for OSPFv2, and enter the <b>show ospf3 neighbor detail</b> command for OSPFv3. |

---

### Example: Configuring a Passive OSPF Interface

This example shows how to configure a passive OSPF interface. A passive OSPF interface advertises its address but does not run the OSPF protocol.

- [Requirements on page 4086](#)
- [Overview on page 4087](#)
- [Configuration on page 4087](#)
- [Verification on page 4088](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.



### Overview

By default, OSPF must be configured on an interface for direct interface addresses to be advertised as interior routes. To advertise the direct interface addresses without actually running OSPF on that interface (adjacencies are not formed and hello packets are not generated), you configure that interface as a passive interface.

Enabling OSPF on an interface (by including the **interface** statement), disabling it (by including the **disable** statement), and not actually having OSPF run on an interface (by including the **passive** statement) are mutually exclusive states.



**NOTE:** If you do not want to see notifications for state changes in a passive OSPF interface, you can disable the OSPF traps for the interface by including the **no-interface-state-traps** statement. The **no-interface-state-traps** statement is supported only for OSPFv2.

In this example, you configure interface **ge-0/2/0** as a passive OSPF interface in area 0.0.0.1 by including the **passive** statement.

### Configuration

#### CLI Quick Configuration

To quickly configure a passive OSPF interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface ge-0/2/0 passive
```

#### Step-by-Step Procedure

To configure a passive OSPF interface:

1. Create an OSPF area.



**NOTE:** For an OSPFv3 interface, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Configure the passive interface.

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface ge-0/2/0 passive
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  area 0.0.0.1 {
    interface ge-0/2/0.0 {
      passive;
    }
  }
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Status of OSPF Interfaces

**Purpose** Verify the status of the OSPF interface. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.

**Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

---

### Example: Configuring OSPFv2 Peer interfaces

This example shows how to configure an OSPFv2 peer interface.

- [Requirements on page 4088](#)
- [Overview on page 4089](#)
- [Configuration on page 4089](#)
- [Verification on page 4089](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4053](#).

- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).
- Configure Generalized MPLS per your network requirements. See *LMP Configuration Overview* in the *Junos OS MPLS Applications Library for Routing Devices*.

### Overview

You can configure an OSPFv2 peer interface for many reasons, including when you configure Generalized MPLS (GMPLS). This example configures a peer interface for GMPLS. GMPLS requires traffic engineering information to be transported through a link separate from the control channel. You establish this separate link by configuring a peer interface. The OSPFv2 peer interface name must match the Link Management Protocol (LMP) peer name. You configure GMPLS and the LMP settings separately from OSPF.

This example assumes that GMPLS and the LMP peer named **oxc1** are already configured, and you need to configure the OSPFv2 peer interface in area 0.0.0.0.

### Configuration

**CLI Quick Configuration** To quickly configure an OSPFv2 peer interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 peer-interface oxc1
```

**Step-by-Step Procedure** To configure a peer OSPFv2 interface used by the LMP:

1. Create an OSPF area.  

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```
2. Configure the peer interface.  

```
[edit protocols ospf area 0.0.0.0]
user@host# set peer-interface oxc1
```
3. If you are done configuring the device, commit the configuration.  

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
  area 0.0.0.0 {
    peer-interface oxc1;
  }
```

### Verification

Confirm that the configuration is working properly.

### ***Verifying the Configured OSPFv2 Peer***

**Purpose** Verify the status of the OSPFv2 peer. When an OSPFv2 peer is configured for GMPLS, the Peer Name field displays the name of the LMP peer that you created for GMPLS, which is also the configured OSPFv2 peer.

**Action** From operational mode, enter the **show link-management** command.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)

## **Example: Configuring Multiple Address Families for OSPFv3**

- [Understanding Multiple Address Families for OSPFv3 on page 4090](#)
- [Example: Configuring Multiple Address Families for OSPFv3 on page 4091](#)

### **Understanding Multiple Address Families for OSPFv3**

---

By default, OSPFv3 supports only unicast IPv6 routes. In Junos OS Release 9.2 and later, you can configure OSPFv3 to support multiple address families, including IPv4 unicast, IPv4 multicast, and IPv6 multicast. This multiple address family support allows OSPFv3 to support both IPv6 and IPv4 nodes. Junos OS maps each address family to a separate realm as defined in Internet draft draft-ietf-ospf-af-alt-06.txt, *Support for Address Families in OSPFv3*. Each realm maintains a separate set of neighbors and link-state database.

When you configure multiple address families for OSPFv3, there is a new instance ID field that allows multiple OSPFv3 protocol instances per link. This allows a single link to belong to multiple areas.

You configure each realm independently. We recommend that you configure an area and at least one interface for each realm.

These are the default import and export routing tables for each of the four address families:

- IPv6 unicast: **inet6.0**
- IPv6 multicast: **inet6.2**
- IPv4 unicast: **inet.0**
- IPv4 multicast: **inet.2**

With the exception of virtual links, all configurations supported for the default IPv6 unicast family are supported for the address families that have to be configured as realms.

### Example: Configuring Multiple Address Families for OSPFv3

This example shows how to configure multiple address families for OSPFv3.

- [Requirements on page 4091](#)
- [Overview on page 4091](#)
- [Configuration on page 4092](#)
- [Verification on page 4093](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4053](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

#### Overview

By default, OSPFv3 supports unicast IPv6 routes, but you can configure OSPFv3 to support multiple address families. To support an address family other than unicast IPv6, you configure a realm that allows OSPFv3 to advertise IPv4 unicast, IPv4 multicast, or IPv6 multicast routes. Junos OS then maps each address family that you configure to a separate realm with its own set of neighbors and link-state database.




**NOTE:** By default, LDP synchronization is only supported for OSPFv2. If you configure an IPv4 unicast or IPv4 multicast realm, you can also configure LDP synchronization. Since LDP synchronization is only supported for IPv4, this support is only available for OSPFv3 if you configure an IPv4 realm.

When configuring OSPFv3 to support multiple address families, consider the following:

- You configure each realm independently. We recommend that you configure an area and at least one interface for each realm.
- OSPFv3 uses IPv6 link-local addresses as the source of hello packets and next hop calculations. As such, you must enable IPv6 on the link regardless of the additional realm you configure.

[Figure 119 on page 4092](#) shows a connection between Routers R0 and R1. In this example, you configure interface **fe-0/1/0** on Router R0 in area 0 to advertise IPv4 unicast routes, in addition to the default unicast IPv6 routes in area 1, by including the **realm ipv4-unicast** statement. Depending on your network requirements, you can also advertise IPv4 multicast routes by including the **realm-ipv4-multicast** statement, and you can advertise IPv6 multicast routes by including the **realm-ipv6-multicast** statement.

Figure 119: IPv4 Unicast Realm

**Configuration****CLI Quick Configuration**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To quickly configure multiple address families for OSPFv3, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 11.1.2.1/24
set interfaces fe-0/1/0 unit 0 family inet6
set protocols ospf3 area 0.0.0.0 interface fe-0/1/0
set protocols ospf3 realm ipv4-unicast area 0.0.0.0 interface fe-0/1/0
```

**Step-by-Step Procedure**

To configure multiple address families for OSPFv3:

1. Configure the device interface participating in OSPFv3.
 

```
[edit]
user@host# set interfaces fe-0/1/0 unit 0 family inet address 11.1.2.1/24
user@host# set interfaces fe-0/1/0 unit 0 family inet6
```
2. Enter OSPFv3 configuration mode.
 

```
[edit ]
user@host# edit protocols ospf3
```
3. Add the interface you configured to the OSPFv3 area.
 

```
[edit protocols ospf3 ]
user@host# set area 0.0.0.0 interface fe-0/1/0
```
4. Configure an IPv4 unicast realm. This allows OSPFv3 to support both IPv4 unicast and IPv6 unicast routes.
 

```
[edit protocols ospf3 ]
user@host# set realm ipv4-unicast area 0.0.0.0 interface fe-0/1/0
```

- If you are done configuring the device, commit the configuration.

```
[edit protocols ospf3 ]
user@host# commit
```



**NOTE:** Repeat this entire configuration on the neighboring device that is part of the realm.

### Results

Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 11.1.2.1/24;
    }
    family inet6;
  }
}

user@host# show protocols ospf3
realm ipv4-unicast {
  area 0.0.0.0 {
    interface fe-0/1/0.0;
  }
}
area 0.0.0.0 {
  interface fe-0/1/0.0;
}
```

### Verification

Confirm that the configuration is working properly.

- [Verifying the Link-State Database on page 4093](#)
- [Verifying the Status of OSPFv3 Interfaces with Multiple Address Families on page 4093](#)

#### Verifying the Link-State Database

**Purpose** Verify the status of the link-state database for the configured realm, or address family.

**Action** From operational mode, enter the **show ospf3 database realm ipv4-unicast** command.

#### Verifying the Status of OSPFv3 Interfaces with Multiple Address Families

**Purpose** Verify the status of the interface for the specified OSPFv3 realm, or address family.

**Action** From operational mode, enter the **show ospf3 interface realm ipv4-unicast** command.

- Related Documentation**
- [OSPF Overview on page 4036](#)
  - [OSPF Configuration Overview](#)

## OSPF Route Control Configuration

---

- [Examples: Configuring OSPF Route Summarization on page 4094](#)
- [Examples: Configuring OSPF Traffic Control on page 4103](#)
- [Example: Configuring OSPF Overload Mode on page 4113](#)

### Examples: Configuring OSPF Route Summarization

- [Understanding OSPF Route Summarization on page 4094](#)
- [Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements on page 4094](#)
- [Example: Limiting the Number of Prefixes Exported to OSPF on page 4100](#)
- [Configuring OSPF Refresh and Flooding Reduction in Stable Topologies on page 4102](#)

#### Understanding OSPF Route Summarization

---

Area border routers (ABRs) send summary link advertisements to describe the routes to other areas. Depending on the number of destinations, an area can get flooded with a large number of link-state records, which can utilize routing device resources. To minimize the number of advertisements that are flooded into an area, you can configure the ABR to coalesce, or summarize, a range of IP addresses and send reachability information about these addresses in a single link-state advertisement (LSA). You can summarize one or more ranges of IP addresses, where all routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

For an OSPF area, you can summarize and filter intra-area prefixes. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place. For an OSPF not-so-stubby area (NSSA), you can only coalesce or filter NSSA external (Type 7) LSAs before they are translated into AS external (Type 5) LSAs and enter the backbone area. All external routes learned within the area that do not fall into the range of one of the prefixes are advertised individually to other areas.

In addition, you can also limit the number of prefixes (routes) that are exported into OSPF. By setting a user-defined maximum number of prefixes, you prevent the routing device from flooding an excessive number of routes into an area.

#### Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements

---

This example shows how to summarize routes sent into the backbone area.

- [Requirements on page 4095](#)
- [Overview on page 4095](#)
- [Configuration on page 4096](#)
- [Verification on page 4099](#)



## Requirements

Before you begin:

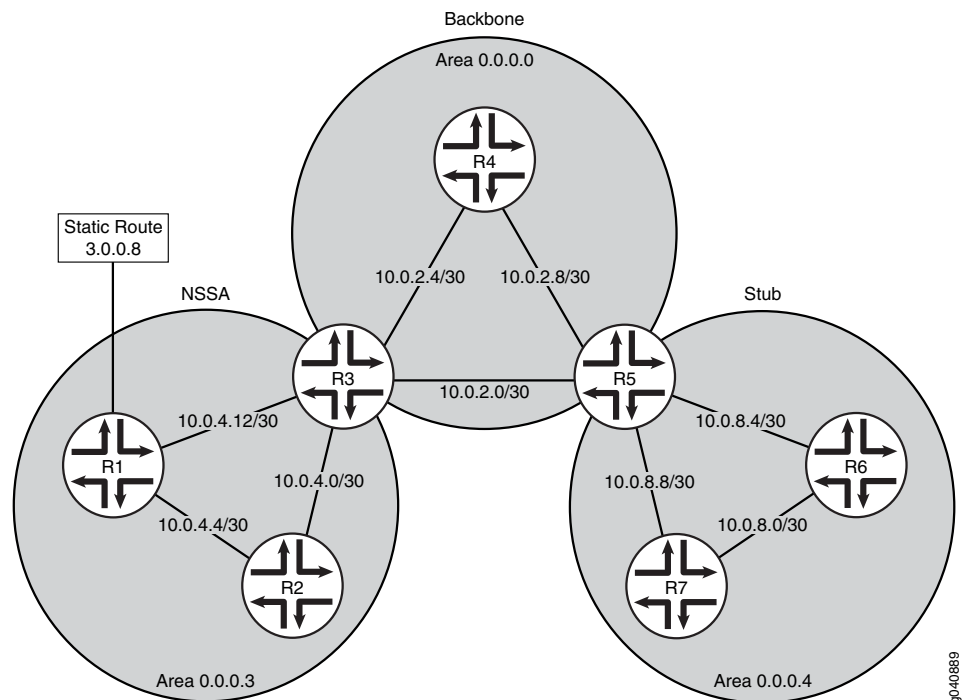
- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#)
- Configure a static route. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Library for Routing Devices*.

## Overview

You can summarize a range of IP addresses to minimize the size of the backbone router's link-state database. All routes that match the specified area range are filtered at the area boundary, and the summary is advertised in their place.

[Figure 120 on page 4095](#) shows the topology used in this example. R5 is the ABR between area 0.0.0.4 and the backbone. The networks in area 0.0.0.4 are 10.0.8.4/30, 10.0.8.0/30, and 10.0.8.8/30, which can be summarized as 10.0.8.0/28. R3 is the ABR between NSSA area 0.0.0.3 and the backbone. The networks in area 0.0.0.3 are 10.0.4.4/30, 10.0.4.0/30, and 10.0.4.12/30, which can be summarized as 10.0.4.0/28. Area 0.0.0.3 also contains external static route 3.0.0.8 that you will prevent from flooding throughout the network.

**Figure 120: Summarizing Ranges of Routes in OSPF**



In this example, you configure the ABRs for route summarization by including the following settings:

- **area-range**—For an area, summarizes a range of IP addresses when sending summary intra-area link advertisements. For an NSSA, summarizes a range of IP addresses when sending NSSA link-state advertisements (Type 7 LSAs). The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas.
- **network/mask-length**—Indicates the summarized IP address range and the number of significant bits in the network mask.
- **restrict**—On the NSSA ABR, prevents the configured summary from being advertised. In this example, we do not want to flood the external route outside of area 0.0.0.3.

### Configuration

#### CLI Quick Configuration

- To quickly configure route summarization for an OSPF area, copy the following commands and paste them into the CLI. The following is the configuration on ABR R5:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3
set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5
set protocols ospf area 0.0.0.4 stub
set protocols ospf area 0.0.0.4 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28
```

- To quickly configure route summarization for an OSPF NSSA, copy the following commands and paste them into the CLI. The following is the configuration on ABR R3:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10
set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1
set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1
set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7
set protocols ospf area 0.0.0.3 interface fe-0/0/1
set protocols ospf area 0.0.0.3 interface fe-0/0/2
set protocols ospf area 0.0.0.0 interface fe-0/0/0
set protocols ospf area 0.0.0.0 interface fe-0/0/4
set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
set protocols ospf area 0.0.0.3 nssa
set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8 restrict
```

#### Step-by-Step Procedure

To summarize routes sent to the backbone area:

1. Configure the interfaces.



**NOTE:** For OSPFv3, include IPv6 addresses.

```
[edit]
user@R5# set interfaces fe-0/0/1 unit 0 family inet address 10.0.8.3
user@R5# set interfaces fe-0/0/2 unit 0 family inet address 10.0.8.4
user@R5# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.3
user@R5# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.5

[edit]
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.10
user@R3# set interfaces fe-0/0/2 unit 0 family inet address 10.0.4.1
user@R3# set interfaces fe-0/0/0 unit 0 family inet address 10.0.2.1
user@R3# set interfaces fe-0/0/4 unit 0 family inet address 10.0.2.7
```

2. Configure the type of OSPF area.



**NOTE:** For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 stub

[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa
```

3. Assign the interfaces to the OSPF areas.

```
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/1
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/2
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R5# set protocols ospf area 0.0.0.0 interface fe-0/0/4

user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/1
user@R3# set protocols ospf area 0.0.0.3 interface fe-0/0/2
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/4
```

4. Summarize the routes that are flooded into the backbone.

```
[edit]
user@R5# set protocols ospf area 0.0.0.4 area-range 10.0.8.0/28

[edit]
user@R3# set protocols ospf area 0.0.0.3 area-range 10.0.4.0/28
```

5. On ABR R3, restrict the external static route from leaving area 0.0.0.3.

```
[edit]
user@R3# set protocols ospf area 0.0.0.3 nssa area-range 3.0.0.0/8 restrict
```

6. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

### Results

Confirm your configuration by entering the `show interfaces` and the `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Configuration on ABR R5:

```
user@R5# show interfaces
fe-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.2.3/32;
    }
  }
}
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.8.3/32;
    }
  }
}
fe-0/0/2 {
  unit 0 {
    family inet {
      address 10.0.8.4/32;
    }
  }
}
fe-0/0/4 {
  unit 0 {
    family inet {
      address 10.0.2.5/32;
    }
  }
}

user@R5# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/0.0;
  interface fe-0/0/4.0;
}
area 0.0.0.4 {
  stub;
  area-range 10.0.8.0/28;
  interface fe-0/0/1.0;
  interface fe-0/0/2.0;
}
```

Configuration on ABR R3:

```
user@R3# show interfaces
fe-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.2.1/32;
    }
  }
}
fe-0/0/1 {
  unit 0 {
    family inet {
```

```

        address 10.0.4.10/32;
    }
}
fe-0/0/2 {
    unit 0 {
        family inet {
            address 10.0.4.1/32;
        }
    }
}
fe-0/0/4 {
    unit 0 {
        family inet {
            address 10.0.2.7/32;
        }
    }
}

user@R3t# show protocols ospf
area 0.0.0.0 {
    interface fe-0/0/0.0;
    interface fe-0/0/4.0;
}
area 0.0.0.3 {
    nssa {
        area-range 3.0.0.0/8 restrict;
    }
    area-range 10.0.4.0/28;
    interface fe-0/0/1.0;
    interface fe-0/0/2.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces** and **show protocols ospf3** commands.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the Summarized Route**

**Purpose** Verify that the routes you configured for route summarization are being aggregated by the ABRs before the routes enter the backbone area. Confirm route summarization by checking the entries of the OSPF link-state database for the routing devices in the backbone.

**Action** From operational mode, enter the **show ospf database** command for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

### Example: Limiting the Number of Prefixes Exported to OSPF

---

This example shows how to limit the number of prefixes exported to OSPF.

- [Requirements on page 4100](#)
- [Overview on page 4100](#)
- [Configuration on page 4100](#)
- [Verification on page 4101](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### Overview

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area.

You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem. If the routing device exceeds the configured prefix export value, the routing device purges the external prefixes and enters into an overload state. This state ensures that the routing device is not overwhelmed as it attempts to process routing information. The prefix export limit number can be a value from 0 through 4,294,967,295.

In this example, you configure a prefix export limit of 100,000 by including the **prefix-export-limit** statement.

#### Configuration

##### CLI Quick Configuration

To quickly limit the number of prefixes exported to OSPF, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf prefix-export-limit 100000
```

**Step-by-Step Procedure** To limit the number of prefixes exported to OSPF:

1. Configure the prefix export limit value.



**NOTE:** For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf prefix-export-limit 100000
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
prefix-export-limit 100000;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Prefix Export Limit

**Purpose** Verify the prefix export counter that displays the number of routes exported into OSPF.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

### Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

---

The OSPF standard requires that every link-state advertisement (LSA) be refreshed every 30 minutes. The Juniper Networks implementation refreshes LSAs every 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes. This requirement can result in traffic overhead that makes it difficult to scale OSPF networks. You can override the default behavior by specifying that the DoNotAge bit be set in self-originated LSAs when they are initially sent by the router or switch. Any LSA with the DoNotAge bit set is reflooded only when a change occurs in the LSA. This feature thus reduces protocol traffic overhead while permitting any changed LSAs to be flooded immediately. Routers or switches enabled for flood reduction continue to send hello packets to their neighbors and to age self-originated LSAs in their databases.

The Juniper implementation of OSPF refresh and flooding reduction is based on RFC 4136, *OSPF Refresh and Flooding Reduction in Stable Topologies*. However, the Juniper implementation does not include the forced-flooding interval defined in the RFC. Not implementing the forced-flooding interval ensures that LSAs with the DoNotAge bit set are reflooded only when a change occurs.

This feature is supported for the following:

- OSPFv2 and OSPFv3 interfaces
- OSPFv3 realms
- OSPFv2 and OSPFv3 virtual links
- OSPFv2 sham links
- OSPFv2 peer interfaces
- All routing instances supported by OSPF
- Logical systems

To configure flooding reduction for an OSPF interface, include the **flood-reduction** statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-id]` hierarchy level.



**NOTE:** If you configure flooding reduction for an interface configured as a demand circuit, the LSAs are not initially flooded, but sent only when their content has changed. Hello packets and LSAs are sent and received on a demand-circuit interface only when a change occurs in the network topology.

---

In the following example, the OSPF interface **so-0/0/1.0** is configured for flooding reduction. As a result, all the LSAs generated by the routes that traverse the specified interface have the DoNotAge bit set when they are initially flooded, and LSAs are refreshed only when a change occurs.

```
[edit]
protocols ospf {
  area 0.0.0.0 {
```



```

interface so-0/0/1.0 {
    flood-reduction;
}
interface lo0.0;
interface so-0/0/0.0;
}

```



**NOTE:** Beginning with Junos OS Release 12.2, you can configure a global default link-state advertisement (LSA) flooding interval in OSPF for self-generated LSAs by including the `lsa-refresh-interval minutes` statement at the `[edit protocols (ospf | ospf3)]` hierarchy level. The Juniper Networks implementation refreshes LSAs every 50 minutes. The range is 25 through 50 minutes. By default, any LSA that is not refreshed expires after 60 minutes.

If you have both the global LSA refresh interval configured for OSPF and OSPF flooding reduction configured for a specific interface in an OSPF area, the OSPF flood reduction configuration takes precedence for that specific interface.

- Related Documentation**
- [OSPF Overview on page 4036](#)
  - [OSPF Configuration Overview](#)

## Examples: Configuring OSPF Traffic Control

- [Understanding OSPF Traffic Control on page 4103](#)
- [Example: Controlling the Cost of Individual OSPF Network Segments on page 4105](#)
- [Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth on page 4109](#)
- [Example: Controlling OSPF Route Preferences on page 4111](#)

### Understanding OSPF Traffic Control

Once a topology is shared across the network, OSPF uses the topology to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest-path-first (SPF) algorithm. Routes with lower total path metrics are preferred over those with higher path metrics.

You can use the following methods to control OSPF traffic:

- Control the cost of individual OSPF network segments
- Dynamically adjust OSPF interface metrics based on bandwidth
- Control OSPF route selection

### ***Controlling the Cost of Individual OSPF Network Segments***

OSPF uses the following formula to determine the cost of a route:

$\text{cost} = \text{reference-bandwidth} / \text{interface bandwidth}$

You can modify the reference-bandwidth value, which is used to calculate the default interface cost. The interface bandwidth value is not user-configurable and refers to the actual bandwidth of the physical interface.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface.

To control the flow of packets across the network, OSPF allows you to manually assign a cost (or metric) to a particular path segment. When you specify a metric for a specific OSPF interface, that value is used to determine the cost of routes advertised from that interface. For example, if all routers in the OSPF network use default metric values, and you increase the metric on one interface to 5, all paths through that interface have a calculated metric higher than the default and are not preferred.



**NOTE:** Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface.

When there are multiple equal-cost routes to the same destination in a routing table, an equal-cost multipath (ECMP) set is formed. If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose one of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. Define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the **[edit policy-options]** hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table.

#### ***Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth***

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, the Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. Junos OS uses the smallest configured bandwidth threshold value that is equal to or greater than the actual interface bandwidth to determine the metric value. If the interface bandwidth is greater than any of the configured bandwidth threshold values, the metric value configured for the interface is used instead of any of the bandwidth-based metric values configured. The ability to recalculate the metric for an interface when its bandwidth changes is especially useful for aggregate interfaces.



**NOTE:** You must also configure a metric for the interface when you enable bandwidth-based metrics.

### Controlling OSPF Route Preferences

You can control the flow of packets through the network using route preferences. Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. Although the default settings are appropriate for most environments, you might want to modify the default settings if all of the routing devices in your OSPF network use the default preference values, or if you are planning to migrate from OSPF to a different interior gateway protocol (IGP). If all of the devices use the default route preference values, you can change the route preferences to ensure that the path through a particular device is selected for the forwarding table any time multiple equal-cost paths to a destination exist. When migrating from OSPF to a different IGP, modifying the route preferences allows you to perform the migration in a controlled manner.

### Example: Controlling the Cost of Individual OSPF Network Segments

This example shows how to control the cost of individual OSPF network segments.

- [Requirements on page 4105](#)
- [Overview on page 4105](#)
- [Configuration on page 4106](#)
- [Verification on page 4108](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.

#### Overview

All OSPF interfaces have a cost, which is a routing metric that is used in the link-state calculation. Routes with lower total path metrics are preferred to those with higher path metrics. In this example, we explore how to control the cost of OSPF network segments.

By default, OSPF assigns a default cost metric of 1 to any link faster than 100 Mbps, and a default cost metric of 0 to the loopback interface (**lo0**). No bandwidth is associated with the loopback interface. This means that all interfaces faster than 100 Mbps have the same default cost metric of 1. If multiple equal-cost paths exist between a source

and destination address, OSPF routes packets along each path alternately, in round-robin fashion.


Having the same default metric might not be a problem if all of the interfaces are running at the same speed. If the interfaces operate at different speeds, you might notice that traffic is not routed over the fastest interface because OSPF equally routes packets across the different interfaces. For example, if your routing device has Fast Ethernet and Gigabit Ethernet interfaces running OSPF, each of these interfaces have a default cost metric of 1.

In the first example, you set the reference bandwidth to 10g (10 Gbps, as denoted by 10,000,000,000 bits) by including the **reference-bandwidth** statement. With this configuration, OSPF assigns the Fast Ethernet interface a default metric of 100, and the Gigabit Ethernet interface a metric of 10. Since the Gigabit Ethernet interface has the lowest metric, OSPF selects it when routing packets. The range is 9600 through 1,000,000,000,000 bits.

Figure 121 on page 4106 shows three routing devices in area 0.0.0.0 and assumes that the link between Device R2 and Device R3 is congested with other traffic. You can also control the flow of packets across the network by manually assigning a metric to a particular path segment. Any value you configure for the metric overrides the default behavior of using the reference-bandwidth value to calculate the route cost for that interface. To prevent the traffic from Device R3 going directly to Device R2, you adjust the metric on the interface on Device R3 that connects with Device R1 so that all traffic goes through Device R1.

In the second example, you set the metric to 5 on interface **fe-1/0/1** on Device R3 that connects with Device R1 by including the **metric** statement. The range is 1 through 65,535.

**Figure 121: OSPF Metric Configuration**



#### Configuration

- [Configuring the Reference Bandwidth on page 4107](#)
- [Configuring a Metric for a Specific OSPF Interface on page 4107](#)

*Configuring the Reference Bandwidth*

**CLI Quick Configuration** To quickly configure the reference bandwidth, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf reference-bandwidth 10g
```

**Step-by-Step Procedure** To configure the reference bandwidth:

1. Configure the reference bandwidth to calculate the default interface cost.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf reference-bandwidth 10g
```



**TIP:** As a shortcut in this example, you enter `10g` to specify 10 Gbps reference bandwidth. Whether you enter `10g` or `10000000000`, the output of `show protocols ospf` command displays 10 Gbps as `10g`, not `10000000000`.

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```



**NOTE:** Repeat this entire configuration on all routing devices in a shared network.

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
reference-bandwidth 10g;
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

*Configuring a Metric for a Specific OSPF Interface*

**CLI Quick Configuration** To quickly configure a metric for a specific OSPF interface, copy the following command and paste it into the CLI.

```
[edit]
```

```
set protocols ospf area 0.0.0.0 interface fe-1/0/1 metric 5
```

**Step-by-Step  
Procedure**

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
```

```
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0 ]
```

```
user@host# set interface fe-1/0/1 metric 5
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
```

```
user@host# commit
```

**Results**

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-1/0/1.0 {
    metric 5;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Configured Metric on page 4108](#)
- [Verifying the Route on page 4109](#)

**Verifying the Configured Metric****Purpose**

Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.

**Action**

From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Verifying the Route**

**Purpose** When choosing paths to a destination, OSPF uses the path with the lowest total cost. Confirm that OSPF is using the appropriate path.

**Action** From operational mode, enter the **show route** command.

**Example: Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth**

This example shows how to dynamically adjust OSPF interface metrics based on bandwidth.

- [Requirements on page 4109](#)
- [Overview on page 4109](#)
- [Configuration on page 4110](#)
- [Verification on page 4111](#)

**Requirements**

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.

**Overview**

You can specify a set of bandwidth threshold values and associated metric values for an OSPF interface. When the bandwidth of an interface changes, the Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value. When you configure bandwidth-based metric values, you typically configure multiple bandwidth and metric values.

In this example, you configure OSPF interface **ae0** for bandwidth-based metrics by including the **bandwidth-based-metrics** statement and the following settings:

- **bandwidth**—Specifies the bandwidth threshold in bits per second. The range is 9600 through 1,000,000,000,000,000.
- **metric**—Specifies the metric value to associate with a specific bandwidth value. The range is 1 through 65,535.

### Configuration

**CLI Quick Configuration** To quickly configure bandwidth threshold values and associated metric values for an OSPF interface, copy the following commands, remove any line breaks, and then paste the commands into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface ae0.0 metric 5
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 1g
metric 60
set protocols ospf area 0.0.0.0 interface ae0.0 bandwidth-based-metrics bandwidth 10g
metric 50
```

To configure the metric for a specific OSPF interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the metric of the OSPF network segment.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface ae0 metric 5
```

3. Configure the bandwidth threshold values and associated metric values.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 1g metric 60
user@host# set interface ae0.0 bandwidth-based-metrics bandwidth 10g metric 50
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface ae0.0 {
    bandwidth-based-metrics {
      bandwidth 1g metric 60;
      bandwidth 10g metric 50;
    }
    metric 5;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.



**Verification**

Confirm that the configuration is working properly.

**Verifying the Configured Metric**

- Purpose** Verify the metric setting on the interface. Confirm that the Cost field displays the interface's configured metric (cost). When choosing paths to a destination, OSPF uses the path with the lowest cost.
- Action** From operational mode, enter the **show ospf interface detail** command for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Example: Controlling OSPF Route Preferences**

This example shows how to control OSPF route selection in the forwarding table. This example also shows how you might control route selection if you are migrating from OSPF to another IGP.

- [Requirements on page 4111](#)
- [Overview on page 4111](#)
- [Configuration on page 4112](#)
- [Verification on page 4113](#)

**Requirements**

This example assumes that OSPF is properly configured and running in your network, and you want to control route selection because you are planning to migrate from OSPF to a different IGP.

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.
- Configure the IGP that you want to migrate to. See the *Junos OS Routing Protocols Library for Routing Devices*.

**Overview**

Route preferences are used to select which route is installed in the forwarding table when several protocols calculate routes to the same destination. The route with the lowest preference value is selected.

By default, internal OSPF routes have a preference value of 10, and external OSPF routes have a preference value of 150. You might want to modify this setting if you are planning to migrate from OSPF to a different IGP. Modifying the route preferences enables you to perform the migration in a controlled manner.

This example makes the following assumptions:

- OSPF is already running in your network.
- You want to migrate from OSPF to IS-IS.

- You configured IS-IS per your network requirements and confirmed it is working properly.

In this example, you increase the OSPF route preference values to make them less preferred than IS-IS routes by specifying 168 for internal OSPF routes and 169 for external OSPF routes. IS-IS internal routes have a preference of either 15 (for Level 1) or 18 (for Level 2), and external routes have a preference of 160 (for Level 1) or 165 (for Level 2). In general, it is preferred to leave the new protocol at its default settings to minimize complexities and simplify any future addition of routing devices to the network. To modify the OSPF route preference values, configure the following settings:

- **preference**—Specifies the route preference for internal OSPF routes. By default, internal OSPF routes have a value of 10. The range is from 0 through 4,294,967,295 ( $2^{32} - 1$ ).
- **external-preference**—Specifies the route preference for external OSPF routes. By default, external OSPF routes have a value of 150. The range is from 0 through 4,294,967,295 ( $2^{32} - 1$ ).

### Configuration

#### CLI Quick Configuration

To quickly configure the OSPF route preference values, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf preference 168 external-preference 169
```

To configure route selection:

1. Enter OSPF configuration mode and set the external and internal routing preferences.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf preference 168 external-preference 169
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

#### Results

Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
preference 168;
external-preference 169;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Route on page 4113](#)

**Verifying the Route**

**Purpose** Verify that the IGP is using the appropriate route. After the new IGP becomes the preferred protocol (in this example, IS-IS), you should monitor the network for any issues. After you confirm that the new IGP is working properly, you can remove the OSPF configuration from the routing device by entering the **delete ospf** command at the **[edit protocols]** hierarchy level.

**Action** From operational mode, enter the **show route** command.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)

**Example: Configuring OSPF Overload Mode**

- [OSPF Overload Function Overview on page 4113](#)
- [Example: Configuring OSPF to Make Routing Devices Appear Overloaded on page 4114](#)

**OSPF Overload Function Overview**

If the time elapsed after the OSPF instance is enabled is less than the specified timeout, overload mode is set.

You can configure the local routing device so that it appears to be overloaded. An overloaded routing device determines it is unable to handle any more OSPF transit traffic, which results in sending OSPF transit traffic to other routing devices. OSPF traffic to directly attached interfaces continues to reach the routing device. You might configure overload mode for many reasons, including:

- If you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic. This could include a routing device that is connected to the network for analysis purposes, but is not considered part of the production network, such as network management routing devices.
- If you are performing maintenance on a routing device in a production network. You can move traffic off that routing device so network services are not interrupted during your maintenance window.

You configure or disable overload mode in OSPF with or without a timeout. Without a timeout, overload mode is set until it is explicitly deleted from the configuration. With a timeout, overload mode is set if the time elapsed since the OSPF instance started is less than the specified timeout.

A timer is started for the difference between the timeout and the time elapsed since the instance started. When the timer expires, overload mode is cleared. In overload mode,

the router link-state advertisement (LSA) is originated with all the transit router links (except stub) set to a metric of 0xFFFF. The stub router links are advertised with the actual cost of the interfaces corresponding to the stub. This causes the transit traffic to avoid the overloaded routing device and to take paths around the routing device. However, the overloaded routing device's own links are still accessible.

The routing device can also dynamically enter the overload state, regardless of configuring the device to appear overloaded. For example, if the routing device exceeds the configured OSPF prefix limit, the routing device purges the external prefixes and enters into an overload state.

In cases of incorrect configurations, the huge number of routes might enter OSPF, which can hamper the network performance. To prevent this, **prefix-export-limit** should be configured which will purge externals and prevent the network from the bad impact.

By allowing any number of routes to be exported into OSPF, the routing device can become overwhelmed and potentially flood an excessive number of routes into an area. You can limit the number of routes exported into OSPF to minimize the load on the routing device and prevent this potential problem.

By default, there is no limit to the number of prefixes (routes) that can be exported into OSPF. To prevent this, **prefix-export-limit** should be configured which will purge externals and prevent the network.

To limit the number of prefixes exported to OSPF:

```
[edit]
set protocols ospf prefix-export-limit number
```

The prefix export limit number can be a value from 0 through 4,294,967,295.

### Example: Configuring OSPF to Make Routing Devices Appear Overloaded

This example shows how to configure a routing device running OSPF to appear to be overloaded.

- [Requirements on page 4114](#)
- [Overview on page 4115](#)
- [Configuration on page 4115](#)
- [Verification on page 4116](#)

#### **Requirements**

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See "[Example: Configuring an OSPF Router Identifier](#)" on page 4048.
- Control OSPF designated router election. See "[Example: Controlling OSPF Designated Router Election](#)" on page 4050

- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4053](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

### Overview

You can configure a local routing device running OSPF to appear to be overloaded, which allows the local routing device to participate in OSPF routing, but not for transit traffic. When configured, the transit interface metrics are set to the maximum value of 65535.

This example includes the following settings:

- **overload**—Configures the local routing device so it appears to be overloaded. You might configure this if you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic, or you are performing maintenance on a routing device in a production network.
- **timeout seconds**—(Optional) Specifies the number of seconds at which the overload is reset. If no timeout interval is specified, the routing device remains in the overload state until the overload statement is deleted or a timeout is set. In this example, you configure 60 seconds as the amount of time the routing device remains in the overload state. By default, the timeout interval is 0 seconds (this value is not configured). The range is from 60 through 1800 seconds.

### Configuration

**CLI Quick Configuration** To quickly configure a local routing device to appear as overloaded, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf overload timeout 60
```

**Step-by-Step Procedure** To configure a local routing device to appear overloaded:

1. Enter OSPF configuration mode.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf
```

2. Configure the local routing device to be overloaded.

```
[edit protocols ospf]
user@host# set overload
```

3. (Optional) Configure the number of seconds at which overload is reset.

```
[edit protocols ospf]
user@host# set overload timeout 60
```

4. (Optional) Configure the limit on the number prefixes exported to OSPF, to minimise the load on the routing device and prevent the device from entering the overload mode.

```
[edit protocols ospf]  
user@host# set prefix-export-limit 50
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]  
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration. The output includes the optional **timeout** and **prefix-export-limit** statements.

```
user@host# show protocols ospf
```

```
prefix-export-limit 50;  
overload timeout 60;
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### ***Verification***

Confirm that the configuration is working properly.

- [Verifying Traffic Has Moved Off Devices on page 4116](#)
- [Verifying Transit Interface Metrics on page 4116](#)
- [Verifying the Overload Configuration on page 4116](#)
- [Verifying the Viable Next Hop on page 4117](#)

### ***Verifying Traffic Has Moved Off Devices***

**Purpose** Verify that the traffic has moved off the upstream devices.

**Action** From operational mode, enter the **show interfaces detail** command.

### ***Verifying Transit Interface Metrics***

**Purpose** Verify that the transit interface metrics are set to the maximum value of 65535 on the downstream neighboring device.

**Action** From operational mode, enter the **show ospf database router detail advertising-router address** command for OSPFv2, and enter the **show ospf3 database router detail advertising-router address** command for OSPFv3.

### ***Verifying the Overload Configuration***

**Purpose** Verify that overload is configured by reviewing the Configured overload field. If the overload timer is also configured, this field also displays the time that remains before it is set to expire.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and the **show ospf3 overview** command for OSPFv3.

#### *Verifying the Viable Next Hop*

**Purpose** Verify the viable next hop configuration on the upstream neighboring device. If the neighboring device is overloaded, it is not used for transit traffic and is not displayed in the output.

**Action** From operational mode, enter the **show route address** command.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)

## OSPF Fault Detection Configuration

- [Example: Configuring OSPF Timers on page 4117](#)
- [Example: Configuring BFD for OSPF on page 4123](#)
- [Example: Configuring BFD Authentication for OSPF on page 4129](#)

### Example: Configuring OSPF Timers

- [OSPF Timers Overview on page 4117](#)
- [Example: Configuring OSPF Timers on page 4118](#)

#### OSPF Timers Overview

OSPF routing devices constantly track the status of their neighbors, sending and receiving hello packets that indicate whether each neighbor still is functioning, and sending and receiving link-state advertisement (LSA) and acknowledgment packets. OSPF sends packets and expects to receive packets at specified intervals.

You configure OSPF timers on the interface of the routing device participating in OSPF. Depending on the timer, the configured interval must be the same on all routing devices on a shared network (area).

You can configure the following OSPF timers:

- Hello interval—Routing devices send hello packets at a fixed interval on all interfaces, including virtual links, to establish and maintain neighbor relationships. The hello interval specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. This interval must be the same on all routing devices on a shared network. By default, the routing device sends hello packets every 10 seconds (broadcast and point-to-point networks) and 30 seconds (nonbroadcast multiple access (NBMA) networks).
- Poll interval—(OSPFv2, Nonbroadcast networks only) Routing devices send hello packets for a longer interval on nonbroadcast networks to minimize the bandwidth required on slow WAN links. The poll interval specifies the length of time, in seconds,

before the routing device sends hello packets out of the interface before establishing adjacency with a neighbor. By default, the routing device sends hello packets every 120 seconds until active neighbors are detected.

Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval.

- **LSA retransmission interval**—When a routing device sends LSAs to its neighbors, the routing device expects to receive an acknowledgment packet from each neighbor within a certain amount of time. The LSA retransmission interval specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting the LSA to an interface's neighbors. By default, the routing device waits 5 seconds for an acknowledgment before retransmitting the LSA.
- **Dead interval**—If a routing device does not receive a hello packet from a neighbor within a fixed amount of time, the routing device modifies its topology database to indicate that the neighbor is nonoperational. The dead interval specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. This interval must be the same on all routing devices on a shared network. By default, this interval is four times the default hello interval, which is 40 seconds (broadcast and point-to-point networks) and 120 seconds (NBMA networks).
- **Transit delay**—Before a link-state update packet is propagated out of an interface, the routing device must increase the age of the packet. The transit delay sets the estimated time required to transmit a link-state update on the interface. By default, the transit delay is 1 second. You should never have to modify the transit delay time.

---

### Example: Configuring OSPF Timers

This example shows how to configure the OSPF timers.

- [Requirements on page 4118](#)
- [Overview on page 4119](#)
- [Configuration on page 4120](#)
- [Verification on page 4123](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050



- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4053](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

### Overview

The default OSPF timer settings are optimal for most networks. However, depending on your network requirements, you might need to modify the timer settings. This example explains why you might need to modify the following timers:

- Hello interval
- Dead interval
- LSA retransmission interval
- Transit delay

### Hello Interval and Dead Interval

The hello interval and the dead interval optimize convergence times by efficiently tracking neighbor status. By lowering the values of the hello interval and the dead interval, you can increase the convergence of OSPF routes if a path fails. These intervals must be the same on all routing devices on a shared network. Otherwise, OSPF cannot establish the appropriate adjacencies.

In the first example, you lower the hello interval to 2 seconds and the dead interval to 8 seconds on point-to-point OSPF interfaces **fe-0/0/1** and **fe-1/0/1** in area 0.0.0.0 by configuring the following settings:

- **hello-interval**—Specifies the length of time, in seconds, before the routing device sends a hello packet out of an interface. By default, the routing device sends hello packets every 10 seconds. The range is from 1 through 255 seconds.
- **dead-interval**—Specifies the length of time, in seconds, that the routing device waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor. By default, the routing device waits 40 seconds (four times the hello interval). The range is 1 through 65,535 seconds.

### LSA Retransmission Interval

The link-state advertisement (LSA) retransmission interval optimizes the sending and receiving of LSA and acknowledgement packets. You must configure the LSA retransmission interval to be equal to or greater than 3 seconds to avoid triggering a retransmit trap because the Junos OS delays LSA acknowledgments by up to 2 seconds. If you have a virtual link, you might find increased performance by increasing the value of the LSA retransmission interval.

In the second example, you increase the LSA retransmission timer to 8 seconds on OSPF interface **fe-0/0/1** in area 0.0.0.1 by configuring the following setting:

- **retransmit-interval**—Specifies the length of time, in seconds, that the routing device waits to receive an LSA packet before retransmitting LSA to an interface's neighbors. By default, the routing device retransmits LSAs to its neighbors every 5 seconds. The range is from 1 through 65,535 seconds.

### Transit Delay

The transit delay sets the time the routing device uses to age a link-state update packet. If you have a slow link (for example, one with an average propagation delay of multiple seconds), you should increase the age of the packet by a similar amount. Doing this ensures that you do not receive a packet back that is younger than the original copy.

In the final example, you increase the transit delay to 2 seconds on OSPF interface **fe-1/0/1** in area 0.0.0.1. By configuring the following setting, this causes the routing device to age the link-state update packet by 2 seconds:

- **transit-delay**—Sets the estimated time required to transmit a link-state update on the interface. You should never have to modify the transit delay time. By default, the routing device ages the packet by 1 second. The range is from 1 through 65,535 seconds.

### Configuration

- [Configuring the Hello Interval and the Dead Interval on page 4120](#)
- [Controlling the LSA Retransmission Interval on page 4121](#)
- [Specifying the Transit Delay on page 4122](#)

### Configuring the Hello Interval and the Dead Interval

#### CLI Quick Configuration

To quickly configure the hello and dead intervals, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-0/0/1 dead-interval 8
set protocols ospf area 0.0.0.0 interface fe-1/0/1 hello-interval 2
set protocols ospf area 0.0.0.0 interface fe-1/0/1 dead-interval 8
```

#### Step-by-Step Procedure

To configure the hello and dead intervals:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

---

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interfaces.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
user@host# set interface fe-1/0/1
```

3. Configure the hello interval.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 hello-interval 2
user@host# set interface fe-1/0/1 hello-interval 2
```

4. Configure the dead interval.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 dead-interval 8
user@host# set interface fe-1/0/1 dead-interval 8
```

5. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



**NOTE:** Repeat this entire configuration on all routing devices in a shared network.

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    hello-interval 2;
    dead-interval 8;
  }
  interface fe-1/0/1.0 {
    hello-interval 2;
    dead-interval 8;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### *Controlling the LSA Retransmission Interval*

**CLI Quick Configuration** To quickly configure the LSA retransmission interval, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-0/0/1 retransmit-interval 8
```

**Step-by-Step Procedure** To configure the LSA retransmission interval:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-0/0/1
```

3. Configure the LSA retransmission interval.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-0/0/1 retransmit-interval 8
```

4. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.1]
user@host# commit
```

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface fe-0/0/1.0 {
    retransmit-interval 8;
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

### *Specifying the Transit Delay*

**CLI Quick Configuration** To quickly configure the transit delay, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface fe-1/0/1 transit-delay 2
```

**Step-by-Step Procedure** To configure the transit delay:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.1
```

2. Specify the interface.  

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface fe-1/0/1
```
3. Configure the transit delay.  

```
[edit protocols ospf area 0.0.0.1 ]
user@host# set interface fe-1/0/1 transit-delay 2
```
4. If you are done configuring the device, commit the configuration.  

```
[edit protocols ospf area 0.0.0.1 ]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface fe-1/0/1.0 {
    transit-delay 2;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the Timer Configuration**

**Purpose** Verify that the interface for OSPF or OSPFv3 has been configured with the applicable timer values. Confirm that the Hello field, the Dead field, and the ReXmit field display the values that you configured.

**Action** From operational mode, enter the **show ospf interface detail** for OSPFv2, and enter the **show ospf3 interface detail** command for OSPFv3.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)

## **Example: Configuring BFD for OSPF**

- [Understanding BFD for OSPF on page 4123](#)
- [Example: Configuring BFD for OSPF on page 4126](#)

### **Understanding BFD for OSPF**

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchange BFD packets. Hello packets are sent

at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the OSPF failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.



**NOTE:** BFD is supported for OSPFv3 in Junos OS Release 9.3 and later.

---

You can configure the following BFD protocol settings:

- **detection-time threshold**—Threshold for the adaptation of the detection time. When the BFD session detection time adapts to a value equal to or greater than the configured threshold, a single trap and a single system log message are sent.
- **full-neighbors-only**—Ability to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors. This setting is available in Junos OS Release 9.5 and later.
- **minimum-interval**—Minimum transmit and receive interval for failure detection. This setting configures both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. Both intervals are in milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. Without NSR, Routing Engine-based sessions can have a minimum interval of 100 ms. In OSPFv3, BFD is always based in the Routing Engine, meaning that BFD is not distributed. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **minimum-receive-interval**—Minimum receive interval for failure detection. This setting configures the minimum receive interval, in milliseconds, after which the routing device expects to receive a hello packet from a neighbor with which it has established a BFD session. You can also specify the minimum receive interval using the **minimum-interval** statement.
- **multiplier**—Multiplier for hello packets. This setting configures the number of hello packets that are not received by a neighbor, which causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down.
- **no-adaptation**—Disables BFD adaption. This setting disables BFD sessions from adapting to changing network conditions. This setting is available in Junos OS Release 9.0 and later.



**NOTE:** We recommend that you do not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

- **transmit-interval minimum-interval**—Minimum transmit interval for failure detection. This setting configures the minimum transmit interval, in milliseconds, at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can also specify the minimum transmit interval using the **minimum-interval** statement.

- **transmit-interval threshold**—Threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The threshold value must be greater than the minimum transmit interval. If you attempt to commit a configuration with a threshold value less than the minimum transmit interval, the routing device displays an error and does not accept the configuration.
- **version**—BFD version. This setting configures the BFD version used for detection. You can explicitly configure BFD version 1, or the routing device can automatically detect the BFD version. By default, the routing device automatically detects the BFD version automatically, which is either 0 or 1.

You can also trace BFD operations for troubleshooting purposes.

---

### Example: Configuring BFD for OSPF

This example shows how to configure the Bidirectional Forwarding Detection (BFD) protocol for OSPF.

- [Requirements on page 4126](#)
- [Overview on page 4126](#)
- [Configuration on page 4128](#)
- [Verification on page 4129](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### Overview

An alternative to adjusting the OSPF hello interval and dead interval settings to increase route convergence is to configure BFD. The BFD protocol is a simple hello mechanism that detects failures in a network. The BFD failure detection timers have shorter timer limits than the OSPF failure detection mechanisms, thereby providing faster detection.



BFD is useful on interfaces that are unable to detect failure quickly, such as Ethernet interfaces. Other interfaces, such as SONET interfaces, already have built-in failure detection. Configuring BFD on those interfaces is unnecessary.

You configure BFD on a pair of neighboring OSPF interfaces. Unlike the OSPF hello interval and dead interval settings, you do not have to enable BFD on all interfaces in an OSPF area.

In this example, you enable failure detection by including the **bfd-liveness-detection** statement on the neighbor OSPF interface **fe-0/1/0** in area 0.0.0.0 and configure the BFD packet exchange interval to 300 milliseconds, configure 4 as the number of missed hello packets that causes the originating interface to be declared down, and configure BFD sessions only for OSPF neighbors with full neighbor adjacency by including the following settings:

- **full-neighbors-only**—In Junos OS Release 9.5 and later, configures the BFD protocol to establish BFD sessions only for OSPF neighbors with full neighbor adjacency. The default behavior is to establish BFD sessions for all OSPF neighbors.
- **minimum-interval**—Configures the minimum interval, in milliseconds, after which the local routing device transmits hello packets as well as the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. You can also specify the minimum transmit and receive intervals separately using the **transmit-interval** **minimum-interval** and **minimum-receive-interval** statements.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

- **multiplier**—Configures the number of hello packets not received by a neighbor that causes the originating interface to be declared down. By default, three missed hello packets cause the originating interface to be declared down. You can configure a value in the range from 1 through 255.

### Configuration

**CLI Quick Configuration** To quickly configure the BFD protocol for OSPF, copy the following commands, remove any line breaks, and then paste the commands into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection multiplier 4
set protocols ospf area 0.0.0.0 interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

**Step-by-Step Procedure** To configure the BFD protocol for OSPF on one neighboring interface:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Specify the interface.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/0/1
```

3. Specify the minimum transmit and receive intervals.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection minimum-interval 300
```

4. Configure the number of missed hello packets that cause the originating interface to be declared down.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection multiplier 4
```

5. Configure BFD sessions only for OSPF neighbors with full neighbor adjacency.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# set interface fe-0/0/1 bfd-liveness-detection full-neighbors-only
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0 ]
user@host# commit
```



**NOTE:** Repeat this entire configuration on the other neighboring interface.

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0 {
    bfd-liveness-detection {
      minimum-interval 300;
      multiplier 4;
      full-neighbors-only;
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### **Verification**

Confirm that the configuration is working properly.

### **Verifying the BFD Sessions**

**Purpose** Verify that the OSPF interfaces have active BFD sessions, and that session components have been configured correctly.

**Action** From operational mode, enter the **show bfd session detail** command.

**Meaning** The output displays information about the BFD sessions.

- The Address field displays the IP address of the neighbor.
- The Interface field displays the interface you configured for BFD.
- The State field displays the state of the neighbor and should show Full to reflect the full neighbor adjacency that you configured.
- The Transmit Interval field displays the time interval you configured to send BFD packets.
- The Multiplier field displays the multiplier you configured.

**Related Documentation**

- *OSPF Configuration Overview*
- [BFD Authentication for OSPF Overview on page 4130](#)

## **Example: Configuring BFD Authentication for OSPF**

- [BFD Authentication for OSPF Overview on page 4130](#)
- [Configuring BFD Authentication for OSPF on page 4131](#)

## BFD Authentication for OSPF Overview

---

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over OSPFv2. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- [BFD Authentication Algorithms on page 4130](#)
- [Security Authentication Keychains on page 4131](#)
- [Strict Versus Loose Authentication on page 4131](#)

### **BFD Authentication Algorithms**

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method,

packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.

- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

### **Security Authentication Keychains**

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### **Strict Versus Loose Authentication**

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Configuring BFD Authentication for OSPF**

Beginning with Junos OS Release 9.6, you can configure authentication for BFD sessions running over OSPFv2. Routing instances are also supported.

The following sections provide instructions for configuring and viewing BFD authentication on OSPF:

- [Configuring BFD Authentication Parameters on page 4131](#)
- [Viewing Authentication Information for BFD Sessions on page 4133](#)

### **Configuring BFD Authentication Parameters**

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the OSPFv2 protocol.
2. Associate the authentication keychain with the OSPFv2 protocol.
3. Configure the related security authentication keychain.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on an OSPF route or routing instance.

[edit]

```
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection  
authentication algorithm keyed-sha-1
```

---



**NOTE:** Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

---

2. Specify the keychain to be used to associate BFD sessions on the specified OSPF route or routing instance with the unique security authentication keychain attributes.

This keychain should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

[edit]

```
user@host# set protocols ospf area 0.0.0.1 interface if2-ospf bfd-liveness-detection  
authentication keychain bfd-ospf
```

---



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

---

3. Specify the unique security authentication information for BFD sessions:
  - The matching keychain name as specified in Step 2.
  - At least one key, a unique integer between 0 and 63. Creating multiple keys enables multiple clients to use the BFD session.
  - The secret data used to allow access to the session.
  - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

[edit security]

```
user@host# authentication-key-chains key-chain bfd-ospf key 53 secret  
$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHM start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host> set protocols ospf interface if2-ospf bfd-liveness-detection authentication
loose-check
```

5. (Optional) View your configuration using the **show bfd session detail** or **show bfd session extensive** command.
6. Repeat the steps in this procedure to configure the other end of the BFD session.



**NOTE:** BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

### *Viewing Authentication Information for BFD Sessions*

You can view the existing BFD authentication configuration using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **if2-ospf** BGP group. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-ospf**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9l.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols ospf]
area 0.0.0.1 {
  interface if2-ospf {
    bfd-liveness-detection {
      authentication {
        algorithm keyed-sha-1;
        key-chain bfd-ospf;
      }
    }
  }
}
[edit security]
authentication key-chains {
  key-chain bfd-ospf {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$9$a5jiKW9l.reP38ny.TszF2/9";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
```

If you commit these updates to your configuration, you see output similar to the following. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured.

## show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3

Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**  
 Session up time 3d 00:34  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Replicated

1 sessions, 1 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

## show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3

Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, **Authenticate**  
**keychain bfd-ospf, algo keyed-md5, mode loose**

Session up time 3d 00:34  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Replicated  
 Min async interval 0.200, min slow interval 1.000  
 Adaptive async tx interval 0.200, rx interval 0.200  
 Local min tx interval 0.200, min rx interval 0.200, multiplier 3  
 Remote min tx interval 0.100, min rx interval 0.100, multiplier 3  
 Threshold transmission interval 0.000, Threshold for detection time 0.000  
 Local discriminator 11, remote discriminator 80  
 Echo mode disabled/inactive  
**Authentication enabled/active, keychain bfd-ospf, algo keyed-sha-1, mode strict**

1 sessions, 1 clients  
 Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

- Related Documentation**
- [OSPF Configuration Overview](#)
  - [Understanding BFD for OSPF on page 4123](#)

## OSPF Redundancy Features Configuration

- [Examples: Configuring Graceful Restart for OSPF on page 4134](#)

### Examples: Configuring Graceful Restart for OSPF

- [Graceful Restart for OSPF Overview on page 4135](#)
- [Example: Configuring Graceful Restart for OSPF on page 4136](#)



- [Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 4140](#)
- [Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4144](#)
- [Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 4147](#)

### Graceful Restart for OSPF Overview

Graceful restart allows a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting device and its neighbors continue forwarding packets without disrupting network performance. Because neighboring devices assist in the restart (these neighbors are called *helper routers*), the restarting device can quickly resume full operation without recalculating algorithms.



**NOTE:** On a broadcast link with a single neighbor, when the neighbor initiates an OSPFv3 graceful restart operation, the restart might be terminated at the point when the local routing device assumes the role of a helper. A change in the LSA is considered a topology change, which terminates the neighbor's restart operation.

Graceful restart is disabled by default. You can globally enable graceful restart for all routing protocols by including the **graceful-restart** statement at the **[edit routing-options]** hierarchy level. To enable graceful restart specifically for OSPF, first you need to globally enable graceful restart for all routing protocols.

This topic describes the following information:

- [Helper Mode for Graceful Restart on page 4135](#)
- [Planned and Unplanned Graceful Restart on page 4136](#)

#### **Helper Mode for Graceful Restart**

When a device enabled for OSPF graceful restart restarts, it retains routes learned before the restart in its forwarding table. The device does not allow new OSPF link-state advertisements (LSAs) to update the routing table. This device continues to forward traffic to other OSPF neighbors (or helper routers), and sends only a limited number of LSAs during the restart period. To reestablish OSPF adjacencies with neighbors, the restarting device must send a grace LSA to all neighbors. In response, the helper routers enter helper mode (the ability to assist a neighboring device attempting a graceful restart) and send an acknowledgment back to the restarting device. If there are no topology changes, the helper routers continue to advertise LSAs as if the restarting device had remained in continuous OSPF operation.



**NOTE:** Helper mode is enabled by default when you start the routing platform, even if graceful restart is not enabled. You can disable helper mode specifically for OSPF.

When the restarting device receives replies from all the helper routers, the restarting device selects routes, updates the forwarding table, and discards the old routes. At this point, full OSPF adjacencies are reestablished and the restarting device receives and processes OSPF LSAs as usual. When the helper routers no longer receive grace LSAs from the restarting device or when the topology of the network changes, the helper routers also resume normal operation.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The Junos OS implementation is based on RFC 4811, *OSPF Out-of-Band Link State Database (LSDB) Resynchronization*, RFC 4812, *OSPF Restart Signaling*, and RFC 4813, *OSPF Link-Local Signaling*. In restart signaling-based helper mode implementations, the restarting device informs its restart status to its neighbors only after the restart is complete. When the restart is complete, the restarting device sends hello messages to its helper routers with the restart signal (RS) bit set in the hello packet header. When a helper router receives a hello packet with the RS bit set in the header, the helper router returns a hello message to the restarting device. The reply hello message from the helper router contains the ResyncState flag and the ResyncTimeout timer that enable the restarting device to keep track of the helper routers that are syncing up with it. When all helpers complete the synchronization, the restarting device exits the restart mode.



**NOTE:** Restart signaling-based graceful restart helper mode is not supported for OSPFv3 configurations.

---

### ***Planned and Unplanned Graceful Restart***

OSPF supports two types of graceful restart: planned and unplanned. During a planned restart, the restarting routing device informs the neighbors before restarting. The neighbors act as if the routing device is still within the network topology, and continue forwarding traffic to the restarting routing device. A grace period is set to specify when the neighbors should consider the restarting routing device as part of the topology. During an unplanned restart, the routing device restarts without warning.

### **Example: Configuring Graceful Restart for OSPF**

---

This example shows how to configure graceful restart specifically for OSPF.

- [Requirements on page 4136](#)
- [Overview on page 4137](#)
- [Configuration on page 4137](#)
- [Verification on page 4140](#)

#### ***Requirements***

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#).
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4053](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

### Overview

Graceful restart enables a routing device undergoing a restart to inform its adjacent neighbors and peers of its condition. During a graceful restart, the restarting routing device and its neighbors continue forwarding packets without disrupting network performance. By default, graceful restart is disabled. You can globally enable graceful restart for all routing protocols by including the **graceful-restart** statement at the **[edit routing-options]** hierarchy level, or you can enable graceful restart specifically for OSPF by including the **graceful-restart** statement at the **[edit protocols (ospf|ospf3)]** hierarchy level.

The first example shows how to enable graceful restart and configure the optional settings for the grace period interval. In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. The grace period interval for OSPF graceful restart is determined as equal to or less than the sum of the **notify-duration** time interval and the **restart-duration** time interval. The grace period is the number of seconds that the routing device's neighbors continue to advertise the routing device as fully adjacent, regardless of the connection state between the routing device and its neighbors.

The **notify-duration** statement configures how long (in seconds) the routing device notifies helper routers that it has completed graceful restart by sending purged grace link-state advertisements (LSAs) over all interfaces. By default, the routing device sends grace LSAs for 30 seconds. The range is from 1 through 3600 seconds.

The **restart-duration** statement configures the amount of time the routing device waits (in seconds) to complete reacquisition of OSPF neighbors from each area. By default, the routing device allows 180 seconds. The range is from 1 through 3600 seconds.

The second example shows how to disable graceful restart for OSPF by including the **disable** statement.

### Configuration

- [Enabling Graceful Restart for OSPF on page 4137](#)
- [Disabling Graceful Restart for OSPF on page 4139](#)

#### Enabling Graceful Restart for OSPF

**CLI Quick Configuration** To quickly enable graceful restart for OSPF, copy the following commands and paste them into the CLI.

[edit]

```
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set routing-options graceful-restart
set protocols ospf graceful-restart restart-duration 190
set protocols ospf graceful-restart notify-duration 40
```

**Step-by-Step Procedure** To enable graceful restart for OSPF:

1. Configure the interfaces.



**NOTE:** For OSPFv3, use IPv6 addresses.

[edit]

```
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
```

```
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

[edit]

```
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
```

```
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Configure graceful restart globally

[edit]

```
user@host# edit routing-options graceful-restart
```

4. Configure OSPF graceful restart.

[edit]

```
user@host# edit protocols ospf graceful-restart
```

5. (Optional) Configure the restart duration time.

[edit protocols ospf graceful-restart]

```
user@host# set restart-duration 190
```

6. (Optional) Configure the notify duration time.

[edit protocols ospf graceful-restart]

```
user@host# set notify-duration 40
```

7. If you are done configuring the device, commit the configuration.

[edit protocols ospf graceful-restart]

```
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces** and **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  restart-duration 190;
  notify-duration 40;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

To confirm an OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

### *Disabling Graceful Restart for OSPF*

**CLI Quick Configuration** To quickly disable graceful restart for OSPF, copy the following command and paste it into the CLI.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

**Step-by-Step Procedure** To disable graceful restart for OSPF:

1. Disable graceful restart for the OSPF protocol only.

This command does not affect the global graceful restart configuration setting.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf graceful-restart disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
graceful-restart disable;
```

To confirm an OSPFv3 configuration, enter the **show protocols ospf3** command.

### ***Verification***

Confirm that the configuration is working properly.

- [Verifying the OSPF Graceful Restart Configuration on page 4140](#)
- [Verifying Graceful Restart Status on page 4140](#)

### ***Verifying the OSPF Graceful Restart Configuration***

**Purpose** Verify information about your OSPF graceful restart configuration.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2. Enter the **show ospf3 overview** command for OSPFv3.

**Meaning** The Restart field displays the status of graceful restart as either enabled or disabled. The Restart duration field displays how much time the restarted routing device requires to complete reacquisition of OSPF neighbors. The Restart grace period field displays how much time the neighbors should consider the restarted routing device as part of the topology.

### ***Verifying Graceful Restart Status***

**Purpose** Verify the status of graceful restart.

**Action** From operational mode, enter the **show route instance detail** command.

**Meaning** The Restart State field displays Pending if the restart has not been completed or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have or have not yet completed graceful restart for the specified routing table.

### **Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart**

This example shows how to disable and reenabling the helper mode capability for OSPFv2 graceful restart.

- [Requirements on page 4141](#)
- [Overview on page 4141](#)

- [Configuration on page 4141](#)
- [Verification on page 4143](#)

### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

### Overview

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. Both the standard and restart signaling-based helper modes are enabled by default.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv2 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the standard OSPFv2 graceful restart helper capability by including the **helper-disable standard** statement. This configuration is useful if you have an environment that contains other vendor equipment that is configured for restart signaling-based graceful restart.



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

The second example shows how to reenabling the standard OSPFv2 restart helper capability that you disabled in the first example.

### Configuration

- [Disabling Helper Mode for OSPFv2 on page 4142](#)
- [Reenabling Helper Mode for OSPFv2 on page 4143](#)

### *Disabling Helper Mode for OSPFv2*

**CLI Quick Configuration** To quickly enable graceful restart for OSPFv2 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart helper-disable standard
```

**Step-by-Step Procedure** To enable graceful restart for OSPFv2 with helper mode disabled:

1. Configure the interfaces.  

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```
2. Configure OSPFv2 on the interfaces  

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```
3. Disable the OSPFv2 graceful restart helper capability.  
If you disable the OSPFv2 graceful restart helper capability, you cannot disable strict LSA checking.  

```
[edit]
user@host# set protocols ospf graceful-restart helper-disable standard
```
4. If you are done configuring the device, commit the configuration.  

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
```



```

user@host# show protocols ospf
graceful-restart {
  helper-disable {
    standard;
  }
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}

```

### *Reenabling Helper Mode for OSPFv2*

#### **CLI Quick Configuration**

To quickly reenabling standard helper-mode for OSPFv2, copy the following command and paste it into the CLI.

```

[edit]
delete protocols ospf graceful-restart helper-disable standard

```



**NOTE:** To reenabling restart signaling-based helper mode, include the `restart-signaling` statement. To reenabling both standard and restart signaling-based helper mode, include the `both` statement.

#### **Step-by-Step Procedure**

To reenabling standard helper mode for OSPFv2:

1. Delete the standard helper-mode statement from the OSPFv2 configuration.

```

[edit]
user@host# delete protocols ospf graceful-restart helper-disable standard

```

2. If you are done configuring the device, commit the configuration.

```

[edit]
user@host# commit

```

#### **Results**

After you reenabling standard helper mode, the `show protocols ospf` command no longer displays the graceful restart configuration.

#### **Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPFv2 Graceful Restart Configuration on page 4143](#)
- [Verifying Graceful Restart Status on page 4144](#)

#### **Verifying the OSPFv2 Graceful Restart Configuration**

#### **Purpose**

Verify information about your OSPFv2 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, the Graceful restart helper mode field displays the status of the standard helper mode capability as enabled or disabled, and the Restart-signaling helper mode field displays the status of the restart

signaling-based helper mode as enabled or disabled. By default, both standard and restart signaling-based helper modes are enabled.

**Action** From operational mode, enter the **show ospf overview** command.

#### ***Verifying Graceful Restart Status***

**Purpose** Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

**Action** From operational mode, enter the **show route instance detail** command.

#### **Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart**

This example shows how to disable and reenabling the helper mode capability for OSPFv3 graceful restart.

- [Requirements on page 4144](#)
- [Overview on page 4144](#)
- [Configuration on page 4145](#)
- [Verification on page 4147](#)

#### ***Requirements***

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### ***Overview***

The OSPF graceful restart helper capability assists a neighboring routing device attempting a graceful restart. By default, the helper capability is globally enabled when you start the routing platform. This means that the helper capability is enabled when you start OSPF, even if graceful restart is not globally enabled or specifically enabled for OSPF. You can further modify your graceful restart configuration to disable the helper capability.

In the first example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPFv3 area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable the OSPFv3 graceful restart helper capability by including the **helper-disable** statement.



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

The second example shows how to reenable the OSPFv3 restart helper capability that you disabled in the first example.

### Configuration

- [Disabling Helper Mode for OSPFv3 on page 4145](#)
- [Reenabling Helper Mode for OSPFv3 on page 4146](#)

### Disabling Helper Mode for OSPFv3

#### CLI Quick Configuration

To quickly enable graceful restart for OSPFv3 with helper mode disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet6 address 2002:0a00:0004::
set interfaces fe-1/1/2 unit 0 family inet6 address 2002:0a00:0005::
set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
set protocols ospf3 graceful-restart helper-disable
```

#### Step-by-Step Procedure

To enable graceful restart for OSPFv3 with helper mode disabled:

1. Configure the interfaces.
 

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet6 address 2002:0a00:0004::
user@host# set interfaces fe-1/1/2 unit 0 family inet6 address 2002:0a00:0005::
```
2. Configure OSPFv3 on the interfaces
 

```
[edit]
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf3 area 0.0.0.0 interface fe-1/1/2
```
3. Disable the OSPFv3 graceful restart helper capability.
 

If you disable the OSPFv3 graceful restart helper capability, you cannot disable strict LSA checking.

```
[edit]
user@host# set protocols ospf3 graceful-restart helper-disable
```
4. If you are done configuring the device, commit the configuration.
 

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces** and the **show protocols ospf3** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet6 {
      address 2002:0a00:0004::/128;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet6 {
      address 2002:0a00:0005::/128;
    }
  }
}
user@host# show protocols ospf3
graceful-restart {
  helper-disable;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

#### *Reenabling Helper Mode for OSPFv3*

**CLI Quick Configuration** To quickly reenable helper-mode for OSPFv3, copy the following command and paste it into the CLI.

```
[edit]
delete protocols ospf3 graceful-restart helper-disable
```

**Step-by-Step Procedure** To reenable helper mode for OSPFv3:

1. Delete the standard helper-mode statement from the OSPFv3 configuration.

```
[edit]
user@host# delete protocols ospf3 graceful-restart helper-disable
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

**Results** After you reenable standard helper mode, the **show protocols ospfs** command no longer displays the graceful restart configuration.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPFv3 Graceful Restart Configuration on page 4147](#)
- [Verifying Graceful Restart Status on page 4147](#)

**Verifying the OSPFv3 Graceful Restart Configuration**

**Purpose** Verify information about your OSPFv3 graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled, and the Helper mode field displays the status of the helper mode capability as either enabled or disabled.

**Action** From operational mode, enter the **show ospf3 overview** command.

**Verifying Graceful Restart Status**

**Purpose** Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

**Action** From operational mode, enter the **show route instance detail** command.

**Example: Disabling Strict LSA Checking for OSPF Graceful Restart**

This example shows how to disable strict link-state advertisement (LSA) checking for OSPF graceful restart.

- [Requirements on page 4147](#)
- [Overview on page 4148](#)
- [Configuration on page 4148](#)
- [Verification on page 4149](#)

**Requirements**

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#)
- Configure a single-area OSPF network. See [“Example: Configuring a Single-Area OSPF Network” on page 4053](#).
- Configure a multiarea OSPF network. See [“Example: Configuring a Multiarea OSPF Network” on page 4055](#).

### Overview

You can disable strict LSA checking to prevent the termination of graceful restart by a helping router. You might configure this option for interoperability with other vendor devices. The OSPF graceful restart helper capability must be enabled if you disable strict LSA checking. By default, LSA checking is enabled.

In this example, interfaces **fe-1/1/1** and **fe-1/1/2** are in OSPF area 0.0.0.0, and you configure those interfaces for graceful restart. You then disable strict LSA checking by including the **no-strict-lsa-checking** statement.



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both statements at the same time, the routing device displays a warning message when you enter the **show protocols ospf** command.

### Configuration

#### CLI Quick Configuration

To quickly enable graceful restart for OSPF with strict LSA checking disabled, copy the following commands and paste them into the CLI.

```
[edit]
set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
set protocols ospf area 0.0.0.0 interface fe-1/1/1
set protocols ospf area 0.0.0.0 interface fe-1/1/2
set protocols ospf graceful-restart no-strict-lsa-checking
```

#### Step-by-Step Procedure

To enable graceful restart for OSPF with strict LSA checking disabled:

1. Configure the interfaces.



**NOTE:** For OSPFv3, use IPv6 addresses.

```
[edit]
user@host# set interfaces fe-1/1/1 unit 0 family inet address 10.0.0.4
user@host# set interfaces fe-1/1/2 unit 0 family inet address 10.0.0.5
```

2. Configure OSPF on the interfaces



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/1
user@host# set protocols ospf area 0.0.0.0 interface fe-1/1/2
```

3. Disable strict LSA checking.

If you disable the strict LSA checking, OSPF graceful restart helper capability must be enabled (which is the default behavior).

```
[edit]
user@host# set protocols ospf graceful-restart no-strict-lsa-checking
```

4. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-1/1/1 {
  unit 0 {
    family inet {
      address 10.0.0.4/32;
    }
  }
}
fe-1/1/2 {
  unit 0 {
    family inet {
      address 10.0.0.5/32;
    }
  }
}
user@host# show protocols ospf
graceful-restart {
  no-strict-lsa-checking;
}
area 0.0.0.0 {
  interface fe-1/1/1.0;
  interface fe-1/1/2.0;
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces** and the **show protocols ospf3** commands.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPF Graceful Restart Configuration on page 4149](#)
- [Verifying Graceful Restart Status on page 4150](#)

### **Verifying the OSPF Graceful Restart Configuration**

**Purpose** Verify information about your OSPF graceful restart configuration. The Restart field displays the status of graceful restart as either enabled or disabled.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** command for OSPFv3.

#### ***Verifying Graceful Restart Status***

**Purpose** Verify the status of graceful restart. The Restart State field displays Pending if the restart has not completed, or Complete if the restart has finished. The Path selection timeout field indicates the amount of time remaining until graceful restart is declared complete. There is a more detailed Restart State field that displays a list of protocols that have completed graceful restart or have not yet completed graceful restart for the specified routing table.

**Action** From operational mode, enter the **show route instance detail** command.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)
- [Graceful Restart Concepts on page 2257](#) in the *Junos OS High Availability Library for Routing Devices*

---

## OSPF Traffic Engineering Configuration

- [Examples: Configuring OSPF Traffic Engineering on page 4150](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode on page 4159](#)

### Examples: Configuring OSPF Traffic Engineering

- [OSPF Support for Traffic Engineering on page 4150](#)
- [Example: Enabling OSPF Traffic Engineering Support on page 4153](#)
- [Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface on page 4157](#)

---

#### OSPF Support for Traffic Engineering

Traffic engineering allows you to control the path that data packets follow, bypassing the standard routing model, which uses routing tables. Traffic engineering moves flows from congested links to alternate links that would not be selected by the automatically computed destination-based shortest path.

To help provide traffic engineering and MPLS with information about network topology and loading, extensions have been added to the Junos OS implementation of OSPF. When traffic engineering is enabled on the routing device, you can enable OSPF traffic engineering support. When you enable traffic engineering for OSPF, the shortest-path-first (SPF) algorithm takes into account the various label-switched paths (LSPs) configured under MPLS and configures OSPF to generate opaque link-state advertisements (LSAs) that carry traffic engineering parameters. The parameters are used to populate the traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. The Constrained



Shortest Path First (CSPF) algorithm uses the traffic engineering database to compute the paths that MPLS LSPs take. RSVP uses this path information to set up LSPs and to reserve bandwidth for them.

By default, traffic engineering support is disabled. To enable traffic engineering, include the **traffic-engineering** statement. You can also configure the following OSPF traffic engineering extensions:

- **advertise-unnnumbered-interfaces**—(OSPFv2 only) Advertises the link-local identifier in the link-local traffic engineering LSA packet. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477, *Signalling Unnumbered Links in Resource Reservation Protocol - Traffic Engineering (RSVP-TE)*.
- **credibility-protocol-preference**—(OSPFv2 only) Assigns a credibility value to OSPF routes in the traffic engineering database. By default, Junos OS prefers IS-IS routes in the traffic engineering database over other interior gateway protocol (IGP) routes even if the routes of another IGP are configured with a lower, that is, more preferred, preference value. The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. In Junos OS Release 9.4 and later, you can configure OSPF to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration.
- **ignore-lsp-metrics**—Ignores RSVP LSP metrics in OSPF traffic engineering shortcut calculations or when you configure LDP over RSVP LSPs. This option avoids mutual dependency between OSPF and RSVP, eliminating the time period when the RSVP metric used for tunneling traffic is not up to date. In addition, If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.
- **multicast-rpf-routes**—(OSPFv2 only) Installs unicast IPv4 routes (not LSPs) in the multicast routing table (**inet.2**) for multicast reverse-path forwarding (RPF) checks. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check if the packet is coming in on an interface that is also sending data back to the packet source.
- **no-topology**—(OSPFv2 only) To disable the dissemination of link-state topology information. If disabled, traffic engineering topology information is no longer distributed within the OSPF area.
- **shortcuts**—Configures OSPF to use MPLS LSPs as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the **inet.3** routing table, and shortcut routes calculated through OSPFv3 are installed in the **inet6.3** routing table.



**NOTE:** Whenever possible, use OSPF IGP shortcuts configured at the `[edit protocols mpls traffic-engineering bgp-igp]` hierarchy level instead of traffic engineering shortcuts configured at the `[edit protocols (ospf | ospf3) traffic-engineering shortcuts]` hierarchy level.

If you configure OSPF IGP shortcuts, `inet.3` routes are moved into the `inet.0` routing table. In addition, you can verify the data path using `ping` or `traceroute` commands since the ping and traceroute packets get tunneled into the LSP. In case of a VPN enabled device, we recommend using `[edit protocols mpls traffic-engineering bgp-igp-both-ribs]` because BGP next-hop resolution for VPN prefixes relies on entries in the `inet.3` table.

If you configure traffic engineering shortcuts, OSPF treats the MPLS LSP as a candidate next hop and installs the routes in the `inet.3` (for OSPFv2) and `inet6.3` (for OSPFv3) routing tables. The only use for these tables is to allow BGP to perform next-hop resolution. In addition, you cannot verify the data path of these routes using `ping` or `traceroute` commands because the ping and traceroute packets get tunneled into the LSP.

- **`lsp-metric-info-summary`**—Advertises the LSP metric in summary LSAs to treat the LSP as a link. This configuration allows other routing devices in the network to use this LSP. To accomplish this, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

When you enable traffic engineering on the routing device, you can also configure an OSPF metric that is used exclusively for traffic engineering. The traffic engineering metric is used for information injected into the traffic engineering database. Its value does not affect normal OSPF forwarding.



**CAUTION:** When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:

- Traffic engineering shortcuts
- IGP shortcuts
- LDP tunneling
- Multiprotocol LSP
- Advertise summary metrics
- Multicast RPF routes

---

### Example: Enabling OSPF Traffic Engineering Support

---

This example shows how to enable OSPF traffic engineering support to advertise the label-switched path (LSP) metric in summary link-state advertisements (LSAs).

- [Requirements on page 4153](#)
- [Overview on page 4153](#)
- [Configuration on page 4154](#)
- [Verification on page 4157](#)

#### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure BGP per your network requirements. See the *Junos OS Routing Protocols Library for Routing Devices*
- Configure MPLS per your network requirements. See the *Junos OS MPLS Applications Library for Routing Devices*.

#### Overview

You can configure OSPF to treat an LSP as a link and have other routing devices in the network use this LSP. To accomplish this, you configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

In this example, there are four routing devices in area 0.0.0.0, and you want OSPF to treat the LSP named R1-to-R4 that goes from the ingress Device R1 to the egress Device R4 as a link.

For OSPF, you enable traffic engineering on all four routing devices in the area by including the **traffic-engineering** statement. This configuration ensures that the shortest-path-first (SPF) algorithm takes into account the LSPs configured under MPLS and configures OSPF to generate LSAs that carry traffic engineering parameters. You further ensure that OSPF uses the MPLS LSP as the next hop and advertises the LSP metric in summary LSAs, by including the optional **shortcuts lsp-metric-into-summary** statement on the ingress Device R1.

For MPLS, you enable traffic engineering so that MPLS performs traffic engineering on both BGP and IGP destinations by including the **traffic-engineering bgp-igp** statement, and you include the LSP named R1-to-R4 by including the **label-switched-path lsp-path-name to address** statement on the ingress Device R1. The address specified in the **to** statement on the ingress Device R1 must match the router ID of the egress Device R4 for the LSP to function as a direct link to the egress routing device and to be used as input to the OSPF SPF calculations. In this example, the router ID of the egress Device R4 is 10.0.0.4.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

**CLI Quick Configuration** To quickly enable OSPF traffic engineering support to advertise the LSP metric in summary LSAs, copy the following commands and paste them into the CLI.

Configuration on R1:

```
[edit]
set routing-options router-id 10.0.0.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

Configuration on R2:

```
[edit]
set routing-options router-id 10.0.0.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R3:

```
[edit]
set routing-options router-id 10.0.0.3
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

Configuration on R4:

```
[edit]
set routing-options router-id 10.0.0.4
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
```

**Step-by-Step Procedure** To enable OSPF traffic engineering support to advertise LSP metrics in summary LSAs:

1. Configure the router ID.

```
[edit]
user@R1# set routing-options router-id 10.0.0.1
```

```
[edit]
user@R2# set routing-options router-id 10.0.0.2
```

```
[edit]
user@R3# set routing-options router-id 10.0.0.3
```

```
[edit]
user@R4# set routing-options router-id 10.0.0.4
```

2. Configure the OSPF area and add the interfaces.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface all
user@R1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface all
user@R2# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R3# set protocols ospf area 0.0.0.0 interface all
user@R3# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface all
user@R4# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

3. Enable OSPF traffic engineering.

```
[edit]
user@R1 set protocols ospf traffic-engineering shortcuts lsp-metric-into-summary
```

```
[edit]
user@R2 set protocols ospf traffic-engineering
```

```
[edit]
user@R3 set protocols ospf traffic-engineering
```

```
[edit]
user@R4 set protocols ospf traffic-engineering
```

4. On Device R1, configure MPLS traffic engineering.

```
[edit ]
user@R1 set protocol mpls traffic-engineering bgp-igp
user@R1 set protocols mpls label-switched-path R1-to-R4 to 10.0.0.4
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the `show routing-options`, `show protocols ospf`, and `show protocols mpls` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@host# show routing-options
router-id 10.0.0.1;

user@host# show protocols ospf
traffic-engineering {
  shortcuts lsp-metric-into-summary;
```

```
}
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}

user@host# show protocols mpls
traffic-engineering bgp-igp;
label-switched-path R1-to-R4 {
  to 10.0.0.4;
}
```

Output for R2:

```
user@host# show routing-options
router-id 10.0.0.2;

user@host# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

Output for R3:

```
user@host# show routing-options
router-id 10.0.0.3;

user@host# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

Output for R4:

```
user@host# show routing-options
router-id 10.0.0.4;

user@host# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show routing-options**, **show protocols ospf3**, and **show protocols mpls** commands.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the Traffic Engineering Capability for OSPF on page 4157](#)
- [Verifying OSPF Entries in the Traffic Engineering Database on page 4157](#)
- [Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF on page 4157](#)

**Verifying the Traffic Engineering Capability for OSPF**

**Purpose** Verify that traffic engineering has been enabled for OSPF. By default, traffic engineering is disabled.

**Action** From operational mode, enter the **show ospf overview** command for OSPFv2, and enter the **show ospf3 overview** for OSPFv3.

**Verifying OSPF Entries in the Traffic Engineering Database**

**Purpose** Verify the OSPF information in the traffic engineering database. The Protocol field displays OSPF and the area from which the information was learned.

**Action** From operational mode, enter the **show ted database** command.

**Verifying That the Traffic Engineering Database Is Learning Node Information from OSPF**

**Purpose** Verify that OSPF is reporting node information. The Protocol name field displays OSPF and the area from which the information was learned.

**Action** From operational mode, enter the **show ted protocol** command.

**Example: Configuring the Traffic Engineering Metric for a Specific OSPF Interface**

This example shows how to configure the OSPF metric value used for traffic engineering.

- [Requirements on page 4157](#)
- [Overview on page 4158](#)
- [Configuration on page 4158](#)
- [Verification on page 4159](#)

**Requirements**

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure OSPF for traffic engineering. See “[Example: Enabling OSPF Traffic Engineering Support](#)” on page 4153

### Overview

You can configure an OSPF metric that is used exclusively for traffic engineering. To modify the default value of the traffic engineering metric, include the **te-metric** statement. The OSPF traffic engineering metric does not affect normal OSPF forwarding. By default, the traffic engineering metric is the same value as the OSPF metric. The range is 1 through 65,535.

In this example, you configure the OSPF traffic engineering metric on OSPF interface **fe-0/1/1** in area 0.0.0.0.

### Configuration

**CLI Quick Configuration** To quickly configure the OSPF traffic engineering metric for a specific interface, copy the following command and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.0 interface fe-0/1/1 te-metric 10
```

**Step-by-Step Procedure** To configure an OSPF traffic engineering metric for a specific interface used only for traffic engineering:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
```

2. Configure the traffic engineering metric of the OSPF network segments.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/1/1 te-metric 10
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf area 0.0.0.0]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.0 {
  interface fe-0/1/1.0 {
    te-metric 10;
  }
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.



**Verification**

Confirm that the configuration is working properly.

**Verifying the Configured Traffic Engineering Metric**

**Purpose** Verify the traffic engineering metric value. Confirm that Metric field displays the configured traffic engineering metric.

**Action** From operational mode, enter the **show ted database extensive** command.

**Related Documentation**

- *OSPF Configuration Overview*
- *Junos OS MPLS Applications Library for Routing Devices*

**Example: Configuring OSPF Passive Traffic Engineering Mode**

- [OSPF Passive Traffic Engineering Mode on page 4159](#)
- [Example: Configuring OSPF Passive Traffic Engineering Mode on page 4159](#)

**OSPF Passive Traffic Engineering Mode**

Ordinarily, interior routing protocols such as OSPF are not run on links between autonomous systems. However, for inter-AS traffic engineering to function properly, information about the inter-AS link—in particular, the address on the remote interface—must be made available inside the autonomous system (AS). This information is not normally included either in the external BGP (EBGP) reachability messages or in the OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include it in the traffic engineering database. OSPF traffic engineering mode allows MPLS label-switched paths (LSPs) to dynamically discover OSPF AS boundary routers and to allow routers to establish a traffic engineering LSP across multiple autonomous systems.

**Example: Configuring OSPF Passive Traffic Engineering Mode**

This example shows how to configure OSPF passive mode for traffic engineering on an inter-AS interface. The AS boundary router link between the EBGP peers must be a directly connected link and must be configured as a passive traffic engineering link.

- [Requirements on page 4159](#)
- [Overview on page 4160](#)
- [Configuration on page 4160](#)
- [Verification on page 4161](#)

**Requirements**

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure BGP per your network requirements. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure the LSP per your network requirements. See the *Junos OS MPLS Applications Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

### Overview

You can configure OSPF passive mode for traffic engineering on an inter-AS interface. The address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGP link. In this example, you configure interface **so-1/1/0** in area 0.0.0.1 as the inter-AS link to distribute traffic engineering information with OSPF within the AS and include the following settings:

- **passive**—Advertises the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.
- **traffic-engineering**—Configures an interface in OSPF passive traffic-engineering mode to enable dynamic discovery of OSPF AS boundary routers. By default, OSPF passive traffic-engineering mode is disabled.
- **remote-node-id**—Specifies the IP address at the far end of the inter-AS link. In this example, the remote IP address is 192.168.207.2.

### Configuration

To quickly configure OSPF passive mode for traffic engineering, copy the following command, remove any line breaks, and paste it into the CLI.

```
[edit]
set protocols ospf area 0.0.0.1 interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

### Step-by-Step Procedure

To configure OSPF passive traffic engineering mode:

1. Create an OSPF area.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# set protocols ospf area 0.0.0.1
```

2. Configure interface `so-1/1/0` as a passive interface configured for traffic engineering, and specify the IP address at the far end of the inter-AS link.

```
[edit protocols ospf area 0.0.0.1]
user@host# set interface so-1/1/0 passive traffic-engineering remote-node-id
192.168.207.2
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf]
user@host# commit
```

**Results** Confirm your configuration by entering the `show protocols ospf` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
area 0.0.0.1 {
  interface so-1/1/0.0 {
    passive {
      traffic-engineering {
        remote-node-id 192.168.207.2;
      }
    }
  }
}
```

To confirm your OSPFv3 configuration, enter the `show protocols ospf3` command.

### Verification

Confirm that the configuration is working properly.

#### Verifying the Status of OSPF Interfaces

**Purpose** Verify the status of OSPF interfaces. If the interface is passive, the Adj count field is 0 because no adjacencies have been formed. Next to this field, you might also see the word Passive.

**Action** From operational mode, enter the `show ospf interface detail` command for OSPFv2, and enter the `show ospf3 interface detail` command for OSPFv3.

**Related Documentation**

- *OSPF Configuration Overview*
- [About OSPF Interfaces on page 4076](#)
- *Junos OS MPLS Applications Library for Routing Devices*

## OSPF Database Protection Configuration

---

- [Example: Configuring OSPF Database Protection on page 4162](#)

### Example: Configuring OSPF Database Protection

- [OSPF Database Protection Overview on page 4162](#)
- [Configuring OSPF Database Protection on page 4163](#)

#### OSPF Database Protection Overview

---

OSPF database protection allows you to limit the number of link-state advertisements (LSAs) not generated by the local router in a given OSPF routing instance, helping to protect the link-state database from being flooded with excessive LSAs. This feature is particularly useful if VPN routing and forwarding is configured on your provider edge and customer edge routers using OSPF as the routing protocol. An overrun link-state database on the customer edge router can exhaust resources on the provider edge router and impact the rest of the service provider network.

When you enable OSPF database protection, the maximum number of LSAs you specify includes all LSAs whose advertising router ID is not equal to the local router ID (nonsystem-generated LSAs). These might include external LSAs as well as LSAs with any scope such as the link, area, and autonomous system (AS).

Once the specified maximum LSA count is exceeded, the database typically enters into the ignore state. In this state, all neighbors are brought down, and nonsystem-generated LSAs are destroyed. In addition, the database sends out hellos but ignores all received packets. As a result, the database does not form any full neighbors, and therefore does not learn about new LSAs. However, if you have configured the **warning-only** option, only a warning is issued and the database does not enter the ignore state but continues to operate as before.

You can also configure one or more of the following options:

- A warning threshold for issuing a warning message before the LSA limit is reached.
- An ignore state time during which the database must remain in the ignore state and after which normal operations can be resumed.
- An ignore state count that limits the number of times the database can enter the ignore state, after which it must enter the isolate state. The isolate state is very similar to the ignore state, but has one important difference: once the database enters the isolate state, it must remain there until you issue a command to clear database protection before it can return to normal operations.
- A reset time during which the database must stay out of the ignore or isolate state before it is returned to a normal operating state.

## Configuring OSPF Database Protection

By configuring OSPF database protection, you can help prevent your OSPF link-state database from being overrun with excessive LSAs that are not generated by the local router. You specify the maximum number of LSAs whose advertising router ID is not the same as the local router ID in an OSPF instance. This feature is particularly useful if your provider edge and customer edge routers are configured with VPN routing and forwarding using OSPF.

OSPF database protection is supported on:

- Logical systems
- All routing instances supported by OSPFv2 and OSPFv3
- OSPFv2 and OSPFv3 topologies
- OSPFv3 realms

To configure OSPF database protection:

1. Include the **database-protection** statement at one of the following hierarchy levels:
  - [edit protocols ospf | ospf3]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3)]
  - [edit routing-instances *routing-instance-name* protocols (ospf | ospf3)]
  - [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-unicast | ipv6-multicast)]
2. Include the **maximum-lsa *number*** statement.



**NOTE:** The **maximum-lsa** statement is mandatory, and there is no default value for it. If you omit this statement, you cannot configure OSPF database protection.

3. (Optional) Include the following statements:
  - **ignore-count *number***—Specify the number of times the database can enter the ignore state before it goes into the isolate state.
  - **ignore-time *seconds***—Specify the time limit the database must remain in the ignore state before it resumes regular operations.
  - **reset-time *seconds***—Specify the time during which the database must operate without being in either the ignore or isolate state before it is reset to a normal operating state.
  - **warning-threshold *percent***—Specify the percent of the maximum LSA number that must be exceeded before a warning message is issued.

4. (Optional) Include the **warning-only** statement to prevent the database from entering the ignore state or isolate state when the maximum LSA count is exceeded.



**NOTE:** If you include the **warning-only** statement, values for the other optional statements at the same hierarchy level are not used when the maximum LSA number is exceeded.

5. Verify your configuration by checking the database protection fields in the output of the **show ospf overview** command.

**Related  
Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)

---

## OSPF Policy Configuration

---

- [Examples: Configuring OSPF Routing Policy on page 4164](#)
- [Examples: Configuring Routing Policy for Network Summaries on page 4180](#)

### Examples: Configuring OSPF Routing Policy

- [Understanding OSPF Routing Policy on page 4164](#)
- [Example: Injecting OSPF Routes into the BGP Routing Table on page 4166](#)
- [Example: Redistributing Static Routes into OSPF on page 4169](#)
- [Example: Configuring an OSPF Import Policy on page 4172](#)
- [Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF on page 4176](#)

---

#### Understanding OSPF Routing Policy

---

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks. Each routing policy name must be unique within a configuration. Once a policy is created and named, it must be applied before it is active.

In the **import** statement, you list the name of the routing policy used to filter OSPF external routes from being installed into the routing tables of OSPF neighbors. You can filter the routes, but not link-state address (LSA) flooding. An external route is a route that is outside the OSPF Autonomous System (AS). The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements.

In the **export** statement, you list the name of the routing policy to be evaluated when routes are being exported from the routing table into OSPF.

By default, if a routing device has multiple OSPF areas, learned routes from other areas are automatically installed into area 0 of the routing table.

To specify more than one policy and create a policy chain, you list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an accept or reject action is executed, the policy chain evaluation ends.

This topic describes the following information:

- [Routing Policy Terms on page 4165](#)
- [Routing Policy Match Conditions on page 4165](#)
- [Routing Policy Actions on page 4166](#)

### ***Routing Policy Terms***

Routing policies are made up of one or more terms. A term is a named structure in which match conditions and actions are defined. You can define one or more terms. The name can contain letters, numbers, and hyphens ( - ) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

### ***Routing Policy Match Conditions***

A match condition defines the criteria that a route must match for an action to take place. You can define one or more match conditions for each term. If a route matches all of the match conditions for a particular term, the actions defined for that term are processed.

Each term can include two statements, **from** and **to**, that define the match conditions:

- In the **from** statement, you define the criteria that an incoming route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The **from** statement is optional. If you omit the **from** and the **to** statements, all routes are considered to match.



**NOTE:** In export policies, omitting the **from** statement from a routing policy term might lead to unexpected results. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

- In the **to** statement, you define the criteria that an outgoing route must match. You can specify one or more match conditions. If you specify more than one, they all must match the route for a match to occur.

The order of the match conditions in a term is not important because a route must match all match conditions in a term for an action to be taken.

For a complete list of match conditions, see *Routing Policy Match Conditions* in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### **Routing Policy Actions**

An action defines what the routing device does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term. If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy.
- Actions that manipulate route characteristics.
- Trace action, which logs route matches.

The **then** statement is optional. If you omit it, one of the following occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the **accept** or **reject** action specified by the default policy is executed.

For a complete list of routing policy actions, see *Actions in Routing Policy Terms* in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### **Example: Injecting OSPF Routes into the BGP Routing Table**

---

This example shows how to create a policy that injects OSPF routes into the BGP routing table.

- [Requirements on page 4166](#)
- [Overview on page 4167](#)
- [Configuration on page 4167](#)
- [Verification on page 4169](#)
- [Troubleshooting on page 4169](#)

#### **Requirements**

Before you begin:

- Configure network interfaces.
- Configure external peer sessions. See [“Example: Configuring External BGP Point-to-Point Peer Sessions” on page 3262](#).



- Configure interior gateway protocol (IGP) sessions between peers.

### Overview

In this example, you create a routing policy called **injectpolicy1** and a routing term called **injectterm1**. The policy injects OSPF routes into the BGP routing table.

### Configuration

- [Configuring the Routing Policy on page 4167](#)
- [Configuring Tracing for the Routing Policy on page 4168](#)

### Configuring the Routing Policy

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 from protocol ospf
set policy-options policy-statement injectpolicy1 term injectterm1 from area 0.0.0.1
set policy-options policy-statement injectpolicy1 term injectterm1 then accept
set protocols bgp export injectpolicy1
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To inject OSPF routes into a BGP routing table:

1. Create the policy term.  

```
[edit policy-options policy-statement injectpolicy1]
user@host# set term injectterm1
```
2. Specify OSPF as a match condition.  

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from protocol ospf
```
3. Specify the routes from an OSPF area as a match condition.  

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set from area 0.0.0.1
```
4. Specify that the route is to be accepted if the previous conditions are matched.  

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# set then accept
```
5. Apply the routing policy to BGP.  

```
[edit]
user@host# set protocols bgp export injectpolicy1
```

**Results** Confirm your configuration by entering the **show policy-options** and **show protocols bgp** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement injectpolicy1 {
  term injectterm1 {
    from {
      protocol ospf;
      area 0.0.0.1;
    }
    then accept;
  }
}

user@host# show protocols bgp
export injectpolicy1;
```

If you are done configuring the device, enter **commit** from configuration mode.

#### *Configuring Tracing for the Routing Policy*

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement injectpolicy1 term injectterm1 then trace
set routing-options traceoptions file ospf-bgp-policy-log
set routing-options traceoptions file size 5m
set routing-options traceoptions file files 5
set routing-options traceoptions flag policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Include a trace action in the policy.

```
[edit policy-options policy-statement injectpolicy1 term injectterm1]
user@host# then trace
```

2. Configure the tracing file for the output.

```
[edit routing-options traceoptions]
user@host# set file ospf-bgp-policy-log
user@host# set file size 5m
user@host# set file files 5
user@host# set flag policy
```

**Results** Confirm your configuration by entering the **show policy-options** and **show routing-options** commands from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
```

```

policy-statement injectpolicy1 {
  term injectterm1 {
    then {
      trace;
    }
  }
}

user@host# show routing-options
traceoptions {
  file ospf-bgp-policy-log size 5m files 5;
  flag policy;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

### Verifying That the Expected BGP Routes Are Present

**Purpose** Verify the effect of the export policy.

**Action** From operational mode, enter the **show route** command.

### Troubleshooting

- [Using the show log Command to Examine the Actions of the Routing Policy on page 4169](#)

### Using the show log Command to Examine the Actions of the Routing Policy

**Problem** The routing table contains unexpected routes, or routes are missing from the routing table.

**Solution** If you configure policy tracing as shown in this example, you can run the **show log ospf-bgp-policy-log** command to diagnose problems with the routing policy. The **show log ospf-bgp-policy-log** command displays information about the routes that the **injectpolicy1** policy term analyzes and acts upon.

### Example: Redistributing Static Routes into OSPF

This example shows how to create a policy that redistributes static routes into OSPF.

- [Requirements on page 4169](#)
- [Overview on page 4170](#)
- [Configuration on page 4170](#)
- [Verification on page 4171](#)

### Requirements

Before you begin:

- Configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure static routes. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Library for Routing Devices*.

### Overview

In this example, you create a routing policy called `exportstatic1` and a routing term called `exportstatic1`. The policy injects static routes into OSPF. This example includes the following settings:

- **policy-statement**—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens ( - ) and be up to 255 characters long.
- **term**—Defines the match condition and applicable actions for the routing policy. The term name can contain letters, numbers, and hyphens ( - ) and be up to 255 characters long. You specify the name of the term and define the criteria that an incoming route must match by including the **from** statement and the action to take if the route matches the conditions by including the **then** statement. In this example you specify the static protocol match condition and the accept action.
- **export**—Applies the export policy you created to be evaluated when routes are being exported from the routing table into OSPF.

### Configuration

#### CLI Quick Configuration

To quickly create a policy that injects static routes into OSPF, copy the following commands and paste them into the CLI.

```
[edit]
set policy-options policy-statement exportstatic1 term exportstatic1 from protocol static
set policy-options policy-statement exportstatic1 term exportstatic1 then accept
set protocols ospf export exportstatic1
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To inject static routes into OSPF:

1. Create the routing policy.  

```
[edit]
user@host# edit policy-options policy-statement exportstatic1
```
2. Create the policy term.  

```
[edit policy-options policy-statement exportstatic1]
user@host# set term exportstatic1
```
3. Specify static as a match condition.  

```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set from protocol static
```

4. Specify that the route is to be accepted if the previous condition is matched.

```
[edit policy-options policy-statement exportstatic1 term exportstatic1]
user@host# set then accept
```

5. Apply the routing policy to OSPF.



**NOTE:** For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@host# set protocols ospf export exportstatic1
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the `show policy-options` and `show protocols ospf` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement exportstatic1 {
  term exportstatic1 {
    from protocol static;
    then accept;
  }
}
```

```
user@host# show protocols ospf
export exportstatic1;
```

To confirm your OSPFv3 configuration, enter the `show policy-options` and the `show protocols ospf3` commands.

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Static Routes Are Present on page 4171](#)
- [Verifying That AS External LSAs Are Added to the Routing Table on page 4172](#)

### Verifying That the Expected Static Routes Are Present

**Purpose** Verify the effect of the export policy.

**Action** From operational mode, enter the `show route` command.

### ***Verifying That AS External LSAs Are Added to the Routing Table***

- Purpose** On the routing device where you configured the export policy, verify that the routing device originates an AS external LSA for the static routes that are added to the routing table.
- Action** From operational mode, enter the **show ospf database** command for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

---

### **Example: Configuring an OSPF Import Policy**

This example shows how to create an OSPF import policy. OSPF import policies apply to external routes only. An external route is a route that is outside the OSPF autonomous system (AS).

- [Requirements on page 4172](#)
- [Overview on page 4172](#)
- [Configuration on page 4173](#)
- [Verification on page 4175](#)

#### ***Requirements***

Before you begin:

- Configure static routes. See *Examples: Configuring Static Routes* in the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.

#### ***Overview***

External routes are learned by AS boundary routers. External routes can be advertised throughout the OSPF domain if you configure the AS boundary router to redistribute the route into OSPF. An external route might be learned by the AS boundary router from a routing protocol other than OSPF, or the external route might be a static route that you configure on the AS boundary router.

For OSPFv3, the link-state advertisement (LSA) is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An area border router (ABR) originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

OSPF import policy allows you to prevent external routes from being added to the routing tables of OSPF neighbors. The import policy does not impact the OSPF database. This means that the import policy has no impact on the link-state advertisements. The filtering

is done only on external routes in OSPF. The intra-area and interarea routes are not considered for filtering. The default action is to accept the route when the route does not match the policy.

This example includes the following OSPF policy settings:

- **policy-statement**—Defines the routing policy. You specify the name of the policy and further define the elements of the policy. The policy name must be unique and can contain letters, numbers, and hyphens ( - ) and be up to 255 characters long.
- **export**—Applies the export policy you created to be evaluated when network summary LSAs are flooded into an area. In this example, the export policy is named `export_static`.
- **import**—Applies the import policy you created to prevent external routes from being added to the routing table. In this example, the import policy is named `filter_routes`.

The devices you configure in this example represent the following functions:

- **R1**—Device R1 is in area 0.0.0.0 and has a direct connection to device R2. R1 has an OSPF export policy configured. The export policy redistributes static routes from R1's routing table into R1's OSPF database. Because the static route is in R1's OSPF database, the route is advertised in an LSA to R1's OSPF neighbor. R1's OSPF neighbor is device R2.
- **R2**—Device R2 is in area 0.0.0.0 and has a direct connection to device R1. R2 has an OSPF import policy configured that matches the static route to the 10.0.16.0/30 network and prevents the static route from being installed in R2's routing table. R2's OSPF neighbor is device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure an OSPF import policy, copy the following commands, removing any line breaks, and then paste the commands into the CLI.

Configuration on Device R1:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
set protocols ospf export export_static
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement export_static from protocol static
set policy-options policy-statement export_static then accept
```

Configuration on Device R2:

```
[edit]
set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
set protocols ospf import filter_routes
set protocols ospf area 0.0.0.0 interface so-0/2/0
set policy-options policy-statement filter_routes from route-filter 10.0.16.0/30 exact
set policy-options policy-statement filter_routes then reject
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure an OSPF import policy:

1. Configure the interfaces.

```
[edit]
user@R1# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.1/30
```

```
[edit]
user@R2# set interfaces so-0/2/0 unit 0 family inet address 10.0.2.2/30
```

2. Enable OSPF on the interfaces.



**NOTE:** For OSPFv3, include the `ospf3` statement at the [edit protocols] hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

```
[edit]
user@R2# set protocols ospf area 0.0.0.0 interface so-0/2/0
```

3. On R1, redistribute the static route into OSPF.

```
[edit]
user@R1# set protocols ospf export export_static
user@R1# set policy-options policy-statement export_static from protocol static
user@R1# set policy-options policy-statement export_static then accept
```

4. On R2, configure the OSPF import policy.

```
[edit]
user@R2# set protocols ospf import filter_routes
user@R2# set policy-options policy-statement filter_routes from route-filter
10.0.16.0/30 exact
user@R2# set policy-options policy-statement filter_routes then reject
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
so-0/2/0 {
  unit 0 {
    family inet {
```



```

        address 10.0.2.1/30;
    }
}

user@R1# show policy-options
policy-statement export_static {
    from protocol static;
    then accept;
}

user@R1# show protocols ospf
export export_static;
area 0.0.0.0 {
    interface so-0/2/0.0;
}

```

Output for R2:

```

user@R2# show interfaces
so-0/2/0 {
    unit 0 {
        family inet {
            address 10.0.2.2/30;
        }
    }
}

user@R2# show policy-options
policy-statement filter_routes {
    from {
        route-filter 10.0.16.0/30 exact;
    }
    then reject;
}

user@R2# show protocols ospf
import filter_routes;
area 0.0.0.0 {
    interface so-0/2/0.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, **show routing-options**, and **show protocols ospf3** commands on the appropriate device.

### Verification

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 4175](#)
- [Verifying the Routing Table on page 4176](#)

### Verifying the OSPF Database

**Purpose** Verify that OSPF is advertising the static route in the OSPF database.

**Action** From operational mode, enter the **show ospf database** for OSPFv2, and enter the **show ospf3 database** command for OSPFv3.

#### *Verifying the Routing Table*

**Purpose** Verify the entries in the routing table.

**Action** From operational mode, enter the **show route** command.

### **Example: Configuring a Route Filter Policy to Specify Priority for Prefixes Learned Through OSPF**

---

This example shows how to create an OSPF import policy that prioritizes specific prefixes learned through OSPF.

- [Requirements on page 4176](#)
- [Overview on page 4176](#)
- [Configuration on page 4177](#)
- [Verification on page 4180](#)

#### **Requirements**

Before you begin:

- Configure the device interfaces.
- Configure the router identifiers for the devices in your OSPF network. See “[Example: Configuring an OSPF Router Identifier](#)” on page 4048.
- Control OSPF designated router election. See “[Example: Controlling OSPF Designated Router Election](#)” on page 4050.
- Configure a single-area OSPF network. See “[Example: Configuring a Single-Area OSPF Network](#)” on page 4053.
- Configure a multiarea OSPF network. See “[Example: Configuring a Multiarea OSPF Network](#)” on page 4055.

#### **Overview**

In a network with a large number of OSPF routes, it can be useful to control the order in which routes are updated in response to a network topology change. In Junos OS Release 9.3 and later, you can specify a priority of high, medium, or low for prefixes included in an OSPF import policy. In the event of an OSPF topology change, high priority prefixes are updated in the routing table first, followed by medium and then low priority prefixes.

OSPF import policy can only be used to set priority or to filter OSPF external routes. If an OSPF import policy is applied that results in a **reject** terminating action for a nonexternal route, then the **reject** action is ignored and the route is accepted anyway. By default, such a route is now installed in the routing table with a priority of low. This behavior prevents traffic black holes, that is, silently discarded traffic, by ensuring consistent routing within the OSPF domain.

In general, OSPF routes that are not explicitly assigned a priority are treated as priority medium, except for the following:

- Summary discard routes have a default priority of low.
- Local routes that are not added to the routing table are assigned a priority of low.
- External routes that are rejected by import policy and thus not added to the routing table are assigned a priority of low.

Any available match criteria applicable to OSPF routes can be used to determine the priority. Two of the most commonly used match criteria for OSPF are the **route-filter** and **tag** statements.

In this example, the routing device is in area 0.0.0.0, with interfaces fe-0/1/0 and fe-1/1/0 connecting to neighboring devices. You configure an import routing policy named **ospf-import** to specify a priority for prefixes learned through OSPF. Routes associated with these prefixes are installed in the routing table in the order of the prefixes' specified priority. Routes matching **200.3.0.0/16 orlonger** are installed first because they have a priority of **high**. Routes matching **200.2.0.0/16 orlonger** are installed next because they have a priority of **medium**. Routes matching **200.1.0.0/16 orlonger** are installed last because they have a priority of **low**. You then apply the import policy to OSPF.



**NOTE:** The priority value takes effect when a new route is installed, or when there is a change to an existing route.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set interfaces fe-0/1/0 unit 0 family inet address 192.168.8.4/30
set interfaces fe-0/2/0 unit 0 family inet address 192.168.8.5/30
set policy-options policy-statement ospf-import term t1 from route-filter 200.1.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t1 then priority low
set policy-options policy-statement ospf-import term t1 then accept
set policy-options policy-statement ospf-import term t2 from route-filter 200.2.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t2 then priority medium
set policy-options policy-statement ospf-import term t2 then accept
set policy-options policy-statement ospf-import term t3 from route-filter 200.3.0.0/16
  orlonger
set policy-options policy-statement ospf-import term t3 then priority high
set policy-options policy-statement ospf-import term t3 then accept
set protocols ospf import ospf-import
set protocols ospf area 0.0.0.0 interface fe-0/1/0.0
set protocols ospf area 0.0.0.0 interface fe-0/2/0.0
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an OSPF import policy that prioritizes specific prefixes:

1. Configure the device interfaces.

```
[edit interfaces]
user@host# set fe-0/1/0 unit 0 family inet address 192.168.8.4/30

user@host# set fe-0/2/0 unit 0 family inet address 192.168.8.5/30
```

2. Enable OSPF on the interfaces.



**NOTE:** For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit protocols ospf area 0.0.0.0]
user@host# set interface fe-0/1/0.0
user@host# set interface fe-0/2/0.0
```

3. Configure the policy to specify the priority for prefixes learned through OSPF.

```
[edit policy-options policy-statement ospf-import]
user@host# set term t1 from route-filter 200.1.0.0/16 orlonger
user@host# set term t1 then priority low
user@host# set term t1 then accept

user@host# set term t2 from route-filter 200.2.0.0/16 orlonger
user@host# set term t2 then priority medium
user@host# set term t2 then accept

user@host# set term t3 from route-filter 200.3.0.0/16 orlonger
user@host# set term t3 then priority high
user@host# set term t3 then accept
```

4. Apply the policy to OSPF.

```
[edit protocols ospf]
user@host# set import ospf-import
```

5. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
fe-0/1/0 {
```

```

    unit 0 {
        family inet {
            address 192.168.8.4/30;
        }
    }
}
fe-0/2/0 {
    unit 0 {
        family inet {
            address 192.168.8.5/30;
        }
    }
}

user@host# show protocols ospf
import ospf-import;
area 0.0.0.0 {
    interface fe-0/1/0.0;
    interface fe-0/2/0.0;
}

user@host# show policy-options
policy-statement ospf-import {
    term t1 {
        from {
            route-filter 200.1.0.0/16 orlonger;
        }
        then {
            priority low;
            accept;
        }
    }
    term t2 {
        from {
            route-filter 200.2.0.0/16 orlonger;
        }
        then {
            priority medium;
            accept;
        }
    }
    term t3 {
        from {
            route-filter 200.3.0.0/16 orlonger;
        }
        then {
            priority high;
            accept;
        }
    }
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show protocols ospf3**, and **show policy-options** commands.

**Verification**

Confirm that the configuration is working properly.

**Verifying the Prefix Priority in the OSPF Routing Table**

**Purpose** Verify the priority assigned to the prefix in the OSPF routing table.

**Action** From operational mode, enter the **show ospf route detail** for OSPFv2, and enter the **show ospf3 route detail** command for OSPFv3.

**Related Documentation**

- [OSPF Overview on page 4036](#)
- [OSPF Configuration Overview](#)
- *Routing Policy Match Conditions in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
- *Actions in Routing Policy Terms in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*

**Examples: Configuring Routing Policy for Network Summaries**

- [Import and Export Policies for Network Summaries Overview on page 4180](#)
- [Example: Configuring an OSPF Export Policy for Network Summaries on page 4181](#)
- [Example: Configuring an OSPF Import Policy for Network Summaries on page 4190](#)

**Import and Export Policies for Network Summaries Overview**

By default, OSPF uses network-summary link-state advertisements (LSAs) to transmit route information across area boundaries. Each area border router (ABR) floods network-summary LSAs to other routing devices in the same area. The ABR also controls which routes from the area are used to generate network-summary LSAs into other areas. Each ABR maintains a separate topological database for each area to which they are connected. In Junos OS Release 9.1 and later, you can configure export and import policies for OSPFv2 and OSPFv3 that enable you to control how network-summary LSAs, which contain information about interarea OSPF prefixes, are distributed and generated. For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area.

The export policy enables you to specify which summary LSAs are flooded into an area. The import policy enables you to control which routes learned from an area are used to generate summary LSAs into other areas. You define a routing policy at the **[edit policy-options policy-statement *policy-name*]** hierarchy level. As with all OSPF export policies, the default for network-summary LSA export policies is to reject everything. Similarly, as with all OSPF import policies, the default for network-summary LSA import policies is to accept all OSPF routes.

### Example: Configuring an OSPF Export Policy for Network Summaries

This example shows how to create an OSPF export policy to control the network-summary (Type 3) LSAs that the ABR floods into an OSPF area.

- [Requirements on page 4181](#)
- [Overview on page 4181](#)
- [Configuration on page 4183](#)
- [Verification on page 4188](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#)

#### Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.




**NOTE:** For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

The following guidelines apply to export network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the **route-filter** policy match condition for these types of policies.
- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

Figure 122 on page 4182 shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

Figure 122: Sample Topology Used for an OSPF Export Network Summary Policy



In this example, you configure R4 with an export network summary policy named `export-policy` that only allows routes that match the `10.0.4.4` prefix from area 3 into area 4. The export policy controls the network-summary LSAs that R4 floods into area 4. This results in only the allowed interarea route to enter area 4, and all other interarea routes to be purged from the OSPF database and the routing table of the devices in area 4. You first define the policy and then apply it to the ABR by including the **network-summary-export** statement for OSPFv2 or the **inter-area-prefix-export** statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface `ge-0/0/1` has an IP address of 10.0.4.13/30 and connects to R3. Interface `ge-0/0/0` has an IP address of 10.0.4.5/30 and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface `ge-0/0/0` has an IP address of 10.0.4.6/30 and connects to R1. Interface `ge-0/0/2` has an IP address of 10.0.4.1 and connects to R3.
- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface `ge-0/0/2` has an IP address of 10.0.4.2/30 and connects to R2. Interface `ge-0/0/1` has an IP address of 10.0.4.14/30 and connects to R1. Interface `ge-0/0/0` has an IP address of 10.0.2.1/30 and connects to R4.
- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface `ge-0/0/0` has an IP address of 10.0.2.4/30 and connects to R3. Interface `ge-0/0/1` has an IP



address of 10.0.8.6/30 and connects to R5. Interface **ge-0/0/2** has an IP address of 10.0.8.9/30 and connects to R6.

- R5—Device R5 is an internal router in area 4. Interface **ge-0/0/1** has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface **ge-0/0/2** has an IP address of 10.0.8.10/30 and connects to R4.

### Configuration

**CLI Quick Configuration** To quickly configure an OSPF export policy for network summaries, copy the following commands, removing any line breaks, and then paste the commands into the CLI.

Configuration on Device R1:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.0.4.5/30
set interfaces ge-0/0/1 unit 0 family inet address 10.0.4.13/30
set protocols ospf area 0.0.0.3 interface ge-0/0/1
set protocols ospf area 0.0.0.3 interface ge-0/0/0
```

Configuration on Device R2:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.0.4.6/30
set interfaces ge-0/0/2 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface ge-0/0/2
set protocols ospf area 0.0.0.3 interface ge-0/0/1
```

Configuration on Device R3:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces ge-0/0/1 unit 0 family inet address 10.0.4.14/30
set interfaces ge-0/0/2 unit 0 family inet address 10.0.2.1/30
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.3 interface ge-0/0/1
set protocols ospf area 0.0.0.3 interface ge-0/0/2
```

Configuration on Device R4:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.0.2.1/30
set interfaces ge-0/0/1 unit 0 family inet address 10.0.8.6/30
set interfaces ge-0/0/2 unit 0 family inet address 10.0.8.9/30
set protocols ospf area 0.0.0.0 interface ge-0/0/0
set protocols ospf area 0.0.0.4 network-summary-export export-policy
set protocols ospf area 0.0.0.4 interface ge-0/0/1
set protocols ospf area 0.0.0.4 interface ge-0/0/2
set policy-options policy-statement export-policy term term1 from route-filter 10.0.4.4/30
  prefix-length-range /30-/30
set policy-options policy-statement export-policy term term1 then accept
```

Configuration on Device R5:

```
[edit]
set interfaces ge-0/0/1 unit 0 family inet address 10.0.8.5/30
```

```
set protocols ospf area 0.0.0.4 interface ge-0/0/1
```

Configuration on Device R6:

```
[edit]
set interfaces ge-0/0/2 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface ge-0/0/2
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure an OSPF export policy for network summaries:

1. Configure the interfaces.



**NOTE:** For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces ge-0/0/0 unit 0 family inet address 10.0.4.5/30
user@R1# set interfaces ge-0/0/1 unit 0 family inet address 10.0.4.13/30

[edit]
user@R2# set interfaces ge-0/0/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces ge-0/0/2 unit 0 family inet address 10.0.4.1/30

[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
user@R3# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30

[edit]
user@R4# set interfaces ge-0/0/0 unit 0 family inet address 10.0.4.2/30
user@R4# set interfaces ge-0/0/1 unit 0 family inet address 10.0.4.14/30
user@R4# set interfaces ge-0/0/2 unit 0 family inet address 10.0.2.1/30

[edit]
user@R5# set interfaces ge-0/0/1 unit 0 family inet address 10.0.8.5/30

[edit]
user@R6# set interfaces ge-0/0/2 unit 0 family inet address 10.0.8.10/30
```

2. Enable OSPF on the interfaces.



**NOTE:** For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@R1# set protocols ospf area 0.0.0.3 interface ge-0/0/1
user@R1# set protocols ospf area 0.0.0.3 interface ge-0/0/0

[edit]
user@R2# set protocols ospf area 0.0.0.3 interface ge-0/0/2
```

```

user@R2# set protocols ospf area 0.0.0.3 interface ge-0/0/1
[edit]
user@R3# set protocols ospf area 0.0.0.3 interface ge-0/0/1
user@R3# set protocols ospf area 0.0.0.3 interface ge-0/0/2
user@R3# set protocols ospf area 0.0.0.0 interface ge-0/0/0

[edit]
user@R4# set protocols ospf area 0.0.0.0 interface ge-0/0/0
user@R4# set protocols ospf area 0.0.0.4 interface ge-0/0/1
user@R4# set protocols ospf area 0.0.0.4 interface ge-0/0/2

[edit]
user@R5# set protocols ospf area 0.0.0.4 interface fe-0/0/1

[edit]
user@R6# set protocols ospf area 0.0.0.4 interface fe-0/0/2

```

3. On R4, configure the export network summary policy.

```

[edit ]
user@R4# set policy-options policy-statement export-policy term term1 from
route-filter 10.0.4.4/30 prefix-length-range /30-/30
user@R4# set policy-options policy-statement export-policy term term1 then accept

```

4. On R4, apply the export network summary policy to OSPF.



**NOTE:** For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```

[edit]
user@R4# set protocols ospf area 0.0.0.4 network-summary-export export-policy

```

5. If you are done configuring the devices, commit the configuration.

```

[edit]
user@host# commit

```

**Results** Confirm your configuration by entering the `show interfaces`, `show policy-options`, and `show protocols ospf` commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```

user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.5/30;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {

```

```
        address 10.0.4.13/30;
    }
}
```

```
user@R1# show protocols ospf
area 0.0.0.3 {
    interface ge-0/0/1.0;
    interface ge-0/0/0.0;
}
```

Output for R2:

```
user@R2# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.0.4.6/30;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.0.4.1/30;
        }
    }
}
```

```
user@R2# show protocols ospf
area 0.0.0.3 {
    interface ge-0/0/2.0;
    interface ge-0/0/1.0;
}
```

Output for R3:

```
user@R3# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.0.2.1/30;
            address 10.0.4.2/30;
        }
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.0.4.14/30;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.0.4.2/30;
            address 10.0.2.1/30;
        }
    }
}
```

```

    }
  }
}

user@R3# show protocols ospf
area 0.0.0.0 {
  interface ge-0/0/0.0;
}
area 0.0.0.3 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
}

```

Output for R4:

```

user@R4# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.0.2.1/30;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.8.6/30;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 10.0.8.9/30;
    }
  }
}

user@R4# show protocols ospf
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
  }
  area 0.0.0.4 {
    network-summary-export export-policy;
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
  }
}

user@R4# show policy-options
policy-statement export-policy {
  term term1 {
    from {
      route-filter 10.0.4.4/30 prefix-length-range /30-/30;
    }
    then accept;
  }
}

```

```
}  
}
```

Output for R5:

```
user@R5# show interfaces  
ge-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.8.5/30;  
    }  
  }  
}  
  
user@R5# show protocols ospf  
ospf {  
  area 0.0.0.4 {  
    interface ge-0/0/1.0;  
  }  
}
```

Output for R6:

```
user@R6# show interfaces  
ge-0/0/2 {  
  unit 0 {  
    family inet {  
      address 10.0.8.10/30;  
    }  
  }  
}  
  
user@R6# show protocols ospf  
area 0.0.0.4 {  
  interface ge-0/0/2.0;  
}
```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands on the appropriate device.

### **Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 4188](#)
- [Verifying the Routing Table on page 4189](#)

### **Verifying the OSPF Database**

**Purpose** Verify that the OSPF database for the devices in area 4 includes the interarea route that we permitted on the ABR R4. The other interarea routes that are not specified should age out or no longer be present in the OSPF database.

**Action** From operational mode, enter the **show ospf database** command.

## Sample Output

```
user@R4>show ospf database
OSPF database, Area 0.0.0.0
  Type      ID            Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.0.2.1        10.0.2.1    0x80000004  911  0x22 0xda1f  36
Router  *10.0.2.2        10.0.2.2    0x80000003  1505 0x22 0xda1d  36
Network *10.0.2.2        10.0.2.2    0x80000002  213  0x22 0x6d97  32
Summary 10.0.4.0        10.0.2.1    0x80000003  1495 0x22 0x60c1  28
Summary 10.0.4.4        10.0.2.1    0x80000002  1490 0x22 0x44d9  28
Summary 10.0.4.12       10.0.2.1    0x80000003  1490 0x22 0xe72e  28
Summary *10.0.8.4    10.0.2.2    0x80000004   644 0x22 0x414  28
Summary *10.0.8.8    10.0.2.2    0x80000003  1503 0x22 0xdd37  28
```

```
OSPF database, Area 0.0.0.4
  Type      ID            Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  *10.0.2.2        10.0.2.2    0x80000004  1508 0x22 0x597  48
Router  10.0.8.5         10.0.8.5    0x80000003  1517 0x22 0x8cc  36
Router  10.0.8.10        10.0.8.10    0x80000003  1514 0x22 0x3090 36
Network 10.0.8.5         10.0.8.5    0x80000001  1517 0x22 0x35b4 32
Network 10.0.8.10      10.0.8.10    0x80000001  1514 0x22 0x17c3 32
Summary *10.0.4.4     10.0.2.2    0x80000001  1492 0x22 0x4ad2 28
```

```
user@R5>show ospf database
OSPF database, Area 0.0.0.4
  Type      ID            Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.0.2.2        10.0.2.2    0x80000004  1479 0x22 0x597  48
Router  *10.0.8.5         10.0.8.5    0x80000003  1486 0x22 0x8cc  36
Router  10.0.8.10        10.0.8.10    0x80000003  1485 0x22 0x3090 36
Network *10.0.8.5         10.0.8.5    0x80000001  1486 0x22 0x35b4 32
Network 10.0.8.10      10.0.8.10    0x80000001  1485 0x22 0x17c3 32
Summary 10.0.4.4       10.0.2.2    0x80000001  1463 0x22 0x4ad2 28
```

```
user@R6>show ospf database
OSPF database, Area 0.0.0.4
  Type      ID            Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.0.2.2        10.0.2.2    0x80000004  2162 0x22 0x597  48
Router  10.0.8.5         10.0.8.5    0x80000003  2171 0x22 0x8cc  36
Router  *10.0.8.10       10.0.8.10    0x80000003  2166 0x22 0x3090 36
Network 10.0.8.5         10.0.8.5    0x80000001  2171 0x22 0x35b4 32
Network *10.0.8.10     10.0.8.10    0x80000001  2166 0x22 0x17c3 32
Summary 10.0.4.4       10.0.2.2    0x80000001  2146 0x22 0x4ad2 28
```

### Verifying the Routing Table

- Purpose** Verify that the routes corresponding to the rejected network summaries are no longer present in R4's, R5's, or R6's routing table.
- Action** From operational mode, enter the **show route protocol ospf** command for both OSPFv2 and OSPFv3.

## Sample Output

```
user@R4> show route protocol ospf
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.4.0/30      * [OSPF/10] 00:37:05, metric 2
                  > to 10.0.2.1 via ge-3/0/2.4
10.0.4.4/30      * [OSPF/10] 00:36:59, metric 3
```

```
10.0.4.12/30      > to 10.0.2.1 via ge-3/0/2.4
                  *[OSPF/10] 00:37:05, metric 2
224.0.0.5/32     > to 10.0.2.1 via ge-3/0/2.4
                  *[OSPF/10] 00:38:05, metric 1
                  MultiRecv
```

user@R5> show route protocol ospf

inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.4.4/30      *[OSPF/10] 00:37:09, metric 4
                  > to 10.0.8.6 via ge-3/0/2.5
10.0.8.8/30      *[OSPF/10] 00:37:30, metric 2
                  > to 10.0.8.6 via ge-3/0/2.5
224.0.0.5/32     *[OSPF/10] 00:38:20, metric 1
                  MultiRecv
```

user@6> show route protocol ospf

inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

```
10.0.4.4/30      *[OSPF/10] 00:38:19, metric 4
                  > to 10.0.8.9 via ge-3/0/2.6
10.0.8.4/30      *[OSPF/10] 00:38:34, metric 2
                  > to 10.0.8.9 via ge-3/0/2.6
224.0.0.5/32     *[OSPF/10] 00:39:34, metric 1
                  MultiRecv
```

---

### Example: Configuring an OSPF Import Policy for Network Summaries

This example shows how to create an OSPF import policy to control the network-summary (Type 3) LSAs that the ABR advertises out of an OSPF area.

- [Requirements on page 4190](#)
- [Overview on page 4190](#)
- [Configuration on page 4192](#)
- [Verification on page 4198](#)

#### Requirements

Before you begin:

- Configure the router identifiers for the devices in your OSPF network. See [“Example: Configuring an OSPF Router Identifier” on page 4048](#).
- Control OSPF designated router election. See [“Example: Controlling OSPF Designated Router Election” on page 4050](#).

#### Overview

OSPF uses network-summary LSAs to transmit route information across area boundaries. Depending on your network environment, you might want to further filter the network-summary LSAs between OSPF areas. For example, if you create OSPF areas to define administrative boundaries, you might not want to advertise internal route information between those areas. To further improve the control of route distribution



between multiple OSPF areas, you can configure network summary policies on the ABR for the area that you want to filter the advertisement of network-summary LSAs.




**NOTE:** For OSPFv3, the LSA is referred to as the interarea prefix LSA and performs the same function as a network-summary LSA performs for OSPFv2. An ABR originates an interarea prefix LSA for each IPv6 prefix that must be advertised into an area. In this topic, the terms network summary policy and network-summary policy are used to describe both OSPFv2 and OSPFv3 functionality.

The following guidelines apply to import network summary policies:

- You should have a thorough understanding of your network before configuring these policies. Incorrect network summary policy configuration might result in an unintended result such as suboptimal routing or dropped traffic.
- We recommend that you use the **route-filter** policy match condition for these types of policies.
- We recommend that you use the **accept** and **reject** routing policy terms for these types of policies.

Figure 123 on page 4191 shows a sample topology with three OSPF areas. R4 generates network summaries for the routes in area 4 and sends them out of area 4 to area 0. R3 generates network summaries for the routes in area 3 and sends them out of area 3 to area 0.

**Figure 123: Sample Topology Used for an OSPF Import Network Summary Policy**



In this example, you configure R3 with an import network summary policy named **import-policy** so R3 only generates network summaries for the route 10.0.4.12/30. The import policy controls the routes and therefore the network summaries that R3 advertises out of area 3, so applying this policy means that R3 only advertises route 10.0.4.12/30 out of area 3. This results in existing network summaries from other interarea routes

getting purged from the OSPF database in area 0 and area 4, as well as the routing tables of the devices in areas 0 and area 4. You first define the policy and then apply it to the ABR by including the **network-summary-import** statement for OSPFv2 or the **inter-area-prefix-import** statement for OSPFv3.

The devices operate as follows:

- R1—Device R1 is an internal router in area 3. Interface **fe-0/1/0** has an IP address of 10.0.4.13/30 and connects to R3. Interface **fe-0/0/1** has an IP address of 10.0.4.5/30 and connects to R2.
- R2—Device R2 is an internal router in area 3. Interface **fe-0/0/1** has an IP address of 10.0.4.6/30 and connects to R1. Interface **fe-1/0/0** has an IP address of 10.0.4.1/30 and connects to R3.
- R3—Device R3 participates in area 3 and area 0. R3 is the ABR between area 3 and area 0, and passes network-summary LSAs between the areas. Interface **fe-1/0/0** has an IP address of 10.0.4.2/30 and connects to R2. Interface **fe-1/1/0** has an IP address of 10.0.4.14/30 and connects to R1. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R4.
- R4—Device R4 participates in area 0 and area 4. R4 is the ABR between area 0 and area 4, and passes network-summary LSAs between the areas. Interface **fe-0/0/1** has an IP address of 10.0.2.1/30 and connects to R3. Interface **fe-1/1/0** has an IP address of 10.0.8.6/30 and connects to R5. Interface **fe-1/0/0** has an IP address of 10.0.8.9/30 and connects to R6.
- R5—Device R5 is an internal router in area 4. Interface **fe-1/1/0** has an IP address of 10.0.8.5/30 and connects to R4.
- R6—Device R6 is an internal router in area 4. Interface **fe-1/0/0** has an IP address of 10.0.8.10/30 and connects to R4.

### *Configuration*

#### **CLI Quick Configuration**

To quickly configure an OSPF import policy for network summaries, copy the following commands, removing any line breaks, and then paste the commands into CLI.

Configuration on Device R1:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-0/0/1
```

Configuration on Device R2:

```
[edit]
set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30
set protocols ospf area 0.0.0.3 interface fe-0/1/0
set protocols ospf area 0.0.0.3 interface fe-1/0/0
```

Configuration on Device R3:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set policy-options policy-statement import-policy term term1 from route-filter 10.0.4.12/30
  prefix-length-range /30-/30
set policy-options policy-statement import-policy term term1 then accept
set protocols ospf area 0.0.0.3 interface fe-1/0/0
set protocols ospf area 0.0.0.3 interface fe-1/1/0
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

Configuration on Device R4:

```
[edit]
set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
set protocols ospf area 0.0.0.0 interface fe-0/0/1
set protocols ospf area 0.0.0.4 interface fe-1/1/0
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

Configuration on Device R5:

```
[edit]
set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
set protocols ospf area 0.0.0.4 interface fe-1/1/0
```

Configuration on Device R6:

```
[edit]
set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30
set protocols ospf area 0.0.0.4 interface fe-1/0/0
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure an OSPF import policy for network summaries:

1. Configure the interfaces.



**NOTE:** For OSPFv3, use IPv6 addresses.

```
[edit]
user@R1# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.13/30
user@R1# set interfaces fe-0/0/1 unit 0 family inet address 10.0.4.5/30

[edit]
user@R2# set interfaces fe-0/1/0 unit 0 family inet address 10.0.4.6/30
user@R2# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.1/30

[edit]
user@R3# set interfaces fe-1/0/0 unit 0 family inet address 10.0.4.2/30
user@R3# set interfaces fe-1/1/0 unit 0 family inet address 10.0.4.14/30
```

```

user@R3#set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
[edit]
user@R4# set interfaces fe-0/0/1 unit 0 family inet address 10.0.2.1/30
user@R4# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.6/30
user@R4# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.9/30
[edit]
user@R5# set interfaces fe-1/1/0 unit 0 family inet address 10.0.8.5/30
[edit]
user@R6# set interfaces fe-1/0/0 unit 0 family inet address 10.0.8.10/30

```

2. Enable OSPF on the interfaces.



**NOTE:** For OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```

[edit]
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R1# set protocols ospf area 0.0.0.3 interface fe-0/0/1
[edit]
user@R2# set protocols ospf area 0.0.0.3 interface fe-0/1/0
user@R2# set protocols ospf area 0.0.0.3 interface fe-1/0/0
[edit]
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/0/0
user@R3# set protocols ospf area 0.0.0.3 interface fe-1/1/0
user@R3# set protocols ospf area 0.0.0.0 interface fe-0/0/1
[edit]
user@R4# set protocols ospf area 0.0.0.0 interface fe-0/0/1
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/1/0
user@R4# set protocols ospf area 0.0.0.4 interface fe-1/0/0
[edit]
user@R5# set protocols ospf area 0.0.0.4 interface fe-1/1/0
[edit]
user@R6# set protocols ospf area 0.0.0.4 interface fe-1/0/0

```

3. On R3, configure the import network summary policy.

```

[edit ]
user@R3# set policy-options policy-statement import-policy term term1 from
route-filter 10.0.4.12/30 prefix-length-range /30-/30
user@R3# set policy-options policy-statement import-policy term term1 then accept

```

4. On R3, apply the import network summary policy to OSPF.



**NOTE:** For OSPFv3, include the `inter-area-prefix-export` statement at the `[edit protocols ospf3 area area-id]` hierarchy level.

```

[edit]

```

```
user@R3# set protocols ospf area 0.0.0.3 network-summary-import import-policy
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

**Results** Confirm your configuration by entering the **show interfaces**, **show policy-options**, and **show protocols ospf** commands on the appropriate device. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Output for R1:

```
user@R1# show interfaces
fe-0/0/1 {
  unit 0 {
    family inet {
      address 10.0.4.5/30;
    }
  }
}
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.13/30;
    }
  }
}

user@R1# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
  interface fe-0/0/1.0;
}
```

Output for R2:

```
user@R2# show interfaces
fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.4.6/30;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.4.1/30;
    }
  }
}

user@R2# show protocols ospf
area 0.0.0.3 {
  interface fe-0/1/0.0;
```

```
    interface fe-1/0/0.0;  
  }
```

Output for R3:

```
user@R3# show interfaces  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.2.1/30;  
    }  
  }  
}  
fe-1/0/0 {  
  unit 0 {  
    family inet {  
      address 10.0.4.2/30;  
    }  
  }  
}  
fe-1/1/0 {  
  unit 0 {  
    family inet {  
      address 10.0.4.14/30;  
    }  
  }  
}  
  
user@R3# show protocols ospf  
area 0.0.0.0 {  
  interface fe-0/0/1.0;  
}  
area 0.0.0.3 {  
  network-summary-import import-policy;  
  interface fe-1/0/0.0;  
  interface fe-1/1/0.0;  
}  
  
user@R3# show policy-options  
policy-statement import-policy {  
  term term1 {  
    from {  
      route-filter 10.0.4.12/30 prefix-length-range /30-/30;  
    }  
    then accept;  
  }  
}
```

Output for R4:

```
user@R4# show interfaces  
fe-0/0/1 {  
  unit 0 {  
    family inet {  
      address 10.0.2.1/30;  
    }  
  }  
}
```

```

fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.9/30;
    }
  }
}
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.6/30;
    }
  }
}

user@R4# show protocols ospf
area 0.0.0.0 {
  interface fe-0/0/1.0;
}
area 0.0.0.4 {
  interface fe-0/1/0.0;
  interface fe-1/0/0.0;
}

```

Output for R5:

```

user@R5# show interfaces
fe-1/1/0 {
  unit 0 {
    family inet {
      address 10.0.8.5/30;
    }
  }
}

user@R5# show protocols ospf
area 0.0.0.4 {
  interface fe-1/1/0.0;
}

```

Output for R6:

```

user@R6# show interfaces
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.0.8.10/30;
    }
  }
}

user@R6# show protocols ospf
area 0.0.0.4 {
  interface fe-1/0/0.0;
}

```

To confirm your OSPFv3 configuration, enter the **show interfaces**, **show policy-options**, and **show protocols ospf3** commands on the appropriate device.

**Verification**

Confirm that the configuration is working properly.

- [Verifying the OSPF Database on page 4198](#)
- [Verifying the Routing Table on page 4198](#)

**Verifying the OSPF Database**

**Purpose** Verify that the OSPF database for the devices in area 4 includes the interarea route that we are advertising from R3. Any other routes from area 3 should not be advertised into area 4, so those entries should age out or no longer be present in the OSPF database.

**Action** From operational mode, enter the **show ospf database netsummary area 0.0.0.4** command for OSPFv2, and enter the **show ospf3 database inter-area-prefix area 0.0.0.4** command for OSPFv3.

**Verifying the Routing Table**

**Purpose** Verify that the specified route is included in R4's, R5's, or R6's routing table. Any other routes from area 3 should not be advertised into area 4.

**Action** From operational mode, enter the **show route protocol ospf** command for both OSPFv2 and OSPFv3.

- Related Documentation**
- [OSPF Overview on page 4036](#)
  - [OSPF Configuration Overview](#)
  - [Routing Policy Match Conditions in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices](#)
  - [Actions in Routing Policy Terms in the Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices](#)

---

## OSPF Monitoring Configuration

- [Example: Configuring OSPF Trace Options on page 4198](#)

### Example: Configuring OSPF Trace Options

- [Tracing OSPF Protocol Traffic on page 4198](#)
- [Example: Tracing OSPF Protocol Traffic on page 4200](#)

---

#### Tracing OSPF Protocol Traffic

Tracing operations record detailed messages about the operation of OSPF. You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace.



You can specify the following OSPF protocol-specific trace options:

- **database-description**—All database description packets, which are used in synchronizing the OSPF topological database
- **error**—OSPF error packets
- **event**—OSPF state transitions
- **flooding**—Link-state flooding packets
- **graceful-restart**—Graceful-restart events
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable
- **ldp-synchronization**—Synchronization events between OSPF and LDP
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database
- **nsr-synchronization**—Nonstop routing synchronization events
- **on-demand**—Trace demand circuit extensions
- **packet-dump**—Dump the contents of selected packet types
- **packets**—All OSPF packets
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events
- **spf**—Shortest path first (SPF) calculations

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



**NOTE:** Use the **detail** flag modifier with caution as it might cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the OSPF protocol using the **traceoptions flag** statement included at the **[edit protocols ospf]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



**NOTE:** Use the trace flag **all** with caution as it might cause the CPU to become very busy.

---

### Example: Tracing OSPF Protocol Traffic

---

This example shows how to trace OSPF protocol traffic.

- [Requirements on page 4200](#)
- [Overview on page 4200](#)
- [Configuration on page 4201](#)
- [Verification on page 4205](#)

#### **Requirements**

This example assumes that OSPF is properly configured and running in your network, and you want to trace OSPF protocol traffic for debugging purposes.

#### **Overview**

You can trace OSPF protocol traffic to help debug OSPF protocol issues. When you trace OSPF protocol traffic, you specify the name of the file and the type of information you want to trace. All files are placed in a directory on the routing device's hard disk. On M Series and T Series routers, trace files are stored in the `/var/log` directory.

This example shows a few configurations that might be useful when debugging OSPF protocol issues. The verification output displayed is specific to each configuration.

---



**TIP:** To keep track of your log files, create a meaningful and descriptive name so it is easy to remember the content of the trace file. We recommend that

you place global routing protocol tracing output in the file `routing-log`, and OSPF tracing output in the file `ospf-log`.

In the first example, you globally enable tracing operations for all routing protocols that are actively running on your routing device to the file `routing-log`. With this configuration, you keep the default settings for the trace file size and the number of trace files. After enabling global tracing operations, you enable tracing operations to provide detailed information about OSPF packets, including link-state advertisements, requests, and updates, database description packets, and hello packets to the file `ospf-log`, and you configure the following options:

- **size**—Specifies the maximum size of each trace file, in KB, MB, or GB. In this example, you configure 10 KB as the maximum size. When the file reaches its maximum size, it is renamed with a .0 extension. When the file again reaches its maximum size, it is renamed with a .1 extension, and the newly created file is renamed with a .0 extension. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option. You specify **k** for KB, **m** for MB, and **g** for GB. By default, the trace file size is 128 KB. The file size range is 10 KB through the maximum file size supported on your system.
- **files**—Specifies the maximum number of trace files. In this example, you configure a maximum of 5 trace files. When a trace file reaches its maximum size, it is renamed with a .0 extension, then a .1 extension, and so on until the maximum number of trace files is reached. When the maximum number of files is reached, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option. By default, there are 10 files. The range is 2 through 1000 files.

In the second example, you trace all SPF calculations to the file `ospf-log` by including the **spf** flag. You keep the default settings for the trace file size and the number of trace files.

In the third example, you trace the creation, receipt, and retransmission of all LSAs to the file `ospf-log` by including the **lsa-request**, **lsa-update**, and **lsa-ack** flags. You keep the default settings for the trace file size and the number of trace files.

### **Configuration**

- [Configuring Global Tracing Operations and Tracing OSPF Packet Information on page 4201](#)
- [Tracing SPF Calculations on page 4203](#)
- [Tracing Link-State Advertisements on page 4204](#)

### **Configuring Global Tracing Operations and Tracing OSPF Packet Information**

#### **CLI Quick Configuration**

To quickly enable global tracing operations for all routing protocols actively running on your routing device and to trace detailed information about OSPF packets, copy the following commands and paste them into the CLI.

[edit]

```
set routing-options traceoptions file routing-log
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions file files 5 size 10k
set protocols ospf traceoptions flag lsa-ack
set protocols ospf traceoptions flag database-description
set protocols ospf traceoptions flag hello
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-request
```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Modifying the Junos OS Configuration* in *CLI User Guide*.

To configure global routing tracing operations and tracing operations for OSPF packets:

1. Configure tracing at the routing options level to collect information about the active routing protocols on your routing device.

```
[edit]
user@host# edit routing-options traceoptions
```

2. Configure the filename for the global trace file.

```
[edit routing-options traceoptions]
user@host# set file routing-log
```

3. Configure the filename for the OSPF trace file.



**NOTE:** To specify OSPFv3, include the `ospf3` statement at the `[edit protocols]` hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

4. Configure the maximum number of trace files.

```
[edit protocols ospf traceoptions]
user@host# set file files 5
```

5. Configure the maximum size of each trace file.

```
[edit protocols ospf traceoptions]
user@host# set file size 10k
```

6. Configure tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-ack
user@host# set flag database-description
user@host# set flag hello
user@host# set flag lsa-update
user@host# set flag lsa-request
```

7. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
```

```
user@host# commit
```

**Results** Confirm your configuration by entering the **show routing-options** and the **show protocols ospf** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
traceoptions {
  file routing-log;
}
```

```
user@host# show protocols ospf
traceoptions {
  file ospf-log size 10k files 5;
  flag lsa-ack;
  flag database-description;
  flag hello;
  flag lsa-update;
  flag lsa-request;
}
```

To confirm your OSPFv3 configuration, enter the **show routing-options** and the **show protocols ospf3** commands.

### *Tracing SPF Calculations*

**CLI Quick Configuration** To quickly trace SPF calculations, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag spf
```

**Step-by-Step Procedure** To configure SPF tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the SPF tracing flag.

```
[edit protocols ospf traceoptions]
user@host# set flag spf
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
  file ospf-log ;
  flag spf;
}
```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### *Tracing Link-State Advertisements*

**CLI Quick Configuration** To quickly trace the creation, receipt, and retransmission of all LSAs, copy the following commands and paste them into the CLI.

```
[edit]
set protocols ospf traceoptions file ospf-log
set protocols ospf traceoptions flag lsa-request
set protocols ospf traceoptions flag lsa-update
set protocols ospf traceoptions flag lsa-ack
```

**Step-by-Step Procedure** To configure link-state advertisement tracing operations for OSPF:

1. Configure the filename for the OSPF trace file.



**NOTE:** To specify OSPFv3, include the **ospf3** statement at the **[edit protocols]** hierarchy level.

```
[edit]
user@host# edit protocols ospf traceoptions
user@host# set file ospf-log
```

2. Configure the link-state advertisement tracing flags.

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-request
user@host# set flag lsa-update
user@host# set flag lsa-ack
```

3. If you are done configuring the device, commit the configuration.

```
[edit protocols ospf traceoptions]
user@host# commit
```

**Results** Confirm your configuration by entering the **show protocols ospf** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols ospf
traceoptions {
  file ospf-log;
```

```

    flag lsa-request;
    flag lsa-update;
    flag lsa-ack;
}

```

To confirm your OSPFv3 configuration, enter the **show protocols ospf3** command.

### Verification

Confirm that the configuration is working properly.

### Verifying Trace Operations

<b>Purpose</b>	Verify that the Trace options field displays the configured trace operations, and verify that the Trace file field displays the location on the routing device where the file is saved, the name of the file to receive the output of the tracing operation, and the size of the file.
<b>Action</b>	From operational mode, enter the <b>show ospf overview extensive</b> command for OSPFv2, and enter the <b>show ospf3 overview extensive</b> command for OSPFv3.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">OSPF Overview on page 4036</a></li> <li>• <a href="#">OSPF Configuration Overview</a></li> <li>• <a href="#">Tracing and Logging Junos OS Operations in the Junos OS Administration Library for Routing Devices</a></li> <li>• <a href="#">Example: Tracing Global Routing Protocol Operations in the Junos OS Routing Protocols Library for Routing Devices</a></li> </ul>

## Configuration Statements

- [area on page 4207](#)
- [area-range on page 4209](#)
- [authentication \(Protocols OSPF\) on page 4211](#)
- [context-identifier \(Protocols OSPF\) on page 4212](#)
- [bfd-liveness-detection \(Protocols OSPF\) on page 4213](#)
- [database-protection on page 4217](#)
- [disable \(OSPF\) on page 4219](#)
- [export \(Protocols OSPF\) on page 4221](#)
- [external-preference \(Protocols OSPF\) on page 4222](#)
- [graceful-restart \(Protocols OSPF\) on page 4223](#)
- [import \(Protocols OSPF\) on page 4225](#)
- [interface \(Protocols OSPF\) on page 4226](#)
- [no-nssa-abr on page 4228](#)
- [no-rfc-1583 on page 4229](#)
- [ospf on page 4230](#)

- [ospf3](#) on page 4230
- [overload \(Protocols OSPF\)](#) on page 4231
- [preference \(Protocols OSPF\)](#) on page 4232
- [prefix-export-limit \(Protocols OSPF\)](#) on page 4233
- [reference-bandwidth \(Protocols OSPF\)](#) on page 4234
- [rib-group \(Protocols OSPF\)](#) on page 4235
- [topology \(OSPF\)](#) on page 4236
- [traceoptions \(Protocols OSPF\)](#) on page 4237
- [traffic-engineering \(OSPF\)](#) on page 4240



## area

<b>Syntax</b>	<pre> area <i>area-id</i> {     interface <i>interface-name</i> {         passive;         topology (ipv4-multicast   <i>name</i>) {             disable;         }     }     virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i> {         topology (ipv4-multicast   <i>name</i>) {             disable;         }     } } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<b>ospf</b>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.</p> <p>Specify multiple <b>area</b> statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas. Use the <b>area-range</b> statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the <b>virtual-link</b> statement.</p> <p>To specify that the routing device is directly connected to the OSPF backbone, include the <b>area 0.0.0.0</b> statement.</p> <p>All routing devices on the backbone must be contiguous. If they are not, use the <b>virtual-link</b> statement to create the appearance of connectivity to the backbone.</p>

You can also configure any interface that belongs to one or more topologies to advertise the direct interface addresses without actually running OSPF on that interface. By default, OSPF must be configured on an interface in order for direct interface addresses to be advertised as interior routes.



**NOTE:** If you configure an interface with the **passive** statement, it applies to all the topologies to which the interface belongs. You cannot configure an interface as passive for only one specific topology and have it remain active for any other topologies to which it belongs.

<b>Options</b>	<b><i>area-id</i></b> —Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number <b>0.0.0.0</b> is reserved for the OSPF backbone area.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">OSPF Areas and Router Functionality Overview on page 4041</a></li><li>• <a href="#">Understanding Multiple Address Families for OSPFv3 on page 4090</a></li><li>• <i>virtual-link</i></li></ul>

## area-range

<b>Syntax</b>	<b>area-range</b> <i>network/mask-length</i> <exact> <override-metric <i>metric</i> > <restrict>;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) area <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) area <i>area-id</i> nssa],</p> <p>[edit routing-instances <i>routing-instance-name</i> realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>(Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple <b>area-range</b> statements.</p> <p>For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple <b>area-range</b> statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.</p>
<b>Default</b>	By default, area border routing devices do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.
<b>Options</b>	<p><b>exact</b>—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.</p> <p><b>mask-length</b>—Number of significant bits in the network mask.</p> <p><b>network</b>—IP address. You can specify one or more IP addresses.</p>

**override-metric *metric***—(Optional) Override the metric for the IP address range and configure a specific metric value.

**restrict**—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

**Range:** 1 through 16,777,215

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Summarizing Ranges of Routes in OSPF Link-State Advertisements on page 4094</a></li></ul>
------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------

## authentication (Protocols OSPF)

<b>Syntax</b>	<pre> authentication {   md5 key-identifier {     key key-value;     start-time YYYY-MM-DD.hh:mm;   }   simple-password key; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> <a href="#">interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <a href="#">interface interface-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf area <i>area-id</i> <a href="#">interface interface-name</a>],</p> <p>[edit protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure an authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding OSPFv2 Authentication</i></li> <li>• <i>Example: Configuring MD5 Authentication for OSPFv2 Exchanges</i></li> <li>• <i>Example: Configuring a Transition of MD5 Keys on an OSPFv2 Interface</i></li> <li>• <i>Example: Configuring Simple Authentication for OSPFv2 Exchanges</i></li> </ul>

## context-identifier (Protocols OSPF)

---

<b>Syntax</b>	context-identifier <i>identifier</i>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> ], [edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure OSPF context-identifier information.
<b>Options</b>	<i>identifier</i> —IPv4 address that defines a protection pair. The context identifier is manually configured on both the primary and protector provider edge (PE) devices.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ospf context-identifier on page 4260</a></li></ul>

## bfd-liveness-detection (Protocols OSPF)

**Syntax**    `bfd-liveness-detection {`  
                   `authentication {`  
                     `algorithm algorithm-name;`  
                     `key-chain key-chain-name;`  
                     `loose-check;`  
                   `}`  
                   `detection-time {`  
                     `threshold milliseconds;`  
                   `}`  
                   `full-neighbors-only`  
                   `minimum-interval milliseconds;`  
                   `minimum-receive-interval milliseconds;`  
                   `multiplier number;`  
                   `no-adaptation;`  
                   `transmit-interval {`  
                     `minimum-interval milliseconds;`  
                     `threshold milliseconds;`  
                   `}`  
                   `version (1 | automatic);`  
                   `}`

**Hierarchy Level**    `[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],`  
                           `[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast |`  
                             `ipv4-multicast | ipv6-multicast) area area-id interface interface-name],`  
                           `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`  
                             `(ospf | ospf3) area area-id interface interface-name],`  
                           `[edit logical-systems logical-system-name routing-instances routing-instance-name protocols`  
                             `ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface`  
                             `interface-name],`  
                           `[edit protocols (ospf | ospf3) area area-id interface interface-name],`  
                           `[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id`  
                             `interface interface-name],`  
                           `[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface`  
                             `interface-name],`  
                           `[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast |`  
                             `ipv4-multicast | ipv6-multicast) area area-id interface interface-name]`

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                               Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                               **detection-time threshold** and **transmit-interval threshold** options added in Junos OS Release 8.2.  
                               Support for logical systems introduced in Junos OS Release 8.3.  
                               **no-adaptation** option introduced in Junos OS Release 9.0.  
                               **no-adaptation** option introduced in Junos OS Release 9.0 for EX Series switches.  
                               Support for OSPFv3 introduced in Junos OS Release 9.3.  
                               Support for OSPFv3 introduced in Junos OS Release 9.3 for EX Series switches.  
                               **full-neighbors-only** option introduced in Junos OS Release 9.5.  
                               **full-neighbors-only** option introduced in Junos OS Release 9.5 for EX Series switches.

**authentication algorithm**, **authentication key-chain**, and **authentication loose-check** options introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Configure bidirectional failure detection timers and authentication for OSPF.

The remaining statements are explained separately.



**Options** **authentication algorithm *algorithm-name***—Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

**authentication key-chain *key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**full-neighbors-only**—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established

a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** **automatic**

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring BFD for OSPF on page 4126</a></li><li>• <a href="#">Example: Configuring BFD Authentication for OSPF on page 4129</a></li></ul>
------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## database-protection

<b>Syntax</b>	<pre>database-protection {   ignore-count <i>number</i>;   ignore-time <i>seconds</i>;   maximum-lsa <i>number</i>;   reset-time <i>seconds</i>;   warning-only;   warning-threshold <i>percent</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-unicast   ipv6-multicast)]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Configure the maximum number of link-state advertisements (LSAs) that are not generated by the router or switch in a given OSPF instance.
<b>Default</b>	By default, OSPF database protection is not enabled.
<b>Options</b>	<p><b>ignore-count <i>number</i></b>—Configure the number of times the database can enter the ignore state. When the ignore count is exceeded, the database enters the isolate state.</p> <p><b>Range:</b> 1 through 32</p> <p><b>Default:</b> 5</p> <p><b>ignore-time <i>seconds</i></b>—Configure the time the database must remain in the ignore state before it resumes regular operations (enters retry state).</p> <p><b>Range:</b> 30 through 3,600 seconds</p> <p><b>Default:</b> 300 seconds</p> <p><b>maximum-lsa <i>number</i></b>—Configure the maximum number of LSAs whose advertising router ID is different from the local router ID in a given OSPF instance. This includes external LSAs as well as LSAs with any scope, such as the link, area, and autonomous system (AS). This value is mandatory.</p> <p><b>Range:</b> 1 through 1,000,000</p> <p><b>Default:</b> None</p> <p><b>reset-time <i>seconds</i></b>—Configure the time period during which the database must operate without being in the ignore or isolate state before it is reset to a normal operating state.</p> <p><b>Range:</b> 60 through 86,400 seconds</p> <p><b>Default:</b> 600 seconds</p>

**warning-only**—Specify that only a warning should be issued when the maximum LSA number is exceeded. If configured, no other action is taken against the database.

**warning-threshold *percent***—Configure the percentage of the maximum number of LSAs to be exceeded before a warning message is logged.

**Range:** 30 through 100 percent

**Default:** 75 percent

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">OSPF Database Protection Overview on page 4162</a></li><li>• <a href="#">Configuring OSPF Database Protection on page 4163</a></li></ul>
------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## disable (OSPF)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <b>area</b> <i>area-id</i> <b>peer-interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf   ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (<b>ospf</b>   ospf3)],</p> <p>[edit protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit protocols (ospf   ospf3) virtual-link],</p> <p>[edit protocols ospf <b>area</b> <i>area-id</i> <b>peer-interface</b> <i>interface-name</i>],</p> <p>[edit protocols ospf <b>area</b> <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf   ospf3) virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast) <b>area</b> <i>area-id</i> <b>interface</b> <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Disable OSPF, an OSPF interface, or an OSPF virtual link.</p> <p>By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that</p>

are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id *router-id* transit-area *area-id* topology *name*]** hierarchy level.



**NOTE:** If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id *router-id* transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

<b>Default</b>	The configured object is enabled (operational) unless explicitly disabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>OSPF Configuration Overview</i></li><li>• <i>Configuring RSVP and OSPF for LMP Peer Interfaces</i></li></ul>

## export (Protocols OSPF)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into OSPF.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding OSPF Routing Policy on page 4164</a></li> <li>• <a href="#">Import and Export Policies for Network Summaries Overview on page 4180</a></li> <li>• <a href="#">import on page 4225</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## external-preference (Protocols OSPF)

---

<b>Syntax</b>	<code>external-preference <i>preference</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code> <code>[edit protocols (<b>ospf</b>   ospf3)],</code> <code>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the <b>realm</b> statement introduced in Junos OS Release 9.2. Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Set the route preference for OSPF external routes.
<b>Options</b>	<b><i>preference</i></b> —Preference value. <b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ ) <b>Default:</b> 150
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Controlling OSPF Route Preferences on page 4111</a></li><li>• <a href="#">preference on page 4232</a></li></ul>



## graceful-restart (Protocols OSPF)

<b>Syntax</b>	<pre> graceful-restart {   disable;   helper-disable (standard   restart-signaling   both);   no-strict-lsa-checking;   notify-duration <i>seconds</i>;   restart-duration <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for the <b>no-strict-lsa-checking</b> statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the helper mode <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure graceful restart for OSPF.</p> <p>Graceful restart allows a routing device to restart with minimal effects to the network, and is enabled for all routing protocols at the [edit routing-options] hierarchy level.</p>
<b>Options</b>	<p><b>disable</b>—Disable graceful restart for OSPF.</p> <p><b>helper-disable (standard   restart-signaling   both)</b>—Disable helper mode for graceful restart. When helper mode is disabled, a device cannot help a neighboring device that is attempting to restart. Beginning with Junos OS Release 11.4, you can configure restart signaling-based helper mode for OSPFv2 graceful restart configurations. The <b>standard</b>, <b>restart-signaling</b>, and <b>both</b> options are only supported for OSPFv2. Specify <b>standard</b> to disable helper mode for standard graceful restart (based on RFC 3623). Specify <b>restart-signaling</b> to disable helper mode for restart signaling-based graceful restart (based on RFC 4811, RFC 4812, and RFC 4813). Specify <b>both</b> to disable helper mode for both standard and restart signaling-based graceful restart. The last committed statement takes precedence over the previously configured statement.</p> <p><b>Default:</b> Helper mode is enabled by default. For OSPFv2, both standard and restart-signaling based helper modes are enabled by default.</p> <p><b>no-strict-lsa-checking</b>—Disable strict OSPF link-state advertisement (LSA) checking to prevent the termination of graceful restart by a helping router. LSA checking is enabled by default.</p>



**NOTE:** The **helper-disable** statement and the **no-strict-lsa-checking** statement cannot be configured at the same time. If you attempt to configure both

statements at the same time, the routing device displays a warning message when you enter the `show protocols (ospf | ospf3)` command.

.....  
**notify-duration seconds**—Estimated time needed to send out purged grace LSAs over all the interfaces.

**Range:** 1 through 3600 seconds

**Default:** 30 seconds

**restart-duration seconds**—Estimated time needed to reacquire a full OSPF neighbor from each area.

**Range:** 1 through 3600 seconds

**Default:** 180 seconds

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---------------------------------------------------------------------------------------------------------------------

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Graceful Restart for OSPF on page 4136</a></li><li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv2 Graceful Restart on page 4140</a></li><li>• <a href="#">Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart on page 4144</a></li><li>• <a href="#">Example: Disabling Strict LSA Checking for OSPF Graceful Restart on page 4147</a></li><li>• <i>Junos OS High Availability Library for Routing Devices</i></li></ul> |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## import (Protocols OSPF)

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Filter OSPF routes from being added to the routing table.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding OSPF Routing Policy on page 4164</a></li> <li>• <a href="#">Import and Export Policies for Network Summaries Overview on page 4180</a></li> <li>• <a href="#">export on page 4221</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> </ul>

## interface (Protocols OSPF)

**Syntax** interface *interface-name* {  
 disable;  
 authentication key <key-id identifier>;  
 bfd-liveness-detection {  
 authentication {  
 algorithm *algorithm-name*;  
 key-chain *key-chain-name*;  
 loose-check;  
 }  
 detection-time {  
 threshold *milliseconds*;  
 }  
 minimum-interval *milliseconds*;  
 minimum-receive-interval *milliseconds*;  
 transmit-interval {  
 threshold *milliseconds*;  
 minimum-interval *milliseconds*;  
 }  
 multiplier *number*;  
 }  
 dead-interval *seconds*;  
 demand-circuit;  
 hello-interval *seconds*;  
 ipsec-sa *name*;  
 interface-type *type*;  
 ldp-synchronization {  
 disable;  
 hold-time *seconds*;  
 }  
 metric *metric*;  
 neighbor *address* <eligible>;  
 no-interface-state-traps;  
 passive;  
 poll-interval *seconds*;  
 priority *number*;  
 retransmit-interval *seconds*;  
 te-metric *metric*;  
 topology (ipv4-multicast | *name*) {  
 metric *metric*;  
 }  
 transit-delay *seconds*;  
}

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols (ospf | ospf3) *area area-id*],  
 [edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast) *area area-id*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 (ospf | ospf3) *area area-id*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
 ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],  
 [edit protocols (ospf | ospf3) *area area-id*],  
 [edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) *area area-id*],

[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) *area area-id*],  
 [edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast |  
 ipv4-multicast | ipv6-multicast) *area area-id*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Support for the **topology** statement introduced in Junos OS Release 9.0.  
 Support for the **topology** statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Support for the **realm** statement introduced in Junos OS Release 9.2.  
 Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.  
 Support for the **no-interface-state-traps** statement introduced in Junos OS Release 10.3.  
 This statement is supported only for OSPFv2.  
 Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Enable OSPF routing on a routing device interface.  
 You must include at least one **interface** statement in the configuration to enable OSPF on the routing device.

**Options** *interface-name*—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify **all**. Specifying a particular interface and **all** produces an invalid configuration.



**NOTE:** For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.

The remaining statements are explained separately.



**NOTE:** You cannot run both OSPF and ethernet-tcc encapsulation between two Juniper Networks routing devices.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

**Related Documentation**

- *OSPF Configuration Overview*
- *Example: Configuring Multitopology Routing Based on Applications*
- *Example: Configuring Multitopology Routing Based on a Multicast Source*
- [Example: Configuring Multiple Address Families for OSPFv3 on page 4091](#)
- *neighbor*

## no-nssa-abr

---

<b>Syntax</b>	no-nssa-abr;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ( <a href="#">ospf</a>   ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ( <a href="#">ospf</a>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit protocols ( <a href="#">ospf</a>   ospf3)], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols ( <a href="#">ospf</a>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for the <b>realm</b> statement introduced in Junos OS Release 9.2. Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Disable exporting Type 7 link-state advertisements into not-so-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring OSPF Not-So-Stubby Areas on page 4064</a></li></ul>

## no-rfc-1583

<b>Syntax</b>	no-rfc-1583;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Disable compatibility with RFC 1583, <i>OSPF Version 2</i> . If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.
<b>Default</b>	Compatibility with RFC 1583 is enabled by default.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control-level—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Disabling OSPFv2 Compatibility with RFC 1583 on page 4074</a></li> </ul>

## ospf

---

<b>Syntax</b>	ospf { ... }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable OSPF routing on the routing device.  You must include the <b>ospf</b> statement to enable OSPF on the routing device.
<b>Default</b>	OSPF is disabled on the routing device.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>OSPF Configuration Overview</i></li><li>• <i>[edit protocols ospf] Hierarchy Level</i></li></ul>


## ospf3

---

<b>Syntax</b>	ospf3 { ... }
<b>Hierarchy Level</b>	[edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable OSPFv3 routing on the routing device.  You must include the <b>ospf3</b> statement to enable OSPFv3.
<b>Default</b>	OSPFv3 is disabled on the routing device.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>OSPF Configuration Overview</i></li><li>• <i>[edit protocols ospf3] Hierarchy Level</i></li></ul>



## overload (Protocols OSPF)

<b>Syntax</b>	<pre>overload {     timeout <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<b>ospf</b>   ospf3)],</p> <p>[edit protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf <b>topology</b> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.</p>
<div>  <p><b>NOTE:</b> Traffic destined to directly attached interfaces continues to reach the routing device.</p> </div>	
<b>Options</b>	<p><b>timeout <i>seconds</i></b>—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the <b>overload</b> statement is deleted or a timeout is set.</p> <p><b>Range:</b> 60 through 1800 seconds</p> <p><b>Default:</b> 0 seconds</p>



**NOTE:** Multitopology Routing does not support the timeout option.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring OSPF to Make Routing Devices Appear Overloaded on page 4114</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source</a></li> </ul>


## preference (Protocols OSPF)

<b>Syntax</b>	<code>preference preference;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Set the route preference for OSPF internal routes.
<b>Options</b>	<p><b>preference</b>—Preference value.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> 10</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Controlling OSPF Route Preferences on page 4111</a></li> <li>• <a href="#">external-preference on page 4222</a></li> </ul>

## prefix-export-limit (Protocols OSPF)

<b>Syntax</b>	<code>prefix-export-limit <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   <a href="#">ospf3</a>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf <a href="#">topology</a> (default   ipv4-multicast   <i>name</i>)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0.</p> <p>Support for Multitopology Routing introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Configure a limit to the number of prefixes exported into OSPF.
<b>Options</b>	<p><b><i>number</i></b>—Prefix limit.</p> <p><b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>)</p> <p><b>Default:</b> None</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Limiting the Number of Prefixes Exported to OSPF on page 4100</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing Based on Applications</a></li> <li>• <a href="#">Example: Configuring Multitopology Routing Based on a Multicast Source</a></li> </ul>

## reference-bandwidth (Protocols OSPF)

<b>Syntax</b>	<code>reference-bandwidth <i>reference-bandwidth</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<b>ospf</b>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:</p> $\text{cost} = \text{ref-bandwidth} / \text{bandwidth}$
<b>Options</b>	<p><b><i>reference-bandwidth</i></b>—Reference bandwidth, in bits per second.</p> <p><b>Range:</b> 9600 through 1,000,000,000,000 bits</p> <p><b>Default:</b> 100 Mbps (100,000,000 bits)</p>
<div>  <p><b>NOTE:</b> The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the metric statement.</p> </div>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Controlling the Cost of Individual OSPF Network Segments on page 4105</a></li> <li>• <i>metric</i></li> </ul>

## rib-group (Protocols OSPF)

<b>Syntax</b>	<code>rib-group group-name;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (<a href="#">ospf</a>   ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.
<b>Options</b>	<i>group-name</i> —Name of the routing table group.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Exporting Specific Routes from One Routing Table Into Another Routing Table</i></li> <li>• <i>Example: Importing Direct and Static Routes Into a Routing Instance</i></li> <li>• <i>Understanding Multiprotocol BGP</i></li> <li>• <a href="#">interface-routes on page 2989</a></li> <li>• <a href="#">rib-group on page 3031</a></li> </ul>

## topology (OSPF)

---

<b>Syntax</b>	<pre>topology (default   ipv4-multicast   <i>name</i>) {     spf-options {         delay <i>milliseconds</i>;         holddown <i>milliseconds</i>;         rapid-runs <i>number</i>;     }     topology-id <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>ospf</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>ospf</b>], [edit protocols <b>ospf</b>], [edit routing-instances <i>routing-instance-name</i> protocols <b>ospf</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Enable a topology for OSPF multitenancy routing. You must first configure one or more topologies under the <b>[edit routing-options]</b> hierarchy level.</p>
<b>Options</b>	<p><b>default</b>—Name of the default topology. This topology is automatically created, and all routes that correspond to it are automatically added to the <b>inet.0</b> routing table. You can modify certain default parameters, such as for the SPF algorithm.</p> <p><b>ipv4-multicast</b>—Name of the topology for IPv4 multicast traffic.</p> <p><b><i>name</i></b>—Name of a topology you configured at the <b>[edit routing-options]</b> hierarchy level to create a topology for a specific type of traffic, such as voice or video.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Multitenancy Routing Based on Applications</i></li><li>• <i>Example: Configuring Multitenancy Routing Based on a Multicast Source</i></li></ul>

## traceoptions (Protocols OSPF)

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   (<b>ospf</b>   ospf3)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit protocols (<b>ospf</b>   ospf3)], [edit protocols ospf3 realm (ipv4-unicast   ipv4-multicast   ipv6-multicast)], [edit routing-instances <i>routing-instance-name</i> protocols (<b>ospf</b>   ospf3)], [edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast     ipv4-multicast   ipv6-multicast)]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2.</p> <p>Support for the <b>realm</b> statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure OSPF protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>



**NOTE:** The **traceoptions** statement is not supported on QFabric systems.

<b>Default</b>	The default OSPF protocol-level tracing options are those inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place OSPF tracing output in the file <b>ospf-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p>

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP.
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis. Specific to the Juniper Networks implementation of OSPF, Junos OS performs LSA analysis before running the shortest-path-first (SPF) algorithm. LSA analysis helps to speed the calculations performed by the SPF algorithm.
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **restart-signaling**—(OSPFv2 only) Restart-signaling graceful restart events.
- **spf**—Shortest-path-first (SPF) calculations.

#### Global Tracing Flags



- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations. If you do not specify this option, only unusual or abnormal operations are traced.
- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Tracing OSPF Protocol Traffic on page 4200</a></li> </ul>

## traffic-engineering (OSPF)

---

<b>Syntax</b>	<pre>traffic-engineering {   &lt;advertise-unnumbered-interfaces&gt;;   &lt;credibility-protocol-preference&gt;;   ignore-lsp-metrics;   multicast-rpf-routes;   no-topology;   shortcuts {     lsp-metric-into-summary;   } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ( <b>ospf</b>   ospf3)], [edit protocols ( <b>ospf</b>   ospf3)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>multicast-rpf-routes</b> option introduced in Junos OS Release 7.5. <b>advertise-unnumbered-interfaces</b> option introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for OSPFv3 ( <b>ospf3</b> ) introduced in Junos OS Release 9.4. Support for OSPFv3 ( <b>ospf3</b> ) introduced in Junos OS Release 9.4 for EX Series switches. <b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4. <b>credibility-protocol-preference</b> statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable the OSPF traffic engineering features.
<b>Default</b>	Traffic engineering support is disabled.
<b>Options</b>	<p><b>advertise-unnumbered-interfaces</b>—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. This statement must be included on both ends of an unnumbered link to allow an ingress LER to update the link in its traffic engineering database and use it for CSPF calculations. The link-local identifier is then used by RSVP to signal unnumbered interfaces as defined in RFC 3477.</p> <p><b>credibility-protocol-preference</b>—(Optional) (OSPFv2 only) Use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior, in which the traffic engineering database prefers IS-IS routes even if OSPF routes are configured with a lower, that is, preferred, preference value. For example, OSPF routes have a default preference value of 10, whereas IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, whereas IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.</p>

**multicast-rpf-routes**—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table. The **inet.2** routing table consists of unicast routes used for multicast RPF lookup. RPF is an antispoofing mechanism used to check whether the packet is coming in on an interface that is also sending data back to the packet source.



**NOTE:** You must enable OSPF traffic engineering shortcuts to use the **multicast-rpf-routes** statement. You must not allow LSP advertisements into OSPF when configuring the **multicast-rpf-routes** statement.

**no-topology**—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.



**CAUTION:** When the OSPF traffic engineering configuration is considerably modified, the routing table entries are deleted and the routing table is recreated. Changes to configuration that can cause this behavior include enabling or disabling:

- Traffic engineering shortcuts
- IGP shortcuts
- LDP tunneling
- Multiprotocol LSP
- Advertise summary metrics
- Multicast RPF routes

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Enabling OSPF Traffic Engineering Support on page 4153</a></li> </ul>
------------------------------	-------------------------------------------------------------------------------------------------------------------------------------



## CHAPTER 49

# Administration

- [Routine Monitoring on page 4243](#)
- [Operational Commands on page 4243](#)

## Routine Monitoring

---

- [Monitoring OSPF Routing Information on page 4243](#)

### Monitoring OSPF Routing Information

**Purpose** Use the monitoring functionality to monitor OSPF routing information on routing devices.

**Action** To view OSPF routing information in the CLI, enter the following CLI commands:

- `show ospf neighbor`
- `show ospf interface`
- `show ospf statistics`

**Related Documentation**

- [show \(ospf | ospf3\) interface on page 4270](#)
- [clear \(ospf | ospf3\) neighbor on page 4250](#)
- [show \(ospf | ospf3\) statistics on page 4298](#)

## Operational Commands

---

- `clear (ospf | ospf3) database`
- `clear (ospf | ospf3) database-protection`
- `clear (ospf | ospf3) io-statistics`
- `clear (ospf | ospf3) neighbor`
- `clear (ospf | ospf3) statistics`
- `clear (ospf | ospf3) overload`
- `show (ospf | ospf3) backup coverage`
- `show (ospf | ospf3) backup neighbor`

- `show ospf context-identifier`
- `show ospf database`
- `show (ospf | ospf3) interface`
- `show (ospf | ospf3) io-statistics`
- `show (ospf | ospf3) log`
- `show (ospf | ospf3) neighbor`
- `show (ospf | ospf3) overview`
- `show (ospf | ospf3) route`
- `show (ospf | ospf3) statistics`

## clear (ospf | ospf3) database

**List of Syntax**    [Syntax on page 4245](#)  
                           [Syntax \(EX Series Switch and QFX Series\) on page 4245](#)

**Syntax**    clear (ospf | ospf3) database  
                   <advertising-router (*router-id* | self) >  
                   <area *area-id* >  
                   <asbrsummary >  
                   <external >  
                   <instance *instance-name* >  
                   <inter-area-prefix >  
                   <inter-area-router >  
                   <intra-area-prefix >  
                   <link-local >  
                   <logical-system (all | *logical-system-name*) >  
                   <lsa-id *lsa-id* >  
                   <netsummary >  
                   <network >  
                   <nssa >  
                   <opaque-area >  
                   <purge >  
                   <realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) >  
                   <router >

**Syntax (EX Series Switch and QFX Series)**    clear (ospf | ospf3) database  
                                                           <advertising-router (*router-id* | self) >  
                                                           <area *area-id* >  
                                                           <asbrsummary >  
                                                           <external >  
                                                           <instance *instance-name* >  
                                                           <inter-area-prefix >  
                                                           <inter-area-router >  
                                                           <intra-area-prefix >  
                                                           <link-local >  
                                                           <lsa-id *lsa-id* >  
                                                           <netsummary >  
                                                           <network >  
                                                           <nssa >  
                                                           <opaque-area >  
                                                           <purge >  
                                                           <router >

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   **advertising-router** *router-id*, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.  
                                   **area** *area-id* option added in Junos OS Release 8.3.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   **realm** option added in Junos OS Release 9.2.  
                                   **advertising-router** (*router-id* | **self**) option added in Junos OS Release 9.5.  
                                   **advertising-router** (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.3 for the QFX Series.

**purge** option (and all options that are dependent on the **purge** option) hidden in Junos OS Release 13.3.

**Description** With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine.



**CAUTION:** You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries. This command is useful only for testing. Use it with care, because it causes significant network disruption.

**Options** **none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

**advertising-router (router-id | self)**—(Hidden) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

**area area-id**—(Optional) Discard entries for the LSAs in the specified area.

**asbrsummary**—(Optional) Discard summary AS boundary router LSA entries.

**external**—(Optional) Discard external LSAs.

**instance instance-name**—(Optional) Delete or discard entries for the specified routing instance only.

**inter-area-prefix**—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

**inter-area-router**—(OSPFv3 only) (Optional) Discard interarea router LSAs.

**intra-area-prefix**—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

**logical-system (all | logical-system-name)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**link-local**—(Optional) Delete link-local LSAs.

**lsa-id lsa-id**—(Optional) Discard the LSA entries with the specified LSA identifier.

**netsummary**—(Hidden) Discard summary network LSAs.

**network**—(Hidden) Discard network LSAs.

**nssa**—(Hidden) Discard not-so-stubby area (NSSA) LSAs.

**opaque-area**—(Hidden) Discard opaque area-scope LSAs.



**purge**—(Hidden) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

**realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)**—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**router**—(Hidden) Discard router LSAs.

**Required Privilege Level**

clear

**Related Documentation**

- [show ospf database on page 4262](#)
- [show ospf3 database](#)

**List of Sample Output** [clear ospf database on page 4247](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ospf database

```
user@host> clear ospf database
```

## clear (ospf | ospf3) database-protection

---

<b>Syntax</b>	clear (ospf   ospf3) database-protection <instance <i>instance-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear the Open Shortest Path First (OSPF) link-state database from its isolated state. Reset the ignore count, ignore timer, and reset timer, and resume normal operations.
<b>Options</b>	<b>instance <i>instance-name</i></b> —(Optional) Clear the OSPF link-state database for the specified routing instance only.
<b>Required Privilege Level</b>	clear
<b>Output Fields</b>	This command produces no output.

### Sample Output

#### clear ospf database-protection

```
user@host> clear ospf database-protection
```

## clear (ospf | ospf3) io-statistics

---

<b>List of Syntax</b>	<a href="#">Syntax on page 4249</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 4249</a>
<b>Syntax</b>	clear (ospf   ospf3) io-statistics <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switch and QFX Series)</b>	clear (ospf   ospf3) io-statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear Open Shortest Path First (OSPF) input and output statistics.
<b>Options</b>	<b>none</b> —Clear OSPF input and output statistics.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear ospf io-statistics on page 4249</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

## clear (ospf | ospf3) neighbor

---

List of Syntax	<a href="#">Syntax on page 4250</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 4250</a>
Syntax	<pre>clear (ospf   ospf3) neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear (ospf   ospf3) neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;neighbor&gt;</pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>realm</b> option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Tear down Open Shortest Path First (OSPF) neighbor connections.
Options	<p><b>none</b>—Tear down OSPF connections with all neighbors for all routing instances.</p> <p><b>area <i>area-id</i></b>—(Optional) Tear down neighbor connections for the specified area only.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Tear down neighbor connections for the specified routing instance only.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Tear down neighbor connections for the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear the state of the specified neighbor only.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">show (ospf   ospf3) neighbor on page 4281</a></li></ul>
List of Sample Output	<a href="#">clear ospf neighbor on page 4251</a>

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

### Sample Output

clear ospf neighbor

```
user@host> clear ospf neighbor
```

## clear (ospf | ospf3) statistics

---

List of Syntax	<a href="#">Syntax on page 4252</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 4252</a>
Syntax	<code>clear (ospf   ospf3) statistics</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</code>
Syntax (EX Series Switch and QFX Series)	<code>clear (ospf   ospf3) statistics</code> <code>&lt;instance <i>instance-name</i>&gt;</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>realm</b> option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) statistics.
Options	<b>none</b> —Clear OSPF statistics.  <b>instance <i>instance-name</i></b> —(Optional) Clear statistics for the specified routing instance only.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b> —(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">show (ospf   ospf3) statistics on page 4298</a></li></ul>
List of Sample Output	<a href="#">clear ospf statistics on page 4252</a>
Output Fields	See <a href="#">show (ospf   ospf3) statistics</a> for an explanation of output fields.

### Sample Output

#### clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

Packet type	Total	Last 5 seconds
-------------	-------	----------------

	Sent	Received	Sent	Received
Hello	3254	2268	3	1
DbD	41	46	0	0
LSReq	8	7	0	0
LSUpdate	212	154	0	0
LSAck	65	98	0	0

DBDs retransmitted	:	3, last 5 seconds	:	0
LSAs flooded	:	12, last 5 seconds	:	0
LSAs flooded high-prio	:	0, last 5 seconds	:	0
LSAs retransmitted	:	0, last 5 seconds	:	0
LSAs transmitted to nbr:	:	3, last 5 seconds	:	0
LSAs requested	:	5, last 5 seconds	:	0
LSAs acknowledged	:	19, last 5 seconds	:	0

Flood queue depth	:	0
Total rexmit entries	:	0
db summaries	:	0
lsreq entries	:	0

Receive errors:

626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

Packet type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	3	1	3	1
DbD	0	0	0	0
LSReq	0	0	0	0
LSUpdate	0	0	0	0
LSAck	0	0	0	0

DBDs retransmitted	:	0, last 5 seconds	:	0
LSAs flooded	:	0, last 5 seconds	:	0
LSAs flooded high-prio	:	0, last 5 seconds	:	0
LSAs retransmitted	:	0, last 5 seconds	:	0
LSAs transmitted to nbr:	:	0, last 5 seconds	:	0
LSAs requested	:	0, last 5 seconds	:	0
LSAs acknowledged	:	0, last 5 seconds	:	0

Flood queue depth	:	0
Total rexmit entries	:	0
db summaries	:	0
lsreq entries	:	0

Receive errors:

None

## clear (ospf | ospf3) overload

---

<b>List of Syntax</b>	<a href="#">Syntax on page 4254</a> <a href="#">Syntax (EX Series Switches) on page 4254</a>
<b>Syntax</b>	<code>clear (ospf   ospf3) overload</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>
<b>Syntax (EX Series Switches)</b>	<code>clear (ospf   ospf3) overload</code> <code>&lt;instance <i>instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).
<b>Options</b>	<b>none</b> —Clear the overload bit and rebuild LSAs for all routing instances.  <b>instance <i>instance-name</i></b> —(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>List of Sample Output</b>	<a href="#">clear ospf overload on page 4254</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ospf overload

```
user@host> clear ospf overload
```



## show (ospf | ospf3) backup coverage

<b>Syntax</b>	<pre>show (ospf   ospf3) backup coverage &lt;instance <i>instance-name</i>&gt; &lt; logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-unicast   ipv6-unicast)&gt; &lt;topology <i>topology-name</i>&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show (ospf   ospf3) backup coverage &lt;instance <i>instance-name</i>&gt; &lt;topology <i>topology-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display information about the level of backup coverage available for all the nodes and prefixes in the network.
<b>Options</b>	<p><b>none</b>—Display information about the level backup coverage for all OSPF routing instances in all logical systems.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.</p> <p><b>realm (ipv4-unicast   ipv6-unicast)</b>—(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.</p> <p><b>topology (default   <i>topology-name</i>)</b>—(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">show (ospf   ospf3) backup lsp</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ospf backup coverage on page 4256</a> <a href="#">show ospf3 backup coverage on page 4256</a>
<b>Output Fields</b>	<p><a href="#">Table 324 on page 4255</a> lists the output fields for the <b>show (ospf   ospf3) backup coverage</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 324: show (ospf | ospf3) backup coverage Output Fields**

Field Name	Field Description
Node Coverage	Information about backup coverage for each OSPF node.
Area	Area number. Area 0.0.0.0 is the backbone.

Table 324: show (ospf | ospf3) backup coverage Output Fields (*continued*)

Field Name	Field Description
<b>Covered Nodes</b>	Number of nodes for which backup coverage is available.
<b>Total Nodes</b>	Total number of OSPF nodes.
<b>Route Coverage</b>	Information about backup coverage for each type of OSPF route.
<b>Path Type</b>	Type of OSPF path: <b>Intra</b> , <b>Inter</b> , <b>Ext1</b> , <b>Ext2</b> , and <b>All</b> .
<b>Covered Routes</b>	For each path type, the number of routes for which backup coverage is available.
<b>Total Routes</b>	For each path type, the total number of configured routes.
<b>Percent Covered</b>	For all nodes and for each path type, the percentage for which backup coverage is available.

## Sample Output

### show ospf backup coverage

```

user@host> show ospf backup coverage
Topology default coverage:

Node Coverage:

Area              Covered  Total  Percent
                  Nodes   Nodes  Covered
0.0.0.0           4        5    80.00%

Route Coverage:

Path Type  Covered  Total  Percent
          Routes Routes  Covered
Intra      8        14    57.14%
Inter      0         0   100.00%
Ext1       0         0   100.00%
Ext2       1         1   100.00%
All        9        15    60.00%

```

### show ospf3 backup coverage

```

user @host > show ospf3 backup coverage
show ospf3 backup coverage
Node Coverage:

Area              Covered  Total  Percent
                  Nodes   Nodes  Covered
0.0.0.0           4        5    80.00%

Route Coverage:

Path Type  Covered  Total  Percent
          Routes Routes  Covered

```

Intra	4	6	66.67%
Inter	0	0	100.00%
Ext1	0	0	100.00%
Ext2	1	1	100.00%
All	5	7	71.43%

## show (ospf | ospf3) backup neighbor

<b>Syntax</b>	<pre>show (ospf   ospf3) backup neighbor &lt;area <i>area-id</i>&gt; &lt;instance (default   <i>instance-name</i>)&gt; &lt;logical-system (default   ipv4-multicast   <i>logical-system-name</i>)&gt; &lt;topology (default   ipv4-multicast   <i>topology-name</i>)&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show (ospf   ospf3) backup neighbor &lt;area <i>area-id</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;topology (default   ipv4-multicast   <i>topology-name</i>)&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display the neighbors through which direct next hops for the backup paths are available.
<b>Options</b>	<p><b>none</b>—Display all neighbors that have direct next hops for backup paths.</p> <p><b>area <i>area-id</i></b>—(Optional) Display the area information.</p> <p><b>instance (default   <i>instance-name</i>)</b>—(Optional) Display information about the default routing instance or a particular routing instance.</p> <p><b>logical-system (default   ipv4-multicast   <i>logical-system-name</i>)</b>—(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.</p> <p><b>topology (default   ipv4-multicast   <i>topology-name</i>)</b>—(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">show (ospf   ospf3) backup spf</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ospf backup neighbor on page 4259</a>
<b>Output Fields</b>	<a href="#">Table 325 on page 4258</a> lists the output fields for the <b>show (ospf   ospf3) backup neighbor</b> command. Output fields are listed in the approximate order in which they appear.

**Table 325: show (ospf | ospf3) backup neighbor Output Fields**

Field Name	Field Description	Level of Output
Neighbor to Self Metric	Metric from the backup neighbor to the OSPF node.	All levels
Self to Neighbor Metric	Metric from the OSPF node to the backup neighbor.	All levels

Table 325: show (ospf |ospf3) backup neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Direct next-hop	Interface and address of the direct next hop.	All levels

## Sample Output

### show ospf backup neighbor

```
user@host> show ospf backup neighbor
Topology default backup neighbors:

Area 0.0.0.5 backup neighbors:

10.0.0.5
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/0/0.111 via 10.0.175.5

10.0.0.6
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```

## show ospf context-identifier

---

List of Syntax	<a href="#">Syntax on page 4260</a> <a href="#">Syntax (EX Series Switches and QFX Series) on page 4260</a>
Syntax	<pre>show ospf context-identifier &lt;brief   detail&gt; &lt;area <i>area-id</i>&gt; &lt;context-id&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf context-identifier &lt;brief   detail&gt; &lt;area <i>area-id</i>&gt; &lt;context-id&gt; &lt;instance <i>instance-name</i>&gt;</pre>
Release Information	Command introduced in Junos OS Release 10.4. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection.
Options	<p><b>none</b>—Display information about all context identifiers.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>area <i>area-id</i></b>—(Optional) Display information about the context identifier for the specified area.</p> <p><b>context-id</b>—(Optional) Display information about the specified context identifier.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the context identifier for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li><i>egress-protection (Layer 2 circuit)</i> in the <i>Junos OS VPNs Library for Routing Devices</i></li><li><i>egress-protection (MPLS)</i> in the <i>Junos OS VPNs Library for Routing Devices</i></li></ul>
List of Sample Output	<a href="#">show ospf context-identifier on page 4261</a> <a href="#">show ospf context-identifier detail on page 4261</a>
Output Fields	<a href="#">Table 326 on page 4261</a> lists the output fields for the <b>show ospf context-identifier</b> command. Output fields are listed in the approximate order in which they appear.

Table 326: show ospf context-identifier Output Fields

Field Name	Field Description	Level of Output
<b>Context</b>	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices.	All levels
<b>Status</b>	State of the path: <b>active</b> or <b>inactive</b> .	All levels
<b>Metric</b>	Advertised OSPF metric.	All levels
<b>Area</b>	OSPF area number.	All levels
<b>Other Advertisements</b>	Other advertisements received by the OSPF node: <ul style="list-style-type: none"> <li>• <b>Advertising router</b>—Address of the device that sent the advertisement.</li> <li>• <b>Type</b>—Type of OSPF path: <b>inter-area</b> and <b>stub</b>.</li> <li>• <b>Metric</b>—Advertised OSPF metric.</li> <li>• <b>None</b>—No additional advertisements were received by the OSPF node.</li> </ul>	<b>detail</b>

## Sample Output

### show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

### show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
Other Advertisements:
Advertising router: 8.8.8.103
Type: stub link
Metric: 65534
```

```
show ospf database
```

<b>List of Syntax</b> <a href="#">Syntax on page 4262</a> <a href="#">Syntax (EX Series Switches and QFX Series) on page 4262</a>	
<b>Syntax</b>	<pre>show ospf database &lt;brief   detail   extensive   summary&gt; &lt;advertising-router (address   self)&gt; &lt;area area-id&gt; &lt;asbrsummary&gt; &lt;external&gt; &lt;instance instance-name&gt; &lt;link-local&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;lsa-id lsa-id&gt; &lt;netsummary&gt; &lt;network&gt; &lt;nssa&gt; &lt;opaque-area&gt; &lt;router&gt;</pre>
<b>Syntax (EX Series Switches and QFX Series)</b>	<pre>show ospf database &lt;brief   detail   extensive   summary&gt; &lt;advertising-router (address   self)&gt; &lt;area area-id&gt; &lt;asbrsummary&gt; &lt;external&gt; &lt;instance instance-name&gt; &lt;link-local&gt; &lt;lsa-id lsa-id&gt; &lt;netsummary&gt; &lt;network&gt; &lt;nssa&gt; &lt;opaque-area&gt; &lt;router&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>advertising-router self (address   self)</b> option introduced in Junos OS Release 9.5.</p> <p><b>advertising-router self (address   self)</b> option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.
<b>Options</b>	<p><b>none</b>—Display standard information about entries in the OSPFv2 link-state database for all routing instances.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>advertising-router (address   self)</b>—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p>



**area** *area-id*—(Optional) Display the LSAs in a particular area.

**asbrsummary**—(Optional) Display summary AS boundary router LSA entries.

**external**—(Optional) Display external LSAs.

**instance** *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

**link-local**—(Optional) Display information about link-local LSAs.

**logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**lsa-id** *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

**netsummary**—(Optional) Display summary network LSAs.

**network**—(Optional) Display information about network LSAs.

**nssa**—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

**opaque-area**—(Optional) Display opaque area-scope LSAs.

**router**—(Optional) Display information about router LSAs.

**Required Privilege Level**

view

**Related Documentation**

- [clear \(ospf | ospf3\) database on page 4245](#)

**List of Sample Output**

[show ospf database on page 4265](#)  
[show ospf database brief on page 4265](#)  
[show ospf database detail on page 4265](#)  
[show ospf database extensive on page 4267](#)  
[show ospf database summary on page 4269](#)

**Output Fields**

[Table 327 on page 4263](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

**Table 327: show ospf database Output Fields**

Field Name	Field Description	Level of Output
<b>area</b>	Area number. Area 0.0.0.0 is the backbone area.	All levels
<b>Type</b>	Type of link advertisement: <b>ASBRSum</b> , <b>Extern</b> , <b>Network</b> , <b>NSSA</b> , <b>OpaqArea</b> , <b>Router</b> , or <b>Summary</b> .	All levels
<b>ID</b>	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
<b>Adv Rtr</b>	Address of the routing device that sent the advertisement.	All levels

Table 327: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Seq</b>	Link sequence number of the advertisement.	All levels
<b>Age</b>	Time elapsed since the LSA was originated, in seconds.	All levels
<b>Opt</b>	Optional OSPF capabilities associated with the LSA.	All levels
<b>Cksum</b>	Checksum value of the LSA.	All levels
<b>Len</b>	Length of the advertisement, in bytes.	All levels
<b>Router</b>	Router link-state advertisement information: <ul style="list-style-type: none"> <li><b>bits</b>—Flags describing the routing device that generated the LSP.</li> <li><b>link count</b>—Number of links in the advertisement.</li> <li><b>id</b>—ID of a routing device or subnet on the link.</li> <li><b>data</b>—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP.</li> <li><b>type</b>—Type of link. It can be <b>PointToPoint</b>, <b>Transit</b>, <b>Stub</b>, or <b>Virtual</b>.</li> <li><b>TOS count</b>—Number of type-of-service (ToS) entries in the advertisement.</li> <li><b>TOS 0 metric</b>—Metric for ToS 0.</li> <li><b>TOS</b>—Type-of-service (ToS) value.</li> <li><b>metric</b>—Metric for the ToS.</li> </ul>	<b>detail extensive</b>
<b>Network</b>	Network link-state advertisement information: <ul style="list-style-type: none"> <li><b>mask</b>—Network mask.</li> <li><b>attached router</b>—ID of the attached neighbor.</li> </ul>	<b>detail extensive</b>
<b>Summary</b>	Summary link-state advertisement information: <ul style="list-style-type: none"> <li><b>mask</b>—Network mask.</li> <li><b>TOS</b>—Type-of-service (ToS) value.</li> <li><b>metric</b>—Metric for the ToS.</li> </ul>	<b>detail extensive</b>
<b>Gen timer</b>	How long until the LSA is regenerated.	<b>extensive</b>
<b>Aging timer</b>	How long until the LSA expires.	<b>extensive</b>
<b>Installed <i>hh:mm:ss</i> ago</b>	How long ago the route was installed.	<b>extensive</b>
<b>expires in <i>hh:mm:ss</i></b>	How long until the route expires.	<b>extensive</b>
<b>sent <i>hh:mm:ss</i> ago</b>	How long ago the LSA was sent.	<b>extensive</b>
<b>Last changed <i>hh:mm:ss</i> ago</b>	How long ago the route was changed.	<b>extensive</b>

Table 327: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Change count	Number of times the route has changed.	extensive
Ours	Indicates that this is a local advertisement.	extensive
Router LSAs	Number of router link-state advertisements in the link-state database.	summary
Network LSAs	Number of network link-state advertisements in the link-state database.	summary
Summary LSAs	Number of summary link-state advertisements in the link-state database.	summary
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	summary

## Sample Output

### show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.70.103  10.255.70.103 0x80000002   215  0x20 0x4112  48
Router     *10.255.71.242 10.255.71.242 0x80000002   214  0x20 0x11b1  48
Summary    *23.1.1.0      10.255.71.242 0x80000002   172  0x20 0x6d72  28
Summary    *24.1.1.0      10.255.71.242 0x80000002   177  0x20 0x607e  28
NSSA       *33.1.1.1      10.255.71.242 0x80000002   217  0x28 0x73bd  36

      OSPF link state database, Area 0.0.0.2
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.71.52   10.255.71.52   0x80000004   174  0x20 0xd021  36
Router     *10.255.71.242 10.255.71.242 0x80000003   173  0x20 0xe191  36
Network    *23.1.1.1      10.255.71.242 0x80000002   173  0x20 0x9c76  32
Summary    *12.1.1.0      10.255.71.242 0x80000001   217  0x20 0xfeec  28
Summary    *24.1.1.0      10.255.71.242 0x80000002   177  0x20 0x607e  28
NSSA       *33.1.1.1      10.255.71.242 0x80000001   222  0x28 0xe047  36

      OSPF link state database, Area 0.0.0.3
  Type      ID            Adv Rtr      Seq          Age  Opt  Cksum  Len
Router     10.255.71.238   10.255.71.238 0x80000003   179  0x20 0x3942  36
Router     *10.255.71.242 10.255.71.242 0x80000003   177  0x20 0xf37d  36
Network    *24.1.1.1      10.255.71.242 0x80000002   177  0x20 0xc591  32
Summary    *12.1.1.0      10.255.71.242 0x80000001   217  0x20 0xfeec  28
Summary    *23.1.1.0      10.255.71.242 0x80000002   172  0x20 0x6d72  28
NSSA       *33.1.1.1      10.255.71.242 0x80000001   222  0x28 0xeb3b  36

```

### show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 4265](#).

### show ospf database detail

```
user@host> show ospf database detail
```

```

    OSPF link state database, Area 0.0.0.1
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.70.103  10.255.70.103  0x80000002  261  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000002  260  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary *23.1.1.0      10.255.71.242  0x80000002  218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0      10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA *33.1.1.1        10.255.71.242  0x80000002  263  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

```

    OSPF link state database, Area 0.0.0.2
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.52   10.255.71.52   0x80000004  220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000003  219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *23.1.1.1      10.255.71.242  0x80000002  219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0      10.255.71.242  0x80000001  263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0      10.255.71.242  0x80000002  223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA *33.1.1.1        10.255.71.242  0x80000001  268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

```

```

    OSPF link state database, Area 0.0.0.3
    Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router 10.255.71.238  10.255.71.238  0x80000003  225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router *10.255.71.242  10.255.71.242  0x80000003  223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1      10.255.71.242  0x80000002  223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242

```

```

    attached router 10.255.71.238
Summary *12.1.1.0      10.255.71.242    0x80000001    263    0x20 0xfeec    28
    mask 255.255.255.0
    TOS 0x0, metric 1
Summary *23.1.1.0      10.255.71.242    0x80000002    218    0x20 0x6d72    28
    mask 255.255.255.0
    TOS 0x0, metric 1
NSSA  *33.1.1.1        10.255.71.242    0x80000001    268    0x28 0xeb3b    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

### show ospf database extensive

```

user@host> show ospf database extensive
    OSPF link state database, Area 0.0.0.1
Type      ID          Adv Rtr      Seq      Age    Opt  Cksum  Len
Router    10.255.70.103    10.255.70.103  0x80000002  286    0x20 0x4112  48
    bits 0x0, link count 2
    id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
    TOS count 0, TOS 0 metric 1
    id 12.1.1.0, data 255.255.255.0, Type Stub (3)
    TOS count 0, TOS 0 metric 1
    Aging timer 00:55:14
    Installed 00:04:43 ago, expires in 00:55:14
    Last changed 00:04:43 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000002    285    0x20 0x11b1  48
    bits 0x3, link count 2
    id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
    TOS count 0, TOS 0 metric 1
    id 12.1.1.0, data 255.255.255.0, Type Stub (3)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:15
    Aging timer 00:55:15
    Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
    Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243    0x20 0x6d72    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:57
    Aging timer 00:55:57
    Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002    248    0x20 0x607e    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1        10.255.71.242    0x80000002    288    0x28 0x73bd    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
    Last changed 00:04:48 ago, Change count: 2, Ours

    OSPF link state database, Area 0.0.0.2
Type      ID          Adv Rtr      Seq      Age    Opt  Cksum  Len
Router    10.255.71.52     10.255.71.52     0x80000004    245    0x20 0xd021  36
    bits 0x0, link count 1

```

```

id 23.1.1.1, data 23.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:55
Installed 00:04:02 ago, expires in 00:55:55
Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 244 0x20 0xe191 36
bits 0x3, link count 1
id 23.1.1.1, data 23.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1 10.255.71.242 0x80000002 244 0x20 0x9c76 32
mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.52
Gen timer 00:45:56
Aging timer 00:55:56
Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0 10.255.71.242 0x80000001 288 0x20 0xfeec 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0 10.255.71.242 0x80000002 248 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA *33.1.1.1 10.255.71.242 0x80000001 293 0x28 0xe047 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
Last changed 00:04:53 ago, Change count: 1, Ours

OSPF link state database, Area 0.0.0.3
Type ID Adv Rtr Seq Age Opt Cksum Len
Router 10.255.71.238 10.255.71.238 0x80000003 250 0x20 0x3942 36
bits 0x0, link count 1
id 24.1.1.1, data 24.1.1.2, Type Transit (2)
TOS count 0, TOS 0 metric 1
Aging timer 00:55:50
Installed 00:04:07 ago, expires in 00:55:50
Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242 10.255.71.242 0x80000003 248 0x20 0xf37d 36
bits 0x3, link count 1
id 24.1.1.1, data 24.1.1.1, Type Transit (2)
TOS count 0, TOS 0 metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1 10.255.71.242 0x80000002 248 0x20 0xc591 32

```

```

mask 255.255.255.0
attached router 10.255.71.242
attached router 10.255.71.238
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288  0x20 0xfeec  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1      10.255.71.242    0x80000001    293  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours

```

#### show ospf database summary

```

user@host> show ospf database summary
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:

```

```
show (ospf | ospf3) interface
```

List of Syntax	<a href="#">Syntax on page 4270</a> <a href="#">Syntax (EX Series Switches and QFX Series) on page 4270</a>
Syntax	<pre>show (ospf   ospf3) interface &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;<i>interface-name</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf   ospf3) interface &lt;brief   detail   extensive&gt; &lt;area <i>area-id</i>&gt; &lt;<i>interface-name</i>&gt; &lt;instance <i>instance-name</i>&gt;</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>area</b> option introduced in Junos OS Release 9.2.</p> <p><b>area</b> option introduced in Junos OS Release 9.2 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the status of OSPF interfaces.
Options	<p><b>none</b>—Display standard information about the status of all OSPF interfaces for all routing instances</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>area <i>area-id</i></b>—(Optional) Display information about the interfaces that belong to the specified area.</p> <p><b><i>interface-name</i></b>—(Optional) Display information for the specified interface.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view



**List of Sample Output** [show ospf interface brief on page 4273](#)  
[show ospf interface detail on page 4273](#)  
[show ospf3 interface detail on page 4273](#)  
[show ospf interface detail\(When Multiarea Adjacency Is Configured\) on page 4273](#)  
[show ospf interface area area-id on page 4275](#)  
[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 4275](#)  
[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 4275](#)

**Output Fields** [Table 328 on page 4271](#) lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

**Table 328: show (ospf | ospf3) interface Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
<b>State</b>	State of the interface: <b>BDR</b> , <b>Down</b> , <b>DR</b> , <b>DRother</b> , <b>Loop</b> , <b>PtToPt</b> , or <b>Waiting</b> .	All levels
<b>Area</b>	Number of the area that the interface is in.	All levels
<b>DR ID</b>	Address of the area's designated router.	All levels
<b>BDR ID</b>	Backup designated router for a particular subnet.	All levels
<b>Nbrs</b>	Number of neighbors on this interface.	All levels
<b>Type</b>	Type of interface: <b>LAN</b> , <b>NBMA</b> , <b>P2MP</b> , <b>P2P</b> , or <b>Virtual</b> .	<b>detail extensive</b>
<b>Address</b>	IP address of the neighbor.	<b>detail extensive</b>
<b>Mask</b>	Netmask of the neighbor.	<b>detail extensive</b>
<b>Prefix-length</b>	(OSPFv3) IPv6 prefix length, in bits.	<b>detail extensive</b>
<b>OSPF3-Intf-Index</b>	(OSPFv3) OSPF version 3 interface index.	<b>detail extensive</b>
<b>MTU</b>	Interface maximum transmission unit (MTU).	<b>detail extensive</b>
<b>Cost</b>	Interface cost (metric).	<b>detail extensive</b>
<b>DR addr</b>	Address of the designated router.	<b>detail extensive</b>
<b>BDR addr</b>	Address of the backup designated router.	<b>detail extensive</b>
<b>Adj count</b>	Number of adjacent neighbors.	<b>detail extensive</b>
<b>Secondary</b>	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	<b>detail extensive</b>

Table 328: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Flood Reduction</b>	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the <b>DoNotAge</b> bit set. As a result, LSAs are refreshed only when a change occurs.	<b>extensive</b>
<b>Priority</b>	Router priority used in designated router (DR) election on this interface.	<b>detail extensive</b>
<b>Flood list</b>	List of link-state advertisements (LSAs) that might be about to flood this interface.	<b>extensive</b>
<b>Ack list</b>	Acknowledgment list. List of pending acknowledgments on this interface.	<b>extensive</b>
<b>Descriptor list</b>	List of packet descriptors.	<b>extensive</b>
<b>Hello</b>	Configured value for the hello timer.	<b>detail extensive</b>
<b>Dead</b>	Configured value for the dead timer.	<b>detail extensive</b>
<b>Auth type</b>	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> <li>• <b>MD5</b>—The MD5 mechanism is configured in accordance with RFC 2328.</li> <li>• <b>None</b>—No authentication method is configured.</li> <li>• <b>Password</b>—A simple password (RFC 2328) is configured.</li> </ul>	<b>detail extensive</b>
<b>Topology</b>	(Multiarea adjacency) Name of topology: <b>default</b> or <b>name</b> .	
<b>LDP sync state</b>	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: <b>in sync</b> , <b>in holddown</b> , and <b>not supported</b> .	<b>extensive</b>
<b>reason</b>	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	<b>extensive</b>
<b>config holdtime</b>	(OSPFv2 and LDP synchronization) Configured value of the hold timer.  If the state is not synchronized, and the hold time is not infinity, the <b>remaining</b> field displays the number of seconds that remain until the configured hold timer expires.	<b>extensive</b>
<b>IPSec SA name</b>	(OSPFv2) Name of the IPSec security association name.	<b>detail extensive</b>
<b>Active key ID</b>	(OSPFv2 and MD5) Number from <b>0</b> to <b>255</b> that uniquely identifies an MD5 key.	<b>detail extensive</b>
<b>Start time</b>	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as <b>Start time 1970 Jan 01 00:00:00 PST</b> .	<b>detail extensive</b>

Table 328: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
ReXmit	Configured value for the Retransmit timer.	detail extensive
Stub, Not Stub, or Stub NSSA	Type of area.	detail extensive

## Sample Output

### show ospf interface brief

```

user@host> show ospf interface brief
Intf          State   Area      DR ID      BDR ID      Nbrs
at-5/1/0.0    PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
ge-2/3/0.0    DR      0.0.0.0    192.168.4.16 192.168.4.15 1
lo0.0         DR      0.0.0.0    192.168.4.16 0.0.0.0     0
so-0/0/0.0    Down    0.0.0.0    0.0.0.0    0.0.0.0     0
so-6/0/1.0    PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
so-6/0/2.0    Down    0.0.0.0    0.0.0.0    0.0.0.0     0
so-6/0/3.0    PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1

```

### show ospf interface detail

```

user@host> show ospf interface detail
Interface      State   Area      DR ID      BDR ID      Nbrs
fe-0/0/1.0     BDR    0.0.0.0    192.168.37.12 10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     0
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa

```

### show ospf3 interface detail

```

user@host> show ospf3 interface so-0/0/3.0 detail
Interface      State   Area      DR-ID      BDR-ID      Nbrs
so-0/0/3.0     PtToPt 0.0.0.0    0.0.0.0    0.0.0.0     1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub

```

### show ospf interface detail (When Multiarea Adjacency Is Configured)

```

user@host> show ospf interface detail
regress@router> show ospf interface detail
Interface      State   Area      DR ID      BDR ID      Nbrs
lo0.0          DR      0.0.0.0    10.255.245.2 0.0.0.0     0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 127.0.0.1, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None

```

```

Topology default (ID 0) -> Cost: 0
lo0.0          DR          0.0.0.0          10.255.245.2    0.0.0.0          0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
DR addr: 10.255.245.2, Adj count: 0, Priority: 128
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 0
so-0/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0          0.0.0.0          0.0.0.0          0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2          0.0.0.0          0.0.0.0          1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary

```

```

Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

#### show ospf interface area area-id

```

user@host> show ospf interface area 1.1.1.1
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1
so-1/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1

```

#### show ospf interface extensive (When Flooding Reduction Is Enabled)

```

user@host> show ospf interface extensive
Interface      State   Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0     PtToPt  0.0.0.0   0.0.0.0    0.0.0.0     0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

#### show ospf interface extensive (When LDP Synchronization Is Configured)

```

user@host> show ospf interface extensive
Interface      State   Area      DR ID      BDR ID
Nbrs
so-1/0/3.0     Down    0.0.0.0   0.0.0.0    0.0.0.0
0
Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddtime: 10 seconds, remaining: 1

```

## show (ospf | ospf3) io-statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 4276</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 4276</a>
<b>Syntax</b>	show (ospf   ospf3) io-statistics <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switch and QFX Series)</b>	show (ospf   ospf3) io-statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display Open Shortest Path First (OSPF) input and output statistics.
<b>Options</b>	none—Display OSPF input and output statistics.  logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear (ospf   ospf3) statistics on page 4252</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ospf io-statistics on page 4277</a>
<b>Output Fields</b>	Table 329 on page 4276 lists the output fields for the <b>show ospf io-statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 329: show (ospf | ospf3) io-statistics Output Fields**

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

## Sample Output

### show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1  
Receive errors:  
None
```

## show (ospf | ospf3) log

**List of Syntax** [Syntax on page 4278](#)  
[Syntax \(EX Series Switch and QFX Series\) on page 4278](#)

**Syntax** `show (ospf | ospf3) log`  
`<instance instance-name>`  
`<logical-system (all | logical-system-name)>`  
`<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>`  
`<topology topology-name>`

**Syntax (EX Series Switch and QFX Series)** `show (ospf | ospf3) log`  
`<instance instance-name>`  
`<topology topology-name>`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
**topology** option introduced in Junos OS Release 9.0.  
**topology** option introduced in Junos OS Release 9.0 for EX Series switches.  
**realm** option introduced in Junos OS Release 9.2.  
Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.

**Options** **none**—Display entries in the OSPF log of SPF calculations for all routing instances.

**instance *instance-name***—(Optional) Display entries for the specified routing instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**topology *topology-name***—(Optional) (OSPFv2 only) Display entries for the specified topology.

**realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)**—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**Required Privilege Level** view

**List of Sample Output** [show ospf log on page 4279](#)  
[show ospf log topology voice on page 4279](#)

**Output Fields** [Table 330 on page 4278](#) lists the output fields for the **show (ospf | ospf3) log** command. Output fields are listed in the approximate order in which they appear.

**Table 330: show (ospf | ospf3) log Output Fields**

Field Name	Field Description
<b>When</b>	Time, in weeks ( <b>w</b> ) and days ( <b>d</b> ), since the SPF calculation was made.



Table 330: show (ospf | ospf3) log Output Fields (*continued*)

Field Name	Field Description
Type	Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field.

## Sample Output

### show ospf log

```

user@host> show ospf log
When          Type          Elapsed
1w4d 17:25:58 Stub          0.000017
1w4d 17:25:58 SPF            0.000070
1w4d 17:25:58 Stub          0.000019
1w4d 17:25:58 Interarea     0.000054
1w4d 17:25:58 External      0.000005
1w4d 17:25:58 Cleanup       0.000203
1w4d 17:25:58 Total        0.000537
1w4d 17:24:48 SPF            0.000125
1w4d 17:24:48 Stub          0.000017
1w4d 17:24:48 SPF            0.000100
1w4d 17:24:48 Stub          0.000016
1w4d 17:24:48 Interarea     0.000056
1w4d 17:24:48 External      0.000005
1w4d 17:24:48 Cleanup       0.000238
1w4d 17:24:48 Total        0.000600
...

```

### show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

    Last instance of each event type
When          Type          Elapsed
00:06:11      SPF            0.000116
00:06:11      Stub          0.000114
00:06:11      Interarea     0.000126
00:06:11      External      0.000067
00:06:11      NSSA          0.000037
00:06:11      Cleanup       0.000186

    Maximum length of each event type
When          Type          Elapsed
00:13:43      SPF            0.000140
00:13:33      Stub          0.000116
00:13:43      Interarea     0.000128
00:13:33      External      0.000075
00:13:38      NSSA          0.000039
00:13:53      Cleanup       0.000657

    Last 100 events

```

When	Type	Elapsed
00:13:53	SPF	0.000090
00:13:53	Stub	0.000041
00:13:53	Interarea	0.000123
00:13:53	External	0.000040
00:13:53	NSSA	0.000038
00:13:53	Cleanup	0.000657
00:13:53	Total	0.001252
.		
.		
00:06:11	SPF	0.000116
00:06:11	Stub	0.000114
00:06:11	Interarea	0.000126
00:06:11	External	0.000067
00:06:11	NSSA	0.000037
00:06:11	Cleanup	0.000186
00:06:11	Total	0.000818

## show (ospf | ospf3) neighbor

**List of Syntax**    [Syntax on page 4281](#)  
                          [Syntax \(EX Series Switches and QFX Series\) on page 4281](#)

**Syntax**    `show (ospf | ospf3) neighbor`  
                  `<brief | detail | extensive>`  
                  `<area area-id>`  
                  `<instance (all | instance-name)>`  
                  `<interface interface-name>`  
                  `<logical-system (all | logical-system-name)>`  
                  `<neighbor>`  
                  `<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>`

**Syntax (EX Series Switches and QFX Series)**    `show (ospf | ospf3) neighbor`  
                  `<brief | detail | extensive>`  
                  `<area area-id>`  
                  `<instance (all | instance-name)>`  
                  `<interface interface-name>`  
                  `<neighbor>`

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                  Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                  **instance all** option introduced in Junos OS Release 9.1.  
                                  **instance all** option introduced in Junos OS Release 9.1 for EX Series switches.  
                                  **area**, **interface**, and **realm** options introduced in Junos OS Release 9.2.  
                                  **area** and **interface** options introduced in Junos OS Release 9.2 for EX Series switches.  
                                  Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Display information about OSPF neighbors.

CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the **show (ospf | ospf3) neighbor** command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the **show (ospf | ospf3) neighbor** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.

**Options**    **none**—Display standard information about all OSPF neighbors for all routing instances.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**area *area-id***—(Optional) Display information about the OSPF neighbors for the specified area.

**instance (all | *instance-name*)**—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.

**interface *interface-name***—(Optional) Display information about OSPF neighbors for the specified logical interface.

**logical-system** (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

**neighbor**—(Optional) Display information about the specified OSPF neighbor.

**realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**Required Privilege Level** view

**Related Documentation** • [clear \(ospf | ospf3\) neighbor on page 4250](#)

**List of Sample Output** [show ospf neighbor brief on page 4284](#)  
[show ospf neighbor detail on page 4284](#)  
[show ospf neighbor extensive on page 4285](#)  
[show ospf3 neighbor detail on page 4286](#)  
[show ospf neighbor area area-id on page 4286](#)  
[show ospf neighbor interface interface-name on page 4286](#)  
[show ospf3 neighbor instance all \(OSPFv3 Multiple Family Address Support Enabled\) on page 4286](#)

**Output Fields** [Table 331 on page 4282](#) lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

**Table 331: show (ospf | ospf3) neighbor Output Fields**

Field Name	Field Description	Level of Output
<b>Address</b>	Address of the neighbor.	All levels
<b>Interface</b>	Interface through which the neighbor is reachable.	All levels

Table 331: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	<p>State of the neighbor:</p> <ul style="list-style-type: none"> <li>• <b>Attempt</b>—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor.</li> <li>• <b>Down</b>—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the <b>Down</b> state, although at a reduced frequency.</li> <li>• <b>Exchange</b>—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged.</li> <li>• <b>ExStart</b>—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number.</li> <li>• <b>Full</b>—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements.</li> <li>• <b>Init</b>—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet.</li> <li>• <b>Loading</b>—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the <b>Exchange</b> state.</li> <li>• <b>2Way</b>—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state <b>2Way</b> or greater.</li> </ul>	All levels
<b>ID</b>	Router ID of the neighbor.	All levels
<b>Pri</b>	Priority of the neighbor to become the designated router.	All levels
<b>Dead</b>	Number of seconds until the neighbor becomes unreachable.	All levels
<b>Link state acknowledgment list</b>	Number of link-state acknowledgments received.	<b>extensive</b>
<b>Link state retransmission list</b>	<p>Total number of link-state advertisements retransmitted. For <b>extensive</b> output only, the following information is also displayed:</p> <ul style="list-style-type: none"> <li>• <b>Type</b>—Type of link advertisement: <b>ASBR</b>, <b>Sum</b>, <b>Extern</b>, <b>Network</b>, <b>NSSA</b>, <b>OpaqArea</b>, <b>Router</b>, or <b>Summary</b>.</li> <li>• <b>LSA ID</b>—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.</li> <li>• <b>Adv rtr</b>—Address of the routing device that sent the advertisement.</li> <li>• <b>Seq</b>—Link sequence number of the advertisement.</li> </ul>	<b>detail extensive</b>

Table 331: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Neighbor-address</b>	(OSPFv3 only) If the neighbor uses virtual links, the <b>Neighbor-address</b> is the site-local, local, or global address. If the neighbor uses a physical interface, the <b>Neighbor-address</b> is an IPv6 link-local address.	detail extensive
<b>area</b>	Area that the neighbor is in.	detail extensive
<b>OSPF3-Intf-Index</b>	(OSPFv3 only) Displays the OSPFv3 interface index.	detail extensive
<b>opt</b>	Option bits received in the hello packets from the neighbor.	detail extensive
<b>DR or DR-ID</b>	Address of the designated router.	detail extensive
<b>BDR or BDR-ID</b>	Address of the backup designated router.	detail extensive
<b>Up</b>	Length of time since the neighbor came up.	detail extensive
<b>adjacent</b>	Length of time since the adjacency with the neighbor was established.	detail extensive

## Sample Output

### show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

### show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2      ge-1/2/0.1     Full       10.5.1.2    128  37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2      ge-1/2/0.10     ExStart    10.5.1.38   128  34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2      ge-1/2/0.11     Full       10.5.1.42   128  38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```

```

10.5.12.2      ge-1/2/0.12      ExStart  10.5.1.46      128   33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

### show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address      Interface      State      ID      Pri  Dead
10.5.1.2      ge-1/2/0.1     Full      10.5.1.2  128   33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.56.0    172.25.27.82 0x8000004d
Router  10.5.1.94    10.5.1.94    0x8000005c
Network 10.5.24.2    10.5.1.94    0x80000036
Summary 10.8.57.0    172.25.27.82 0x80000024
Extern  1.10.90.0   10.8.1.2     0x80000041
Extern  1.4.109.0    10.6.1.2     0x80000041
Router  10.5.1.190   10.5.1.190   0x8000005f
Network 10.5.48.2    10.5.1.190   0x8000003d
Summary 10.8.58.0    172.25.27.82 0x8000004d
Extern  1.10.91.0    10.8.1.2     0x80000041
Extern  1.4.110.0    10.6.1.2     0x80000041
Router  10.5.1.18    10.5.1.18    0x8000005f
Network 10.5.5.2     10.5.1.18    0x80000033
Summary 10.8.59.0    172.25.27.82 0x8000003a
Summary 10.8.62.0    172.25.27.82 0x80000025

10.5.10.2     ge-1/2/0.10     ExStart  10.5.1.38      128   38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec
10.5.11.2     ge-1/2/0.11     Full      10.5.1.42      128   33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
Summary 10.8.58.0    172.25.27.82 0x8000004d

```

Extern	1.10.91.0	10.8.1.2	0x80000041
Extern	1.1.247.0	10.5.1.2	0x8000003f
Extern	1.4.110.0	10.6.1.2	0x80000041
Router	10.5.1.18	10.5.1.18	0x8000005f
Network	10.5.5.2	10.5.1.18	0x80000033
Summary	10.8.59.0	172.25.27.82	0x8000003a

**show ospf3 neighbor detail**

```

user@host> show ospf3 neighbor detail
ID          Interface          State    Pri    Dead
10.255.71.13 fe-0/0/2.0          Full     128    30
Neighbor-address fe80::290:69ff:fe9b:e002
area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43

```

**show ospf neighbor area area-id**

```

user@host >show ospf neighbor area 1.1.1.1
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0          Full     10.255.245.5 128    37
Area 1.1.1.1

```

**show ospf neighbor interface interface-name**

```

user@host >show ospf neighbor interface so-0/0/0.0
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0          Full     10.255.245.4 128    32
Area 2.2.2.2

```

**show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)**

```

user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID          Interface          State    Pri    Dead
100.1.1.1    fe-0/0/2.0          Full     128    37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID          Interface          State    Pri    Dead
100.1.2.1    fe-0/0/2.1          Full     128    33
Neighbor-address fe80::217:cb00:c97c:8c03

```



## show (ospf | ospf3) overview

<b>List of Syntax</b>	<a href="#">Syntax on page 4287</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 4287</a>
<b>Syntax</b>	show (ospf   ospf3) overview <brief   extensive> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )> <realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)>
<b>Syntax (EX Series Switch and QFX Series)</b>	show (ospf   ospf3) overview <brief   extensive> <instance <i>instance-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>realm</b> option introduced in Junos OS Release 9.2. Database protection introduced in Junos 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display Open Shortest Path First (OSPF) overview information.
<b>Options</b>	<b>none</b> —Display standard information about all OSPF neighbors for all routing instances.  <b>brief   extensive</b> —(Optional) Display the specified level of output.  <b>instance <i>instance-name</i></b> —(Optional) Display all OSPF interfaces under the named routing instance.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b> —(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ospf overview on page 4289</a> <a href="#">show ospf overview (With Database Protection) on page 4290</a> <a href="#">show ospf3 overview (With Database Protection) on page 4290</a> <a href="#">show ospf overview extensive on page 4290</a>
<b>Output Fields</b>	<a href="#">Table 211 on page 2362</a> lists the output fields for the <b>show ospf overview</b> command. Output fields are listed in the approximate order in which they appear.

Table 332: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: <b>Current</b> , <b>Warning</b> (threshold), and <b>Allowed</b> .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: <b>Current</b> and <b>Allowed</b> .	All levels
Restart	Graceful restart capability: <b>enabled</b> or <b>disabled</b> .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 332: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): <b>enabled</b> or <b>disabled</b> .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): <b>enabled</b> or <b>disabled</b> .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: <b>enabled</b> or <b>disabled</b> .	All levels
Trace options	OSPF-specific trace options.	<b>extensive</b>
Trace file	Name of the file to receive the output of the tracing operation.	<b>extensive</b>
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: <b>Normal Stub</b> , <b>Not Stub</b> , or <b>Not so Stubby Stub</b> .	All levels
Authentication Type	Type of authentication: <b>None</b> , <b>Password</b> , or <b>MD5</b> .  <b>NOTE:</b> The <b>Authentication Type</b> field refers to the authentication configured at the <b>[edit protocols ospf area area-id]</b> level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

## Sample Output

### show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

### show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

### show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

### show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
```

```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

```
show (ospf | ospf3) route
```

- |                                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                           | <a href="#">Syntax on page 4292</a><br><a href="#">Syntax (EX Series Switch and QFX Series) on page 4292</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Syntax</b>                                   | <pre>show (ospf   ospf3) route &lt;brief   detail   extensive&gt; &lt;abr   asbr   extern   inter   intra&gt; &lt;destination&gt; &lt;instance (default   ipv4-multicast   instance-name)&gt; &lt;logical-system (default   ipv4-multicast   logical-system-name)&gt; &lt;network&gt; &lt;no-backup-coverage&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt; &lt;router&gt; &lt;topology (default   ipv4-multicast   topology-name)&gt; &lt;transit&gt;</pre>                                                                                                                                             |
| <b>Syntax (EX Series Switch and QFX Series)</b> | <pre>show (ospf   ospf3) route &lt;brief   detail   extensive&gt; &lt;abr   asbr   extern   inter   intra&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;network&gt; &lt;no-backup-coverage&gt; &lt;router&gt; &lt;topology (default   ipv4-multicast   topology-name)&gt; &lt;transit&gt;</pre>                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                      | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>topology</b> option introduced in Junos OS Release 9.0.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                              | Display the entries in the Open Shortest Path First (OSPF) routing table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                  | <p><b>none</b>—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.</p> <p><b>destination</b>—Display routes to the specified IP address (with optional destination prefix length).</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>abr</b>—(Optional) Display routes to area border routers.</p> <p><b>asbr</b>—(Optional) Display routes to autonomous system border routers.</p> <p><b>extern</b>—(Optional) Display external routes.</p> <p><b>inter</b>—(Optional) Display interarea routes.</p> |

**intra**—(Optional) Display intra-area routes.

**instance** (**default** | **ipv4-multicast** | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

**logical-system** (**default** | **ipv4-multicast** | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

**network**—(Optional) Display routes to networks.

**no-backup-coverage**—(Optional) Display routes with no backup coverage.

**realm** (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

**router**—(Optional) Display routes to all routers.

**topology** (**default** | **ipv4-multicast** | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

**transit**—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

**Required Privilege Level**

view

**List of Sample Output**

[show ospf route on page 4295](#)  
[show ospf route detail on page 4295](#)  
[show ospf3 route on page 4295](#)  
[show ospf3 route detail on page 4296](#)  
[show ospf route topology voice on page 4296](#)

**Output Fields**

[Table 333 on page 4293](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

**Table 333: show (ospf | ospf3) route Output Fields**

Field Name	Field Description	Output Level
<b>Topology</b>	Name of the topology.	All levels
<b>Prefix</b>	Destination of the route.	All levels

Table 333: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
<b>Path type</b>	How the route was learned: <ul style="list-style-type: none"> <li>• <b>Inter</b>—Interarea route</li> <li>• <b>Ext1</b>—External type 1 route</li> <li>• <b>Ext2</b>—External type 2 route</li> <li>• <b>Intra</b>—Intra-area route</li> </ul>	All levels
<b>Route type</b>	The type of routing device from which the route was learned: <ul style="list-style-type: none"> <li>• <b>AS BR</b>—Route to AS border router.</li> <li>• <b>Area BR</b>—Route to area border router.</li> <li>• <b>Area/AS BR</b>—Route to router that is both an <b>Area BR</b> and <b>AS BR</b>.</li> <li>• <b>Network</b>—Network router.</li> <li>• <b>Router</b>—Route to a router that is neither an <b>Area BR</b> nor an <b>AS BR</b>.</li> <li>• <b>Transit</b>—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link.</li> <li>• <b>Discard</b>—Route to a summary discard.</li> </ul>	All levels
<b>NH Type</b>	Next-hop type: <b>LSP</b> or <b>IP</b> .	All levels
<b>Metric</b>	Route's metric value.	All levels
<b>NH-interface</b>	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
<b>NH-addr</b>	(OSPFv3 only) IPv6 address of the next hop.	All levels
<b>NextHop Interface</b>	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
<b>Nexthop addr/label</b>	(OSPFv2 only) If the <b>NH Type</b> is <b>IP</b> , then it is the address of the next hop. If the <b>NH Type</b> is <b>LSP</b> , then it is the name of the label-switched path.	All levels
<b>Area</b>	Area ID of the route.	detail
<b>Origin</b>	Router from which the route was learned.	detail
<b>Type 7</b>	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	detail
<b>P-bit</b>	Route was learned through NSSA LSA and the propagate bit was set.	detail
<b>Fwd NZ</b>	Forwarding address is nonzero. <b>Fwd NZ</b> is only displayed if the route is learned through an NSSA LSA.	detail



Table 333: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
<b>optional-capability</b>	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when <b>Route Type</b> is <b>Area BR</b> , <b>AS BR</b> , <b>Area/AS BR</b> , or <b>Router</b> ), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> <li>• <b>0x4 (V)</b>—Routing device is at the end of a virtual active link.</li> <li>• <b>0x2 (E)</b>—Routing device is an autonomous system boundary router.</li> <li>• <b>0x1 (B)</b>—Routing device is an area border router.</li> </ul>	<b>detail</b>
<b>priority</b>	The priority assigned to the prefix: <ul style="list-style-type: none"> <li>• <b>high</b></li> <li>• <b>medium</b></li> <li>• <b>low</b></li> </ul> <p><b>NOTE:</b> The <b>priority</b> field applies only to routes of type <b>Network</b>.</p>	<b>detail</b>

## Sample Output

### show ospf route

```
user@host> show ospf route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.71.12     Intra Router   IP    1       fe-0/0/2.0   192.16.22.86
10.255.71.13/32  Intra Network IP    0       lo0.0
192.168.222.84/30 Intra Network LSP   1       fe-0/0/2.0   1sp-ab
```

### show ospf route detail

```
user@host> show ospf route detail
Topology default Route Table:

Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
10.255.14.174    Inter AS BR   IP    210     t1-3/0/1.0
area 0.0.0.2, origin 10.255.14.185
10.255.14.178    Intra Router   IP    200     t3-3/1/3.0
area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30    Intra Network IP    10      t3-3/1/2.0
area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32     Inter Network IP    210     t1-3/0/1.0
area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24     Ext2  Network   IP    0       t1-3/0/1.0
area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30     Inter Network IP    220     t1-3/0/1.0
area 0.0.0.2, origin 10.255.14.185, priority high
```

### show ospf3 route

```
user@host> show ospf3 route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type   Type    Type                Interface    addr/label
```

```

10.255.71.13      Intra Router      IP      1
NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1      Intra Router      IP      40 fxp1.1      192.168.36.17

    area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3      Intra AS BR      IP      1 fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32   Intra Network    IP      40 fxp1.1      192.168.36.17

    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32   Intra Network    IP      0 lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32   Intra Network    IP      1 fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3, priority low
    Intra Transit    IP      1
NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network    IP      1
NH-interface fe-0/0/2.0
abcd::71:12/128   Intra Network    IP      0
NH-interface lo0.0
abcd::71:13/128   Intra Network    LSP     1
NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

#### show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix
Path      Route      NH      Metric
type      type
10.255.14.174 Intra Area/AS BR IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178 Intra Router IP 200
NH-interface t3-3/1/3.0
Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2 Intra Transit IP 200
NH-interface t1-3/0/1.0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128 Inter Network IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48 Ext1 Network IP 110
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48 Ext2 Network IP 0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48 Ext2 Network IP 0
NH-interface so-1/2/2.0
Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128 Intra Network IP 0
NH-interface lo0.0
Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

#### show ospf route topology voice

```
user@host show ospf route topology voice
```

Topology voice Route Table:

Prefix	Path Type	Route Type	NH Type	Metric	NextHop Interface	Nexthop addr/label
10.255.8.2	Intra	Router	IP	1	so-0/2/0.0	
10.255.8.3	Intra	Router	IP	2	so-0/2/0.0	
10.255.8.1/32	Intra	Network	IP	0	lo0.0	
10.255.8.2/32	Intra	Network	IP	1	so-0/2/0.0	
10.255.8.3/32	Intra	Network	IP	2	so-0/2/0.0	
192.168.8.0/29	Intra	Network	IP	2	so-0/2/0.0	
192.168.8.44/30	Intra	Network	IP	2	so-0/2/0.0	
192.168.8.46/32	Intra	Network	IP	1	so-0/2/0.0	
192.168.8.48/30	Intra	Network	IP	1	so-0/2/1.0	
192.168.8.52/30	Intra	Network	IP	2	so-0/2/0.0	
192.168.9.44/30	Intra	Network	IP	1	so-0/2/0.0	
192.168.9.45/32	Intra	Network	IP	2	so-0/2/0.0	

## show (ospf | ospf3) statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 4298</a> <a href="#">Syntax (EX Series Switch and QFX Series) on page 4298</a>
<b>Syntax</b>	<pre>show (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)&gt;</pre>
<b>Syntax (EX Series Switch and QFX Series)</b>	<pre>show (ospf   ospf3) statistics &lt;instance <i>instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>realm</b> option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display OSPF statistics.
<b>Options</b>	<p><b>none</b>—Display OSPF statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display all statistics for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>realm (ipv4-multicast   ipv4-unicast   ipv6-multicast)</b>—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the <b>realm</b> option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear (ospf   ospf3) statistics on page 4252</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ospf statistics on page 4300</a> <a href="#">show ospf statistics logical-system all on page 4300</a> <a href="#">show ospf3 statistics on page 4301</a>
<b>Output Fields</b>	<p><a href="#">Table 334 on page 4298</a> lists the output fields for the <b>show (ospf   ospf3) statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 334: show (ospf | ospf3) statistics Output Fields**

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.

Table 334: show (ospf | ospf3) statistics Output Fields (*continued*)

Field Name	Field Description
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	<p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.
Receive errors	<p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> <li>• mtu mismatches</li> <li>• no interface found</li> <li>• no virtual link found</li> <li>• nssa mismatches</li> <li>• stub area mismatches</li> <li>• subnet mismatches</li> </ul> <p>If there are no receive errors, the output displays <b>none</b>.</p>

## Sample Output

### show ospf statistics

```

user@host> show ospf statistics
Packet type          Total
                   Sent      Received
Hello                 31         14
DbD                   9          10
LSReq                 2           2
LSUpdate              8          16
LSAck                 9           9

                   Last 5 seconds
                   Sent      Received
Hello                 2           2
DbD                   0           0
LSReq                 0           0
LSUpdate              0           0
LSAck                 0           0

DBDs retransmitted   :          3, last 5 seconds :          0
LSAs flooded         :         12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested       :          5, last 5 seconds :          0
LSAs acknowledged    :         19, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
  862 no interface found
 115923 no virtual link found

```

### show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B
Packet type          Total
                   Sent      Received
Hello              313740      313653
DbD                 3           2
LSReq               1           1
LSUpdate            2752      1825
LSAck               1821      2747

                   Last 5 seconds
                   Sent      Received
Hello                1           0
DbD                  0           0
LSReq                0           0
LSUpdate             0           0
LSAck                0           0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        2741, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          2, last 5 seconds :          0
LSAs requested       :          1, last 5 seconds :          0
LSAs acknowledged    :       1831, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:

```

```

None
-----

logical-system: A

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
    Hello           313698      313695         0         0
      DbD              2         3         0         0
    LSReq             1         1         0         0
  LSUpdate           1825      2752         0         0
    LSAck            2747      1821         0         0

DBDs retransmitted   :           0, last 5 seconds :           0
LSAs flooded         :        1825, last 5 seconds :           0
LSAs flooded high-prio :         10, last 5 seconds :           0
LSAs retransmitted   :           0, last 5 seconds :           0
LSAs transmitted to nbr:         1, last 5 seconds :           0
LSAs requested       :           2, last 5 seconds :           0
LSAs acknowledged   :        2748, last 5 seconds :           0

Flood queue depth    :           0
Total rexmit entries :           0
db summaries         :           0
lsreq entries        :           0

Receive errors:
None
-----

```

### show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
    Hello              0         0         0         0
      DbD              0         0         0         0
    LSReq              0         0         0         0
  LSUpdate             0         0         0         0
    LSAck              0         0         0         0

DBDs retransmitted   :           0, last 5 seconds :           0
LSAs flooded         :           0, last 5 seconds :           0
LSAs flooded high-prio :           0, last 5 seconds :           0
LSAs retransmitted   :           0, last 5 seconds :           0
LSAs transmitted to nbr:           0, last 5 seconds :           0
LSAs requested       :           0, last 5 seconds :           0
LSAs acknowledged   :           0, last 5 seconds :           0

Flood queue depth    :           0
Total rexmit entries :           0
db summaries         :           0
lsreq entries        :           0

Receive errors:
None

```





## PART 14

# Routing Information Protocol

- [Overview on page 4305](#)
- [Configuration on page 4311](#)
- [Administration on page 4399](#)



## CHAPTER 50

# Overview

- [RIP Overview on page 4305](#)

## RIP Overview

---

- [RIP Overview on page 4305](#)

## RIP Overview

RIP is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using the hop count as the metric.

In a RIP network, each router's forwarding table is distributed among the nodes through the flooding of routing table information. Because topology changes are flooded throughout the network, every node maintains the same list of destinations. Packets are then routed to these destinations based on path-cost calculations done at each node in the network.



NOTE: In general, the term *RIP* refers to RIP version 1 and RIP version 2.

This topic contains the following sections:

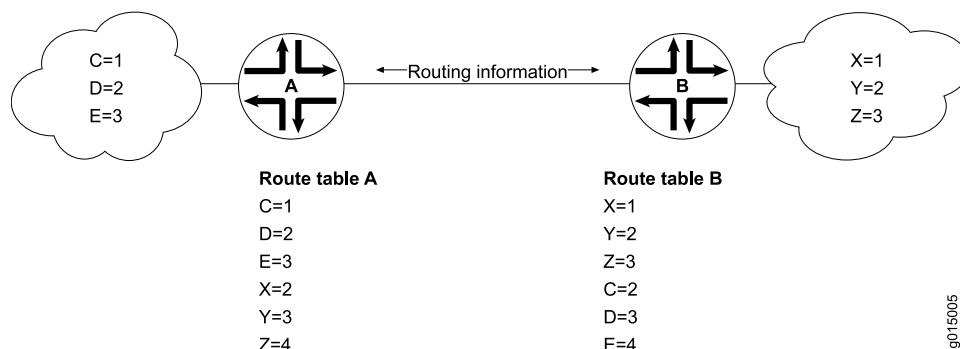
- [Distance-Vector Routing Protocols on page 4305](#)
- [RIP Protocol Overview on page 4306](#)
- [RIP Packets on page 4307](#)
- [Maximizing Hop Count on page 4308](#)
- [Split Horizon and Poison Reverse Efficiency Techniques on page 4308](#)
- [Limitations of Unidirectional Connectivity on page 4309](#)

## Distance-Vector Routing Protocols

---

Distance-vector routing protocols transmit routing information that includes a distance vector, typically expressed as the number of hops to the destination. This information is flooded out all protocol-enabled interfaces at regular intervals (every 30 seconds in the case of RIP) to create a network map that is stored in each node's local topology database. [Figure 124 on page 4306](#) shows how distance-vector routing works.

Figure 124: Distance-Vector Protocol



In Figure 124 on page 4306, Routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors Routers C, D, and E, which are 1, 2, and 3 hops away, respectively. Router B has known RIP neighbors Routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each router floods its entire routing table information out all RIP-enabled interfaces. In this case, flooding exchanges routing table information across the RIP link.

When Router A receives routing information from Router B, it adds 1 to the hop count to determine the new hop count. For example, Router X has a hop count of 1, but when Router A imports the route to X, the new hop count is 2. The imported route also includes information about where the route was learned, so that the original route is imported as a route to Router X through Router B with a hop count of 2.

When multiple routes to the same host are received, RIP uses the distance-vector algorithm to determine which path to import into the forwarding table. The route with the smallest hop count is imported. If there are multiple routes with the same hop count, all are imported into the forwarding table, and traffic is sent along the paths in round-robin fashion.

### RIP Protocol Overview

The RIP IGP uses the Bellman-Ford, or *distance-vector*, algorithm to determine the best route to a destination. RIP uses the hop count as the metric. RIP enables hosts and routers to exchange information for computing routes through an IP-based network. RIP is intended to be used as an IGP in reasonably homogeneous networks of moderate size.

The Junos® operating system (Junos OS) supports RIP versions 1 and 2.



**NOTE:** RIP is not supported for multipoint interfaces.

RIP version 1 packets contain the minimal information necessary to route packets through a network. However, this version of RIP does not support authentication or subnetting.

RIP uses User Datagram Protocol (UDP) port 520.

RIP has the following architectural limitations:

- The longest network path cannot exceed 15 hops (assuming that each network, or hop, has a cost of 1).
- RIP depends on counting to infinity to resolve certain unusual situations—When the network consists of several hundred routers, and when a routing loop has formed, the amount of time and network bandwidth required to resolve a next hop might be great.
- RIP uses only a fixed metric to select a route. Other IGPs use additional parameters, such as measured delay, reliability, and load.

### RIP Packets

RIP packets contain the following fields:

- Command—Indicates whether the packet is a request or response message. Request messages seek information for the router's routing table. Response messages are sent periodically and also when a request message is received. Periodic response messages are called *update messages*. Update messages contain the command and version fields and 25 destinations (by default), each of which includes the destination IP address and the metric to reach that destination.



**NOTE:** Beginning with Junos OS Release 11.1, three additional command field types are available to support RIP demand circuits. When you configure an interface for RIP demand circuits, the command field indicates whether the packet is an update request, update response, or update acknowledge message. Neighbor interfaces send updates on demand, not periodically. These command field types are only valid on interfaces configured for RIP demand circuits. For more detailed information, see *RIP Demand Circuits Overview*.

- Version number—Version of RIP that the originating router is running.
- Address family identifier—Address family used by the originating router. The family is always IP.
- Address—IP address included in the packet.
- Metric—Value of the metric advertised for the address.
- Mask—Mask associated with the IP address (RIP version 2 only).
- Next hop—IP address of the next-hop router (RIP version 2 only).

Routing information is exchanged in a RIP network by RIP request and RIP response packets. A router that has just booted can broadcast a RIP request on all RIP-enabled interfaces. Any routers running RIP on those links receive the request and respond by sending a RIP response packet immediately to the router. The response packet contains the routing table information required to build the local copy of the network topology map.

In the absence of RIP request packets, all RIP routers broadcast a RIP response packet every 30 seconds on all RIP-enabled interfaces. The RIP broadcast is the primary way in which topology information is flooded throughout the network.

Once a router learns about a particular destination through RIP, it starts a timer. Every time it receives a new response packet with information about the destination, the router resets the timer to zero. However, if the router receives no updates about a particular destination for 180 seconds, it removes the destination from its RIP routing table.

In addition to the regular transmission of RIP packets every 30 seconds, if a router detects a new neighbor or detects that an interface is unavailable, it generates a triggered update. The new routing information is immediately broadcast out all RIP-enabled interfaces, and the change is reflected in all subsequent RIP response packets.

### Maximizing Hop Count

The successful routing of traffic across a RIP network requires that every node in the network maintain the same view of the topology. Topology information is broadcast between RIP neighbors every 30 seconds. If Router A is many hops away from a new host, Router B, the route to B might take significant time to propagate through the network and be imported into Router A's routing table. If the two routers are 5 hops away from each other, Router A cannot import the route to Router B until 2.5 minutes after Router B is online (30 seconds per hop). For large numbers of hops, the delay becomes prohibitive. To help prevent this delay from growing arbitrarily large, RIP enforces a maximum hop count of 15 hops. Any prefix that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity. This maximum hop count is called the *network diameter*.


### Split Horizon and Poison Reverse Efficiency Techniques

Because RIP functions by periodically flooding the entire routing table out to the network, it generates a lot of traffic. The split horizon and poison reverse techniques can help reduce the amount of network traffic originated by RIP hosts and make the transmission of routing information more efficient.

If a router receives a set of route advertisements on a particular interface, RIP determines that those advertisements do not need to be retransmitted out the same interface. This technique, known as *split horizon*, helps limit the amount of RIP routing traffic by eliminating information that other neighbors on that interface have already learned.

[Figure 125 on page 4308](#) shows an example of the split horizon technique.


**Figure 125: Split Horizon Example**



In [Figure 125 on page 4308](#), Router A advertises routes to Routers C, D, and E to Router B. In this example, Router A can reach Router C in 2 hops. When Router A advertises the route to Router B, Router B imports it as a route to Router C through Router A in 3 hops. If Router B then readvertised this route to Router A, Router A would import it as a route to Router C through Router B in 4 hops. However, the advertisement from Router B to Router A is unnecessary, because Router A can already reach the route in 2 hops. The split horizon technique helps reduce extra traffic by eliminating this type of route advertisement.

Similarly, the poison reverse technique helps to optimize the transmission of routing information and improve the time to reach network convergence. If Router A learns about unreachable routes through one of its interfaces, it advertises those routes as unreachable (hop count of 16) out the same interface. [Figure 126 on page 4309](#) shows an example of the poison reverse technique.


**Figure 126: Poison Reverse Example**



In [Figure 126 on page 4309](#), Router A learns through one of its interfaces that routes to Routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs Router B that Routers C, D, and E are definitely not reachable through Router A.

### Limitations of Unidirectional Connectivity

Because RIP processes routing information based solely on the receipt of routing table updates, it cannot ensure bidirectional connectivity. As [Figure 127 on page 4310](#) shows, RIP networks are limited by their unidirectional connectivity.

**Figure 127: Limitations of Unidirectional Connectivity**

In [Figure 127 on page 4310](#), Routers A and D flood their routing table information to Router B. Because the path to Router E has the fewest hops when routed through Router A, that route is imported into Router B's forwarding table. However, suppose that Router A can transmit traffic but is not receiving traffic from Router B because of an unavailable link or invalid routing policy. If the only route to Router E is through Router A, any traffic destined for Router A is lost, because bidirectional connectivity was never established.

OSPF establishes bidirectional connectivity with a three-way handshake.

#### Related Documentation

- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)
- [RIP Configuration Overview](#)
- [Example: Configuring RIP on page 4311](#)



## CHAPTER 51

# Configuration

- [RIP Configuration Tasks on page 4311](#)
- [RIP Configuration Statements on page 4373](#)

### RIP Configuration Tasks

---

- [Example: Configuring RIP on page 4311](#)
- [Example: Configuring Authentication for RIP Routes on page 4318](#)
- [Example: Configuring BFD for RIP on page 4324](#)
- [Example: Configuring BFD Authentication for RIP on page 4330](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4338](#)
- [Examples: Controlling Traffic with Metrics in a RIP Network on page 4344](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4352](#)
- [Example: Redistributing Routes Among RIP Instances on page 4356](#)
- [Example: Configuring RIP Timers on page 4361](#)
- [Example: Tracing RIP Protocol Traffic on page 4368](#)

### Example: Configuring RIP

- [Understanding Basic RIP Routing on page 4311](#)
- [Example: Configuring a Basic RIP Network on page 4312](#)

#### Understanding Basic RIP Routing

---

RIP is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). By default, RIP does not advertise the subnets that are directly connected through the device's interfaces. For traffic to pass through a RIP network, you must create a routing policy to export these routes. Advertising only the direct routes propagates the routes to the immediately adjacent RIP-enabled router only. To propagate all routes through the entire RIP network, you must configure the routing policy to export the routes learned through RIP.

### Example: Configuring a Basic RIP Network

This example shows how to configure a basic RIP network.

- [Requirements on page 4312](#)
- [Overview on page 4312](#)
- [Configuration on page 4312](#)
- [Verification on page 4315](#)

#### Requirements

No special configuration beyond device initialization is required before configuring this example.

#### Overview


In this example, you configure a basic RIP network, create a RIP group called **rip-group**, and add the directly connected interfaces to the RIP group. Then you configure a routing policy to advertise direct routes using policy statement **advertise-routes-through-rip**.

By default, Junos OS does not advertise RIP routes, not even routes that are learned through RIP. To advertise RIP routes, you must configure and apply an export routing policy that advertises RIP-learned and direct routes.

In Junos OS, you do not need to configure the RIP version. RIP version 2 is used by default.

To use RIP on the device, you must configure RIP on all of the RIP interfaces within the network. [Figure 128 on page 4312](#) shows the topology used in this example.

**Figure 128: Sample RIP Network Topology**



“CLI Quick Configuration” on [page 4312](#) shows the configuration for all of the devices in [Figure 128 on page 4312](#). The section “Step-by-Step Procedure” on [page 4313](#) describes the steps on Device R1.

#### Configuration

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

##### Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
```

```

set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a basic RIP network:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

```

```

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32

```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```

[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1

```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Verification**

Confirm that the configuration is working properly.

- [Checking the Routing Table on page 4315](#)
- [Looking at the Routes That Device R1 Is Advertising to Device R2 on page 4315](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 4316](#)
- [Verifying the RIP-Enabled Interfaces on page 4316](#)
- [Verifying the Exchange of RIP Messages on page 4316](#)
- [Verifying Reachability of All Hosts in the RIP Network on page 4317](#)

**Checking the Routing Table**

**Purpose** Verify that the routing table is populated with the expected routes..

**Action** From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[RIP/100] 00:59:15, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32    *[RIP/100] 02:52:48, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32    *[RIP/100] 00:45:05, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32   *[RIP/100] 02:52:48, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32   *[RIP/100] 00:45:05, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32     *[RIP/100] 00:45:09, metric 1
                 MultiRecv
```

**Meaning** The output shows that the routes have been learned from Device R2 and Device R3.

If you were to delete the **from protocol rip** condition in the routing policy on Device R2, the remote routes from Device R3 would not be learned on Device R1.

**Looking at the Routes That Device R1 Is Advertising to Device R2**

**Purpose** Verify that Device R1 is sending the expected routes.

**Action** From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R1> show route advertising-protocol rip 10.0.0.1
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32    *[Direct/0] 05:18:26
                 > via lo0.1
192.168.1.1/32   *[Direct/0] 05:18:25
                 > via lo0.1
```

**Meaning** Device R1 is sending routes to its directly connected networks.

**Looking at the Routes That Device R1 Is Receiving from Device R2**

**Purpose** Verify that Device R1 is receiving the expected routes.

**Action** From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30          * [RIP/100] 02:31:22, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32       * [RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32       * [RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32      * [RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32      * [RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

**Meaning** Device R1 is receiving from Device R2 all of Device R2's directly connected networks. Device R1 is also receiving from Device R2 all of Device R3's directly connected networks, which Device R2 learned from Device R3 through RIP.

**Verifying the RIP-Enabled Interfaces**

**Purpose** Verify that all RIP-enabled Interfaces are available and active.

**Action** From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
fe-1/2/0.1	Up	10.0.0.1	224.0.0.9	mcast	both	1

**Meaning** The output shows that the RIP-enabled interface on Device R1 is operational.

In general for this command, the output shows a list of the RIP neighbors that are configured on the device. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Local State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

**Verifying the Exchange of RIP Messages**

**Purpose** Verify that RIP messages are being sent and received on all RIP-enabled interfaces.

**Action** From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter                Total    Last 5 min  Last minute
-----
Updates Sent                2669         10         2
Triggered Updates Sent      2          0          0
Responses Sent              0          0          0
Bad Messages                0          0          0
RIPv1 Updates Received      0          0          0
RIPv1 Bad Route Entries     0          0          0
RIPv1 Updates Ignored       0          0          0
RIPv2 Updates Received     2675        11         2
RIPv2 Bad Route Entries     0          0          0
RIPv2 Updates Ignored       0          0          0
Authentication Failures     0          0          0
RIP Requests Received       0          0          0
RIP Requests Ignored        0          0          0
none                        0          0          0

```

**Meaning** The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

#### *Verifying Reachability of All Hosts in the RIP Network*

**Purpose** Use the **traceroute** command on each loopback address in the network to verify that all hosts in the RIP network are reachable from each Juniper Networks device.

**Action** From operational mode, enter the **traceroute** command.

```

user@R1> traceroute 192.168.3.3
traceroute to 192.168.3.3 (192.168.3.3), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  1.094 ms  1.028 ms  0.957 ms
 2  192.168.3.3 (192.168.3.3)  1.344 ms  2.245 ms  2.125 ms

```

**Meaning** Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

**Related  
Documentation**

- *Example: Configuring Point-to-Multipoint RIP Networks*

## Example: Configuring Authentication for RIP Routes

- [Understanding RIP Authentication on page 4318](#)
- [Example: Configuring Route Authentication for RIP on page 4318](#)
- [Enabling Authentication with Plain-Text Passwords \(CLI Procedure\) on page 4323](#)
- [Enabling Authentication with MD5 Authentication \(CLI Procedure\) on page 4323](#)

### Understanding RIP Authentication

---

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. Authentication provides an additional layer of security on the network beyond the other security features. By default, this authentication is disabled.

Authentication keys can be specified in either plain-text or MD5 form. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

This type of authentication is not supported on RIPv1 networks.

### Example: Configuring Route Authentication for RIP

---

This example shows how to configure authentication for a RIP network.

- [Requirements on page 4318](#)
- [Overview on page 4318](#)
- [Configuration on page 4319](#)
- [Verification on page 4322](#)

#### **Requirements**

No special configuration beyond device initialization is required before configuring this example.

#### **Overview**

You can configure the router to authenticate RIP route queries. By default, authentication is disabled. You can use one of the following authentication methods:




- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication—Creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum.

This example shows MD5 authentication.

Figure 129 on page 4319 shows the topology used in this example.

**Figure 129: RIP Authentication Network Topology**



"CLI Quick Configuration" on page 4319 shows the configuration for all of the devices in Figure 129 on page 4319. The section "Step-by-Step Procedure" on page 4320 describes the steps on Device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2

```

```
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$Lf1Xds2gJDHmoJCu1hKvoJGUjq"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip authentication-type md5
set protocols rip authentication-key "$9$G.UkP5T39tOz3K87V4oz36/Cu"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure RIP authentication:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
```

```
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Require MD5 authentication for RIP route queries received on an interface.

The passwords must match on neighboring RIP routers. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Do not enter the password as shown here. The password shown here is the encrypted password that is displayed in the configuration after the actual password is already configured.

```
[edit protocols rip]
user@R1# set authentication-type md5
user@R1# set authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"
```

6. Configure tracing operations to track authentication.

```
[edit protocols rip traceoptions]
user@R1# set file rip-authentication-messages
user@R1# set flag auth
user@R1# set flag packets
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  traceoptions {
    file rip-authentication-messages;
    flag auth;
```

```

        flag packets;
    }
    authentication-type md5;
    authentication-key "$9$ONLRBhreK87dsM8i.5FAtM8XxNb"; ## SECRET-DATA
    group rip-group {
        export advertise-routes-through-rip;
        neighbor fe-1/2/0.1;
    }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking for Authentication Failures on page 4322](#)
- [Verifying That MD5 Authentication Is Enabled in RIP Update Packets on page 4323](#)

### Checking for Authentication Failures

**Purpose** Verify that there are no authentication failures.

**Action** From operational mode, enter the **show rip statistics** command.

```

user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval
30s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          2669         10          2
Triggered Updates Sent      2          0          0
Responses Sent          0          0          0
Bad Messages           0          0          0
RIPv1 Updates Received     0          0          0
RIPv1 Bad Route Entries    0          0          0
RIPv1 Updates Ignored       0          0          0
RIPv2 Updates Received    2675         11          2
RIPv2 Bad Route Entries     0          0          0
RIPv2 Updates Ignored       0          0          0
Authentication Failures      0          0          0
RIP Requests Received       0          0          0
RIP Requests Ignored        0          0          0
none                       0          0          0

```

**Meaning** The output shows that there are no authentication failures.

### Verifying That MD5 Authentication Is Enabled in RIP Update Packets

**Purpose** Use tracing operations to verify that MD5 authentication is enabled in RIP updates.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log rip-authentication-messages | match md5
Feb 15 15:45:13.969462      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:45:43.229867      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:13.174410      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:42.716566      sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:47:11.425076      sending msg 0xb9a8c04, 3 rtes (needs MD5)
...
```

**Meaning** The **(needs MD5)** output shows that all route updates require MD5 authentication.

### Enabling Authentication with Plain-Text Passwords (CLI Procedure)

To configure authentication that requires a plain-text password to be included in the transmitted packet, enable simple authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 335 on page 4323](#).
3. If you are finished configuring the router, commit the configuration.

**Table 335: Configuring Simple RIP Authentication**

Task	CLI Configuration Editor
Navigate to <b>Rip</b> level in the configuration hierarchy.	From the <b>[edit]</b> hierarchy level, enter  <b>edit protocols rip</b>
Set the authentication type to <b>simple</b> .	Set the authentication type to <b>simple</b> :  <b>set authentication-type simple</b>
Set the authentication key to a simple-text password.  The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.	Set the authentication key to a simple-text password:  <b>set authentication-key <i>password</i></b>

### Enabling Authentication with MD5 Authentication (CLI Procedure)

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in [Table 336 on page 4324](#).
3. If you are finished configuring the router, commit the configuration.

Table 336: Configuring MD5 RIP Authentication

Task	CLI Configuration Editor
Navigate to <b>Rip</b> level in the configuration hierarchy.	From the <b>[edit]</b> hierarchy level, enter  <b>edit protocols rip</b>
Set the authentication type to <b>MD5</b> .	Set the authentication type to <b>md5</b> :  <b>set authentication-type md5</b>
Set the MD5 authentication key (password).  The key can be from 1 through 16 contiguous characters long and can include any ASCII strings.	Set the MD5 authentication key:  <b>set authentication-key password</b>

**Related Documentation**

- [Example: Configuring RIP on page 4311](#)

### Example: Configuring BFD for RIP

- [Understanding BFD for RIP on page 4324](#)
- [Example: Configuring BFD for RIP on page 4325](#)

#### Understanding BFD for RIP

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. BFD failure detection times are shorter than RIP detection times, providing faster reaction times to various kinds of failures in the network. Instead of waiting for the routing protocol neighbor timeout, BFD provides rapid detection of link failures. BFD timers are adaptive and can be adjusted to be more or less aggressive. For example, a timer can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured.

BFD enables quick failover between a primary and a secondary routed path. The protocol tests the operational status of the interface multiple times per second. BFD provides for configuration timers and thresholds for failure detection. For example, if the minimum interval is set for 50 milliseconds and the threshold uses the default value of three missed messages, a failure is detected on an interface within 200 milliseconds of the failure.

Intervening devices (for example, an Ethernet LAN switch) hide link-layer failures from routing protocol peers, such as when two routers are connected by way of a LAN switch, where the local interface status remains up even when a physical fault happens on the remote link. Link-layer failure detection times vary, depending on the physical media and the Layer 2 encapsulation. BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore, you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

### Example: Configuring BFD for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- [Requirements on page 4325](#)
- [Overview on page 4325](#)
- [Configuration on page 4327](#)
- [Verification on page 4329](#)

#### Requirements

No special configuration beyond device initialization is required before configuring this example.

#### Overview

To enable failure detection, include the **bfd-liveness-detection** statement:

```

bfd-liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  version (1 | automatic);
}

```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.



**NOTE:** BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.

To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.


You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.



In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 130 on page 4327 shows the topology used in this example.

Figure 130: RIP BFD Network Topology



"CLI Quick Configuration" on page 4327 shows the configuration for all of the devices in Figure 130 on page 4327. The section "Step-by-Step Procedure" on page 4328 describes the steps on Device R1.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip

```

```
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step  
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Configure tracing operations to track BFD messages.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

**Results**

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
```

```

unit 1 {
    family inet {
        address 10.0.0.1/30;
    }
}

user@R1# show protocols
bfd {
    traceoptions {
        file bfd-trace;
        flag all;
    }
}
rip {
    group rip-group {
        export advertise-routes-through-rip;
        bfd-liveness-detection {
            minimum-interval 600;
        }
        neighbor fe-1/2/0.1;
    }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Up on page 4329](#)
- [Checking the BFD Trace File on page 4330](#)

### Verifying That the BFD Sessions Are Up

**Purpose** Make sure that the BFD sessions are operating.

**Action** From operational mode, enter the **show bfd session** command.

```
user@R1> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

1 sessions, 1 clients

Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

**Meaning** The output shows that there are no authentication failures.

### Checking the BFD Trace File

**Purpose** Use tracing operations to verify that BFD packets are being exchanged.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

**Meaning** The output shows the normal functioning of BFD.

**Related Documentation**

- [Example: Configuring RIP on page 4311](#)
- [Example: Configuring Authentication for RIP Routes on page 4318](#)
- [Example: Configuring Point-to-Multipoint RIP Networks](#)

### Example: Configuring BFD Authentication for RIP

- [Understanding BFD Authentication for RIP on page 4330](#)
- [Example: Configuring BFD Authentication for RIP on page 4332](#)

---

#### Understanding BFD Authentication for RIP

BFD enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over RIP. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and the level of authentication that can be configured:

- [BFD Authentication Algorithms on page 4331](#)
- [Security Authentication Keychains on page 4331](#)
- [Strict Versus Loose Authentication on page 4332](#)

### ***BFD Authentication Algorithms***

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.



**NOTE:** Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

### ***Security Authentication Keychains***

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and

associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

### ***Strict Versus Loose Authentication***

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

### **Example: Configuring BFD Authentication for RIP**

---

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for a RIP network.

- [Requirements on page 4332](#)
- [Overview on page 4332](#)
- [Configuration on page 4333](#)
- [Verification on page 4337](#)

### ***Requirements***

No special configuration beyond device initialization is required before configuring this example.

The devices must be running Junos OS Release 9.6 or later.


### ***Overview***

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the RIP protocol.
2. Associate the authentication keychain with the RIP protocol.
3. Configure the related security authentication keychain.

[Figure 131 on page 4333](#) shows the topology used in this example.

Figure 131: RIP BFD Authentication Network Topology



"CLI Quick Configuration" on page 4333 shows the configuration for all of the devices in Figure 131 on page 4333. The section "Step-by-Step Procedure" on page 4334 describes the steps on Device R1.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
    "$9$dlV2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
    "2012-2-16.12:00:00 -0800"

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct

```

```
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
  "$9$d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
  keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret
  "$9$d1V2aZGi.fzDiORSeXxDikqmT"
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD authentication:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```



5. Enable BFD.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection minimum-interval 600
```

6. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.



**NOTE:** Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication algorithm keyed-md5
```

7. Specify the keychain to be used to associate BFD sessions on RIP with the unique security authentication keychain attributes.

The keychain you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication key-chain bfd-rip
```

8. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication loose-check
```

9. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 7.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-rip]
user@R1# set key 53 secret "$9$d1V2aZGi.fzDiORSeXxDikqmT"
user@R1# set key 53 start-time "2012-2-16.12:00:00 -0800"
```

10. Configure tracing operations to track BFD authentication.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}

user@R1# show protocols
bfd {
  traceoptions {
    file bfd-trace;
    flag all;
  }
}
rip {
  group rip-group {
    export advertise-routes-through-rip;
    bfd-liveness-detection {
      minimum-interval 600;
    }
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

user@R1# show security
authentication-key-chains {
  key-chain bfd-rip {
    key 53 {
      secret "$9$d1V2aZGi.fzDiORSeXxDikqmT"; ## SECRET-DATA
      start-time "2012-2-16.12:00:00 -0800";
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Verification**

Confirm that the configuration is working properly.

- [Verifying That the BFD Sessions Are Authenticated on page 4337](#)
- [Viewing Extensive Information About the BFD Authentication on page 4337](#)
- [Checking the BFD Trace File on page 4338](#)

**Verifying That the BFD Sessions Are Authenticated**

**Purpose** Make sure that the BFD sessions are authenticated.

**Action** From operational mode, enter the **show bfd session detail** command.

```
user@R1> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**  
 Session up time 01:39:34  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Logical system 6, routing table index 53

1 sessions, 1 clients  
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

**Meaning** **Authenticate** is displayed to indicate that BFD authentication is configured.

**Viewing Extensive Information About the BFD Authentication**

**Purpose** View the keychain name, the authentication algorithm and mode for each client in the session, and the BFD authentication configuration status.

**Action** From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.0.0.2	Up	fe-1/2/0.1	1.800	0.600	3

Client RIP, TX interval 0.600, RX interval 0.600, **Authenticate**  
**keychain bfd-rip, algo keyed-md5, mode loose**  
 Session up time 01:46:29  
 Local diagnostic None, remote diagnostic None  
 Remote state Up, version 1  
 Logical system 6, routing table index 53  
 Min async interval 0.600, min slow interval 1.000  
 Adaptive async TX interval 0.600, RX interval 0.600  
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 3  
 Remote min TX interval 0.600, min RX interval 0.600, multiplier 3  
 Local discriminator 225, remote discriminator 226  
 Echo mode disabled/inactive  
**Authentication enabled/active, keychain bfd-rip, algo keyed-md5, mode loose**  
 Session ID: 0x300501

1 sessions, 1 clients  
 Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps

**Meaning** The output shows the keychain name, the authentication algorithm and mode for the client in the session, and the BFD authentication configuration status.

#### *Checking the BFD Trace File*

**Purpose** Use tracing operations to verify that BFD packets are being exchanged.

**Action** From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32   IfIndex (3) len 4: 0
Feb 16 10:26:32   Protocol (1) len 1: BFD
Feb 16 10:26:32   Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
78 6d 69 74 20 6f
...
```

**Meaning** The output shows the normal functioning of BFD.

- Related Documentation**
- [Example: Configuring BFD for RIP on page 4324](#)
  - [Example: Configuring Authentication for RIP Routes on page 4318](#)
  - [Example: Configuring RIP on page 4311](#)

## Example: Applying Policies to RIP Routes Imported from Neighbors

- [Understanding RIP Import Policy on page 4338](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4338](#)

### Understanding RIP Import Policy

The default RIP import policy is to accept all received RIP routes that pass a sanity check. To filter routes being imported by the local routing device from its neighbors, include the **import** statement, and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local routing device does not import any routes.

### Example: Applying Policies to RIP Routes Imported from Neighbors

This example shows how to configure an import policy in a RIP network.

- [Requirements on page 4339](#)
- [Overview on page 4339](#)

- [Configuration on page 4339](#)
- [Verification on page 4342](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.


### Overview

In this example, Device R1 has an import policy that accepts the 10/8 and 192.168/16 RIP routes and rejects all other RIP routes. This means that the 172.16/16 RIP routes are excluded from Device R1's routing table.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 132 on page 4339](#) shows the topology used in this example.

**Figure 132: RIP Import Policy Network Topology**



[“CLI Quick Configuration” on page 4339](#) shows the configuration for all of the devices in [Figure 132 on page 4339](#). The section [“Step-by-Step Procedure” on page 4340](#) describes the steps on Device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device R1
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip import rip-import
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set policy-options policy-statement rip-import term 1 from protocol rip
set policy-options policy-statement rip-import term 1 from route-filter 10.0.0.0/8 orlonger
set policy-options policy-statement rip-import term 1 from route-filter 192.168.0.0/16
  orlonger
```

```
set policy-options policy-statement rip-import term 1 then accept
set policy-options policy-statement rip-import term 2 then reject
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP import policy:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP group and add the interface.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled.

You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
```

```

user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept

```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```

[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip

```

5. Configure the import policy.

```

[edit policy-options policy-statement rip-import]
user@R1# set term 1 from protocol rip
user@R1# set term 1 from route-filter 10.0.0.0/8 orlonger
user@R1# set term 1 from route-filter 192.168.0.0/16 orlonger
user@R1# set term 1 then accept
user@R1# set term 2 then reject

```

6. Apply the import policy.

```

[edit protocols rip]
user@R1# set import rip-import

```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  import rip-import;
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {

```

```
term 1 {
    from protocol [ direct rip ];
    then accept;
}
}
policy-statement rip-import {
    term 1 {
        from {
            protocol rip;
            route-filter 10.0.0.0/8 orlonger;
            route-filter 192.168.0.0/16 orlonger;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

- [Looking at the Routes That Device R2 Is Advertising to Device R1 on page 4342](#)
- [Looking at the Routes That Device R1 Is Receiving from Device R2 on page 4343](#)
- [Checking the Routing Table on page 4343](#)
- [Testing the Import Policy on page 4343](#)

### **Looking at the Routes That Device R2 Is Advertising to Device R1**

**Purpose** Verify that Device R2 is sending the expected routes.

**Action** From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R2> show route advertising-protocol rip 10.0.0.2
```

```
inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.4/30      *[Direct/0] 2d 01:17:44
                  >   via fe-1/2/0.5
172.16.2.2/32    *[Direct/0] 2d 04:09:52
                  >   via lo0.2
172.16.3.3/32    *[RIP/100] 23:40:02, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
192.168.2.2/32   *[Direct/0] 2d 04:09:52
                  >   via lo0.2
192.168.3.3/32   *[RIP/100] 23:40:02, metric 2, tag 0
                  > to 10.0.0.6 via fe-1/2/0.5
```

**Meaning** Device R2 is sending 172.16/16 routes to Device R1.



*Looking at the Routes That Device R1 Is Receiving from Device R2*

**Purpose** Verify that Device R1 is receiving the expected routes.

**Action** From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[RIP/100] 01:06:03, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32  *[RIP/100] 01:06:03, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32  *[RIP/100] 01:06:03, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
```

**Meaning** The output shows that the 172.16/16 routes are excluded.

*Checking the Routing Table*

**Purpose** Verify that the routing table is populated with the expected routes.

**Action** From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip

inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30      *[RIP/100] 00:54:34, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32  *[RIP/100] 00:54:34, metric 2, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32  *[RIP/100] 00:54:34, metric 3, tag 0
                 > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32    *[RIP/100] 00:49:00, metric 1
                 MultiRecv
```

**Meaning** The output shows that the routes have been learned from Device R2 and Device R3.

If you delete or deactivate the import policy, the routing table contains the 172.16/16 routes.

*Testing the Import Policy*

**Purpose** By using the **test policy** command, monitor the number of rejected prefixes.

**Action** From operational mode, enter the **test policy rip-import 172.16/16** command.

```
user@R1> test policy rip-import 172.16/16
Policy rip-import: 0 prefix accepted, 1 prefix rejected
```

**Meaning** The output shows that the policy rejected one prefix.

**Related Documentation** • [Example: Configuring RIP on page 4311](#)

## Examples: Controlling Traffic with Metrics in a RIP Network

- [Understanding Traffic Control with Metrics in a RIP Network on page 4344](#)
- [Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 4345](#)
- [Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 4346](#)
- [Example: Configuring the Metric Value Added to Imported RIP Routes on page 4348](#)

---

### Understanding Traffic Control with Metrics in a RIP Network

To tune a RIP network and to control traffic flowing through the network, you increase or decrease the cost of the paths through the network. RIP provides two ways to modify the path cost: an incoming metric and an outgoing metric, which are each set to 1 by default. In other words, by default, the metric of routes that RIP imports from a neighbor or exports to a neighbor is incremented by 1. These routes include those learned from RIP as well as those learned from other protocols. The metrics are attributes that specify the cost of any route advertised through a host. By increasing or decreasing the metrics—and thus the cost—of links throughout the network, you can control packet transmission across the network.

The incoming metric modifies the cost of an individual segment when a route across the segment is imported into the routing table. For example, if you set the incoming metric on the segment to **3**, the individual segment cost along the link is changed from 1 to **3**. The increased cost affects all route calculations through that link. Other routes that were previously excluded because of a high hop count might now be included in the router's forwarding table.

The outgoing metric modifies the path cost for all the routes advertised out of a particular interface. Unlike the incoming metric, the outgoing metric modifies the routes that other routers are learning and thereby controls the way they send traffic.

If an exported route was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured in the outgoing metric for that group.

You might want to increase the metric of routes to decrease the likelihood that a particular route is selected and installed in the routing table. This process is sometimes referred to as *route poisoning*. Some reasons that you might want to poison a route are that the route is relatively expensive to use, or it has relatively low bandwidth.

A route with a higher metric than another route becomes the active route only when the lower-metric route becomes unavailable. In this way, the higher-metric route serves as a backup path.

One way to increase the metric of imported routes is to configure an import policy. Another way is to include the **metric-in** statement in the RIP neighbor configuration. One way to increase the metric of export routes is to configure an export policy. Another way is to include the **metric-out** statement in the RIP neighbor configuration.

### Example: Controlling Traffic in a RIP Network with an Incoming Metric

This example shows how to control traffic with an incoming metric.

- [Requirements on page 4345](#)
- [Overview on page 4345](#)
- [Configuration on page 4346](#)
- [Verification on page 4346](#)


#### Requirements

Before you begin, define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through the RIP routing exchanges. See “[Example: Configuring a Basic RIP Network](#)” on page 4312.

#### Overview

In this example, routes to Router D are received by Router A across both of its RIP-enabled interfaces as shown in [Figure 133 on page 4345](#). Because the route through Router B and the route through Router C have the same number of hops, both routes are imported into the forwarding table. However, because the T3 link from Router B to Router D has a higher bandwidth than the T1 link from Router C to Router D, you want traffic to flow from Router A through Router B to Router D.

Figure 133: Controlling Traffic in a RIP Network with the Incoming Metric



To force this flow, you can modify the route metrics as they are imported into Router A's routing table. By setting the incoming metric on the interface from Router A to Router C, you modify the metric on all routes received through that interface. Setting the incoming route metric on Router A changes only the routes in Router A's routing table, and affects only how Router A sends traffic to Router D. Router D's route selection is based on its own routing table, which, by default, includes no adjusted metric values.

In the example, Router C receives a route advertisement from Router D and readvertises the route to Router A. When Router A receives the route, it applies the incoming metric on the interface. Instead of incrementing the metric by 1 (the default), Router A increments it by 3 (the configured incoming metric), giving the route from Router A to Router D

through Router C a total path metric of 4. Because the route through Router B has a metric of 2, it becomes the preferred route for all traffic from Router A to Router D.

This example uses a RIP group called **alpha 1** on interface **ge-0/0/0**.

### Configuration

#### Step-by-Step Procedure

To control traffic with an incoming metric:

1. Enable RIP on the interface.  

```
[edit protocols rip]  
user@host# set group alpha1 neighbor ge-0/0/0
```
2. Set the incoming metric.  

```
[edit protocols rip]  
user@host# set metric-in 3
```
3. If you are done configuring the device, commit the configuration.  

```
[edit]  
user@host# commit
```

### Verification

To verify that the configuration is working properly, enter the **show route protocols rip** command.

---

### Example: Controlling Traffic in a RIP Network with an Outgoing Metric

This example shows how to control traffic with an outgoing metric.

- [Requirements on page 4346](#)
- [Overview on page 4346](#)
- [Configuration on page 4347](#)
- [Verification on page 4347](#)

### Requirements

Before you begin:


- Define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See [“Example: Configuring a Basic RIP Network” on page 4312](#).
- Control traffic with an incoming metric. See [“Example: Controlling Traffic in a RIP Network with an Incoming Metric” on page 4345](#).

### Overview

In this example, each route from Router A to Router D has two hops as shown in [Figure 134 on page 4347](#). However, because the link from Router A to Router B in the RIP group has a higher bandwidth than the link from Router A to Router C in RIP group Alpha 1, you want traffic from Router D to Router A to flow through Router B. To control the

way Router D sends traffic to Router A, you can alter the routes that Router D receives by configuring the outgoing metric on Router A's interfaces in the Alpha 1 RIP group.

**Figure 134: Controlling Traffic in a RIP Network with the Outgoing Metric**



If the outgoing metric for the Alpha 1 RIP group—the A-to-C link—is changed to 3, Router D calculates the total path metric from Router A through Router C as 4. In contrast, the unchanged default total path metric to Router A through Router B in the RIP group is 2. The fact that Router A's interfaces belong to two different RIP groups allows you to configure two different outgoing metrics on its interfaces, because you configure path metrics at the group level.

By configuring the outgoing metric, you control the way Router A sends traffic to Router D. By configuring the outgoing metric on the same router, you control the way Router D sends traffic to Router A.

This example uses an outgoing metric of 3.

### Configuration

#### Step-by-Step Procedure

To control traffic with an outgoing metric:

1. Set the outgoing metric.
 

```
[edit protocols rip group alpha1]
user@host# set metric-out 3
```
2. If you are done configuring the device, commit the configuration.
 

```
[edit]
user@host# commit
```

### Verification

To verify that the configuration is working properly, enter the **show protocols rip** command.

### Example: Configuring the Metric Value Added to Imported RIP Routes

This example shows how to change the default metric to be added to incoming routes to control the route selection process.

- [Requirements on page 4348](#)
- [Overview on page 4348](#)
- [Configuration on page 4348](#)
- [Verification on page 4351](#)

#### Requirements


No special configuration beyond device initialization is required before configuring this example.

#### Overview

Normally, when multiple routes are available, RIP selects the route with the lowest hop count. Changing the default metric enables you to control the route selection process such that a route with a higher hop count can be preferred over of a route with a lower hop count.

[Figure 135 on page 4348](#) shows the topology used in this example.

**Figure 135: RIP Incoming Metrics Network Topology**



Device R1 has two potential paths to reach 172.16.2.2/32. The default behavior is to send traffic out the 0.1/30 interface facing Device R2. Suppose, though, that the path through Device R3 is less expensive to use or has higher bandwidth links. This example shows how to use the **metric-in** statement to ensure that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. “[CLI Quick Configuration on page 4348](#)” shows the configuration for all of the devices in [Figure 135 on page 4348](#). The section “[Step-by-Step Procedure on page 4349](#)” describes the steps on Device R1.

#### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces ge-1/2/1 unit 10 description to-R3

```

```

set interfaces ge-1/2/1 unit 10 family inet address 10.0.0.10/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group primary export advertise-routes-through-rip
set protocols rip group primary neighbor ge-1/2/1.10
set protocols rip group secondary export advertise-routes-through-rip
set protocols rip group secondary neighbor fe-1/2/0.1 metric-in 4
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor ge-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R3**

```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces ge-1/2/1 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group neighbor ge-1/2/1.9
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP metrics:

1. Configure the network interfaces.

```

[edit interfaces]
user@R1# set fe-1/2/0 unit 1 description to-R2
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set ge-1/2/1 unit 10 description to-R3
user@R1# set ge-1/2/1 unit 10 family inet address 10.0.0.10/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32

```

```
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **metric-in 4** setting causes this route to be less likely to be chosen as the active route.

```
[edit protocols rip]
user@R1# set group primary neighbor ge-1/2/1.10
user@R1# set group secondary neighbor fe-1/2/0.1 metric-in 4
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip]
user@R1# set group primary export advertise-routes-through-rip
user@R1# set group secondary export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-R2;
    family inet {
      address 10.0.0.1/30;
    }
  }
}
ge-1/2/1 {
  unit 10 {
    description to-R3;
    family inet {
      address 10.0.0.10/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}
```



```

    }
  }
user@R1# show protocols
rip {
  group primary {
    export advertise-routes-through-rip;
    neighbor ge-1/2/1.10;
  }
  group secondary {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      metric-in 4;
    }
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying That the Expected Route Is Active on page 4351](#)
- [Removing the metric-in Statement on page 4351](#)

### Verifying That the Expected Route Is Active

**Purpose** Make sure that to reach 172.16.2.2/32, Device R1 uses the path through Device R3.

**Action** From operational mode, enter the **show route 172.16.2.2** command.

```

user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      *[RIP/100] 00:15:46, metric 3, tag 0
                   > to 10.0.0.9 via ge-1/2/1.10

```

**Meaning** The **to 10.0.0.9 via ge-1/2/1.10** output shows that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. The metric for this route is 3.

### Removing the metric-in Statement

**Purpose** Delete or deactivate the **metric-in** statement to see what happens to the 172.16.2.2/32 route.

**Action** 1. From configuration mode, deactivate the **metric-in** statement.

```
[edit protocols rip group secondary neighbor fe-1/2/0.1]
user@R1# deactivate metric-in
user@R1# commit
```

2. From operational mode, enter the **show route 172.16.2.2** command.

```
user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      * [RIP/100] 00:00:06, metric 2, tag 0
> to 10.0.0.2 via fe-1/2/0.1
```

**Meaning** The **to 10.0.0.2 via fe-1/2/0.1** output shows that Device R1 uses the path through Device R2 to reach 172.16.2.2/32. The metric for this route is 2.

**Related Documentation**

- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4338](#)

## Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

- [Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4352](#)
- [Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4352](#)

### Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets

---

RIP version 1 (RIPv1) and RIP version 2 (RIPv2) can run simultaneously. This might make sense when you are migrating a RIPv1 network to a RIPv2 network. This also allows interoperation with a device that supports RIPv1 but not RIPv2.

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets. You can configure this behavior by including the [send](#) and [receive](#) statements in the RIP configuration.

### Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

---

This example shows how to configure whether the RIP update messages conform to RIP version 1 (RIPv1) only, to RIP version 2 (RIPv2) only, or to both versions. You can also disable the sending or receiving of update messages.

- [Requirements on page 4352](#)
- [Overview on page 4352](#)
- [Configuration on page 4353](#)
- [Verification on page 4355](#)

#### Requirements


No special configuration beyond device initialization is required before configuring this example.

#### Overview

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets.

Figure 136 on page 4353 shows the topology used in this example.

**Figure 136: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology**



In this example, Device R1 is configured to receive only RIPv2 packets.

“CLI Quick Configuration” on page 4353 shows the configuration for all of the devices in Figure 136 on page 4353. The section “Step-by-Step Procedure” on page 4354 describes the steps on Device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

- Device R1**
- ```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1 receive version-2
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```
- Device R2**
- ```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```
- Device R3**
- ```

set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip

```

```
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP packet versions that can be received:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

```
user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

For the interface that is facing Device R2, the **receive version-2** setting causes this interface to accept only RIPv2 packets.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1 receive version-2
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
```

```

    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      receive version-2;
    }
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

#### Verifying That the Receive Mode Is Set to RIPv2 Only

**Purpose** Make sure that the interfacing Device R2 is configured to receive only RIPv2 packets, instead of both RIPv1 and RIPv2 packets.

**Action** From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
```

| Neighbor   | Local<br>State | Source<br>Address | Destination<br>Address | Send<br>Mode | Receive<br>Mode | In<br>Met |
|------------|----------------|-------------------|------------------------|--------------|-----------------|-----------|
| fe-1/2/0.1 | Up             | 10.0.0.1          | 224.0.0.9              | mcast        | v2 only         | 1         |

**Meaning** In the output, the **Receive Mode** field displays **v2 only**. The default **Receive Mode** is **both**.

**Related Documentation**

- [Example: Configuring RIP on page 4311](#)

## Example: Redistributing Routes Among RIP Instances

- [Understanding Route Redistribution Among RIP instances on page 4356](#)
- [Example: Redistributing Routes Between Two RIP Instances on page 4357](#)

### Understanding Route Redistribution Among RIP instances

---

You can redistribute routes among RIP processes. Another way to say this is to export RIP routes from one RIP instance to other RIP instances.

In Junos OS, route redistribution among routing instances is accomplished by using routing table groups, also called RIB groups. Routing table groups allow you to import and export routes from a protocol within one routing table into another routing table.



**NOTE:** In contrast, the policy-based import and export functions allow you import and export routes between different protocols within the same routing table.

---

Consider the following partial example:

```
protocols {
  rip {
    rib-group inet-to-voice;
  }
}
routing-instances {
  voice {
    protocols {
      rip {
        rib-group voice-to-inet;
      }
    }
  }
}
routing-options {
  rib-groups {
    inet-to-voice {
      import-rib [ inet.0 voice.inet.0 ];
    }
    voice-to-inet {
      import-rib [ voice.inet.0 inet.0 ];
    }
  }
}
```

The way to read the **import-rib** statement is as follows. Take the routes from the protocol (RIP, in this case), and import them into the primary (or local) routing table and also into any other routing tables listed after this. The primary routing table is the routing table where the routing table group is being used. That would be either **inet.0** if used in the main routing instance or **voice.inet.0** if used within the routing instance. In the **inet-to-voice** routing table group, **inet.0** is listed first because this routing table group is used in the

main routing instance. In the **voice-to-inet** routing table group, **voice.inet.0** is listed first because this routing table group is used in the voice routing instance.

### Example: Redistributing Routes Between Two RIP Instances

This example shows how to configure a RIP routing instance and control the redistribution of RIP routes between the routing instance and the master instance.

- [Requirements on page 4357](#)
- [Overview on page 4357](#)
- [Configuration on page 4357](#)
- [Verification on page 4361](#)

#### Requirements


No special configuration beyond device initialization is required before configuring this example.

#### Overview

When you create a routing instance called voice, Junos OS creates a routing table called **voice.inet.0**. The example shows how to install routes learned through the master RIP instance into the **voice.inet.0** routing table. The example also shows how to install routes learned through the voice routing instance into **inet.0**. This is done by configuring routing table groups. RIP routes are installed into each routing table that belongs to a routing table group.

[Figure 137 on page 4357](#) shows the topology used in this example.

**Figure 137: Redistributing Routes Between RIP Instances Network Topology**



[“CLI Quick Configuration” on page 4357](#) shows the configuration for all of the devices in [Figure 137 on page 4357](#). The section [“Step-by-Step Procedure” on page 4358](#) describes the steps on Device R2.

#### Configuration

##### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32

```

```
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip rib-group inet-to-voice
set protocols rip group to-R3 export advertise-routes-through-rip
set protocols rip group to-R3 neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set routing-instances voice protocols rip group to-R1 export advertise-routes-through-rip
set routing-instances voice interface fe-1/2/0.2
set routing-instances voice protocols rip rib-group voice-to-inet
set routing-instances voice protocols rip group to-R1 neighbor fe-1/2/0.2
set routing-options rib-groups inet-to-voice import-rib inet.0
set routing-options rib-groups inet-to-voice import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib inet.0
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To redistribute RIP routes between routing instances:

1. Configure the network interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30

user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
```



```
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Create the routing instance, and add one or more interfaces to the routing instance.

```
[edit routing-instances voice]
user@R2# set interface fe-1/2/0.2
```

3. Create the RIP groups and add the interfaces.

```
[edit protocols rip group to-R3]
user@R2# set neighbor fe-1/2/1.5

[edit routing-instances voice protocols rip group to-R1]
user@R2# set neighbor fe-1/2/0.2
```

4. Create the routing table groups.

```
[edit routing-options rib-groups]
user@R2# set inet-to-voice import-rib inet.0
user@R2# set inet-to-voice import-rib voice.inet.0

user@R2# set voice-to-inet import-rib voice.inet.0
user@R2# set voice-to-inet import-rib inet.0
```

5. Apply the routing table groups.

```
[edit protocols rip]
user@R2# set rib-group inet-to-voice

[edit routing-instances voice protocols rip]
user@R2# set rib-group voice-to-inet
```

6. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

7. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group to-R3]
user@R2# set export advertise-routes-through-rip

[edit routing-instances voice protocols rip group to-R1]
user@R2# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
```

```
        address 10.0.0.2/30;
    }
}
fe-1/2/1 {
    unit 5 {
        family inet {
            address 10.0.0.5/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.2.2/32;
            address 172.16.2.2/32;
        }
    }
}

user@R2# show protocols
rip {
    rib-group inet-to-voice;
    group to-R3 {
        export advertise-routes-through-rip;
        neighbor fe-1/2/1.5;
    }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}

user@R2# show routing-instances
voice {
    interface fe-1/2/0.2;
    protocols {
        rip {
            rib-group voice-to-inet;
            group to-R1 {
                export advertise-routes-through-rip;
                neighbor fe-1/2/0.2;
            }
        }
    }
}

user@R2# show routing-options
rib-groups {
    inet-to-voice {
        import-rib [ inet.0 voice.inet.0 ];
    }
    voice-to-inet {
```

```

import-rib [ voice.inet.0 inet.0 ];
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

### Checking the Routing Tables

**Purpose** Make sure that the routing tables contain the expected routes.

**Action** From operational mode, enter the **show route protocol rip** command.

```

user@R2> show route protocol rip
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32    * [RIP/100] 01:58:14, metric 2, tag 0
                 > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32    * [RIP/100] 02:06:03, metric 2, tag 0
                 > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32   * [RIP/100] 01:58:14, metric 2, tag 0
                 > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32   * [RIP/100] 02:06:03, metric 2, tag 0
                 > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32     * [RIP/100] 01:44:13, metric 1
                 MultiRecv

voice.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32    * [RIP/100] 02:06:03, metric 2, tag 0
                 > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32    * [RIP/100] 01:58:14, metric 2, tag 0
                 > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32   * [RIP/100] 02:06:03, metric 2, tag 0
                 > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32   * [RIP/100] 01:58:14, metric 2, tag 0
                 > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32     * [RIP/100] 01:44:13, metric 1
                 MultiRecv

```

**Meaning** The output shows that both routing tables contain all of the RIP routes.

**Related Documentation**

- [Example: Configuring RIP on page 4311](#)
- [Example: Applying Policies to RIP Routes Imported from Neighbors on page 4338](#)

## Example: Configuring RIP Timers

- [Understanding RIP Timers on page 4362](#)
- [Example: Configuring RIP Timers on page 4362](#)

## Understanding RIP Timers

---

RIP uses several timers to regulate its operation.

The update interval is the interval at which routes that are learned by RIP are advertised to neighbors. This timer controls the interval between routing updates. The update interval is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can occur if all routing devices update their neighbors simultaneously.

To configure the update time interval, include the **update-interval** statement:

**update-interval** *seconds*;

*seconds* can be a value from 10 through 60.

You can set a route timeout interval. If a route is not refreshed after being installed in the routing table by the specified time interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

To configure the route timeout for RIP, include the **route-timeout** statement:

**route-timeout** *seconds*;

*seconds* can be a value from 30 through 360. The default value is 180 seconds.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a specified period so that neighbors can be notified that the route has been dropped. This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

To configure the hold-down timer for RIP, include the **holddown** statement:

**holddown** *seconds*;

*seconds* can be a value from 10 through 180. The default value is 120 seconds.



**NOTE:** In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

---

Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. The route timeout should be at least three times the update interval. Normally, the default values are best left in effect for standard operations.

## Example: Configuring RIP Timers

---

This example shows how to configure the RIP update interval and how to monitor the impact of the change.

- [Requirements on page 4363](#)
- [Overview on page 4363](#)

- [Configuration on page 4363](#)
- [Verification on page 4366](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

### Overview


In this example, Device R2 has an update interval of 60 seconds for its neighbor, Device R1, and an update interval of 10 seconds for its neighbor, Device R3.

This example is not necessarily practical, but it is shown for demonstration purposes. Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. Normally, the default values are best left in effect for standard operations.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

[Figure 138 on page 4363](#) shows the topology used in this example.

**Figure 138: RIP Timers Network Topology**



“CLI Quick Configuration” on [page 4363](#) shows the configuration for all of the devices in [Figure 138 on page 4363](#). The section “Step-by-Step Procedure” on [page 4364](#) describes the steps on Device R2.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2 update-interval 60
set protocols rip group rip-group neighbor fe-1/2/1.5 update-interval 10
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
```

```
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30
```

```
user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30
```

```
user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
```

```
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Configure different update intervals for the two RIP neighbors.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
```

```
user@R2# set neighbor fe-1/2/0.2 update-interval 60
```

```
user@R2# set neighbor fe-1/2/1.5 update-interval 10
```

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
```

```
user@R2# set from protocol direct
```

```
user@R2# set from protocol rip
```

```
user@R2# set then accept
```

4. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R2# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet {
      address 10.0.0.2/30;
    }
  }
}
fe-1/2/1 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
lo0 {
  unit 2 {
    family inet {
      address 192.168.2.2/32;
      address 172.16.2.2/32;
    }
  }
}

user@R2# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.2 {
      update-interval 60;
    }
    neighbor fe-1/2/1.5 {
      update-interval 10;
    }
  }
}

user@R2# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

```
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Checking the RIP Updates Sent by Device R2 on page 4366](#)
- [Checking the RIP Updates Received by Device R2 on page 4367](#)
- [Checking the RIP Updates Received by Device R3 on page 4367](#)

### Checking the RIP Updates Sent by Device R2

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R2> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

```
    rts learned   rts held down   rqsts dropped   resps dropped
          4             2             0             0
```

```
fe-1/2/0.2: 2 routes learned; 5 routes advertised; timeout 180s; update interval
60s
```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| Updates Sent            | 123   | 5          | 1           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 244   | 10         | 2           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |
| none                    | 0     | 0          | 0           |

```
fe-1/2/1.5: 2 routes learned; 5 routes advertised; timeout 180s; update interval
10s
```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| Updates Sent            | 734   | 32         | 6           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 245   | 11         | 2           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |



|                      |   |   |   |
|----------------------|---|---|---|
| RIP Requests Ignored | 0 | 0 | 0 |
| none                 | 0 | 0 | 0 |

**Meaning** The **update interval** field shows that the interval is 60 seconds for Neighbor R1 and 10 seconds for Neighbor R3. The **Updates Sent** field shows that Device R2 is sending updates to Device R1 at roughly 1/6 of the rate that it is sending updates to Device R3.

### *Checking the RIP Updates Received by Device R2*

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

|             |               |               |               |
|-------------|---------------|---------------|---------------|
| rts learned | rts held down | rqsts dropped | resps dropped |
| 5           | 0             | 0             | 0             |

```
fe-1/2/0.1: 5 routes learned; 2 routes advertised; timeout 180s; update interval 30s
```

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 312   | 10         | 2           |
| Triggered Updates Sent  | 2     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 181   | 5          | 1           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 1     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |
| none                    | 0     | 0          | 0           |

**Meaning** The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

### *Checking the RIP Updates Received by Device R3*

**Purpose** Make sure that the RIP update packets are sent at the expected interval.

**Action** From operational mode, enter the **show rip statistics** command.

```
user@R3> show rip statistics
```

```
RIPv2 info: port 520; holddown 120s.
```

|             |               |               |               |
|-------------|---------------|---------------|---------------|
| rts learned | rts held down | rqsts dropped | resps dropped |
| 5           | 0             | 0             | 0             |

```
fe-1/2/0.6: 5 routes learned; 2 routes advertised; timeout 180s; update interval 30s
```

| Counter                | Total | Last 5 min | Last minute |
|------------------------|-------|------------|-------------|
| -----                  | ----- | -----      | -----       |
| Updates Sent           | 314   | 11         | 2           |
| Triggered Updates Sent | 1     | 0          | 0           |
| Responses Sent         | 0     | 0          | 0           |

|                         |     |    |   |
|-------------------------|-----|----|---|
| Bad Messages            | 0   | 0  | 0 |
| RIPv1 Updates Received  | 0   | 0  | 0 |
| RIPv1 Bad Route Entries | 0   | 0  | 0 |
| RIPv1 Updates Ignored   | 0   | 0  | 0 |
| RIPv2 Updates Received  | 827 | 31 | 6 |
| RIPv2 Bad Route Entries | 0   | 0  | 0 |
| RIPv2 Updates Ignored   | 0   | 0  | 0 |
| Authentication Failures | 0   | 0  | 0 |
| RIP Requests Received   | 0   | 0  | 0 |
| RIP Requests Ignored    | 0   | 0  | 0 |
| none                    | 0   | 0  | 0 |

**Meaning** The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

- Related Documentation**
- [Example: Configuring RIP on page 4311](#)
  - [Example: Configuring RIP Demand Circuits](#)

## Example: Tracing RIP Protocol Traffic

- [Understanding RIP Trace Operations on page 4368](#)
- [Example: Tracing RIP Protocol Traffic on page 4369](#)

### Understanding RIP Trace Operations

---

You can trace various types of RIP protocol traffic to help debug RIP protocol issues.

To trace RIP protocol traffic, include the **traceoptions** statement at the **[edit protocols rip]** hierarchy level:

```
traceoptions {  
  file filename <files number> <size size> <world-readable | no-world-readable>;  
  flag flag <flag-modifier> <disable>;  
}
```

You can specify the following RIP protocol-specific trace options using the **flag** statement:

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop active routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted



**NOTE:** Use the **detail** flag modifier with caution as this may cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the RIP protocol using the **traceoptions flag** statement included at the **[edit protocols rip]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing



**NOTE:** Use the trace flag **all** with caution because this may cause the CPU to become very busy.

### Example: Tracing RIP Protocol Traffic

This example shows how to trace RIP protocol operations.

- [Requirements on page 4369](#)
- [Overview on page 4370](#)
- [Configuration on page 4370](#)
- [Verification on page 4372](#)

#### Requirements

No special configuration beyond device initialization is required before configuring this example.


### Overview

In this example, Device R1 is set to trace routing information updates.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

Figure 139 on page 4370 shows the topology used in this example.

**Figure 139: RIP Trace Operations Network Topology**



“CLI Quick Configuration” on page 4370 shows the configuration for all of the devices in Figure 139 on page 4370. The section “Step-by-Step Procedure” on page 4371 describes the steps on Device R1.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

**Device R1**

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip traceoptions file rip-trace-file
set protocols rip traceoptions flag route
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept

```

**Device R2**

```

set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip

```

```
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Device R3**

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

This example shows multiple loopback interface addresses to simulate attached networks.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Configure the RIP group, and add the interface to the group.

To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

```
[edit protocols rip group rip-group]
user@R1# set neighbor fe-1/2/0.1
```

3. Configure RIP tracing operations.

```
[edit protocols rip traceoptions]
user@R1# set file rip-trace-file
user@R1# set flag route
```

4. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

5. Apply the routing policy.

In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}

user@R1# show protocols
rip {
  traceoptions {
    file rip-trace-file;
    flag route;
  }
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}

user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

### **Verification**

Confirm that the configuration is working properly.

### **Checking the Log File**

**Purpose** Make sure that the RIP route updates are logged in the configured log file.

**Action** 1. Deactivate the extra loopback interface address on Device R3.

```
[edit interfaces lo0 unit 3 family inet]
user@R3# deactivate address 172.16.3.3/32
user@R3# commit
```

- From operational mode on Device R1, enter the **show log rip-trace-file** command with the **| match 172.16.3.3** option.

```

user@R1> show log rip-trace-file | match 172.16.3.3
Mar  1 11:39:53.975192 Setting RIPv2 rtbit on route 172.16.3.3/32, tsi =
0xbb69228
Mar  1 11:39:59.847118 172.16.3.3/32: metric-in: 16, change: 3 -> 16; # gw:
1, pkt_upd_src 10.0.0.2, inx: 0, rte_upd_src 10.0.0.2
Mar  1 11:39:59.847568 CHANGE 172.16.3.3/32      nhid 591 gw 10.0.0.2
RIP      pref 100/0 metric 3/0 fe-1/2/0.1 <Delete Int>
Mar  1 11:39:59.847629 Best route to 172.16.3.3/32 got deleted. Doing route calculation
on the stored rte-info

```

**Meaning** The output shows that the route to 172.16.3.3/32 was deleted.

**Related Documentation**


- [Example: Configuring RIP on page 4311](#)

## RIP Configuration Statements

- [any-sender on page 4374](#)
- [authentication-key \(Protocols RIP\) on page 4375](#)
- [authentication-type \(Protocols RIP\) on page 4376](#)
- [bfd-liveness-detection \(Protocols RIP\) on page 4377](#)
- [check-zero on page 4380](#)
- [export \(Protocols RIP\) on page 4381](#)
- [group \(Protocols RIP\) on page 4382](#)
- [holddown \(Protocols RIP\) on page 4384](#)
- [import \(Protocols RIP\) on page 4385](#)
- [message-size on page 4386](#)
- [metric-in \(Protocols RIP\) on page 4387](#)
- [metric-out \(Protocols RIP\) on page 4388](#)
- [neighbor \(Protocols RIP\) on page 4389](#)
- [preference \(Protocols RIP\) on page 4390](#)
- [receive \(Protocols RIP\) on page 4391](#)
- [rib-group \(Protocols RIP\) on page 4392](#)
- [rip on page 4392](#)
- [route-timeout \(Protocols RIP\) on page 4393](#)
- [send \(Protocols RIP\) on page 4394](#)
- [traceoptions \(Protocols RIP\) on page 4395](#)
- [update-interval \(Protocols RIP\) on page 4398](#)

## any-sender

---

|  |   |
|--|---|
| <b>Syntax</b>  | any-sender;   |
| <b>Hierarchy Level</b>   | [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i> ] |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 8.0.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.   |
| <b>Description</b>   | <p>Disable strict sender address checks.</p> <p>If the sender of a RIP message does not belong to the subnet of the interface, the message is discarded. This situation might cause problems with dropped packets when RIP is running on point-to-point interfaces, or when the addresses on the interfaces do not fall in the same subnet. You can resolve this by disabling strict address checks on the RIP traffic.</p>   |
| <div> <b>NOTE:</b> The <b>any-sender</b> statement is supported only for peer-to-peer interfaces.</div> |   |
| <b>Required Privilege Level</b>  | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4311</a></li></ul>   |



## authentication-key (Protocols RIP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>authentication-key password;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Require authentication for RIP route queries received on an interface.   |
| <b>Options</b>                  | <p><b>password</b>—Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Route Authentication for RIP on page 4318</a></li> </ul>   |

## authentication-type (Protocols RIP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>authentication-type type;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure the type of authentication for RIP route queries received on an interface.  |
| <b>Default</b>                  | If you do not include this statement and the <b>authentication-key</b> statement, RIP authentication is disabled.   |
| <b>Options</b>                  | <b>type</b> —Authentication type: <ul style="list-style-type: none"><li>• <b>md5</b>—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.</li><li>• <b>none</b>—Disable authentication. If <b>none</b> is configured, the configured authentication key is ignored.</li><li>• <b>simple</b>—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.</li></ul>  |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Route Authentication for RIP on page 4318</a></li><li>• <a href="#">authentication-key on page 4375</a></li></ul>  |

## bfd-liveness-detection (Protocols RIP)

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre> bfd-liveness-detection {     authentication {         algorithm <i>algorithm-name</i>;         key-chain <i>key-chain-name</i>;         loose-check;     }     detection-time {         threshold <i>milliseconds</i>;     }     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     multiplier <i>number</i>;     no-adaptation;     transmit-interval {         minimum-interval <i>milliseconds</i>;         threshold <i>milliseconds</i>;     }     version (1   automatic); } </pre>   |
| <b>Hierarchy Level</b>     | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],<br/> [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols<br/> rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],<br/> [edit protocols rip <b>group</b> <i>group-name</i>],<br/> [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b><br/> <i>neighbor-name</i>]</p>   |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.0.</p> <p>Options <b>detection-time threshold</b> and <b>transmit-interval threshold</b> introduced in Junos OS Release 8.2.</p> <p>Support for logical systems introduced in Junos OS Release 8.3.</p> <p>Option <b>no-adaptation</b> introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6.</p> <p>Options <b>authentication algorithm</b>, <b>authentication key-chain</b>, and <b>authentication loose-check</b> introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> |
| <b>Description</b>         | Configure bidirectional failure detection timers and authentication.   |
| <b>Options</b>             | <p><b>authentication algorithm <i>algorithm-name</i></b> —Configure the algorithm used to authenticate the specified BFD session: <b>simple-password</b>, <b>keyed-md5</b>, <b>keyed-sha-1</b>, <b>meticulous-keyed-md5</b>, or <b>meticulous-keyed-sha-1</b>.</p> <p><b>authentication key-chain <i>key-chain-name</i></b>—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the <b>authentication-key-chains key-chain</b> statement at the <b>[edit security]</b> hierarchy level.</p>  |

**authentication loose-check**—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

**detection-time threshold *milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

**minimum-interval *milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval**, **minimum-interval**, and **minimum-receive-interval** statements.

**Range:** 1 through 255,000 milliseconds

**minimum-receive-interval *milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

**Range:** 1 through 255,000 milliseconds

**multiplier *number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

**Range:** 1 through 255

**Default:** 3

**no-adaptation**—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

**transmit-interval threshold *milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )

**transmit-interval minimum-interval *milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

**Range:** 1 through 255,000

**version**—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

**Default:** automatic

|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |


- Related Documentation**
- [Example: Configuring BFD for RIP on page 4325](#)
  - [Example: Configuring BFD Authentication for RIP on page 4332](#)

## check-zero

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | (check-zero   no-check-zero);   |
| <b>Hierarchy Level</b>          | <pre>[edit logical-systems <i>logical-system-name</i> protocols <i>rip</i>], [edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>rip</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit protocols <i>rip</i>], [edit protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>rip</i>], [edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <i>neighbor</i> <i>neighbor-name</i>]</pre>   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Some of the reserved fields in RIP version 1 packets must be zero, whereas in RIP version 2 packets, most of these reserved fields can contain nonzero values. By default, RIP discards version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</p> <p>If you find that you are receiving RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero, you can configure RIP to receive these packets even though they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</p> <p>Check whether the reserved fields in a RIP packet are zero:</p> <ul style="list-style-type: none"><li>• <b>check-zero</b>—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.</li><li>• <b>no-check-zero</b>—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.</li></ul> |
| <b>Default</b>                  | check-zero  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4311</a></li></ul>   |

## export (Protocols RIP)

|   |   |
|---|---|
| <b>Syntax</b>   | <code>export [ <i>policy-names</i> ];</code>  |
| <b>Hierarchy Level</b>  | <p>[edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit protocols rip <b>group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>]</p>  |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>  | <p>Apply a policy to routes being exported to the neighbors.</p> <p>By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.</p> <p>If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the <b>metric-in</b> and <b>metric-out</b> statements.</p> |
| <div>  <b>NOTE:</b> The export policy on RIP does not support manipulating routing information of the next hop.         </div> |   |
| <b>Options</b>  | <i>policy-names</i> —Name of one or more policies.  |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 4311</a></li> <li>• <a href="#">import on page 4385</a></li> </ul>  |

## group (Protocols RIP)

---

```
Syntax  group group-name {  
        bfd-liveness-detection {  
            authentication {  
                algorithm algorithm-name;  
                key-chain key-chain-name;  
                loose-check;  
            }  
            detection-time {  
                threshold milliseconds;  
            }  
            minimum-interval milliseconds;  
            minimum-receive-interval milliseconds;  
            transmit-interval {  
                threshold milliseconds;  
                minimum-interval milliseconds;  
            }  
            multiplier number;  
            version (0 | 1 | automatic);  
        }  
        demand-circuit;  
        export policy;  
        max-retrans-time seconds;  
        metric-out metric;  
        preference number;  
        route-timeout seconds;  
        update-interval seconds;  
        neighbor neighbor-name {  
            authentication-key password;  
            authentication-type type;  
            bfd-liveness-detection {  
                authentication {  
                    algorithm algorithm-name;  
                    key-chain key-chain-name;  
                    loose-check;  
                }  
                detection-time {  
                    threshold milliseconds;  
                }  
                minimum-interval milliseconds;  
                minimum-receive-interval milliseconds;  
                transmit-interval {  
                    threshold milliseconds;  
                    minimum-interval milliseconds;  
                }  
                multiplier number;  
                version (0 | 1 | automatic);  
            }  
            (check-zero | no-check-zero);  
            demand-circuit;  
            import policy-name;  
            max-retrans-time seconds;  
            message-size number;
```



```

metric-in metric;
metric-out metric;
receive receive-options;
route-timeout seconds;
send send-options;
update-interval seconds;
}
}

```

|                                 |  |
|---------------------------------|--|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.   |
| <b>Options</b>                  | <p><b><i>group-name</i></b>—Name of a group, up to 16 characters long.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 4311</a></li> </ul>  |

## holddown (Protocols RIP)


---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>holddown seconds;</code>   |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | <p>Configure how long the expired route is retained in the routing table before being removed.</p> <p>When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.</p> |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 180 seconds<br><b>Default:</b> 120 seconds  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 4362</a></li><li>• <a href="#">RIP Demand Circuits Overview</a></li></ul>  |

## import (Protocols RIP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>import [ <i>policy-names</i> ];</code>   |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Apply one or more policies to routes being imported by the local routing device from neighbors.  |
| <b>Options</b>                  | <i>policy-names</i> —Name of one or more policies.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Applying Policies to RIP Routes Imported from Neighbors on page 4338</a></li> <li>• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i></li> <li>• <a href="#">export on page 4381</a></li> </ul>  |

## message-size

|  |  |
|--|--|
| <b>Syntax</b>  | <code>message-size <i>number</i>;</code>   |
| <b>Hierarchy Level</b>   | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>   | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement for SRX Series devices introduced in Junos OS Release 9.5.</p> <p>Statement for J Series platform introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>   | Specify the number of route entries to be included in every RIP update message.  |
| <div>  <p><b>TIP:</b> To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message. Do not change the default number of route entries in a RIP update message.</p> </div> |  |
| <b>Options</b>   | <p><i>number</i>—Number of route entries per update message.</p> <p><b>Range:</b> 25 through 255 entries</p> <p><b>Default:</b> 25 entries</p>   |
| <b>Required Privilege Level</b>  | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP on page 4311</a></li> </ul>  |

## metric-in (Protocols RIP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>metric-in <i>metric</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.   |
| <b>Options</b>                  | <p><i>metric</i>—Metric value.</p> <p><b>Range:</b> 1 through 16</p> <p><b>Default:</b> 1</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Metric Value Added to Imported RIP Routes on page 4348</a></li> </ul>  |

## metric-out (Protocols RIP)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>metric-out <i>metric</i>;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <b>neighbor</b> <i>neighbor-name</i>]</code>  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | <p>Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.</p> <p>If you have included the <b>export</b> statement, RIP exports routes it has learned to the neighbors configured by including the <b>neighbor</b> statement.</p> <p>The metric associated with a RIP route (unless modified by an export policy) is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with a <b>metric-in</b> value of 2 is advertised with a combined metric of 7 when advertised to RIP neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the <b>metric-out</b> statement.</p> <p>The metric for a route can be modified with an export policy. That metric is seen when the route is exported to the next hop.</p> <p>To increase the metric for routes advertised outside a group, include the <b>metric-out</b> statement.</p> |
| <b>Options</b>                  | <b><i>metric</i></b> —Metric value.<br><b>Range:</b> 1 through 16<br><b>Default:</b> 1   |
| <b>Required Privilege Level</b> | <b>routing</b> —To view this statement in the configuration.<br><b>routing-control</b> —To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Examples: Controlling Traffic with Metrics in a RIP Network on page 434450</a></li></ul>   |

## neighbor (Protocols RIP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre> neighbor <i>neighbor-name</i> {   authentication-key <i>password</i>;   authentication-type <i>type</i>;   bfd-liveness-detection {     authentication {       algorithm <i>algorithm-name</i>;       key-chain <i>key-chain-name</i>;       loose-check;     }     detection-time {       threshold <i>milliseconds</i>;     }     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     transmit-interval {       threshold <i>milliseconds</i>;       minimum-interval <i>milliseconds</i>;     }     multiplier <i>number</i>;     version (0   1   automatic);   }   (check-zero   no-check-zero);   demand-circuit;   import <i>policy-name</i>;   max-retrans-time <i>seconds</i>;   message-size <i>number</i>;   metric-in <i>metric</i>;   metric-out <i>metric</i>;   receive <i>receive-options</i>;   route-timeout <i>seconds</i>;   send <i>send-options</i>;   update-interval <i>seconds</i>; } </pre> |
| <b>Hierarchy Level</b>     | <pre> [edit logical-systems <i>logical-system-name</i> protocols rip <b>group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   rip <b>group</b> <i>group-name</i>], [edit protocols rip <b>group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols rip <b>group</b> <i>group-name</i>] </pre>   |
| <b>Release Information</b> | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>  |
| <b>Description</b>         | Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.   |
| <b>Options</b>             | <p><b><i>neighbor-name</i></b>—Name of an interface over which a routing device communicates to its neighbors.</p> <p>The remaining statements are explained separately.</p>  |

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring RIP on page 4311](#)

---

## preference (Protocols RIP)

---

**Syntax** `preference preference;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
rip **group** *group-name*],  
[edit protocols rip **group** *group-name*],  
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.

**Description** Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

**Options** *preference*—Preference value. A lower value indicates a more preferred route.  
**Range:** 0 through 4,294,967,295 ( $2^{32} - 1$ )  
**Default:** 100

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Route Preferences Overview](#)



## receive (Protocols RIP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>receive receive-options;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor neighbor-name</a>]</p> |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Configure RIP receive options.   |
| <b>Options</b>                  | <p><i>receive-options</i>—One of the following:</p> <ul style="list-style-type: none"> <li>• <b>both</b>—Accept both RIP version 1 and version 2 packets.</li> <li>• <b>none</b>—Do not receive RIP packets.</li> <li>• <b>version-1</b>—Accept only RIP version 1 packets.</li> <li>• <b>version-2</b>—Accept only RIP version 2 packets.</li> </ul> <p><b>Default:</b> <b>both</b></p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4352</a></li> <li>• <a href="#">send on page 4394</a></li> </ul>  |

## rib-group (Protocols RIP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>rib-group group-name;</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Install RIP routes into multiple routing tables by configuring a routing table group.   |
| <b>Options</b>                  | <i>group-name</i> —Name of the routing table group.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Redistributing Routes Between Two RIP Instances on page 4357</a></li></ul>   |

## rip

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>rip {...}</code>  |
| <b>Hierarchy Level</b>          | [edit logical-systems <i>logical-system-name</i> protocols],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols],<br>[edit protocols],<br>[edit routing-instances <i>routing-instance-name</i> protocols] |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Enable RIP routing on the routing device.   |
| <b>Default</b>                  | RIP is disabled on the routing device.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP on page 4311</a></li></ul>   |

## route-timeout (Protocols RIP)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>route-timeout seconds;</code>  |
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group group-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>],</p> <p>[edit protocols <a href="#">rip</a>],</p> <p>[edit protocols rip <a href="#">group group-name</a>],</p> <p>[edit protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols rip <a href="#">group group-name</a> neighbor <i>neighbor-name</i>]</p> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.   |
| <b>Options</b>                  | <p><b>seconds</b>—Estimated time to wait before making updates to the routing table.</p> <p><b>Range:</b> 30 through 360 seconds</p> <p><b>Default:</b> 180 seconds</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring RIP Timers on page 4362</a></li> <li>• <a href="#">RIP Demand Circuits Overview</a></li> </ul>   |

## send (Protocols RIP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>send <i>send-options</i>;</code>  |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>  rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols rip group <i>group-name</i> <a href="#">neighbor</a> <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols rip group <i>group-name</i> <a href="#">neighbor</a></code><br><code>  <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure RIP send options.   |
| <b>Options</b>                  | <i>send-options</i> —One of the following: <ul style="list-style-type: none"><li>• <b>broadcast</b>—Broadcast RIP version 2 packets (RIP version 1 compatible).</li><li>• <b>multicast</b>—Multicast RIP version 2 packets. This is the default.</li><li>• <b>none</b>—Do not send RIP updates.</li><li>• <b>version-1</b>—Broadcast RIP version 1 packets.</li></ul> <b>Default:</b> multicast   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 4352</a></li><li>• <a href="#">receive on page 4391</a></li></ul>   |

## traceoptions (Protocols RIP)

|                            |   |
|----------------------------|---|
| <b>Syntax</b>              | <pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>   |
| <b>Hierarchy Level</b>     | [edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> ],<br>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ],<br>[edit protocols <a href="#">rip</a> ],<br>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> ] |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>         | Set RIP protocol-level tracing options.   |



**NOTE:** The `traceoptions` statement is not supported on QFabric systems.

**Default** The default RIP protocol-level trace options are inherited from the global `traceoptions` statement.

**Options** **disable**—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

**file *filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file `/var/log/rip-log`.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

### RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets

- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

#### Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the log file.

|                              |  |
|------------------------------|--|
| <b>Required Privilege</b>    | routing—To view this statement in the configuration.   |
| <b>Level</b>                 | routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <a href="#">Example: Tracing RIP Protocol Traffic on page 4369</a></li> </ul> |

## update-interval (Protocols RIP)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>update-interval seconds;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code><br><code>    <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit protocols <a href="#">rip</a>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit protocols <a href="#">rip</a> group <i>group-name</i> neighbor <i>neighbor-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i>],</code><br><code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">rip</a> group <i>group-name</i> neighbor</code><br><code>    <i>neighbor-name</i>]</code> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 7.6.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.   |
| <b>Options</b>                  | <b>seconds</b> —Estimated time to wait before making updates to the routing table.<br><b>Range:</b> 10 through 60 seconds<br><b>Default:</b> 30 seconds   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring RIP Timers on page 4362</a></li></ul>  |



## CHAPTER 52

# Administration

- [Routine Monitoring on page 4399](#)
- [RIP Operational Commands on page 4399](#)

## Routine Monitoring

---

- [Monitoring RIP Routing Information on page 4399](#)

### Monitoring RIP Routing Information

**Purpose** Use the monitoring functionality to monitor RIP routing on routing devices.

**Action** To view RIP routing information in the CLI, enter the following CLI commands:

- **show rip statistics**
- **show rip neighbor**

**Related Documentation**

- [show rip neighbor on page 4404](#)
- [show rip statistics on page 4406](#)

## RIP Operational Commands

---

- [clear rip general-statistics](#)
- [clear rip statistics](#)
- [show rip general-statistics](#)
- [show rip neighbor](#)
- [show rip statistics](#)

## clear rip general-statistics

---

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4400</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4400</a>  |
| <b>Syntax</b>                                     | clear rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip general-statistics  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.                |
| <b>Description</b>                                | Clear RIP general statistics.   |
| <b>Options</b>                                    | <b>none</b> —Clear RIP general statistics.<br><br><b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | clear   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">show rip general-statistics on page 4402</a></li></ul>  |
| <b>List of Sample Output</b>                      | <a href="#">clear rip general-statistics on page 4400</a>   |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.   |

## Sample Output

### clear rip general-statistics

```
user@host> clear rip general-statistics
```

## clear rip statistics

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4401</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4401</a>   |
| <b>Syntax</b>                                     | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><logical-system (all   <i>logical-system-name</i> )><br><neighbor><br><peer (all   <i>address</i> )>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | clear rip statistics<br><instance (all   <i>instance-name</i> )><br><neighbor>   |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>                                | Clear RIP statistics.  |
| <b>Options</b>                                    | <p><b>none</b>—Reset RIP counters for all neighbors for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear RIP statistics for the specified neighbor only.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Clear RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | clear  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li>• <a href="#">show rip statistics on page 4406</a></li> </ul>   |
| <b>List of Sample Output</b>                      | <a href="#">clear rip statistics on page 4401</a>  |
| <b>Output Fields</b>                              | When you enter this command, you are provided feedback on the status of your request.  |

## Sample Output

### clear rip statistics

```
user@host> clear rip statistics
```

## show rip general-statistics

|   |  |
|---|--|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4402</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4402</a>   |
| <b>Syntax</b>                                     | show rip general-statistics<br><logical-system (all   <i>logical-system-name</i> )>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | show rip general-statistics  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series. |
| <b>Description</b>                                | Display brief RIP statistics.  |
| <b>Options</b>                                    | none—Display brief RIP statistics.<br><br>logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system. |
| <b>Required Privilege Level</b>                   | view   |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"> <li><a href="#">clear rip general-statistics on page 4400</a></li> </ul>  |
| <b>List of Sample Output</b>                      | <a href="#">show rip general-statistics on page 4402</a>   |
| <b>Output Fields</b>                              | Table 337 on page 4402 lists the output fields for the <b>show rip general-statistics</b> command. Output fields are listed in the approximate order in which they appear.               |

Table 337: show rip general-statistics Output Fields

| Field Name  | Field Description                                      |
|-------------|--|
| bad msgs    | Number of invalid messages received.                   |
| no rcv intf | Number of packets received with no matching interface. |
| curr memory | Amount of memory currently used by RIP.                |
| max memory  | Most memory used by RIP.                               |

## Sample Output

### show rip general-statistics

```
user@host> show rip general-statistics
```

```
RIPv2 I/O info:
  bad msgs      :      0
  no recv intf  :      0
  curr memory   :      0
  max memory    :      0
```

## show rip neighbor

---

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4404</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4404</a>  |
| <b>Syntax</b>                                     | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;name&gt;</pre>   |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show rip neighbor &lt;instance (all   <i>instance-name</i>)&gt; &lt;name&gt;</pre>   |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>                                | Display information about RIP neighbors.  |
| <b>Options</b>                                    | <p><b>none</b>—Display information about all RIP neighbors for all instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name</b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> |
| <b>Required Privilege Level</b>                   | view  |
| <b>List of Sample Output</b>                      | <a href="#">show rip neighbor on page 4405</a><br><a href="#">show rip neighbor (With Demand Circuits Configured) on page 4405</a>  |
| <b>Output Fields</b>                              | <a href="#">Table 338 on page 4405</a> lists the output fields for the <b>show rip neighbor</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 338: show rip neighbor Output Fields

| Field Name                 | Field Description   |
|----------------------------|---|
| <b>Neighbor</b>            | <p>Name of the RIP neighbor.</p> <p><b>NOTE:</b> Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.</p> <p>If you configure demand circuits at the <b>[edit protocols rip group group-name neighbor neighbor-name]</b> hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the <b>[edit protocols rip group group-name]</b> hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.</p> |
| <b>State</b>               | State of the connection: <b>Up</b> or <b>Dn</b> (Down).   |
| <b>Source Address</b>      | Address of the port on the local router.  |
| <b>Destination Address</b> | Address of the port on the remote router.   |
| <b>Send Mode</b>           | Send options: <b>broadcast</b> , <b>multicast</b> , <b>none</b> , or <b>version 1</b> .   |
| <b>Receive Mode</b>        | Type of packets to accept: <b>both</b> , <b>none</b> , <b>version 1</b> , or <b>version 2</b> .   |
| <b>In Met</b>              | Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.  |

## Sample Output

### show rip neighbor

```

user@host> show rip neighbor
Neighbor      Local  Source      Destination  Send  Receive  In
-----      -
ge-2/3/0.0    Up    192.168.9.105  192.168.9.107  bcast  both      1
at-5/1/1.42    Dn    (null)        (null)        mcast  v2 only   3
at-5/1/0.42    Dn    (null)        (null)        mcast  both      3
at-5/1/0.0     Up    20.0.0.1      224.0.0.9     mcast  both      3
so-0/0/0.0     Up    192.168.9.97  224.0.0.9     mcast  both      3

```

### show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor      Local  Source      Destination  Send  Receive  In
-----      -
so-0/1/0.0(DC) Up    10.10.10.2   224.0.0.9     mcast  both      1
so-0/2/0.0(DC) Up    13.13.13.2   224.0.0.9     mcast  both      1

```

## show rip statistics

---

|   |   |
|---|---|
| <b>List of Syntax</b>                             | <a href="#">Syntax on page 4406</a><br><a href="#">Syntax (EX Series Switches and QFX Series) on page 4406</a>  |
| <b>Syntax</b>                                     | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>name</i>&gt; &lt;peer (all   <i>address</i>)&gt;</pre>  |
| <b>Syntax (EX Series Switches and QFX Series)</b> | <pre>show rip statistics &lt;instance (all   <i>instance-name</i>)&gt; &lt;<i>name</i>&gt;</pre>  |
| <b>Release Information</b>                        | Command introduced before Junos OS Release 7.4.<br>Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>                                | Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.  |
| <b>Options</b>                                    | <p><b>none</b>—Display RIP statistics for all routing instances.</p> <p><b>instance (all   <i>instance-name</i>)</b>—(Optional) Display RIP statistics for all instances or for only the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>name</i></b>—(Optional) Display detailed information about only the specified RIP neighbor.</p> <p><b>peer (all   <i>address</i>)</b>—(Optional) Display RIP statistics for a single peer or all peers.</p> |
| <b>Required Privilege Level</b>                   | view  |
| <b>Related Documentation</b>                      | <ul style="list-style-type: none"><li>• <a href="#">clear rip statistics on page 4401</a></li></ul>   |
| <b>List of Sample Output</b>                      | <a href="#">show rip statistics on page 4407</a>  |
| <b>Output Fields</b>                              | <a href="#">Table 339 on page 4407</a> lists the output fields for the <b>show rip statistics</b> command. Output fields are listed in the approximate order in which they appear.  |



Table 339: show rip statistics Output Fields

| Field Name               | Field Description  |
|--------------------------|--|
| <b>RIP info</b>          | <p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> <li>• <b>port</b>—UDP port number used for RIP.</li> <li>• <b>update interval</b>—Interval between routing table updates, in seconds.</li> <li>• <b>holddown</b>—Hold-down interval, in seconds.</li> <li>• <b>timeout</b>—Timeout interval, in seconds.</li> <li>• <b>restart in progress</b>—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart.</li> <li>• <b>restart time</b>—Estimated time for the graceful restart to finish, in seconds.</li> <li>• <b>restart will complete in</b>—Remaining time for the graceful restart to finish, in seconds.</li> <li>• <b>rts learned</b>—Number of routes learned through RIP.</li> <li>• <b>rts held down</b>—Number of routes held down by RIP.</li> <li>• <b>rqsts dropped</b>—Number of received request packets that were dropped.</li> <li>• <b>resps dropped</b>—Number of received response packets that were dropped.</li> </ul>   |
| <b>logical-interface</b> | <p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> <li>• <b>routes learned</b>—Number of routes learned on the logical interface.</li> <li>• <b>routes advertised</b>—Number of routes advertised by the logical interface.</li> </ul>  |
| <b>Counter</b>           | <p>List of counter types:</p> <ul style="list-style-type: none"> <li>• <b>Updates Sent</b>—Number of update messages sent.</li> <li>• <b>Triggered Updates Sent</b>—Number of triggered update messages sent.</li> <li>• <b>Responses Sent</b>—Number of response messages sent.</li> <li>• <b>Bad Messages</b>—Number of invalid messages received.</li> <li>• <b>RIPv1 Updates Received</b>—Number of RIPv1 update messages received.</li> <li>• <b>RIPv1 Bad Route Entries</b>—Number of RIPv1 invalid route entry messages received.</li> <li>• <b>RIPv1 Updates Ignored</b>—Number of RIPv1 update messages ignored.</li> <li>• <b>RIPv2 Updates Received</b>—Number of RIPv2 update messages received.</li> <li>• <b>RIPv2 Bad Route Entries</b>—Number of RIPv2 invalid route entry messages received.</li> <li>• <b>RIPv2 Updates Ignored</b>—Number of RIPv2 update messages ignored.</li> <li>• <b>Authentication Failures</b>—Number of received update messages that failed authentication.</li> <li>• <b>RIP Requests Received</b>—Number of RIP request messages received.</li> <li>• <b>RIP Requests Ignored</b>—Number of RIP request messages ignored.</li> </ul> |
| <b>Total</b>             | Total number of packets for the selected counter.  |
| <b>Last 5 min</b>        | Number of packets for the selected counter in the most recent 5-minute period.   |
| <b>Last minute</b>       | Number of packets for the selected counter in the most recent 1-minute period.   |

## Sample Output

### show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s  
restart in progress: restart time 60s; restart will complete in 55s  
      rts learned  rts held down  rqsts dropped  resps dropped  
                  0              0              0              0

so-0/0/0.0: 0 routes learned; 501 routes advertised

| Counter                 | Total | Last 5 min | Last minute |
|-------------------------|-------|------------|-------------|
| -----                   | ----- | -----      | -----       |
| Updates Sent            | 0     | 0          | 0           |
| Triggered Updates Sent  | 0     | 0          | 0           |
| Responses Sent          | 0     | 0          | 0           |
| Bad Messages            | 0     | 0          | 0           |
| RIPv1 Updates Received  | 0     | 0          | 0           |
| RIPv1 Bad Route Entries | 0     | 0          | 0           |
| RIPv1 Updates Ignored   | 0     | 0          | 0           |
| RIPv2 Updates Received  | 0     | 0          | 0           |
| RIPv2 Bad Route Entries | 0     | 0          | 0           |
| RIPv2 Updates Ignored   | 0     | 0          | 0           |
| Authentication Failures | 0     | 0          | 0           |
| RIP Requests Received   | 0     | 0          | 0           |
| RIP Requests Ignored    | 0     | 0          | 0           |

## PART 15

# MPLS Applications

- [Overview on page 4411](#)
- [Configuration on page 4433](#)
- [Administration on page 4565](#)
- [Troubleshooting on page 4733](#)



## CHAPTER 53

# Overview

- [MPLS Overview on page 4411](#)
- [MPLS Features on page 4423](#)
- [Introduction to LDP for QFX5100 on page 4426](#)

## MPLS Overview

---

- [MPLS Overview on page 4411](#)
- [Understanding MPLS Components on page 4412](#)
- [Understanding MPLS Label Operations on page 4415](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on page 4422](#)

## MPLS Overview

You can configure Multiprotocol Label Switching (MPLS) to increase transport efficiency in the network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as voice over IP (VoIP) and other business-critical functions.

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports can be assigned different labels.
- A packet arriving at a particular provider edge (PE) switch can be assigned a label that is different from that of the same packet entering the network at a different PE switch. As a result, forwarding decisions that depend on the ingress PE switch can be easily made.
- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than letting it follow the route chosen by the normal dynamic routing algorithm as the packet travels through the network. In MPLS, a label can be used to represent the route so that the packet need not carry the identity of the explicit route.

### Related Documentation

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [Understanding MPLS Components on page 4412](#)

- [Understanding MPLS Label Operations on page 4415](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- *Junos OS MPLS Applications Library for Routing Devices*

## Understanding MPLS Components

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

- [Provider Edge Switches on page 4412](#)
- [Provider Switch on page 4413](#)
- [Components Required for All Switches in the MPLS Network on page 4413](#)

### Provider Edge Switches

---

To implement MPLS on a network, you must configure two provider edge (PE) switches—that is, an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

- [MPLS Protocol and Label-Switched Paths on page 4412](#)
- [IP Over MPLS for Customer Edge Interfaces on page 4412](#)
- [BGP Layer 3 VPN Configuration on page 4413](#)
- [Routing Instances for Layer 3 VPN on page 4413](#)

#### ***MPLS Protocol and Label-Switched Paths***

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the **[edit protocols mpls]** hierarchy level.

#### ***IP Over MPLS for Customer Edge Interfaces***

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See [“Configuring MPLS on Provider Edge Switches” on page 4468](#).

### **BGP Layer 3 VPN Configuration**

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

### **Routing Instances for Layer 3 VPN**

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### **Provider Switch**

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

### **Components Required for All Switches in the MPLS Network**

The following MPLS components are configured on both the PE switches and the provider switches:

- [Interior Gateway Protocol on page 4413](#)
- [MPLS Protocol on page 4414](#)
- [RSVP on page 4414](#)
- [Family mpls on page 4414](#)

### **Interior Gateway Protocol**

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.



**NOTE:** The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to family mpls, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

### **MPLS Protocol**

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

### **RSVP**

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

### **Family mpls**

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.



**NOTE:** You can enable **family mpls** on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

#### **Related Documentation**

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on page 4422](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)



- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [Configuring MPLS on Provider Switches on page 4471](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

## Understanding MPLS Label Operations

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network header is analyzed and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters a label-switched path (LSP), the ingress provider edge (PE) switch examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label. The packet is then forwarded to the next provider switch in the LSP. This switch and all subsequent switches in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next switch in the path. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is forwarded based on its IP routing information.

This topic describes:


- [MPLS Label-Switched Paths and MPLS Labels on page 4415](#)
- [Reserved Labels on page 4416](#)
- [MPLS Label Operations on page 4416](#)
- [Penultimate-Hop Popping and Ultimate-Hop Popping on page 4418](#)

### MPLS Label-Switched Paths and MPLS Labels

When a packet enters the MPLS network, it is assigned to an LSP. Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the MPLS label (32 bits). Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.

[Figure 140 on page 4416](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 140: Label Encoding



g040575

### Reserved Labels

Labels range from 0 through 1,048,575. Labels 0 through 999,999 are for internal use.

Some of the reserved labels (in the range 0 through 15) have well-defined meanings. The following reserved labels are used by QFX Series and EX4600 devices:

- 0, IPv4 Explicit Null label—This value is valid only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 4 (IPv4) packet.
- 1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.
- 3, Implicit Null label—This label is used in the signaling protocol (RSVP) only to request label popping by the downstream switch. It never actually appears in the encapsulation. Labels with a value of 3 must not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.

### MPLS Label Operations

QFX Series and EX4600 devices support the following MPLS label operations:

- Push
- Pop
- Swap



**NOTE:** There is a limit with regard to the number of labels that QFX and EX4600 devices can affix (push operations) to the label stack or remove (pop operations) from the label stack.

- For Push operations—As many as three labels are supported.
- For Pop operations—As many as two labels are supported.

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that already has an MPLS label.

The pop operation removes a label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the underlying IP packet is forwarded as a native IP packet.

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

- Incoming interface
- Label
- Label forwarding table

Figure 141 on page 4417 shows an IP packet without a label arriving on the customer edge interface (ge-0/0/1) of the ingress PE switch. The ingress PE switch examines the packet and identifies that packet's destination as the egress PE switch. The ingress PE switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (ge-0/0/5). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface ge-0/0/5 with label 100. The provider switch swaps label 100 with label 200 and forwards the MPLS packet through its core interface (ge-0/0/7) to the next hop on the tunnel, which is the egress PE switch. The egress PE switch receives the MPLS packet through its core interface (ge-0/0/7), removes the MPLS label, and sends the IP packet out of its customer edge interface (ge-0/0/1) to a destination that is beyond the tunnel.

**Figure 141: MPLS Label Swapping**




Figure 141 on page 4417 shows the path of a packet as it passes in one direction from the ingress PE switch to the egress PE switch. However, the MPLS configuration also allows traffic to travel in the reverse direction. Thus, each PE switch operates as both an ingress switch and an egress switch.

### Penultimate-Hop Popping and Ultimate-Hop Popping

The switches enable penultimate-hop popping (PHP) by default with IP over MPLS configurations. With PHP, the penultimate provider switch is responsible for popping the MPLS label and forwarding the traffic to the egress PE switch. The egress PE switch then performs an IP route lookup and forwards the traffic. This reduces the processing load on the egress PE switch, because it is not responsible for popping the MPLS label.

- The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop switch removes the label and sends the packet to the egress PE switch.
- If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised and the egress PE switch of the LSP removes the label.

#### **Related Documentation**

- [Understanding MPLS Components on page 4412](#)
- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [Configuring MPLS on Provider Switches on page 4471](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

## Understanding CoS MPLS EXP Classifiers and Rewrite Rules

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. (For information about DSCP and IEEE 802.1p classifiers and general information about classifiers, see [“Understanding CoS Classifiers” on page 5810](#). For information about DSCP and IEEE 802.1p rewrite rules, see [“Understanding CoS Rewrite Rules” on page 5914](#).)

When a packet enters a customer-edge interface on the ingress provider edge (PE) switch, the switch associates the packet with a particular CoS servicing level before placing the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the classifier is translated and encoded in the MPLS header by means of the experimental (EXP) bits.

EXP classifiers map incoming MPLS packets to a forwarding class and a loss priority, and assign MPLS packets to output queues based on the forwarding class mapping. EXP classifiers are behavior aggregate (BA) classifiers.

EXP rewrite rules change (rewrite) the CoS value of the EXP bits in outgoing packets on the egress queues of the switch so that the new (rewritten) value matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** There is no default EXP classifier. There is no default EXP rewrite rule. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. If you want to rewrite the EXP bit value at the egress interface, you must configure EXP rewrite rules and apply them to logical interfaces.

EXP classifiers and rewrite rules are applied only to interfaces that are configured as **family mpls** (for example, set interfaces xe-0/0/35 unit 0 family mpls.)

This topic includes:

- [EXP Classifiers on page 4419](#)
- [EXP Rewrite Rules on page 4420](#)
- [Schedulers on page 4421](#)

### EXP Classifiers

Unlike DSCP and IEEE 802.1p BA classifiers, EXP classifiers are global to the switch and apply to all switch interfaces that are configured as **family mpls**. When you configure and apply an EXP classifier, MPLS traffic on all **family mpls** interfaces uses the EXP classifier, even on interfaces that also have a fixed classifier. If an interface has both an EXP classifier and a fixed classifier, the EXP classifier is applied to MPLS traffic and the fixed classifier is applied to all other traffic.

Also unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. If you want to classify MPLS traffic based on the EXP bits, you must explicitly configure an EXP classifier and apply it to the switch interfaces. Each EXP classifier has eight entries that correspond to the eight EXP CoS values (0 through 7, which correspond to bits 000 through 111).

You can configure as many EXP classifiers as you want. However, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All switch interfaces use the global EXP classifier to classify MPLS traffic.

Only one EXP classifier can be configured as the global EXP classifier at any time. If you want to change the global EXP classifier, delete the global EXP classifier configuration (use the **user@switch# delete class-of-service system-defaults classifiers exp** configuration statement), then configure the new global EXP classifier.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.



**NOTE:** The switch uses only the outermost label of incoming EXP packets for classification.

---



**NOTE:** MPLS packets with 802.1Q tags are not supported.

---

### EXP Rewrite Rules

---

As MPLS packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. EXP rewrite rules set the value of the EXP CoS bits within the header of the outgoing MPLS packet on **family mpls** interfaces. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes that CoS value into the packet header, replacing the old CoS value. EXP rewrite rules apply only to MPLS traffic.

EXP rewrite rules apply only to logical interfaces. You cannot apply EXP rewrite rules to physical interfaces.

There are no default EXP rewrite rules. If you want to rewrite the EXP value in MPLS packets, you must configure EXP rewrite rules and apply them to logical interfaces. If no rewrite rules are applied, all MPLS labels that are pushed have a value of zero (0). The EXP value remains unchanged on MPLS labels that are swapped.

You can configure as many EXP rewrite rules as you want, but you can only apply 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

You can apply an EXP rewrite rule to an interface that has a DSCP, DSCP IPv6, or IEEE 802.1p rewrite rule. Only MPLS traffic uses the EXP rewrite rule. MPLS traffic does not use DSCP or DSCP IPv6 rewrite rules.

If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

## Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on the switch. Default schedulers are provided only for the best-effort, fcoe, no-loss, and network-control forwarding classes. If you configure a custom forwarding class for MPLS traffic, you need to configure a scheduler to support that forwarding class and provide bandwidth to that forwarding class. See “[Understanding CoS Output Queue Schedulers](#)” on page 5868 and “[Example: Configuring Queue Schedulers](#)” on page 6081 for more information.

### Related Documentation

- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring CoS Bits for an MPLS Network on page 4478](#)

## Understanding Using MPLS-Based Layer 3 VPNs

On the QFX Series and on EX4600, you can use MPLS-based Layer 3 virtual private networks (VPNs) to securely connect geographically diverse sites across an MPLS network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as voice over IP (VoIP) and other business-critical functions.

A VPN uses a public telecommunications infrastructure, such as the Internet, to provide remote offices or individual users with secure access to their organization's network. VPNs are designed to provide the same level of performance and security as privately owned or leased networks but without the attendant costs.

This topic describes:

- [MPLS-Based Layer 3 VPNs on page 4422](#)

---

### MPLS-Based Layer 3 VPNs

---

In Junos OS, Layer 3 VPNs are based on RFC 4364, [BGP/MPLS IP Virtual Private Networks](#). RFC 4364 defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

Customer networks, because they are private, can use either public or private addresses, as defined in RFC 1918, [Address Allocation for Private Internets](#). When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the same private addresses used by other network users. BGP/MPLS VPNs solve this problem by adding a VPN identifier prefix to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and on the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only. Two different VPNs can use overlapping addresses. Each route within a VPN is assigned an MPLS label (for example, MPLS-ARCH, MPLS-BGP, or MPLS-ENCAPS). When BGP distributes a VPN route, it also distributes an MPLS label for that route. Before a customer data packet travels across the service provider's backbone, it is encapsulated along with the MPLS label that corresponds to the route within the customer's VPN that is the best match based on the packet's destination address. This MPLS packet is further encapsulated with another MPLS label or with an IP, so that it gets tunneled across the backbone to the egress provider edge (PE) switch. Thus, the backbone core switches do not need to know the VPN routes.

QFX5100 switches also support interprovider VPNs, and carrier-of-carriers VPNs. For more information, see *Interprovider and Carrier-of-Carriers VPNs*

#### Related Documentation

- [Understanding MPLS Label Operations on page 4415](#)
- [Understanding MPLS Components on page 4412](#)
- [Junos OS VPNs Library for Routing Devices](#)



- *Junos OS MPLS Applications Library for Routing Devices*

## MPLS Features

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [Supported MPLS Scaling Values on page 4426](#)

### MPLS Feature Support on the QFX Series and EX4600 Switch Overview

This topic describes the major MPLS features that are supported and not supported on the QFX Series and on the EX4600 switch..



**NOTE:** The command-line interface (CLI) on QFX Series devices and on the EX4600 switch displays even the MPLS related configuration statements that are not supported. However, configuring the unsupported statements on a device will have no effect on the operation of the device. See the following topics for the list of supported MPLS related configuration statements on QFX Series devices and on the EX4600 switch:

- [“\[edit protocols mpls\] Hierarchy Level” on page 4509](#) for the list of supported configuration statements at the [edit protocols mpls] hierarchy level
- [“\[edit protocols rsvp\] Hierarchy Level” on page 4513](#) for the list of supported configuration statements at the [edit protocols rsvp] hierarchy level

- [Supported MPLS Features on page 4423](#)
- [Unsupported MPLS Features on page 4425](#)

### Supported MPLS Features

[Table 340 on page 4423](#) lists the major MPLS features that are supported on the QFX Series and on the EX4600 switch, and the Juniper Networks Junos operating system (Junos OS) release in which they were introduced.

**Table 340: MPLS Features on the QFX Series and on the EX4600 Switch**

| Feature  | QFX Series           | EX4600                        |
|--|----------------------|-------------------------------|
| QFX standalone switch or EX4600 switch as an MPLS provider edge (PE) switch or provider switch | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| QFX standalone switch or EX4600 switch as a route reflector for BGP labeled routes             | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| BGP labeled unicast  | Junos OS 12.2X50-D10 | Junos OS Junos OS 13.2X51-D25 |
| Carrier-over-carrier BGP inter- autonomous systems (AS) L3VPN implementations                  | Junos OS 14.1X53-D10 | Junos OS 14.1X53-D10          |

Table 340: MPLS Features on the QFX Series and on the EX4600 Switch (*continued*)

| Feature  | QFX Series           | EX4600                        |
|--|----------------------|-------------------------------|
| Classifiers for MPLS firewall filters  | Junos OS 12.3X50-D10 | Junos OS Junos OS 13.2X51-D25 |
| Class of service (CoS) for MPLS traffic  | Junos OS 12.3X50-D10 | Junos OS 13.2X51-D25          |
| Ethernet-over-MPLS (L2 Circuit)  | Junos OS 14.1X53-D10 | Junos OS 14.1X53-D10          |
| Fast reroute<br><b>NOTE:</b> The <b>include-all</b> and <b>include-any</b> options are not supported.  | Junos OS 14.1X53-D10 | Junos OS 14.1X53-D10          |
| Graceful restart for OSPF and RSVP protocols   | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| IP over MPLS label-switched paths (LSPs)   | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| IPv6 tunneling for MPLS to tunnel IPv6 traffic over an MPLS-based IPv4 network (6PE)   | Junos OS 12.3X50-D10 | Junos OS 13.2X51-D25          |
| IS-IS (TE)   | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| IS-IS as an interior gateway protocol (IGP) for MPLS   | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| LDP-based signaling<br><b>NOTE:</b> These LDP features are not supported by QFX: multipoint, link protection, bidirectional forwarding detection (BFD), operation administration and management (OAM), multicast-only fast reroute (MoFRR), and equal-cost multipath (ECMP). | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| LDP tunneling (LDP over RSVP)  | Junos OS 12.3X50-D10 | Junos OS 13.2X51-D25          |
| Maximum transmission unit (MTU) discovery for MPLS paths<br><b>NOTE:</b> This is supported only at the control plane, not at the interface level.  | Junos OS 12.3X50-D10 | Junos OS 13.2X51-D25          |
| MPLS-based Layer 3 virtual private networks (VPNs)   | Junos OS 12.3X50-D10 | Junos OS 13.2X51-D25          |
| MPLS firewall filters  | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| MPLS OAM-LSP ping and traceroute   | Junos OS 12.3X50-D10 | Junos OS 13.2X51-D25          |
| MPLS RSVP auto bandwidth   | Junos OS 13.2X51-D15 | Junos OS 13.2X51-D25          |
| MPLS traffic engineering   | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25          |
| Node protection, link protection, and fast reroute   | Junos OS 14.1X53-D10 | Junos OS 14.1X53-D10          |

Table 340: MPLS Features on the QFX Series and on the EX4600 Switch (*continued*)

| Feature  | QFX Series           | EX4600               |
|--|----------------------|----------------------|
| OSPF (TE)  | Junos OS 14.1X53-D10 | Junos OS 14.1X53-D10 |
| OSPF version 2 (OSPFv2) as an interior gateway protocol (IGP) for MPLS | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25 |
| Per VRF Label support  | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25 |
| RSVP as a signaling protocol for MPLS                                  | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25 |
| SNMP MIB support   | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25 |
| Static LSPs  | Junos OS 12.2X50-D10 | Junos OS 13.2X51-D25 |

### Unsupported MPLS Features

The following major MPLS features are not supported on the QFX Series or on the EX4600 switch:

- Auto-policer
- Bidirectional Forwarding Detection (BFD) for MPLS LSPs
- ECMP for incoming MPLS packets
- L2 circuit-based local switching
- Link coloring using administrative groups
- MPLS-based circuit cross-connects (CCC)
- MPLS-based Layer 2 virtual private networks (VPNs)
- MPLS over routed VLAN interfaces (RVIs) and Layer 3 subinterfaces
- Point-to-multipoint LSP support
- Port mirroring on MPLS interfaces
- Virtual Private LAN Service (VPLS)

### Related Documentation

- *Carrier-of-Carriers VPNs*
- [MPLS Configuration Guidelines on page 4433](#)
- [Supported MPLS Scaling Values on page 4426](#)
- [Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch on page 4733](#)
- *Interprovider and Carrier-of-Carriers VPNs*
- [Understanding MPLS Components on page 4412](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)

- *Junos OS MPLS Applications Library for Routing Devices*

## Supported MPLS Scaling Values

This topic lists the MPLS scaling values supported on QFX Series switches.

[Table 341 on page 4426](#) lists the MPLS scaling values supported on Juniper QFX switches and on the EX4600 switch.

**Table 341: MPLS Scaling Values**

| Feature  | QFX3500 Scaling Value        | QFX5100 and EX4600 Scaling Value                               |
|--|------------------------------|--|
| Maximum number of MPLS labels in a packet's label stack                  | 3 labels for Push operations | 3 labels for Push operations                                   |
|  | 2 labels for Pop operations  | 2 labels for Pop operations                                    |
|  | 1 label for Swap operations  | 1 label for Swap operations                                    |
| Maximum number of MPLS labels on provider switches                       | 4096                         | 16386  |
| Maximum number of tunnel (combination of routes and LSPs) initiations    | Ingress LSPs: 1024           | Ingress LSPs: 1024   |
|  | Transit LSPs: 4000           | Transit LSPs: 16386  |
| Maximum number of unique next-hops on egress provider edge (PE) switches | 512                          | 512  |
| Maximum number of MPLS firewall filters                                  | 768                          | 1536   |
| Virtual Routing and Forwarding (VRF)                                     | 1K                           | 1K   |
| Layer 3 Host   | IPV4: 8K                     | See "Understanding the Unified Forwarding Table" on page 1545. |
| Layer 3 Longest Prefix Match (LPM)                                       | IPV4: 16K                    | See "Understanding the Unified Forwarding Table" on page 1545. |
|  | IPV6: 4K                     |  |

### Related Documentation

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [MPLS Configuration Guidelines on page 4433](#)

## Introduction to LDP for QFX5100

- [LDP Introduction on page 4427](#)
- [Junos OS LDP Protocol Implementation on page 4427](#)
- [LDP Operation on page 4427](#)

- [Tunneling LDP LSPs in RSVP LSPs on page 4428](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4428](#)
- [Label Operations on page 4428](#)
- [LDP Message Types on page 4430](#)
- [Discovery Messages on page 4430](#)
- [Session Messages on page 4430](#)
- [Advertisement Messages on page 4430](#)
- [Notification Messages on page 4431](#)
- [LDP Session Protection on page 4431](#)
- [LDP Graceful Restart on page 4431](#)

## LDP Introduction

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

## Junos OS LDP Protocol Implementation

The Junos OS implementation of LDP supports LDP version 1. The Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

## LDP Operation

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routers through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *Logical Interfaces*.

**Related Documentation**

- [Logical Interfaces](#)

## Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4428](#)
- [Label Operations on page 4428](#)

## Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.

## Label Operations

[Figure 142 on page 4429](#) depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see *Label Description*.) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an

LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 142: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs





Figure 143 on page 4429 depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

Figure 143: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs



## LDP Message Types

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- [Discovery Messages on page 4430](#)
- [Session Messages on page 4430](#)
- [Advertisement Messages on page 4430](#)
- [Notification Messages on page 4431](#)

## Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- Basic discovery—A router periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Extended discovery—LDP sessions between routers not directly connected are supported by LDP extended discovery. A router periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router decides whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

## Session Messages

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers are LDP peers and can exchange advertisement messages.

## Advertisement Messages

Advertisement messages create, change, and delete label mappings for forwarding equivalence classes (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.



## Notification Messages

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.

## LDP Session Protection

LDP session protection is based on the LDP targeted hello functionality defined in RFC 5036, *LDP Specification*, and is supported by the Junos OS as well as the LDP implementations of most other vendors. It involves sending unicast User Datagram Protocol (UDP) hello packets to a remote neighbor address and receiving similar packets from the neighbor router.

If you configure LDP session protection on a router, the LDP sessions are maintained as follows:

1. An LDP session is established between a router and a remote neighboring router.
2. If all of the direct links between the routers go down, the LDP session remains up so long as there is IP connectivity between the routers based on another connection over the network.
3. When the direct link between the routers is reestablished, the LDP session is not restarted. The routers simply exchange LDP hellos with each other over the direct link. They can then begin forwarding LDP-signaled MPLS packets using the original LDP session.

By default, LDP targeted hellos are set to the remote neighbor so long as the LDP session is up, even if there are no more link neighbors to that router. You can also specify the duration you would like to maintain the remote neighbor connection in the absence of link neighbors. When the last link neighbor for a session goes down, the Junos OS starts an LDP session protection timer. If this timer expires before any of the link neighbors come back up, the remote neighbor connection is taken down and the LDP session is terminated. If you configure a different value for the timer while it is currently running, the Junos OS updates the timer to the specified value without disrupting the current state of the LDP session.

## LDP Graceful Restart

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart is disabled by default. However, helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

# Configuration

- [Configuration Guidelines on page 4433](#)
- [LDP Configuration Guidelines for QFX5100 on page 4434](#)
- [Configuration Examples on page 4445](#)
- [Configuration Tasks on page 4467](#)
- [Configuration Statements on page 4508](#)
- [LDP Configuration Statements for QFX5100 on page 4523](#)

## Configuration Guidelines

---

- [MPLS Configuration Guidelines on page 4433](#)

### MPLS Configuration Guidelines

When configuring MPLS on QFX Series devices or on EX4600, note that the number of IP prefixes supported depends on the specific platform being used. See the scale specifications in the data sheet of your device for additional information.

- We recommend the following:
  - If your ingress provider edge (PE) switch needs to support more than 8000 external IP prefixes, use a larger capacity device as an ingress PE switch.
  - If you use a switch as a route reflector for BGP labeled routes, use it as a dedicated route reflector (that is, the switch must not participate in managing data traffic).
  - If you use a switch as a PE switch or as a route reflector for BGP labeled routes, configure routing policies on the PE switch and the route reflector to filter external IP routes from the routing table.

The configuration example for a routing policy named `fib_policy` (at the **[edit policy-options]** and **[edit routing-options]** hierarchy levels) to filter BGP labeled routes from the `inet.0` routing table is given below:

```
user@switch# show policy-options
policy-statement fib_policy {
  from {
    protocol bgp;
    rib inet.0;
  }
}
```

```
        then reject;
    }

    user@switch# show routing-options
    forwarding-table {
        export fib_policy;
    }
```

- Packet fragmentation using the **allow-fragmentation** statement at the **[edit protocols mpls path-mtu]** hierarchy level is not supported on QFX Series devices or on the EX4600 switch. Therefore, you must ensure that the maximum transmission unit (MTU) values configured on every MPLS interface is sufficient to handle MPLS packets. The packets whose size exceeds the MTU value of an interface will be dropped.

#### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [Configuring MPLS on Provider Switches on page 4471](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)

---

## LDP Configuration Guidelines for QFX5100

- [Minimum LDP Configuration on page 4434](#)
- [Enabling and Disabling LDP on page 4435](#)
- [Enabling Strict Targeted Hello Messages for LDP on page 4435](#)
- [Filtering Inbound LDP Label Bindings on page 4435](#)
- [Filtering Outbound LDP Label Bindings on page 4437](#)
- [Specifying the Transport Address Used by LDP on page 4439](#)
- [Collecting LDP Statistics on page 4440](#)
- [Tracing LDP Protocol Traffic on page 4442](#)

### Minimum LDP Configuration

To enable LDP on a single interface, include the **ldp** statement and specify the interface using the **interface** statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {
    interface interface-name;
}
```

To enable LDP on all interfaces, specify **all** for *interface-name*.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

## Enabling and Disabling LDP

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {
  interface interface-name;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

To enable LDP on all interfaces, specify **all** for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the **interface** statement with the **disable** option:

```
interface interface-name {
  disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

## Enabling Strict Targeted Hello Messages for LDP

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

```
LDP: Ignoring targeted hello from 10.0.0.1
```

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

**import** [ *policy-names* ];

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 342 on page 4436](#) lists the only **from** operators that apply to LDP received-label filtering.

**Table 342: from Operators That Apply to LDP Received-Label Filtering**

| from Operator       | Description  |
|---------------------|--|
| <b>interface</b>    | Matches on bindings received from a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>     | Matches on bindings received from the specified LDP router ID                              |
| <b>next-hop</b>     | Matches on bindings received from a neighbor advertising the specified interface address   |
| <b>route-filter</b> | Matches on bindings with the specified prefix  |

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels; so if multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface. When a router has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
  ldp {
    import only-32;
    ...
  }
}
policy-options {
  policy-statement only-32 {
    term first {
      from {
        route-filter 0.0.0.0/0 upto /31;
      }
      then reject;
    }
    then accept;
  }
}
```

Accept 131.108/16 or longer from router ID 10.10.255.2 and accept all prefixes from all other neighbors:

```
[edit]
protocols {
  ldp {
    import nosy-neighbor;
    ...
  }
}
policy-options {
  policy-statement nosy-neighbor {
    term first {
      from {
        neighbor 10.10.255.2;
        route-filter 131.108.0.0/16 orlonger accept;
        route-filter 0.0.0.0/0 orlonger reject;
      }
    }
    then accept;
  }
}
```

### Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being advertised to neighboring routers. To configure outbound label filtering, include the **export** statement:

```
export [policy-name];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in [Table 343 on page 4438](#).

**Table 343: to Operators for LDP Outbound-Label Filtering**

| to Operator      | Description  |
|------------------|--|
| <b>interface</b> | Matches on bindings sent to a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>  | Matches on bindings sent to the specified LDP router ID                              |
| <b>next-hop</b>  | Matches on bindings sent to a neighbor advertising the specified interface address   |

If a binding is filtered, the binding is not advertised to the neighboring router, but it can be installed as part of an LSP on the local router. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels. If multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Examples: Filtering Outbound LDP Label Bindings

Block transmission of the route for **10.10.255.6/32** to any neighbors:

```
[edit protocols]
ldp {
  export block-one;
```



```

}
policy-options {
  policy-statement block-one {
    term first {
      from {
        route-filter 10.10.255.6/32 exact;
      }
      then reject;
    }
    then accept;
  }
}

```

Send only **131.108/16** or longer to router ID **10.10.255.2**, and send all prefixes to all other routers:

```

[edit protocols]
ldp {
  export limit-lsps;
}
policy-options {
  policy-statement limit-lsps {
    term allow-one {
      from {
        route-filter 131.108.0.0/16 orlonger;
      }
      to {
        neighbor 10.10.255.2;
      }
      then accept;
    }
    term block-the-rest {
      to {
        neighbor 10.10.255.2;
      }
      then reject;
    }
    then accept;
  }
}

```

## Specifying the Transport Address Used by LDP

Routers must first establish a TCP session between each other before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

To configure the LDP transport address, include the `transport-address` statement:

**transport-address** (router-id | interface);

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify the **router-id** option, the address of the router identifier is used as the transport address (unless otherwise configured, the router identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the router identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

Related Documentation

- [transport-address on page 4563](#)

## Collecting LDP Statistics

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the **traffic-statistics** statement at the **[edit protocols ldp]** hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the **interval** option. The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise, LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {  
  file filename <files number> <size size> <world-readable | no-world-readable>;  
  interval interval;  
  no-penultimate-hop;  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

This section includes the following topics:

- [LDP Statistics Output on page 4440](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router on page 4441](#)
- [LDP Statistics Limitations on page 4442](#)

---

### LDP Statistics Output

The following sample output is from an LDP statistics file:

| FEC               | Type    | Packets | Bytes | Shared |
|-------------------|---------|---------|-------|--------|
| 10.255.350.448/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.450/32 | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.451/32 | Transit | 0       | 0     | No     |

```

                Ingress      0      0      No
220.220.220.1/32 Transit      0      0      Yes
                Ingress      0      0      No
220.220.220.2/32 Transit      0      0      Yes
                Ingress      0      0      No
220.220.220.3/32 Transit      0      0      Yes
                Ingress      0      0      No
May 28 15:02:05, read 12 statistics in 00:00:00 seconds

```

The LDP statistics file includes the following columns of data:

- **read**—Number of bytes of data passed by the FEC since its LSP came up.
- **read**—FEC for which LDP traffic statistics are collected.
- **read**—Number of packets passed by the FEC since its LSP came up.
- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).

### Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement. For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```

traffic-statistics {
    no-penultimate-hop;
}

```

For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.



**NOTE:** When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following sample output is from an LDP statistics file showing routers on which the **no-penultimate-hop** option is configured:

```

FEC              Type      Packets      Bytes      Shared
10.255.245.218/32 Transit      0           0         No
                  Ingress      4          246         No

```

|                   |         |                     |
|-------------------|---------|---------------------|
| 10.255.245.221/32 | Transit | statistics disabled |
|                   | Ingress | statistics disabled |
| 13.1.1.0/24       | Transit | statistics disabled |
|                   | Ingress | statistics disabled |
| 13.1.3.0/24       | Transit | statistics disabled |
|                   | Ingress | statistics disabled |

---

### LDP Statistics Limitations

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.
- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

## Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 4442](#)
- [Tracing LDP Protocol Traffic Within FECs on page 4443](#)
- [Examples: Tracing LDP Protocol Traffic on page 4444](#)

---

### Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

To trace LDP protocol traffic, you can specify options in the global **traceoptions** statement at the **[edit routing-options]** hierarchy level, and you can specify LDP-specific options by including the **traceoptions** statement:

```
traceoptions {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place LDP-tracing output in the file **ldp-log**.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.
- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packets**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

You can also configure the **filter** flag modifier with the **match-on address** sub-option for the **packets** flag. This allows you to trace based on the source and destination addresses of the packets.

- **path**—Trace label-switched path operations.
- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

### Tracing LDP Protocol Traffic Within FECs

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers.

You can trace LDP protocol traffic within a specific FEC and filter LDP trace statements based on an FEC. This is useful when you want to trace or troubleshoot LDP protocol traffic associated with an FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on an FEC:

```
[edit protocols ldp traceoptions]
set flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.

- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).
- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

---

### Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag path;
    }
  }
}
```

Trace all LDP outgoing messages:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag packets;
    }
  }
}
```

Trace all LDP error conditions:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag error;
    }
  }
}
```

Trace all LDP incoming messages and all label-binding operations:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5 world-readable;
      flag packets receive;
      flag binding;
    }
  }
}
```

```

        interface all {
        }
    }
}

```

Trace LDP protocol traffic for an FEC associated with the LSP:

```

[edit]
protocols {
  ldp {
    traceoptions {
      flag route filter match-on fec policy filter-policy-for-ldp-fec;
    }
  }
}

```

## Configuration Examples

- [Example: Configuring MPLS-Based Layer 3 VPNs on page 4445](#)
- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 4454](#)
- [Example: Configuring LDP Downstream on Demand on page 4462](#)

### Example: Configuring MPLS-Based Layer 3 VPNs

You can implement an MPLS-based Layer 3 virtual private network (VPN) on QFX3500 switches to interconnect sites for customers who want the service provider to handle all the Layer 3 routing functions. To support an MPLS-based Layer 3 VPN, you need to add components of the Layer 3 VPN to the configuration of the two provider edge (PE) switches. You do not need to change the configuration of the provider switches.

This example shows how to configure an MPLS-based Layer 3 VPN spanning two corporate sites:

- [Requirements on page 4446](#)
- [Overview and Topology on page 4446](#)
- [Configuring the Local PE Switch on page 4449](#)
- [Configuring the Remote PE Switch on page 4451](#)

## Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.3 or later for the QFX Series
- Three QFX3500 switches

Before you configure the Layer 3 VPN components, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches” on page 4468](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches” on page 4471](#).

## Overview and Topology

Layer 3 VPNs allow customers to leverage the service provider’s technical expertise to ensure efficient site-to-site routing. The customer’s customer edge (CE) switch uses a routing protocol such as BGP or OSPF to communicate with the service provider’s provider edge (PE) switch to carry IP prefixes across the network. MPLS-based Layer 3 VPNs use only IP over MPLS; other protocol packets are not supported. This example includes two PE switches, PE1 and PE2.


In the basic MPLS configuration of the PE switches using IP over MPLS, the PE switches were configured to use OSPF as the routing protocol between the MPLS switches and RSVP as the signaling protocol. Traffic engineering was enabled. A label-switched path (LSP) was configured.

The following components must be added to the PE switches for an MPLS-based Layer 3 VPN:

- BGP group with **family inet-vpn unicast**
- Routing instance with instance type **vrf**

[Figure 144 on page 4446](#) illustrates the topology of this MPLS-based Layer 3 VPN.

**Figure 144: MPLS-Based Layer 3 VPN**



[Table 344 on page 4447](#) shows the settings of the customer edge interface on the local CE switch.



Table 344: Local CE Switch in the MPLS-Based Layer 3 VPN Topology

| Property                 | Settings  | Description                         |
|--------------------------|---|-------------------------------------|
| Local CE switch hardware | QFX3500 switch  | CE1                                 |
| Customer edge interface  | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.14/16</b> | Interface that connects CE1 to PE1. |

Table 345 on page 4447 shows the settings of the customer edge interface on the remote CE switch.

Table 345: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology

| Property                  | Settings  | Description                         |
|---------------------------|---|-------------------------------------|
| Remote CE switch hardware | QFX3500 switch  | CE2                                 |
| Customer edge interface   | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.1/16</b> | Interface that connects CE2 to PE2. |

Table 346 on page 4447 shows the Layer 3 VPN components of the local PE switch.

Table 346: Layer 3 VPN Components of the Local PE Switch

| Property                 | Settings  | Description   |
|--------------------------|---|---|
| Local PE switch hardware | QFX3500 switch  | PE1   |
| Customer edge interface  | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.1/16</b>            | Connects PE1 to CE1.<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface           | <b>xe-0/0/6 unit 0</b><br><b>family inet address 60.0.0.60/16</b><br><b>family mpls</b> | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.  |

Table 346: Layer 3 VPN Components of the Local PE Switch (*continued*)

| Property           | Settings   | Description  |
|--------------------|--|--|
| Loopback interface | <b>lo0 unit 0</b><br><b>family inet address 21.21.21.21/32</b> | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration. |
| BGP                | <b>bgp</b>   | Added for the Layer 3 VPN configuration.   |
| Routing instance   | <b>L3VPN-1</b>   | Added for the Layer 3 VPN configuration.   |

Table 347 on page 4448 shows the Layer 3 VPN components of the remote PE switch.

Table 347: Layer 3 VPN Components of the Remote PE Switch

| Property                  | Settings   | Description  |
|---------------------------|--|--|
| Remote PE switch hardware | QFX3500 switch   | PE2  |
| Customer edge interface   | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.14/16</b><br><b>family mpls</b> | Connects PE2 to CE2.<br><br>For the Layer 3 VPN configuration, added <b>family mpls</b> .<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface            | <b>xe-0/0/6 unit 0</b><br><b>family inet address 60.2.0.60/16</b><br><b>family mpls</b>              | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.   |
| Loopback interface        | <b>lo0 unit 0</b><br><b>family inet address 22.22.22.22/32</b>                                       | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.   |
| BGP                       | <b>bgp</b>   | Added for the Layer 3 VPN configuration.   |
| Routing instances         | <b>L3VPN-1</b>   | Added for the Layer 3 VPN configuration.   |

### Configuring the Local PE Switch

- CLI Quick Configuration** To quickly configure the Layer 3 VPN components on the local PE switch, copy the following commands and paste them into the switch terminal window of PE1:
- ```
[edit]
set protocols bgp local-address 21.21.21.21 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 22.22.22.22
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-0/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label
set routing-options router-id 21.21.21.21
set routing-options autonomous-system 10
```
- Step-by-Step Procedure** To configure the Layer 3 VPN components on the local PE switch:
1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:
 

```
[edit protocols bgp]
user@switchPE1# set local-address 21.21.21.21 family inet-vpn unicast
```
  2. Configure the BGP group, specifying the group name and type:
 

```
[edit protocols bgp]
user@switchPE1# set group PE1-PE2 type internal
```
  3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:
 

```
[edit protocols bgp]
user@switchPE1# set neighbor 22.22.22.22
```
  4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:
 

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 instance-type vrf
```
  5. Configure a description for this routing instance:
 

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
```
  6. Configure the routing instance to use a route distinguisher:
 

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 route-distinguisher 21:21
```



**NOTE:** Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances require a route distinguisher to allow BGP to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. If you configure different VPN routing instances with the same route distinguisher, the commit fails.

7. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 vrf-target target:21:21
```



**NOTE:** You can create more complex policies by explicitly configuring VRF import and export policies using the import and export options. See the *Junos OS VPNs Library for Routing Devices*.

8. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 vrf-table-label
```

9. Configure the router ID and autonomous system (AS):



**NOTE:** We recommend that you explicitly configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

```
[edit routing-options]
user@switchPE1# set router-id 21.21.21.21 autonomous-system 10
```

**Results** Display the results of the configuration:

```
user@switchPE1> show configuration
```

```
interfaces {
  ge-0/0/14 {
    unit 0 {
      family inet {
        address 51.51.0.1/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 21.21.21.21/32;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      family inet {
        address 60.0.0.60/16;
      }
      family mpls;
    }
  }
}
```

```

protocols {
  mpls {
    label-switched-path 21-22 {
      from 21.21.21.21;
      to 22.22.22.22;
      no-cspf;
    }
    interface xe-0/0/6.0;
    interface lo0.0;
  }
  bgp {
    local-address 21.21.21.21;
    family inet-vpn {
      unicast;
    }
    group PE1-PE2 {
      type internal;
      neighbor 22.22.22.22;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface ge-0/0/14.0;
      interface lo0.0;
      interface xe-0/0/6.0;
    }
  }
}
routing-instances {
  L3VPN-1 {
    instance-type vrf;
    description "BETWEEN PE1 AND PE2";
    route-distinguisher 21:21;
    vrf-target target:21:21;
    vrf-table-label;
  }
}
routing-options {
  router-id 21.21.21.21;
  autonomous-system 10;
}

```

### Configuring the Remote PE Switch

#### CLI Quick Configuration

To quickly configure the Layer 3 VPN components on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE2:

```

[edit]
set protocols bgp local-address 22.22.22.22 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 21.21.21.21
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-0/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label;
set routing-options router-id 22.22.22.22

```

**set routing-options autonomous-system 10****Step-by-Step  
Procedure**

To configure Layer 3 VPN components on the remote PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:  

```
[edit protocols bgp]
user@switchPE2# set local-address 22.22.22.22 family inet-vpn unicast
```
2. Configure the BGP group, specifying the group name and type:  

```
[edit protocols bgp]
user@switchPE2# set group PE1-PE2 type internal
```
3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:  

```
[edit protocols bgp]
user@switchPE2# set neighbor 21.21.21.21
```
4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 instance-type vrf
```
5. Configure a description for this routing instance:  

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
```
6. Configure the routing instance to apply to the customer edge interface:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 interface ge-0/0/14.0
```
7. Configure the routing instance to use a route distinguisher, using the format *ip-address:number*:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 route-distinguisher 21:21
```
8. Configure the VPN routing and forwarding (VRF) target of the routing instance:  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 vrf-target target:21:21
```
9. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header.  

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 vrf-table-label
```
10. Configure the router ID and autonomous system (AS):  

```
[edit routing-options]
user@switchPE2# set router-id 22.22.22.22 autonomous-system 10
```

**Results** Display the results of the configuration:

```
user@switchPE2> show configuration

interfaces {
  ge-0/0/14 {
    unit 0 {
      family inet {
        address 11.22.26.14/16;
      }
    }
  }
}
```

```

    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 22.22.22.22/32;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      family inet {
        address 60.2.0.60/16;
      }
      family mpls;
    }
  }
  protocols {
    mpls {
      label-switched-path 22-21 {
        from 22.22.22.22;
        to 21.21.21.21;
        no-cspf;
      }
      interface xe-0/0/6.0;
      interface lo0.0;
    }
    bgp {
      local-address 22.22.22.22;
      family inet-vpn {
        unicast;
      }
      group PE1-PE2 {
        type internal;
        neighbor 21.21.21.21;
      }
    }
    ospf {
      traffic-engineering;
      area 0.0.0.0 {
        interface ge-0/0/14.0;
        interface lo0.0;
        interface xe-0/0/6.0;
      }
    }
  }
  routing-instances {
    L3VPN-1 {
      instance-type vrf;
      description "BETWEEN PE1 AND PE2";
      route-distinguisher 21:21;
      vrf-target target:21:21;
      vrf-table-label;
    }
  }
  routing-options {
    router-id 22.22.22.22;
    autonomous-system 10;
  }

```

- Related Documentation**
- [Configuring MPLS on Provider Edge Switches on page 4468](#)
  - [Configuring MPLS on Provider Switches on page 4471](#)

## Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

This example shows how to configure Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. External BGP (EBGP) is used between the customer edge (CE) and provider edge (PE) devices. The remote CE devices have different AS numbers for loop detection.

- [Requirements on page 4454](#)
- [Overview on page 4454](#)
- [Configuration on page 4457](#)
- [Verification on page 4462](#)

### Requirements

---

No special configuration beyond device initialization is required before you configure this example.

### Overview

---

Detailed information about the Juniper Networks implementation of IPv6 over MPLS is described in the following Internet drafts:

- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)
- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)


These Internet drafts are available on the IETF website at <http://www.ietf.org/>.

This example shows you how to interconnect a two IPv6 networks over an IPv4-based network core, giving you the ability to provide IPv6 service without having to upgrade the routers in your core network. Multiprotocol Border Gateway Protocol (MP-BGP) is configured to exchange routes between the IPv6 networks, and data is tunneled between these IPv6 networks by means of IPv4-based MPLS.

In [Figure 145 on page 4455](#), PE1 and PE2 are dual-stack BGP routers or switches, meaning they have both IPv4 and IPv6 stacks. The PE devices link the IPv6 networks through the customer edge (CE) routers or switches to the IPv4 core network. The CE devices and the PE devices connect through a link layer that can carry IPv6 traffic. The PE devices use IPv6 on the CE router-facing interfaces and use IPv4 and MPLS on the core-facing interfaces. Note that one of the connected IPv6 networks could be the global IPv6 Internet.



Figure 145: IPv6 Networks Linked by MPLS IPv4 Tunnels



The two PE devices are linked through an MP-BGP session using IPv4 addresses. They use the session to exchange IPv6 routes with an IPv6 (value 2) address family indicator (AFI) and a subsequent AFI (SAFI) (value 4). Each PE router sets the next hop for the IPv6 routes advertised on this session to its own IPv4 address. Because MP-BGP requires the BGP next hop to correspond to the same address family as the network layer reachability information (NLRI), this IPv4 address needs to be embedded within an IPv6 format.

The PE devices can learn the IPv6 routes from the CE devices connected to them using MP-BGP or through static configuration. Note that if BGP is used as the PE-router-to-CE-router protocol, the MP-BGP session between the PE device and CE device could occur over an IPv4 or IPv6 Transmission Control Protocol (TCP) session. Also, the BGP routes exchanged on that session would have SAFI unicast. You must configure an export policy to pass routes between IBGP and EBGp, and between BGP and any other protocol.

The PE routers have MPLS LSPs routed to each others' IPv4 addresses. IPv4 provides signaling for the LSPs by means of RSVP. These LSPs are used to resolve the next-hop addresses of the IPv6 routes learned from MP-BGP. The next hops use IPv4-mapped IPv6 addresses, while the LSPs use IPv4 addresses.

The PE devices always advertise IPv6 routes to each other using a label value of 2, the explicit null label for IPv6 as defined in RFC 3032, *MPLS Label Stack Encoding*. As a consequence, each of the forwarding next hops for the IPv6 routes learned from remote PE routers normally push two labels. The inner label is 2 (this label could be different if the advertising PE device is not a Juniper Networks routing or switching platform), and the outer label is the LSP label. If the LSP is a single-hop LSP, then only Label 2 is pushed.

It is also possible for the PE devices to exchange plain IPv6 routes using SAFI unicast. However, there is one major advantage in exchanging labeled IPv6 routes. The

penultimate-hop router for an MPLS LSP can pop the outer label and then send the packet with the inner label as an MPLS packet. Without the inner label, the penultimate-hop router would need to discover whether the packet is an IPv4 or IPv6 packet to set the protocol field in the Layer 2 header correctly.

When the PE1 device in [Figure 145 on page 4455](#) receives an IPv6 packet from the CE1 device, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 device, then no labels need to be pushed and the packet is simply sent to the CE2 device. If the destination matches a prefix that was learned from the PE2 device, then the PE1 router pushes two labels onto the packet and sends it to the Provider router. The inner label is 2 and the outer label is the LSP label for the PE2 router.

Each provider router in the service provider's network handles the packet as it would any MPLS packet, swapping labels as it passes from provider router to provider router. The penultimate-hop provider router for the LSP pops the outer label and sends the packet to the PE2 router. When the PE2 router receives the packet, it recognizes the IPv6 explicit null label on the packet (Label 2). It pops this label and treats it as an IPv6 packet, performing a lookup in the IPv6 forwarding table and forwarding the packet to the CE3 router.

This example includes the following settings:

- In addition to configuring the **family inet6** statement on all the CE router-facing interfaces, you must also configure the statement on all the core-facing interfaces running MPLS. Both configurations are necessary because the router must be able to process any IPv6 packets it receives on these interfaces. You should not see any regular IPv6 traffic arrive on these interfaces, but you will receive MPLS packets tagged with Label 2. Even though Label 2 MPLS packets are sent in IPv4, these packets are treated as native IPv6 packets.
- You enable IPv6 tunneling by including the **ipv6-tunneling** statement in the configuration for the PE routers. This statement allows IPv6 routes to be resolved over an MPLS network by converting all routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.



**NOTE:** BGP automatically runs its import policy even when copying routes from a primary routing table group to a secondary routing table group. If IPv4 labeled routes arrive from a BGP session (for example, when you have configured the **labeled-unicast** statement at the `[edit protocols bgp family inet]` hierarchy level on the PE router), the BGP neighbor's import policy also accepts IPv6 routes, since the neighbor's import policy is run while doing the copy operation to the inet6.3 routing table.

---

- When you configure MP-BGP to carry IPv6 traffic, the IPv4 MPLS label is removed at the destination PE router. The remaining IPv6 packet without a label can then be forwarded to the IPv6 network. To enable this, include the **explicit-null** statement in the BGP configuration.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Device PE1
set interfaces xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
set interfaces xe-0/0/5 unit 2 family mpls
set interfaces xe-0/0/6 unit 5 family inet address 10.1.1.5/30
set interfaces xe-0/0/6 unit 5 family inet6
set interfaces xe-0/0/6 unit 5 family mpls
set interfaces lo0 unit 2 family inet address 1.1.1.2/32
set protocols mpls ipv6-tunneling
set protocols mpls interface xe-0/0/5.2
set protocols mpls interface xe-0/0/6.5
set protocols bgp group toCE1 type external
set protocols bgp group toCE1 local-address ::10.1.1.2
set protocols bgp group toCE1 family inet6 unicast
set protocols bgp group toCE1 export send-bgp6
set protocols bgp group toCE1 peer-as 1
set protocols bgp group toCE1 neighbor ::10.1.1.1
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 1.1.1.2
set protocols bgp group toPE2 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE2 export next-hop-self
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface xe-0/0/6.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols rsvp interface xe-0/0/6.5
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 2

Device PE2
set interfaces xe-0/0/5 unit 10 family inet address 10.1.1.10/30
set interfaces xe-0/0/5 unit 10 family inet6
set interfaces xe-0/0/5 unit 10 family mpls
set interfaces xe-0/0/6 unit 13 family inet6 address ::10.1.1.13/126
set interfaces xe-0/0/6 unit 13 family mpls
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set protocols mpls ipv6-tunneling
set protocols mpls interface xe-0/0/5.10
set protocols mpls interface xe-0/0/6.13
set protocols bgp group toPE1 type internal
set protocols bgp group toPE1 local-address 1.1.1.4
set protocols bgp group toPE1 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE1 export next-hop-self

```

```
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 neighbor 1.1.1.2
set protocols bgp group toCE3 type external
set protocols bgp group toCE3 local-address ::10.1.1.13
set protocols bgp group toCE3 family inet6 unicast
set protocols bgp group toCE3 export send-bgp6
set protocols bgp group toCE3 peer-as 3
set protocols bgp group toCE3 neighbor ::10.1.1.14
set protocols ospf area 0.0.0.0 interface xe-0/0/5.10
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols rsvp interface xe-0/0/5.10
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 2
```

Device P

```
set interfaces xe-0/0/5 unit 6 family inet address 10.1.1.6/30
set interfaces xe-0/0/5 unit 6 family inet6
set interfaces xe-0/0/5 unit 6 family mpls
set interfaces xe-0/0/6 unit 9 family inet address 10.1.1.9/30
set interfaces xe-0/0/6 unit 9 family inet6
set interfaces xe-0/0/6 unit 9 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface xe-0/0/5.6
set protocols mpls interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface xe-0/0/5.6
set protocols ospf area 0.0.0.0 interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols rsvp interface xe-0/0/5.6
set protocols rsvp interface xe-0/0/6.9
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2
```

Device CE1

```
set interfaces xe-0/0/5 unit 1 family inet6 address ::10.1.1.1/126
set interfaces xe-0/0/5 unit 1 family mpls
set interfaces lo0 unit 1 family inet6 address ::1.1.1.1/128
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 local-address ::10.1.1.1
set protocols bgp group toPE1 family inet6 unicast
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 peer-as 2
set protocols bgp group toPE1 neighbor ::10.1.1.2
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1
```

Device CE3

```
set interfaces xe-0/0/5 unit 14 family inet6 address ::10.1.1.14/126
```

```

set interfaces xe-0/0/5 unit 14 family mpls
set interfaces lo0 unit 5 family inet6 address ::1.1.1.5/128
set protocols bgp group toPE2 type external
set protocols bgp group toPE2 local-address ::10.1.1.14
set protocols bgp group toPE2 family inet6 unicast
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 peer-as 2
set protocols bgp group toPE2 neighbor ::10.1.1.13
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 3

```

### Configuring Device PE1

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
user@PE1# set xe-0/0/5 unit 2 family mpls

user@PE1# set xe-0/0/6 unit 5 family inet address 10.1.1.5/30
user@PE1# set xe-0/0/6 unit 5 family inet6
user@PE1# set xe-0/0/6 unit 5 family mpls

user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32

```

2. Configure MPLS on the interfaces.

```

[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set interface xe-0/0/5.2
user@PE1# set interface xe-0/0/6.5

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE1# set group toCE1 type external
user@PE1# set group toCE1 local-address ::10.1.1.2
user@PE1# set group toCE1 family inet6 unicast
user@PE1# set group toCE1 export send-bgp6
user@PE1# set group toCE1 peer-as 1
user@PE1# set group toCE1 neighbor ::10.1.1.1

user@PE1# set group toPE2 type internal
user@PE1# set group toPE2 local-address 1.1.1.2
user@PE1# set group toPE2 family inet6 labeled-unicast explicit-null
user@PE1# set group toPE2 export next-hop-self
user@PE1# set group toPE2 export send-v6

```

```
user@PE1# set group toPE2 neighbor 1.1.1.4
```

4. Configure OSPF

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface xe-0/0/6.5
user@PE1# set interface lo0.2 passive
```

5. Configure a signaling protocol.

```
[edit protocols]
user@PE1# set rsvp interface xe-0/0/6.5
```

6. Configure the routing policies.

```
[edit policy-options]
user@PE1# set policy-statement next-hop-self then next-hop self
```

```
user@PE1# set policy-statement send-bgp6 from family inet6
user@PE1# set policy-statement send-bgp6 from protocol bgp
user@PE1# set policy-statement send-bgp6 then accept
```

```
user@PE1# set policy-statement send-v6 from family inet6
user@PE1# set policy-statement send-v6 from protocol bgp
user@PE1# set policy-statement send-v6 from protocol direct
user@PE1# set policy-statement send-v6 then accept
```

7. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set router-id 1.1.1.2
user@PE1# set autonomous-system 2
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
xe-0/0/5 {
  unit 2 {
    family inet6 {
      address ::10.1.1.2/126;
    }
    family mpls;
  }
}
xe-0/0/6 {
  unit 5 {
    family inet {
      address 10.1.1.5/30;
    }
    family inet6;
    family mpls;
  }
}
lo0 {
```

```

    unit 2 {
        family inet {
            address 1.1.1.2/32;
        }
    }
}

user@R1# show policy-options
policy-statement next-hop-self {
    then {
        next-hop self;
    }
}
policy-statement send-bgp6 {
    from {
        family inet6;
        protocol bgp;
    }
    then accept;
}
policy-statement send-v6 {
    from {
        family inet6;
        protocol [ bgp direct ];
    }
    then accept;
}

user@R1# show protocols
mpls {
    ipv6-tunneling;
    interface xe-0/0/5.2;
    interface xe-0/0/6.5;
}
bgp {
    group toCE1 {
        type external;
        local-address ::10.1.1.2;
        family inet6 {
            unicast;
        }
        export send-bgp6;
        peer-as 1;
        neighbor ::10.1.1.1;
    }
    group toPE2 {
        type internal;
        local-address 1.1.1.2;
        family inet6 {
            labeled-unicast {
                explicit-null;
            }
        }
        export [ next-hop-self send-v6 ];
        neighbor 1.1.1.4;
    }
}

```

```
ospf {
  area 0.0.0.0 {
    interface xe-0/0/6.5;
    interface lo0.2 {
      passive;
    }
  }
}
rsvp {
  interface xe-0/0/6.5;
}

user@R1# show routing-options
router-id 1.1.1.2;
autonomous-system 2;
```

If you are done configuring the device, enter **commit** from configuration mode. Configure the other devices in the topology, as shown in “[CLI Quick Configuration](#)” on [page 4457](#).

---

### Verification

Confirm that the configuration is working properly.

#### *Verifying That the CE Devices Have Connectivity*

**Purpose** Make sure that the tunnel is operating.

**Action** From operational mode, enter the **ping** command.

```
user@CE1> ping ::10.1.1.14
PING6(56=40+8+8 bytes) ::10.1.1.1 --> ::10.1.1.14
16 bytes from ::10.1.1.14, icmp_seq=0 hlim=61 time=10.687 ms
16 bytes from ::10.1.1.14, icmp_seq=1 hlim=61 time=9.239 ms
16 bytes from ::10.1.1.14, icmp_seq=2 hlim=61 time=1.842 ms

user@CE3> ping ::10.1.1.1
PING6(56=40+8+8 bytes) ::10.1.1.14 --> ::10.1.1.1
16 bytes from ::10.1.1.1, icmp_seq=0 hlim=61 time=1.484 ms
16 bytes from ::10.1.1.1, icmp_seq=1 hlim=61 time=1.338 ms
16 bytes from ::10.1.1.1, icmp_seq=2 hlim=61 time=1.351 ms
```

**Meaning** The IPv6 CE devices can communicate over the core IPv4 network.

**Related  
Documentation**

### Example: Configuring LDP Downstream on Demand

This example shows how to configure LDP downstream on demand. LDP is commonly configured using downstream unsolicited advertisement mode, meaning label advertisements for all routes are received from all LDP peers. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between the access and aggregation networks and to reduce the processing requirements for the control plane.



Downstream nodes could potentially receive tens of thousands of label bindings from upstream aggregation nodes. Instead of learning and storing all label bindings for all possible loopback addresses within the entire MPLS network, the downstream aggregation node can be configured using LDP downstream on demand to only request the label bindings for the FECs corresponding to the loopback addresses of those egress nodes on which it has services configured.

- [Requirements on page 4463](#)
- [Overview on page 4463](#)
- [Configuration on page 4463](#)
- [Verification on page 4466](#)

### Requirements

This example uses the following hardware and software components:

- M Series router
- Junos OS 12.2

### Overview

You can enable LDP downstream on demand label advertisement for an LDP session by including the **downstream-on-demand** statement at the **[edit protocols ldp session]** hierarchy level. If you have configured downstream on demand, the Juniper Networks router advertises the downstream on demand request to its peer routers. For a downstream on demand session to be established between two routers, both have to advertise downstream on demand mode during LDP session establishment. If one router advertises downstream unsolicited mode and the other advertises downstream on demand, downstream unsolicited mode is used.

### Configuration

#### *Configuring LDP Downstream on Demand*

**Step-by-Step Procedure** To configure a LDP downstream on demand policy and then configure that policy and enable LDP downstream on demand on the LDP session:

1. Configure the downstream on demand policy (DOD-Request-Loopbacks in this example).

This policy causes the router to forward label request messages only to the FECs that are matched by the DOD-Request-Loopbacks policy.

[edit policy-options]

```
user@host# set prefix-list Request-Loopbacks 10.1.1/32
```

```
user@host# set prefix-list Request-Loopbacks 10.1.1.2/32
```

```
user@host# set prefix-list Request-Loopbacks 10.1.1.3/32
```

```
user@host# set prefix-list Request-Loopbacks 10.1.1.4/32
```

```
user@host# set policy-statement DOD-Request-Loopbacks term 1 from prefix-list
Request-Loopbacks
```

```
user@host# set policy-statement DOD-Request-Loopbacks term 1 then accept
```

2. Specify the DOD-Request-Loopbacks policy using the **dod-request-policy** statement at the **[edit protocols ldp]** hierarchy level.

The policy specified with the **dod-request-policy** statement is used to identify the prefixes to send label request messages. This policy is similar to an egress policy or an import policy. When processing routes from the inet.0 routing table, the Junos OS software checks for routes matching the **DOD-Request-Loopbacks** policy (in this example). If the route matches the policy and the LDP session is negotiated with DOD advertisement mode, label request messages are sent to the corresponding downstream LDP session.

```
[edit protocols ldp]
user@host# set dod-request-policy DOD-Request-Loopbacks
```

3. Include the **downstream-on-demand** statement in the configuration for the LDP session to enable downstream on demand distribution mode.

```
[edit protocols ldp]
user@host# set session 1.1.1.1 downstream-on-demand
```

#### *Distributing LDP Downstream on Demand Routes into Labeled BGP*

**Step-by-Step Procedure** To distribute LDP downstream on demand routes into labeled BGP, use a BGP export policy.

1. Configure the LDP route policy (**redistribute\_ldp** in this example).

```
[edit policy-options]
user@host# set policy-statement redistribute_ldp term 1 from protocol ldp
user@host# set policy-statement redistribute_ldp term 1 from tag 1000
user@host# set policy-statement redistribute_ldp term 1 then accept
```

2. Include the LDP route policy, **redistribute\_ldp** in the BGP configuration (as a part of the BGP group configuration **ebgp-to-abr** in this example).

BGP forwards the LDP routes based on the **redistribute\_ldp** policy to the remote PE router

```
[edit protocols bgp]
user@host# set group ebgp-to-abr type external
user@host# set group ebgp-to-abr local-address 192.168.0.1
user@host# set group ebgp-to-abr peer-as 65319
user@host# set group ebgp-to-abr local-as 65320
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet unicast
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet labeled-unicast
rib inet.3
user@host# set group ebgp-to-abr neighbor 192.168.6.1 export redistribute_ldp
```

**Step-by-Step Procedure** To restrict label propagation to other routers configured in downstream unsolicited mode (instead of downstream on demand), configure the following policies:

1. Configure the **dod-routes** policy to accept routes from LDP.

```
user@host# set policy-options policy-statement dod-routes term 1 from protocol
ldp
```

- ```

user@host# set policy-options policy-statement dod-routes term 1 from tag
1145307136
user@host# set policy-options policy-statement dod-routes term 1 then accept

```
2. Configure the **do-not-propagate-du-sessions** policy to not forward routes to neighbors 1.1.1.1, 2.2.2.2, and 3.3.3.3.
 

```

user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 1.1.1.1
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 2.2.2.2
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 3.3.3.3
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 then reject

```
  3. Configure the **filter-dod-on-du-sessions** policy to prevent the routes examined by the **dod-routes** policy from being forwarded to the neighboring routers defined in the **do-not-propagate-du-sessions** policy.
 

```

user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 from policy dod-routes
user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 to policy do-not-propagate-du-sessions

```
  4. Specify the **filter-dod-routes-on-du-session** policy as the export policy for BGP group **ebgp-to-abr**.
 

```

[edit protocols bgp]
user@host# set group ebgp-to-abr neighbor 192.168.6.2 export
filter-dod-routes-on-du-sessions

```

**Results** From configuration mode, confirm your configuration by entering the **show policy-options** and **show protocols ldp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host#
show policy-options
prefix-list Request-Loopbacks {
  10.1.1.1/32;
  10.1.1.2/32;
  10.1.1.3/32;
  10.1.1.4/32;
}
policy-statement DOD-Request-Loopbacks {
  term 1 {
    from {
      prefix-list Request-Loopbacks;
    }
    then accept;
  }
}
policy-statement redistribute_ldp {
  term 1 {
    from {
      protocol ldp;
      tag 1000;
    }
  }
}

```

```
        then accept;
    }
}

user@host#
show protocols ldp
dod-request-policy DOD-Request-Loopbacks;
session 1.1.1.1 {
    downstream-on-demand;
}

user@host#
show protocols bgp
group ebgp-to-abr {
    type external;
    local-address 192.168.0.1;
    peer-as 65319;
    local-as 65320;
    neighbor 192.168.6.1 {
        family inet {
            unicast;
            labeled-unicast {
                rib {
                    inet.3;
                }
            }
        }
    }
    export redistribute_ldp;
}
}
```

---

## Verification

### *Verifying Label Advertisement Mode*

**Purpose** Confirm that the configuration is working properly.

Use the **show ldp session** command to verify the status of the label advertisement mode for the LDP session.

**Action** Issue the **show ldp session** and **show ldp session detail** commands:

- The following command output for the **show ldp session** command indicates that the **Adv. Mode** (label advertisement mode) is **DOD** (meaning the LDP downstream on demand session is operational):

```
user@host> show ldp session
  Address          State      Connection    Hold time  Adv. Mode
  1.1.1.2          Operational Open          22         DOD
```

- The following command output for the **show ldp session detail** command indicates that the **Local Label Advertisement mode** is **Downstream unsolicited**, the default value (meaning downstream on demand is not configured on the local session). Conversely, the **Remote Label Advertisement mode** and the **Negotiated Label Advertisement mode** both indicate that **Downstream on demand** is configured on the remote session

```
user@host> show ldp session detail
Address: 1.1.1.2, State: Operational, Connection: Open, Hold time: 24
Session ID: 1.1.1.1:0--1.1.1.2:0
Next keepalive in 4 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: configured-tunneled
Keepalive interval: 10, Connect retry interval: 1
Local address: 1.1.1.1, Remote address: 1.1.1.2
Up for 17:54:52
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: disabled, Helper mode: enabled,
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream on demand
Negotiated Label Advertisement mode: Downstream on demand
Nonstop routing state: Not in sync
Next-hop addresses received:
  1.1.1.2
```

## Configuration Tasks

- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [Configuring MPLS on Provider Switches on page 4471](#)
- [Configuring Static Label Switched Paths for MPLS on page 4472](#)
- [Configuring MPLS Firewall Filters and Policers on page 4474](#)
- [Configuring CoS Bits for an MPLS Network on page 4478](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring MPLS to Gather Statistics on page 4481](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 4482](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics on page 4489](#)
- [Configuring the LDP Timer for Hello Messages on page 4492](#)

- [Configuring the Delay Before LDP Neighbors Are Considered Down on page 4493](#)
- [Configuring the Interval for LDP Keepalive Messages on page 4494](#)
- [Configuring the LDP Keepalive Timeout on page 4495](#)
- [Configuring LDP Route Preferences on page 4495](#)
- [Configuring LDP Graceful Restart on page 4495](#)
- [Configuring the Prefixes Advertised into LDP from the Routing Table on page 4498](#)
- [Configuring LDP LSP Traceroute on page 4498](#)
- [Configuring Miscellaneous LDP Properties on page 4500](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 4506](#)

## Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. [Configuring the Ingress PE Switch on page 4468](#)
2. [Configuring the Egress PE Switch on page 4470](#)

### Configuring the Ingress PE Switch

---

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.10.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure OSPF traffic engineering:

```
[edit protocols ospf]
```

- ```

user@switch# set traffic-engineering

```
4. Configure RSVP on the loopback interface and the core interfaces:
 

```

[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0

```
  5. Configure MPLS traffic engineering.
 

```

[edit protocols mpls]
user@switch# set traffic-engineering

```
  6. Configure MPLS on the core interfaces:
 

```

[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0

```
  7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:
 

```

[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls

```
  8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:
 

```

[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 121.100.10.1/16

```
  9. Configure this Layer 3 customer edge interface for the routing protocol:
 

```

[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0

```
  10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):
 

```

[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1

```
  11. Disable constrained-path LSP computation for this LSP:
 

```

[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf

```
  12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:
 

```

[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve

```

## Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0
user@switch# set rsvp interface xe-0/0/6.0
```

4. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:



```
[edit protocols mpls]
```

```
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
```

```
user@switch# set static route 121.121.121.0/24 next-hop 192.168.12.1
```

```
user@switch# set static route 121.121.121.0/24 resolve
```

#### Related Documentation

- [MPLS Configuration Guidelines on page 4433](#)
- [Configuring MPLS on Provider Switches on page 4471](#)
- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [Understanding MPLS Components on page 4412](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)

## Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
```

```
user@switch# set area 0.0.0.0 interface lo0.0
```

```
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
```

```
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

```
user@switch# set area 0.0.0.0 interface ae0
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
```

```
user@switch# set interface xe-0/0/5.0
```

```
user@switch# set interface xe-0/0/6.0
```

```
user@switch# set interface ae0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
```

```
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

4. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
user@switch# set ae0 unit 0 family mpls
```



**NOTE:** You can configure **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

#### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [MPLS Configuration Guidelines on page 4433](#)
- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [Understanding MPLS Components on page 4412](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)

## Configuring Static Label Switched Paths for MPLS

Configuring static label-switched paths (LSPs) for MPLS is similar to configuring static routes on individual switches. As with static routes, there is no error reporting, liveness detection, or statistics reporting.

To configure static LSPs, configure the ingress PE switch and each provider switch along the path up to and including the egress PE switch.

For the ingress PE switch, configure which packets to tag (based on the packet's destination IP address), configure the next switch in the LSP, and the tag to apply to the packet. Manually assigned labels can have values from 0 through 1,048,575.

For the transit switches in the path, configure the next switch in the path and the tag to apply to the packet. Manually assigned labels can have values from 1,000,000 through 1,048,575.

The egress PE switch removes the label and forwards the packet to the IP destination. However, if the previous switch removed the label, the egress switch examines the packet's IP header and forwards the packet toward its IP destination.

Before you configure a static LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See “Configuring MPLS on Provider Edge Switches” on page 4468.



**NOTE:** Do not configure LSPs at the [edit protocols mpls label-switched-path] hierarchy level on the PE switches.

- Configure one or more provider switches. See “Configuring MPLS on Provider Switches” on page 4471.

This topic describes how to configure an ingress PE switch, one or more provider switches, and an egress PE switch for static LSP:

1. [Configuring the Ingress PE Switch on page 4473](#)
2. [Configuring the Provider and the Egress PE Switch on page 4474](#)

### Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for every core interface:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family inet address address
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure the name associated with the static LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name
```

3. Configure the next hop switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress next-hop address-of-next-hop
```

4. Specify the address of the egress switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress to address-of-egress-switch
```

5. Configure the new label that you want to add to the top of the label stack:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress push out-label
```

## Configuring the Provider and the Egress PE Switch

To configure a static LSP for MPLS on the provider and egress PE switch:

1. Configure a transit static LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name transit incoming-label
```

2. Configure the next hop switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name transit incoming-label next-hop
address-of-next-hop
```

3. Only for provider switches, remove the label at the top of the label stack and replace it with the specified label:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name transit incoming-label swap out-label
```

4. Only for the egress PE switch, remove the label at the top of the label stack:



**NOTE:** If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name transit incoming-label pop
```

### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [Configuring MPLS on Provider Switches on page 4471](#)
- [Understanding MPLS Label Operations on page 4415](#)

## Configuring MPLS Firewall Filters and Policers

You can configure firewall filters to filter MPLS traffic. To use an MPLS firewall filter, you must first configure the filter and then apply it to an interface you have configured for forwarding MPLS traffic. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached.



**NOTE:** You can configure ingress MPLS firewall filters only. Egress MPLS firewall filters are not supported. You cannot apply MPLS firewall filters to loopback interfaces.

When you configure an MPLS firewall filter, you define filtering criteria (terms, with match conditions) for the packets and an action (action, or action modifier) for the switch to take if the packets match the filtering criteria.

- [Table 348 on page 4475](#) describes the match conditions you can configure for MPLS firewall filters at the `[edit firewall family mpls filter filter-name term term-name from]` hierarchy level.



**NOTE:** If a packet has multiple MPLS labels, the filter applies the match conditions to only the bottom label in the label stack.

**Table 348: Supported Match Conditions for MPLS Firewall Filters**

Match Condition	Description
<b>exp <i>number</i></b>	<p>Experimental (EXP) bit number or range of bit numbers in the MPLS header of a packet.</p> <p>For <i>number</i>, you can specify one or more values from 0 through 7 in binary, decimal or hexadecimal format, as given below:</p> <ul style="list-style-type: none"> <li>• A single EXP bit—for example, <b>exp 3</b></li> <li>• Several EXP bits—for example, <b>exp 0,4</b></li> <li>• A range of EXP bits—for example, <b>exp [0-5]</b></li> </ul>
<b>label <i>number</i></b>	<p>MPLS label value or range of label values in the MPLS header of a packet.</p> <p>For <i>number</i>, you can specify one or more values from 0 through 1048575 in decimal or hexadecimal format, as given below:</p> <ul style="list-style-type: none"> <li>• A single label—for example, <b>label 3</b></li> <li>• Several labels—for example, <b>label 0,4</b></li> <li>• A range of labels—for example, <b>label [0-5]</b></li> </ul>

- [Table 349 on page 4475](#) describes the actions you can configure for MPLS firewall filters at the **[edit firewall family mpls filter *filter-name* term *term-name* then]** hierarchy level.

**Table 349: Supported Actions for MPLS Firewall Filters**

Action	Description
<b>accept</b>	Accept a packet
<b>count <i>counter-name</i></b>	<p>Count the number of packets that pass this filter or term.</p> <p><b>NOTE:</b> We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.</p>
<b>discard</b>	Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message
<b>policer</b>	Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer.

Table 349: Supported Actions for MPLS Firewall Filters (*continued*)

Action	Description
<b>three-color-policer</b>	Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a three-color policer.

- [Configuring an MPLS Firewall Filter on page 4476](#)
- [Applying an MPLS Firewall Filter to an MPLS Interface on page 4476](#)
- [Configuring Policers for LSPs on page 4477](#)

### Configuring an MPLS Firewall Filter

To configure an MPLS firewall filter:

1. Configure the filter name, term name, and at least one match condition—for example, match on MPLS packets with EXP bits set to either 0 or 4:

```
[edit firewall family mpls]
user@switch# set filter ingress-exp-filter term term-one from exp 0,4
```

2. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term—for example, count MPLS packets with EXP bits set to either 0 or 4:

```
[edit firewall family mpls filter ingress-exp-filter term term-one then]
user@switch# set count counter0
user@switch# set accept
```

### Applying an MPLS Firewall Filter to an MPLS Interface

To apply the MPLS firewall filter to an interface you have configured for forwarding MPLS traffic (using the **family mpls** statement at the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level):



**NOTE:** You can apply firewall filters only to filter MPLS packets that enter an interface.

1. Apply the firewall filter to an MPLS interface—for example, apply the firewall filter to interface xe-0/0/5:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls filter input ingress-exp-filter
```

2. Review your configuration and issue the **commit** command:

```
[edit interfaces]
user@switch# commit
commit complete
```

---

## Configuring Policers for LSPs

---

Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer or three-color policer. MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.

### Related Documentation

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)
- [Supported MPLS Scaling Values on page 4426](#)
- [Overview of Policers on page 5241](#)

## Configuring CoS Bits for an MPLS Network

When traffic enters a labeled-switch path (LSP) tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The *Class of Service Feature Guide for Routing Devices* explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP:

1. Specify a class of service value for the LSP:



**NOTE:** The CoS value set using the `class-of-service` statement at the `[edit protocols mpls]` hierarchy level supersedes the CoS value set at the `[edit class-of-service]` hierarchy level for an interface. Effectively, the CoS value configured for an LSP overrides the CoS value set for an interface.

```
[edit protocols mpls]
user@switch# set class-of-service cos-value
```

### Related Documentation

- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Defining CoS Rewrite Rules on page 6182](#)



## Configuring a Global MPLS EXP Classifier

EXP packet classification associates incoming packets with a particular MPLS CoS servicing level. EXP behavior aggregate (BA) classifiers examine the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. EXP BA classifiers allow you to set the forwarding class and loss priority of an MPLS packet based on the incoming CoS value.

You can configure as many EXP classifiers as you want, however, the switch uses only one MPLS EXP classifier as a global classifier, which is applied only on interfaces configured as **family mpls**. All **family mpls** switch interfaces use the global EXP classifier to classify MPLS traffic.

If an EXP classifier is configured, MPLS traffic on **family mpls** interfaces uses the EXP classifier. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.



**NOTE:** There is no default MPLS EXP classifier. If you want to use an MPLS EXP classifier, you must configure it. The MPLS EXP classifier is global and applies only to all **family mpls** interfaces on the switch. You can configure as many MPLS EXP classifiers as you want, but you can only use one MPLS EXP classifier on switch interfaces at any time.

To configure a unicast MPLS EXP classifier using the CLI:

1. Create an EXP classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1 | exp) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the EXP classifier to the switch interfaces:

```
[edit class-of-service]
user@switch# set system-defaults classifiers exp classifier-name
```

### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)

## Configuring Rewrite Rules for MPLS EXP Classifiers

You configure EXP rewrite rules to alter CoS values in outgoing MPLS packets on the outbound **family mpls** interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure an EXP CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on a logical **family mpls** interface. EXP rewrite rules can only be enabled on logical **family mpls** interfaces, not on physical interfaces or on interfaces of other family types. You can also apply an existing EXP rewrite rule on a logical interface.



**NOTE:** There are no default rewrite rules.

You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch. On a given **family mpls** logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

To create an EXP rewrite rule for MPLS traffic and enable it on a logical interface:

1. Create an EXP rewrite rule:

```
user@switch# set class-of-service rewrite-rules exp rewrite-rule-name forwarding-class forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

For example, to configure an EXP rewrite rule named **exp-rr-1** for a forwarding class named **mpls-1** with a loss priority of **low** that rewrites the EXP code point value to **001**:

```
user@switch# set class-of-service rewrite-rules exp exp-rr-1 forwarding-class mpls-1 loss-priority low code-points 001
```

2. Apply the rewrite rule to a logical interface:

```
user@switch # set class-of-service interfaces interface-name unit logical-unit rewrite-rules exp rewrite-rule-name
```

For example, to apply a rewrite rule named **exp-rr-1** to logical interface **xe-0/0/10.0**:

```
user@switch# set class-of-service interfaces xe-0/0/10 unit 0 rewrite-rules exp exp-rr-1
```



**NOTE:** In this example, all forwarding classes assigned to port xe-0/0/10 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same interface.

#### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Monitoring CoS Rewrite Rules on page 6292](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)

## Configuring MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions, by configuring the **statistics** statement. You must configure the **statistics** statement if you want to collect MPLS traffic statistics using SNMP polling of MPLS Management Information Bases (MIBs).

To enable or disable MPLS statistics collection, include the **statistics** statement:

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
}
```

You can configure these statements at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

The default interval is 300 seconds.

If you configure the **file** option, the statistics are placed in a file, with one entry per LSP. During the specified interval, the following information is recorded in this file:

- The number of packets, number of bytes, packets per second, and bytes per second transmitted by each LSP. Feature parity for the display of packet and byte statistics for sub-LSPs of a point-to-multipoint LSP on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.
- The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. If no bandwidth is configured for an LSP, 0 percent is recorded in the percentage column.

At the end of each periodic report, a summary shows the current time, total number of sessions, number of sessions read, number of sessions ignored, and read errors, if any. Ignored sessions are typically those not in the up state or those with a reserved (0 through 15) incoming label (typically the egress point of an LSP). The reason for a read error appears on the same line as the entry for the LSP on which the error occurred. Gathering statistics is an unreliable process; occasional read errors might affect their accuracy. Sample output follows:

lsp6	0 pkt	0 Byte	0 pps	0 Bps	0
lsp5	0 pkt	0 Byte	0 pps	0 Bps	0
lsp6.1	34845 pkt	2926980 Byte	1049 pps	88179 Bps	132
lsp5.1	0 pkt	0 Byte	0 pps	0 Bps	0
lsp4	0 pkt	0 Byte	0 pps	0 Bps	0

Dec 7 17:28:38 Total 6 sessions: 5 success, 0 fail, 1 ignored

**Related  
Documentation**

- [Configuring Automatic Bandwidth Allocation for LSPs on page 4482](#)

## Configuring Automatic Bandwidth Allocation for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

At the end of the automatic bandwidth allocation time interval, the current maximum average bandwidth usage is compared with the allocated bandwidth for the LSP. If the LSP needs more bandwidth, an attempt is made to set up a new path where bandwidth is equal to the current maximum average usage. If the attempt is successful, the LSP's traffic is routed through the new path and the old path is removed. If the attempt fails, the LSP continues to use its current path.

If you have configured link and node protection for the LSP and traffic has been switched to the bypass LSP, the automatic bandwidth allocation feature continues to operate and take bandwidth samples from the bypass LSP. For the first bandwidth adjustment cycle, the maximum average bandwidth usage taken from the original link and node-protected LSP is used to resignal the bypass LSP if more bandwidth is needed. (Link and node protection are not supported on QFX Series switches.)

If you have configured fast-reroute for the LSP, you might not be able to use this feature to adjust the bandwidth. Because the LSPs use a fixed filter (FF) reservation style, when a new path is signaled, the bandwidth might be double-counted. Double-counting can prevent a fast-reroute LSP from ever adjusting its bandwidth when automatic bandwidth allocation is enabled. (Fast reroute is not supported on QFX Series switches.)

To configure automatic bandwidth allocation, complete the steps in the following sections:

- [Configuring Automatic Bandwidth Allocation on LSPs on page 4483](#)
- [Requesting Automatic Bandwidth Allocation Adjustment on page 4488](#)

## Configuring Automatic Bandwidth Allocation on LSPs

To enable automatic bandwidth allocation on an LSP, include the **auto-bandwidth** statement:

```
auto-bandwidth {
  adjust-interval seconds;
  adjust-threshold percent;
  adjust-threshold-overflow-limit number;
  adjust-threshold-underflow-limit number;
  maximum-bandwidth bps;
  minimum-bandwidth bps;
  minimum-bandwidth-adjust-interval
  minimum-bandwidth-adjust-threshold-change
  minimum-bandwidth-adjust-threshold-value
  monitor-bandwidth;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

If an LSP has an automatic bandwidth configuration, you can disable automatic bandwidth adjustments on a particular path (either primary or secondary) by configuring a static bandwidth value and by disabling the CSPF computation (using the **no-cspf** statement).

For example:

```
user@host> show protocols mpls
label-switched-path primary-path {
  to 192.168.0.1;
  ldp-tunneling;
  optimize-timer 3571;
  least-fill;
  link-protection;
  adaptive;
  auto-bandwidth {
    adjust-interval 7177;
    adjust-threshold 5;
    minimum-bandwidth 1m;
    maximum-bandwidth 2500000000;
    adjust-threshold-overflow-limit 2;
    resignal-minimum-bandwidth;
  }
  primary primary-path;
  secondary secondary-path {
    bandwidth 0;
    no-cspf;
    priority 0 0;
  }
}
```

The statements configured at the **[edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth]** hierarchy level are optional and explained in the following sections:

- [Configuring the Automatic Bandwidth Allocation Interval on page 4484](#)
- [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 4484](#)
- [Configuring the Automatic Bandwidth Adjustment Threshold on page 4485](#)
- [Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 4485](#)
- [Configuring Passive Bandwidth Utilization Monitoring on page 4487](#)

#### ***Configuring the Automatic Bandwidth Allocation Interval***

At the end of the automatic bandwidth allocation interval, the automatic bandwidth computation and new path setup process is triggered.



**NOTE:** To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (**interval** statement at the **[edit protocols mpls statistics]** hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (**adjust-interval** statement at the **[edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth]** hierarchy level). See also [“Configuring Reporting of Automatic Bandwidth Allocation Statistics” on page 4489](#).

To specify the bandwidth reallocation interval in seconds for a specific LSP, include the **adjust-interval** statement:

```
adjust-interval seconds;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]**
- **[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* auto-bandwidth]**

#### ***Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth***

You can maintain the LSP's bandwidth between minimum and maximum bounds by specifying values for the **minimum-bandwidth** and **maximum-bandwidth** statements.

To specify the minimum amount of bandwidth allocated for a specific LSP, include the **minimum-bandwidth** statement:

```
minimum-bandwidth bps;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name* auto-bandwidth]**

- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* [auto-bandwidth](#)]

To specify the maximum amount of bandwidth allocated for a specific LSP, include the **maximum-bandwidth** statement:

```
maximum-bandwidth bps;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* [auto-bandwidth](#)]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* [auto-bandwidth](#)]

### ***Configuring the Automatic Bandwidth Adjustment Threshold***

Use the **adjust-threshold** statement to specify the sensitivity of the automatic bandwidth adjustment of an LSP to changes in bandwidth utilization. You can set the threshold for when to trigger automatic bandwidth adjustments. When configured, bandwidth demand for the current interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the specified **adjust-threshold** percentage, the LSP's bandwidth is adjusted to the current bandwidth demand.

For example, assume that the current bandwidth allocation is 100 megabits per second (Mbps) and that the percentage configured for the **adjust-threshold** statement is 15 percent. If the bandwidth demand increases to 110 Mbps, the bandwidth allocation is not adjusted. However, if the bandwidth demand increases to 120 Mbps (20 percent over the current allocation) or decreases to 80 Mbps (20 percent under the current allocation), the bandwidth allocation is increased to 120 Mbps or decreased to 80 Mbps, respectively.

To configure the threshold for automatic bandwidth adjustment, include the **adjust-threshold** statement:

```
adjust-threshold percent;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* [auto-bandwidth](#)]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* [auto-bandwidth](#)]

### ***Configuring a Limit on Bandwidth Overflow and Underflow Samples***

The automatic bandwidth adjustment timer is a periodic timer which is triggered every adjust interval to determine whether any bandwidth adjustments are required on the LSP's active path. This interval is typically configured as a long period of time, usually hours. If, at the end of adjust interval, the change in bandwidth is above a certain adjust threshold, the LSP is resigaled with the new bandwidth.

During the automatic bandwidth adjustment interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing

congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Every statistics interval, the router samples the average bandwidth utilization of an LSP and if this has exceeded the current maximum average bandwidth utilization, the maximum average bandwidth utilization is updated.

During each sample period, the following conditions are also checked:

- Is the current average bandwidth utilization above the active bandwidth of the path?
- Has the difference between the average bandwidth utilization and the active bandwidth exceeded the adjust threshold (bandwidth utilization has changed significantly)?

If these conditions are true, it is considered to be one bandwidth overflow sample. Using the **adjust-threshold-overflow-limit** statement, you can define a limit on the number of bandwidth overflow samples such that when the limit is reached, the current automatic bandwidth adjustment timer is expired and a bandwidth adjustment is triggered. Once this adjustment is complete, the normal automatic bandwidth adjustment timer is reset to expire after the periodic adjustment interval.

To specify a limit on the number of bandwidth overflow samples before triggering an automatic bandwidth allocation adjustment, configure the **adjust-threshold-overflow-limit** statement:

**adjust-threshold-overflow-limit** *number*;

Similarly, if the current average bandwidth utilization is below the active bandwidth of the path by the configured adjusted threshold (meaning that bandwidth utilization has gone down significantly), the sample is considered to be an underflow sample. The adjusted (new signaling) bandwidth after an adjustment due to underflow is the maximum average bandwidth among the underflow samples. You can specify a limit on the number of bandwidth underflow samples before triggering an automatic bandwidth allocation adjustment by configuring the **adjust-threshold-underflow-limit** statement:

**adjust-threshold-underflow-limit** *number*;

These statements can be configured at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* **auto-bandwidth**]

You must configure the **adjust-threshold** and **minimum-bandwidth** statements whenever you configure the **adjust-threshold-underflow-limit** statement. You must configure the **adjust-threshold** and **maximum-bandwidth** statements whenever you configure the **adjust-threshold-overflow-limit** statement

- You must configure a nonzero value for the **adjust-threshold** statement if you configure the **adjust-threshold-overflow-limit** or **adjust-threshold-underflow-limit** statement.



- Any bandwidth increase or decrease below the value configured for the **adjust-threshold** statement does not constitute an overflow or underflow condition.
- To prevent unlimited increases in LSP bandwidth (to limit overflow beyond a certain bandwidth), you must also configure the **maximum-bandwidth** statement when you configure the **adjust-threshold-overflow-limit** statement.

The following describes the other aspects of the **adjust-threshold-overflow-limit** statement:

- It only applies to bandwidth overflows. If the bandwidth is decreasing, the normal automatic bandwidth adjustment interval is used.
- It does not affect manually triggered automatic bandwidth adjustment.
- It applies to single-class DiffServ-TE LSPs.
- Because the **adjust-threshold-overflow-limit** statement can trigger a bandwidth adjustment, it cannot be enabled at the same time as the **monitor-bandwidth** statement (for information about that statement, see [“Configuring Passive Bandwidth Utilization Monitoring” on page 4487](#)).
- You cannot configure automatic bandwidth adjustments to occur more often than every 300 seconds. The **adjust-threshold-overflow-limit** statement is subject to the same minimum value with regard to the minimum frequency of adjustment allowed. Overflow condition based adjustments can occur no sooner than 300 seconds from the start of the overflow condition. Therefore it is required that:

**sample interval x adjust-threshold-overflow-limit >= 300s**

These values are checked during the commit operation. An error is returned if the value is less than 300 seconds.

- If you change the value of the **adjust-threshold-overflow-limit** statement on a working router, you can expect the following behavior:
  - If you increase the current value of the **adjust-threshold-overflow-limit** statement, the old value is replaced with the new one.
  - If you decrease the current value of the **adjust-threshold-overflow-limit** statement and the current bandwidth overflow count is less than the new value, the old value is replaced with the new one.
  - If you decrease the current value of the **adjust-threshold-overflow-limit** statement and the current bandwidth overflow count is greater than the new value, the adjustment timer is immediately expired and a bandwidth adjustment is initiated.

### ***Configuring Passive Bandwidth Utilization Monitoring***

Use the **monitor-bandwidth** statement to switch to a passive bandwidth utilization monitoring mode. In this mode, no automatic bandwidth adjustments are made, but the maximum average bandwidth utilization is continuously monitored and recorded.

To configure passive bandwidth utilization monitoring, include the **monitor-bandwidth** statement:

**monitor-bandwidth;**

You can include this statement at the following hierarchy levels:

- [edit protocols mpls label-switched-path *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* **auto-bandwidth**]

If you have configured an LSP with primary and secondary paths, the automatic bandwidth allocation statistics are carried over to the secondary path if the primary path fails. For example, consider a primary path whose adjustment interval is half complete and whose maximum average bandwidth usage is currently calculated as 50 Mbps. If the primary path suddenly fails, the time remaining for the next adjustment and the maximum average bandwidth usage are carried over to the secondary path.

### Requesting Automatic Bandwidth Allocation Adjustment

---

For MPLS LSP automatic bandwidth allocation adjustment, the minimum value for the adjustment interval is 5 minutes (300 seconds). You might find it necessary to trigger a bandwidth allocation adjustment manually, for example in the following circumstances:

- When you are testing automatic bandwidth allocation in a network lab.
- When the LSP is configured for automatic bandwidth allocation in monitor mode (the **monitor-bandwidth** statement is included in the configuration as described in [“Configuring Passive Bandwidth Utilization Monitoring” on page 4487](#)), and want to initiate an immediate bandwidth adjustment.

To use the **request mpls lsp adjust-autobandwidth** command, the following must be true:

- Automatic bandwidth allocation must be enabled on the LSP.
- The criteria required to trigger a bandwidth adjustment have been met (the difference between the adjust bandwidth and the current LSP path bandwidth is greater than the threshold limit).

A manually triggered bandwidth adjustment operates only on the active LSP path. Also, if you have enabled periodic automatic bandwidth adjustment, the periodic automatic bandwidth adjustment parameters (the adjustment interval and the maximum average bandwidth) are not reset after a manual adjustment.

For example, suppose the periodic adjust interval is 10 hours and there are currently 5 hours remaining before an automatic bandwidth adjustment is triggered. If you initiate a manual adjustment with the **request mpls lsp adjust-autobandwidth** command, the adjust timer is not reset and still has 5 hours remaining.

To manually trigger a bandwidth allocation adjustment, you need to use the **request mpls lsp adjust-autobandwidth** command. You can trigger the command for all affected LSPs on the router, or you can specify a particular LSP:

```
user@host> request mpls lsp adjust-autobandwidth
```

Once you execute this command, the automatic bandwidth adjustment validation process is triggered. If all the criteria for adjustment are met, the LSP's active path bandwidth is adjusted to the adjusted bandwidth value determined during the validation process.

**Related  
Documentation**

- [Configuring MPLS to Gather Statistics on page 4481](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics on page 4489](#)
- [request mpls lsp adjust-autobandwidth on page 4598](#)
- [show mpls lsp on page 4665](#)

## Configuring Reporting of Automatic Bandwidth Allocation Statistics

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure the device to collect statistics related to automatic bandwidth allocation by completing the following steps:

1. To collect statistics related to automatic bandwidth allocation, configure the **auto-bandwidth** option for the **statistics** statement at the **[edit protocols mpls]** hierarchy level. These settings apply to all LSPs configured on the router on which you have also configured the **auto-bandwidth** statement at the **[edit protocols mpls label-switched-path label-switched-path-name]** hierarchy level.
 

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
}
```
2. Specify the **filename** for the files used to store the MPLS trace operation output using the **file** option. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.
3. Specify the maximum number of trace files using the **files number** option. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
4. Specify the interval for calculating the average bandwidth usage by configuring a time in seconds using the **interval** option. You can also set the adjustment interval on a specific LSP by configuring the **interval** option at the **[edit protocols mpls label-switch-path label-switched-path-name statistics]** hierarchy level.



**NOTE:** To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (interval statement at the [edit protocols mpls statistics] hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (adjust-interval statement at the [edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth] hierarchy level).

- To trace automatic bandwidth allocation, include the **autobw-state** flag for the MPLS **traceoptions** statement at the [edit protocols mpls] hierarchy level.

The following configuration enables the MPLS traceoptions for automatic bandwidth allocation. The trace records are stored in a file called **auto-band-trace** (the filename is user configurable):

```
[edit protocols mpls]
traceoptions {
  file auto-band-trace size 10k files 10 world-readable;
  flag autobw-state;
}
```

- Using the **show log** command, you can display the automatic bandwidth allocation statistics file generated when you configure the **auto-bandwidth** statement. The following shows sample log file output taken from an MPLS statistics file named **auto-band-stats** on a router configured with an LSP named **E-D**. The log file shows that LSP **E-D** is operating over its reserved bandwidth limit initially. Before **Oct 30 17:14:57**, the router triggered an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By **Oct 30 17:16:57**, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```
user@host> show log auto-band-stats
E-D          (LSP ID 5, Tunnel ID 6741)          209 pkt          17094 Byte          1 pps          90 Bps Util
 240.01% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:13:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          241 pkt          19737 Byte          1 pps          88 Bps Util
 234.67% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:14:27 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          276 pkt          22607 Byte          1 pps          95 Bps Util
 253.34% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:14:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)           0 pkt           0 Byte           0 pps           0 Bps Util
  0.00% Reserved Bw          37 Bps
E-D          (LSP ID 6, Tunnel ID 6741)           0 pkt           0 Byte           0 pps           0 Bps Util
  0.00% Reserved Bw          101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:15:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)           0 pkt           0 Byte           0 pps           0 Bps Util
```

```

0.00% Reserved Bw      37 Bps
E-D      (LSP ID 6, Tunnel ID 6741)      33 pkt      2695 Byte      1 pps      89 Bps Util
87.69% Reserved Bw      101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:15:57 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 5, Tunnel ID 6741)      0 pkt      0 Byte      0 pps      0 Bps Util
0.00% Reserved Bw      37 Bps
E-D      (LSP ID 6, Tunnel ID 6741)      65 pkt      5338 Byte      1 pps      88 Bps Util
86.70% Reserved Bw      101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:16:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 6, Tunnel ID 6741)      97 pkt      7981 Byte      1 pps      88 Bps Util
86.70% Reserved Bw      101 Bps
decr nh 0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:16:57 Total 1 sessions: 1
success, 0 fail, 0 ignored

```

7. Issue the `show mpls lsp autobandwidth` command to display current information about automatic bandwidth allocation. The following shows sample output from the `show mpls lsp autobandwidth` command taken at about the same time as the log file shown previously:

```

user@host> show mpls lsp autobandwidth
Lspname      Last      Requested      Reserved      Highwater      AdjustTime LastAdjust
BW           BW           BW           mark           Left (sec)
E-D          300bps      812.005bps    812bps        1.56801kbps 294 sec      Wed Oct 30 17:15:26 2013

```

8. Issue the `file show` command to display the MPLS trace file. You need to specify the file location and file name (the file is located in `/var/log/`). The following shows sample trace file output is taken from an MPLS trace file named `auto-band-trace.0.gz` on a router configured with an LSP named `E-D`. The trace file shows that LSP `E-D` is operating over its reserved bandwidth limit initially. At `Oct 30 17:15:26`, the router triggers an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By `Oct 30 17:15:57`, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```

user@host> file show /var/log/auto-band-trace.0.gz
Oct 30 17:13:57 trace_on: Tracing to "/var/log/E/auto-band-trace" started
Oct 30 17:13:57.466825 LSP E-D (id 5) new bytes arrived      2714 in 29
sec
Oct 30 17:14:27.466713 E-D      (LSP ID 5, Tunnel ID 6741)      241
pkt      19737 Byte      1 pps      88 Bps Util 234.67% Reserved Bw
      37 Bps
Oct 30 17:14:27.466962 LSP E-D (id 5, old id 5); sampled bytes      19737 >
bytes recorded      17094
Oct 30 17:14:27.467035 LSP E-D (id 5) new bytes arrived      2643 in 29
sec
Oct 30 17:14:57.466599 E-D      (LSP ID 5, Tunnel ID 6741)      276
pkt      22607 Byte      1 pps      95 Bps Util 253.34% Reserved Bw
      37 Bps
Oct 30 17:14:57.466758 LSP E-D (id 5, old id 5); sampled bytes      22607 >
bytes recorded      19737
Oct 30 17:14:57.466825 LSP E-D (id 5) new bytes arrived      2870 in 29
sec
Oct 30 17:15:26.265816 Adjust Autobw: LSP E-D (id 5) curr adj bw 300bps updated
with 812.005bps
Oct 30 17:15:26.266064 mpls LSP E-D Autobw change 512.005bps >= threshold 75bps
Oct 30 17:15:26.363372 Autobw Success: LSP E-D () (old id 5 new id 6) update
prev active bw 300 bps with 812 bps

```

```

Oct 30 17:15:26.363686 RPD_MPLS_PATH_BANDWIDTH_CHANGE: MPLS path (lsp E-D)
bandwidth changed, path bandwidth 812 bps
Oct 30 17:15:27.364751 RPD_MPLS_LSP_BANDWIDTH_CHANGE: MPLS LSP E-D bandwidth
changed, lsp bandwidth 812 bps
Oct 30 17:15:27.466849 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:27.467050 E-D (LSP ID 6, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
101 Bps
Oct 30 17:15:57.466858 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:57.467106 E-D (LSP ID 6, Tunnel ID 6741) 33
pkt 2695 Byte 1 pps 89 Bps Util 87.69% Reserved Bw
101 Bps
Oct 30 17:15:57.467201 LSP E-D (id 6, old id 5); LSP up after autobw adjustment
and active for 30 sec
Oct 30 17:15:57.467398 LSP E-D (id 6) psb bytes 2695 < bytes recorded
22607 total bytes 2695 in 30 sec
Oct 30 17:15:57.467461 First sample of the adjust interval after automatic bw
adjustment
Oct 30 17:15:57.467594 Update curr max avg bw 0bps of LSP E-D with new bw
716.225bps
Oct 30 17:16:27.466830 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:16:27.467079 E-D (LSP ID 6, Tunnel ID 6741) 65
pkt 5338 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:27.467171 LSP E-D (id 6, old id 6); sampled bytes 5338 >
bytes recorded 2695
Oct 30 17:16:27.467237 LSP E-D (id 6) new bytes arrived 2643 in 29
sec
Oct 30 17:16:57.466712 E-D (LSP ID 6, Tunnel ID 6741) 97
pkt 7981 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:57.466870 LSP E-D (id 6, old id 6); sampled bytes 7981 >
bytes recorded 5338

```

- Related Documentation**
- [Configuring Automatic Bandwidth Allocation for LSPs on page 4482](#)
  - [show mpls lsp autobandwidth on page 4681](#)

## Configuring the LDP Timer for Hello Messages

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions

between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 4493](#).

### Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the **hello-interval** statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the **hello-interval** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
  hello-interval seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Configuring the Delay Before LDP Neighbors Are Considered Down

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.



**NOTE:** By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see [“Configuring the LDP Timer for Hello Messages” on page 4492](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an LDP peer router advertises a shorter hold time than the value you have configured, the peer router's advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

If the local LDP hold time is not shortened during LDP peer negotiation, the user-configured keepalive interval is left unchanged. However, if the local hold time is reduced during peer negotiation, the keepalive interval is recalculated. If the LDP hold time has been reduced during peer negotiation, the keepalive interval is reduced to one-third of the new hold time value. For example, if the new hold-time value is 45 seconds, the keepalive interval is set to 15 seconds.

This automated keepalive interval calculation can cause different keepalive intervals to be configured on each peer router. This enables the routers to be flexible in how often they send keepalive messages, because the LDP peer negotiation ensures they are sent more frequently than the LDP hold time.

When you reconfigure the hold-time interval, changes do not take effect until after the session is reset. The hold time is negotiated when the LDP peering session is initiated and cannot be renegotiated as long as the session is up (required by RFC 5036, *LDP Specification*). To manually force the LDP session to reset, issue the **clear ldp session** command.

---

### Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

### Configuring the LDP Hold Time for Targeted Hello Messages

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {  
  hold-time seconds;  
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Configuring the Interval for LDP Keepalive Messages

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds. The minimum value is 1 second.



The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 4493](#).

To modify the keepalive interval, include the **keepalive-interval** statement:

```
keepalive-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the LDP Keepalive Timeout

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

```
keepalive-timeout seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

## Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

```
preference preference;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring LDP Graceful Restart

When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router is to assist neighboring routers attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- [Enabling Graceful Restart on page 4496](#)
- [Disabling LDP Graceful Restart or Helper Mode on page 4496](#)
- [Configuring Reconnect Time on page 4497](#)
- [Configuring Recovery Time and Maximum Recovery Time on page 4497](#)

---

### Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router. To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

---

### Disabling LDP Graceful Restart or Helper Mode

To disable LDP graceful restart and recovery, include the **disable** statement:

```
ldp {  
  graceful-restart {  
    disable;  
  }  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the **helper-disable** statement:

```
ldp {  
  graceful-restart {  
    helper-disable;  
  }  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in the initialization message. The router behaves as a router that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

### Configuring Reconnect Time

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router to resume sending LDP messages. After the wait period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

Suppose that Router A and Router B are LDP neighbors. Router A is the restarting Router. The reconnect time is the time that Router A tells Router B to wait after Router B detects that Router A restarted.

To configure the reconnect time, include the **reconnect-time** statement:

```
graceful-restart {
  reconnect-time seconds;
}
```

You can set the reconnect time to a value in the range from 30 through 300 seconds. By default, it is 60 seconds.

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

### Configuring Recovery Time and Maximum Recovery Time

The recovery time is the amount of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router from being adversely affected if it receives a false value for the recovery time from the restarting router, you can configure the maximum recovery time on the neighboring router. A neighboring router maintains its state for the shorter of the two times. For example, Router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring Router B. However, Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```
graceful-restart {  
  maximum-neighbor-recovery-time seconds;  
  recovery-time seconds;  
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring the Prefixes Advertised into LDP from the Routing Table

You can control the set of prefixes that are advertised into LDP and cause the router to be the egress router for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

```
egress-policy policy-name;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

The named policy (configured at the **[edit policy-options]** or **[edit logical-systems logical-system-name policy-options]** hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

---

### Example: Configuring the Prefixes Advertised into LDP

---

Advertise all connected routes into LDP:

```
[edit protocols]  
ldp {  
  egress-policy connected-only;  
}  
policy-options {  
  policy-statement connected-only {  
    from {  
      protocol direct;  
    }  
    then accept;  
  }  
}
```

## Configuring LDP LSP Traceroute

You can trace the route followed by an LDP-signaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This

feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the **oam** statement configured at the **[edit protocols ldp]** hierarchy level. To configure periodic LDP LSP traceroute, include the **periodic-traceroute** statement:

```
periodic-traceroute {
  disable;
  exp exp-value;
  fanout fanout-value;
  frequency minutes;
  paths number-of-paths;
  retries retry-attempts;
  source address;
  ttl ttl-value;
  wait seconds;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit protocols ldp oam fec *address*]**

You can configure the **periodic-traceroute** statement by itself or with any of the following options:

- **exp**—Specify the class of service to use when sending probes.
- **fanout**—Specify the maximum number of next hops to search per node.
- **frequency**—Specify the interval between traceroute attempts.
- **paths**—Specify the maximum number of paths to search.
- **retries**—Specify the number of attempts to send a probe to a specific node before giving up.
- **source**—Specify the IPv4 source address to use when sending probes.
- **ttl**—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- **wait**—Specify the wait interval before resending a probe packet.

## Configuring Miscellaneous LDP Properties

The following sections describe how to configure a number of miscellaneous LDP properties:

- [Configuring LDP to Use the IGP Route Metric on page 4500](#)
- [Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 4500](#)
- [Multiple-Instance LDP and Carrier-of-Carriers VPNs on page 4501](#)
- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 4501](#)
- [Enabling LDP over RSVP-Established LSPs on page 4501](#)
- [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 4502](#)
- [Configuring the TCP MD5 Signature for LDP Sessions on page 4502](#)
- [Configuring LDP Session Protection on page 4503](#)
- [Disabling SNMP Traps for LDP on page 4504](#)
- [Configuring LDP Synchronization with the IGP on LDP Links on page 4504](#)
- [Configuring LDP Synchronization with the IGP on the Router on page 4505](#)
- [Configuring the Label Withdrawal Timer on page 4505](#)
- [Ignoring the LDP Subnet Check on page 4505](#)

---

### Configuring LDP to Use the IGP Route Metric

Use the **track-igp-metric** statement if you want the interior gateway protocol (IGP) route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

To use the IGP route metric, include the **track-igp-metric** statement:

```
track-igp-metric;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

### Preventing Addition of Ingress Routes to the inet.0 Routing Table

By configuring the **no-forwarding** statement, you can prevent ingress routes from being added to the inet.0 routing table instead of the inet.3 routing table even if you enabled the **traffic-engineering bgp-igp** statement at the **[edit protocols mpls]** or the **[edit logical-systems *logical-system-name* protocols mpls]** hierarchy level. By default, the **no-forwarding** statement is disabled.

To omit ingress routes from the inet.0 routing table, include the **no-forwarding** statement:

```
no-forwarding;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Multiple-Instance LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a service provider provider edge (PE) router to a customer carrier customer edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 or Layer 3 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Multiple Instances for Label Distribution Protocol Feature Guide*.

### Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping, include the **explicit-null** statement:

```
explicit-null;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see *Label Description* and *Label Allocation*.

### Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the lo0.0 interface (see “[Enabling and Disabling LDP](#)” on page 4435). You must also configure the LSPs over which you want LDP to operate by including the **ldp-tunneling** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from source;
      to destination;
      ldp-tunneling;
    }
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

**Related  
Documentation**

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 4428](#)

---

### Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This might require that you manually configure the RSVP metric when deploying LDP tunneling over RSVP LSPs in heterogeneous networks.

When a Juniper Networks router is linked to another vendor's router through an RSVP tunnel, and LDP tunneling is also enabled, by default the Juniper Networks router might not use the RSVP tunnel to route traffic to the LDP destinations downstream of the other vendor's egress router if the RSVP path has a metric of 1 larger than the physical OSPF path.

To ensure that LDP tunneling functions properly in heterogeneous networks, you can configure OSPF to ignore the RSVP LSP metric by including the **ignore-lsp-metrics** statement:

```
ignore-lsp-metrics;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ospf traffic-engineering shortcuts]**
- **[edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]**

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section [“Enabling LDP over RSVP-Established LSPs” on page 4501](#).

---

### Configuring the TCP MD5 Signature for LDP Sessions

You can configure an MD5 signature for an LDP TCP connection to protect against the introduction of spoofed TCP segments into LDP session connection streams.

A router using the MD5 signature option is configured with a password for each peer for which authentication is required. The password is stored encrypted.

LDP hello adjacencies can still be created even when peering interfaces are configured with different security signatures. However, the TCP session cannot be authenticated and is never established.

To configure an MD5 signature for an LDP TCP connection, include the **session** and **authentication-key** statement:

```
session address {  
  authentication-key md5-authentication-key;  
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary section for the **session** statement.



Use the **session** statement to configure the address for the remote end of the LDP session.

The **md5-authentication-key** (password) can be up to 69 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks.

You can also configure an authentication key update mechanism for the LDP routing protocol. This mechanism allows you to update authentication keys without interrupting associated routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).

To configure the authentication key update mechanism, include the **key-chain** statement at the **[edit security authentication-key-chains]** hierarchy level, and specify the **key** option to create a keychain consisting of several authentication keys.

```
[edit security authentication-key-chains]
key-chain key-chain-name {
  key key {
    secret secret-data;
    start-time yyyy-mm-dd.hh:mm:ss;
  }
}
```

To configure the authentication key update mechanism for the LDP routing protocol, include the **authentication-key-chain** statement at the **[edit protocols ldp]** hierarchy level to associate the protocol with the **[edit security authentication-key-chains]** authentication keys.

```
[edit protocols ldp]
group group-name {
  neighbor address {
    authentication-key-chain key-chain-name;
  }
}
```

For more information about the authentication key update feature, see *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*.

### Configuring LDP Session Protection

An LDP session is normally created between a pair of routers that are connected by one or more links. The routers form one hello adjacency for every link that connects them and associate all the adjacencies with the corresponding LDP session. When the last hello adjacency for an LDP session goes away, the LDP session is terminated. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished.

You can configure the Junos OS to leave the LDP session between two routers up even if there are no hello adjacencies on the links connecting the two routers by configuring the **session-protection** statement. You can optionally specify a time in seconds using the **timeout** option. The session remains up for the duration specified as long as the routers maintain IP network connectivity.

```
session-protection {
  timeout seconds;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

### Disabling SNMP Traps for LDP

---

Whenever an LDP LSP makes a transition from up to down, or down to up, the router sends an SNMP trap. However, it is possible to disable the LDP SNMP traps on a router, logical system, or routing instance.

For information about the LDP SNMP traps and the proprietary LDP MIB, see the *SNMP MIBs and Traps Reference* and *Interpreting the Enterprise-Specific LDP MIB*.

To disable SNMP traps for LDP, specify the **trap disable** option for the **log-updown** statement:

```
log-updown {  
    trap disable;  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring LDP Synchronization with the IGP on LDP Links

---

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the IGP. If synchronization between LDP and the IGP is not maintained, the LSP goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for synchronization, include the **ldp-synchronization** statement:

```
ldp-synchronization {  
    disable;  
    hold-time seconds;  
}
```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

### Configuring LDP Synchronization with the IGP on the Router

You can configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

To configure the time the LDP waits before informing the IGP that the LDP neighbor and session are operational, include the **igp-synchronization** statement and specify a time in seconds for the **holddown-interval** option:

```
igp-synchronization holddown-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

### Configuring the Label Withdrawal Timer

The label withdrawal timer delays sending a label withdrawal message for a FEC to a neighbor. When an IGP link to a neighbor fails, the label associated with the FEC has to be withdrawn from all the upstream routers if the neighbor is the next hop for the FEC. After the IGP converges and a label is received from a new next hop, the label is readvertised to all the upstream routers. This is the typical network behavior. By delaying label withdrawal by a small amount of time (for example, until the IGP converges and the router receives a new label for the FEC from the downstream next hop), the label withdrawal and sending a label mapping soon could be avoided. The **label-withdrawal-delay** statement allows you to configure this delay time. By default, the delay is 60 seconds.

If the router receives the new label before the timer runs out, the label withdrawal timer is canceled. However, if the timer runs out, the label for the FEC is withdrawn from all of the upstream routers.

By default, LDP waits for 60 seconds before withdrawing labels to avoid resignaling LSPs multiple times while the IGP is reconverging. To configure the label withdrawal delay time in seconds, include the **label-withdrawal-delay** statement:

```
label-withdrawal-delay seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

### Ignoring the LDP Subnet Check

In Junos OS Release 8.4 and later releases, an LDP source address subnet check is performed during the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. This causes an interoperability issue with some other vendors' equipment.

To disable the subnet check, include the **allow-subnet-mismatch** statement:

```
allow-subnet-mismatch;
```

This statement can be included at the following hierarchy levels:

- **[edit protocols ldp interface interface-name]**

- `[edit logical-systems logical-system-name protocols ldp interface interface-name]`

## Configuring Ethernet over MPLS (L2 Circuit)

To implement Ethernet over MPLS, you must configure a Layer 2 circuit on the provider edge (PE) switches. No special configuration is required on the customer edge (CE) switches. The provider switches require MPLS and LDP to be configured on the interfaces that will be receiving and transmitting MPLS packets.



**NOTE:** A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two PE routers. In contrast, each CCC requires a dedicated LSP

This topic describes how to configure the PE switches to support Ethernet over MPLS. You must configure interfaces and protocols on the local PE (PE1) and the remote PE (PE2). The configuration of the interfaces varies depending upon whether the Layer 2 circuit is port-based or VLAN-based.



**NOTE:** This topic refers to the local PE switch as PE1 and the remote PE switch as PE2. It also uses interface names rather than variables to help clarify the connections between the switches. The loopback addresses of the switches are configured as follows:

- PE1: 1.1.1.1
- PE2: 4.4.4.4

- [Configuring the Local PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) on page 4506](#)
- [Configuring the Remote PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) on page 4507](#)
- [Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit on page 4507](#)
- [Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit on page 4508](#)

### Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)

---

To configure the local PE switch (PE1) for a port-based layer 2 circuit (pseudo-wire):

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```

2. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch# set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
```

```
user@switch#set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

### Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)

To configure the remote PE switch (PE2) for a port-based layer 2 circuit:

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```

2. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch#set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
user@switch#set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

### Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit

To configure the local PE switch (PE1) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch#set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
```

```
user@switch#set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

### Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit

---

To configure the remote PE switch (PE2) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch#set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
user@switch#set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

**Related Documentation**

- [Understanding Ethernet-over-MPLS \(L2 Circuit\)](#)

### Configuration Statements

---

- [\[edit protocols mpls\] Hierarchy Level](#) on page 4509
- [\[edit protocols rsvp\] Hierarchy Level](#) on page 4513
- [auto-bandwidth \(MPLS Statistics\)](#) on page 4515
- [auto-bandwidth](#) on page 4516
- [adjust-interval](#) on page 4517
- [adjust-threshold](#) on page 4517

- [adjust-threshold-overflow-limit](#) on page 4518
- [adjust-threshold-underflow-limit](#) on page 4518
- [exp](#) on page 4519
- [maximum-bandwidth \(Protocols MPLS\)](#) on page 4520
- [minimum-bandwidth](#) on page 4520
- [minimum-bandwidth-adjust-interval](#) on page 4521
- [minimum-bandwidth-adjust-threshold-change](#) on page 4521
- [minimum-bandwidth-adjust-threshold-value](#) on page 4522
- [monitor-bandwidth](#) on page 4522
- [system-defaults](#) on page 4523

## [edit protocols mpls] Hierarchy Level

This topic lists the supported configuration statements at the [\[edit protocols mpls\]](#) hierarchy level on the QFX Series and on the EX4600 switch. For more information about these statements, see the *Junos OS MPLS Applications Library for Routing Devices*.



**NOTE:** The command-line interface (CLI) on QFX Series devices and on the EX4600 switch displays even the MPLS related configuration statements that are not supported. However, configuring the unsupported statements on a device will have no effect on the operation of the device.

```
protocols {
  mpls {
    admin-down;
    advertisement-hold-time seconds;
    class-of-service cos-value;
    diffserv-te {
      bandwidth-model {
        extended-mam;
        mam;
        rdm;
      }
      te-class-matrix {
        tnumber {
          priority priority;
          traffic-class {
            ctnumber priority priority;
          }
        }
      }
    }
    disable;
    exclude-srlg;
    explicit-null;
    hop-limit number;
    interface (interface-name | all) {
      disable;
```

```
}
ipv6-tunneling;
label-switched-path lsp-name {
  adaptive;
  admin-down;
  associate-backup-pe-groups;
  associate-lsp lsp-name {
    from from-ip-address;
  }
  auto-bandwidth {
    adjust-interval seconds;
    adjust-threshold percentage;
    maximum-bandwidth bps;
    minimum-bandwidth bps;
    monitor-bandwidth;
  }
  backup;
  bandwidth bps {
    ct0 bps;
    ct1 bps;
    ct2 bps;
    ct3 bps;
  }
  class-of-service cos-value;
  corouted-bidirectional;
  corouted-bidirectional-passive;
  description text;
  disable;
  exclude-srlg;
  from address;
  hop-limit number;
  install {
    destination-prefix/prefix-length <active>;
  }
  inter-domain;
  ldp-tunneling;
  lsp-attributes {
    encoding-type (ethernet | packet | pdh | sonet-sdh);
    gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 |
      pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
    signal-bandwidth type;
    switching-type (fiber | lambda | psc-1 | tdm);
  }
  metric metric;
  no-cspf;
  no-decrement-ttl;
  no-install-to-address;
  no-record;
  oam {
    lsp-ping-interval seconds;
    mpls-tp-mode seconds;
    traceoptions {
      file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
}
```



```

    }
  }
  optimize-hold-dead-delay seconds;
  optimize-timer seconds;
  p2mp lsp-name;
  policing {
    filter filter-name;
    no-auto-policing;
  }
  preference preference;
  primary path-name {
    adaptive;
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    (record | no-record);
    select (manual | unconditional);
    standby;
  }
  (record | no-record);
  retry-limit number;
  retry-timer seconds;
  revert-timer seconds;
  secondary path-name {
    adaptive;
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    (record | no-record);
    select (manual | unconditional);
    standby;
  }
  standby;
  jtemplate;
  to address;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
log-updown {
  no-trap {
    mpls-lsp-traps;
    rfc3812-traps;
  }
}

```

```
}
(syslog | no-syslog);
trap;
trap-path-down;
trap-path-up;
}
mib-mpls-show-p2mp;
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
no-record;
oam{
    lsp-ping-interval seconds;
    mpls-tp-mode seconds;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
optimize-aggressive;
optimize-hold-dead-delay;
optimize-switchover-delay;
optimize-timer;
path path-name {
    (address | hostname) <loose | strict>;
}
path-mtu {
    rsvp {
        mtu-signaling;
    }
}
preference;
record;
revert-timer;
rsvp-error-hold-time;
smart-optimize-timer;
standby;
static-label-switched-path lsp-name {
    bypass bypass-name {
        description string;
        next-hop (address | interface-name | address/interface-name);
        push out-label;
        to address;
    }
    ingress {
        class-of-service cos-value;
        description string;
        install {
            destination-prefix <active>;
        }
        metric metric;
        next-hop (address | interface-name | address/interface-name);
        no-install-to-address;
        policing {
```

```

        filter filter-name;
        no-auto-policing;
    }
    preference preference;
    push out-label;
    to address;
}
transit incoming-label {
    description string;
    next-hop (address | interface-name | address/interface-name);
    pop;
    swap out-label;
}
statistics {
    auto-bandwidth;
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    interval seconds;
}
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag;
}
traffic-engineering;
transit-lsp-association transit-association-lsp-group-name {
    from-1 address-of-associated-lsp-1;
    from-2 address-of-associated-lsp-2;
    lsp-name-1 name-of-associated-lsp-1;
    lsp-name-2 name-of-associated-lsp-2;
}
}
}

```

**Related  
Documentation**

- *Junos OS MPLS Applications Library for Routing Devices*

## [edit protocols rsvp] Hierarchy Level

This topic lists the supported configuration statements at the **[edit protocols rsvp]** hierarchy level on the QFX Series. For more information about these statements, see the *Junos OS MPLS Applications Library for Routing Devices*.

```

protocols {
    rsvp {
        disable;
        graceful-deletion-timeout seconds;
        graceful-restart {
            disable;
            helper-disable;
            maximum-helper-recovery-time seconds;
            maximum-helper-restart-time seconds;
        }
        hello-acknowledgements;
        interface interface-name {

```

```
(aggregate | no-aggregate);
authentication-key key;
bandwidth bps;
disable;
hello-interval seconds;
(reliable | no-reliable);
subscription {
    percentage;
    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
}
update-threshold percentage;
}
keep-multiplier number;
load-balance bandwidth;
no-interface-hello;
no-node-id-subobject;
no-p2mp-sublsp;
node-hello
preemption {
    (aggressive | disabled | normal);
    soft-preemption cleanup-timer seconds;
}
refresh-time seconds;
setup-protection;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-services {
    devices device-names;
}
}
}
```

**Related Documentation**

- [Junos OS MPLS Applications Library for Routing Devices](#)

---

## auto-bandwidth (MPLS Statistics)

---

<b>Syntax</b>	auto-bandwidth;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls statistics], [edit protocols mpls statistics]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Collect statistics related to automatic bandwidth.
<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Automatic Bandwidth Allocation for LSPs on page 4482</a></li><li>• <a href="#">Configuring MPLS to Gather Statistics on page 4481</a></li><li>• <i>statistics</i></li></ul>

## auto-bandwidth

---

<b>Syntax</b>	<pre>auto-bandwidth {   adjust-interval <i>seconds</i>;   adjust-threshold <i>percent</i>;   adjust-threshold-activate-bandwidth <i>bps</i>   adjust-threshold-overflow-limit <i>number</i>;   adjust-threshold-underflow-limit <i>number</i>;   maximum-bandwidth <i>bps</i>;   minimum-bandwidth <i>bps</i>;   minimum-bandwidth-adjust-interval   minimum-bandwidth-adjust-threshold-change   minimum-bandwidth-adjust-threshold-value   monitor-bandwidth; }</pre>
<b>Hierarchy Level</b>	[edit protocols mpls label-switched-path <i>lsp-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Automatic Bandwidth Allocation for LSPs on page 4482</a></li><li>• <a href="#">request mpls lsp adjust-autobandwidth on page 4598</a></li></ul>

## adjust-interval

<b>Syntax</b>	<code>adjust-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the bandwidth reallocation interval.
<b>Options</b>	<b><i>seconds</i></b> —Bandwidth reallocation interval, in seconds. <b>Range:</b> 300 through 315,360,000 seconds <b>Default:</b> 86,400 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Automatic Bandwidth Allocation Interval on page 4484</a></li> </ul>

## adjust-threshold

<b>Syntax</b>	<code>adjust-threshold <i>percent</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.
<b>Options</b>	<b><i>percent</i></b> —Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP's bandwidth is adjusted to the current bandwidth demand.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Automatic Bandwidth Adjustment Threshold on page 4485</a></li> </ul>

## adjust-threshold-overflow-limit

---

<b>Syntax</b>	adjust-threshold-overflow-limit <i>number</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.
<b>Options</b>	<i>number</i> —Number of consecutive bandwidth overflow samples. <b>Range:</b> 1 through 65,535 <b>Default:</b> This feature is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 4485</a></li></ul>

## adjust-threshold-underflow-limit

---

<b>Syntax</b>	adjust-threshold-underflow-limit <i>number</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.
<b>Options</b>	<i>number</i> —Number of consecutive bandwidth underflow samples. <b>Range:</b> 1 through 65,535 <b>Default:</b> This feature is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 4485</a></li></ul>



## exp

Syntax	<pre>exp classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level {             code-points [ aliases ] [ bit-patterns ];         }     } }</pre>
Rewrite Rule Configuration	<pre>exp rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {         loss-priority level {             code-point [ aliases ] [ bit-patterns ];         }     } }</pre>
Global Classifier Association with Interfaces	exp classifier-name;
Hierarchy Level	<pre>[edit class-of-service classifiers], [edit class-of-service rewrite-rules] [edit class-of-service system-defaults classifiers]</pre>
Release Information	Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Define the EXP code point mapping that is applied to MPLS packets. EXP classifiers are not applied to any traffic except MPLS traffic. EXP classifiers are applied only to interfaces that are configured as <b>family mpls</b> (for example, <b>set interfaces xe-0/0/35 unit 0 family mpls</b>.)</p> <p>You can configure as many EXP classifiers as you want. However, the switch uses only one EXP classifier as a global MPLS classifier on all interfaces. You specify the global EXP classifier in the <b>[edit class-of-service system-defaults]</b> hierarchy.</p>
Options	<b>classifier-name</b> —Name of the EXP classifier.
Required Privilege Level	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Global MPLS EXP Classifier on page 4479</a></li> <li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480</a></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419</a></li> <li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li> </ul>

## maximum-bandwidth (Protocols MPLS)

---

<b>Syntax</b>	maximum-bandwidth <i>bps</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the maximum amount of bandwidth in bits per second (bps).
<b>Options</b>	<i>bps</i> —Maximum amount of bandwidth.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 4484</a></li></ul>

## minimum-bandwidth

---

<b>Syntax</b>	minimum-bandwidth <i>bps</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.
<b>Options</b>	<i>bps</i> —Minimum bandwidth for the LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 4484</a></li></ul>

## minimum-bandwidth-adjust-interval

<b>Syntax</b>	<code>minimum-bandwidth-adjust-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the duration (in seconds) for which minimum bandwidth is frozen.
<b>Options</b>	<b><i>seconds</i></b> —Minimum bandwidth reallocation interval, in seconds. <b>Range:</b> 300 through 31,536,000 seconds.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 4484</a></li> </ul>

## minimum-bandwidth-adjust-threshold-change

<b>Syntax</b>	<code>minimum-bandwidth-adjust-threshold-change <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the percentage change in maximum average bandwidth to freeze the minimum bandwidth.
<b>Options</b>	<b><i>percentage</i></b> —Percentage change in maximum average bandwidth. <b>Range:</b> Range: 0 through 100 percent.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 4484</a></li> </ul>

## minimum-bandwidth-adjust-threshold-value

---

<b>Syntax</b>	minimum-bandwidth-adjust-threshold-value <i>bps</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the value in bits per second (bps) to freeze the minimum bandwidth if the maximum average bandwidth falls below this value.
<b>Options</b>	<i>bps</i> —Threshold value for minimum bandwidth if the maximum average bandwidth falls below the specified value.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 4484</a></li></ul>

## monitor-bandwidth

---

<b>Syntax</b>	monitor-bandwidth;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Passive Bandwidth Utilization Monitoring on page 4487</a></li></ul>

## system-defaults

---

<b>Syntax</b>	<pre>system-defaults {   classifiers exp classifier-name; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Configure the global EXP classifier used on all interfaces to classify MPLS traffic.</p> <p>Although you can configure as many EXP classifiers as you want, the switch uses only one EXP classifier as a global MPLS classifier on all interfaces. All switch interfaces use the EXP classifier specified as the system default to classify MPLS traffic.</p>
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Global MPLS EXP Classifier on page 4479</a></li> <li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480</a></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419</a></li> <li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li> </ul>

## LDP Configuration Statements for QFX5100

---

- [allow-subnet-mismatch on page 4525](#)
- [authentication-algorithm on page 4526](#)
- [authentication-key \(Protocols LDP\) on page 4527](#)
- [authentication-key-chain \(Protocols LDP\) on page 4528](#)
- [deaggregate on page 4529](#)
- [disable \(Protocols LDP\) on page 4530](#)
- [dod-request-policy on page 4531](#)
- [downstream-on-demand on page 4531](#)
- [egress-policy on page 4532](#)
- [explicit-null \(Protocols LDP\) on page 4532](#)
- [export \(Protocols LDP\) on page 4533](#)
- [fec on page 4534](#)
- [graceful-restart \(Protocols LDP\) on page 4535](#)
- [hello-interval \(Protocols LDP\) on page 4536](#)

- [helper-disable \(LDP\) on page 4537](#)
- [hold-time \(Protocols LDP\) on page 4538](#)
- [ignore-lsp-metrics on page 4539](#)
- [igp-synchronization on page 4539](#)
- [import \(Protocols LDP\) on page 4540](#)
- [interface \(Protocols LDP\) on page 4541](#)
- [keepalive-interval on page 4542](#)
- [keepalive-timeout on page 4543](#)
- [l2-smart-policy on page 4543](#)
- [label-withdrawal-delay on page 4544](#)
- [ldp on page 4545](#)
- [ldp-synchronization on page 4548](#)
- [log-updown \(Protocols LDP\) on page 4549](#)
- [maximum-neighbor-recovery-time on page 4550](#)
- [no-forwarding on page 4551](#)
- [policing \(Protocols LDP\) on page 4552](#)
- [preference \(Protocols LDP\) on page 4553](#)
- [reconnect-time on page 4554](#)
- [recovery-time on page 4554](#)
- [session \(ldp\) on page 4555](#)
- [session-protection on page 4556](#)
- [strict-targeted-hellos on page 4556](#)
- [targeted-hello on page 4557](#)
- [traceoptions \(Protocols LDP\) on page 4558](#)
- [track-igp-metric on page 4560](#)
- [traffic-statistics \(Protocols LDP\) on page 4561](#)
- [transport-address on page 4563](#)

---

## allow-subnet-mismatch

---

<b>Syntax</b>	allow-subnet-mismatch;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ], [edit protocols ldp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.
<b>Default</b>	The source address in the LDP link hello packet is matched against the interface address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Ignoring the LDP Subnet Check on page 4505</a></li></ul>

## authentication-algorithm

**Syntax** authentication-algorithm *algorithm*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols bgp],  
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],  
 [edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],  
 [edit logical-systems *logical-system-name* protocols ldp session *session-address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *session-address*],  
 [edit logical-systems *logical-system-name* routing-options **bmp**],  
 [edit logical-systems *logical-system-name* routing-options bmp **station** *station-name*],  
 [edit protocols bgp],  
 [edit protocols bgp **group** *group-name*],  
 [edit protocols bgp group *group-name* **neighbor** *address*],  
 [edit protocols ldp session *session-address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],  
 [edit routing-instances *routing-instance-name* protocols ldp session *session-address*],  
 [edit routing-options **bmp**],  
 [edit routing-options bmp **station** *station-name*]

**Release Information** Statement introduced in Junos OS Release 7.6.  
 Statement introduced for BGP in Junos OS Release 8.0.  
 Statement introduced in Junos OS Release 9.0 for EX Series switches.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.  
 Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.  
 Statement introduced for BMP in Junos OS Release 13.3.

**Description** Configure an authentication algorithm type.

**Options** *algorithm*—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

**Default:** hmac-sha-1-96



**NOTE:** The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.



**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Route Authentication for BGP on page 3572](#)
- [Configuring BGP Monitoring Protocol Version 3 on page 3307](#)

## authentication-key (Protocols LDP)

**Syntax** authentication-key *md5-authentication-key*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp session *address*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *address*],  
[edit protocols ldp session *address*],  
[edit routing-instances *routing-instance-name* protocols ldp session *address*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the TCP MD5 Signature for LDP Sessions on page 4502](#)

## authentication-key-chain (Protocols LDP)

---

<b>Syntax</b>	authentication-key-chain <i>key-chain</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>name</i> protocols ldp session <i>address</i> ], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols ldp session <i>address</i> ], [edit protocols ldp session <i>address</i> ], [edit routing-instances <i>instance-name</i> protocols ldp session <i>address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the <b>0.0.0.0/allow</b> statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
<b>Options</b>	<i>key-chain</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li><li>• <a href="#">Configuring Miscellaneous LDP Properties on page 4500</a></li></ul>

## deaggregate

---

<b>Syntax</b>	deaggregate   no-deaggregate;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control forwarding equivalence class (FEC) deaggregation on the router. The use of the <b>deaggregate</b> statement in LDP is a standard practice that we recommend for LDP deployments.
<b>Default</b>	Deaggregation is disabled on the router.
<b>Options</b>	<b>deaggregate</b> —Deaggregate FECs. <b>no-deaggregate</b> —Aggregate FECs.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring FEC Deaggregation</i></li> </ul>

## disable (Protocols LDP)

---

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit protocols ldp graceful-restart], [edit protocols ldp interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.
<b>Default</b>	LDP is enabled on interfaces configured with the LDP <b>interface</b> statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the <b>[edit routing-options]</b> hierarchy level.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling and Disabling LDP on page 4435</a></li><li>• <a href="#">Configuring LDP Graceful Restart on page 4495</a></li></ul>

## dod-request-policy

---

<b>Syntax</b>	<code>dod-request-policy <i>dod-request-policy-name</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy.
<b>Options</b>	<i>dod-request-policy-name</i> —Specify the name of the downstream on demand request policy.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring LDP Downstream on Demand on page 4462</a></li> </ul>

## downstream-on-demand

---

<b>Syntax</b>	<code>downstream-on-demand;</code>
<b>Hierarchy Level</b>	[edit logical systems <i>logical-system-name</i> protocols ldp session <i>session-address</i> ], [edit protocols ldp session <i>session-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring LDP Downstream on Demand on page 4462</a></li> </ul>

## egress-policy

---

<b>Syntax</b>	<code>egress-policy [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control the prefixes advertised into LDP.
<b>Default</b>	Only the loopback address is advertised.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Prefixes Advertised into LDP from the Routing Table on page 4498</a></li></ul>

## explicit-null (Protocols LDP)

---

<b>Syntax</b>	<code>explicit-null;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Advertise label 0 to the egress router of a label-switched path (LSP).
<b>Default</b>	If you do not include the <b>explicit-null</b> statement in the MPLS configuration, label 3 (implicit null) is advertised.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 4501</a></li></ul>

---

## export (Protocols LDP)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Filtering Outbound LDP Label Bindings on page 4437</a></li></ul>

## fec

<b>Syntax</b>	<pre> fec <i>fec-address</i> {     bfd-liveness-detection {         detection-time threshold <i>milliseconds</i>;         ecmp;         failure-action {             remove-nexthop;             remove-route;         }         holddown-interval <i>milliseconds</i>;         ingress-policy <i>ingress-policy-name</i>;         minimum-interval <i>milliseconds</i>;         minimum-receive-interval <i>milliseconds</i>;         minimum-transmit-interval <i>milliseconds</i>;         multiplier <i>detection-time-multiplier</i>;         no-adaptation;         transmit-interval {             minimum-interval <i>milliseconds</i>;             threshold <i>milliseconds</i>;         }         version (0   1   automatic);     }     no-bfd-liveness-detection;     periodic-traceroute {         disable;         exp <i>exp-value</i>;         fanout <i>fanout-value</i>;         frequency <i>minutes</i>;         paths <i>number-of-paths</i>;         retries <i>retry-attempts</i>;         source <i>address</i>;         ttl <i>ttl-value</i>;         wait <i>seconds</i>;     } } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-systems-name</i> protocols ldp oam], [edit protocols ldp oam]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.2 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).
<b>Options</b>	<p><b><i>fec-address</i></b>—Specify the FEC address.</p> <p>The other statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.



Related Documentation • [Configuring BFD for LDP LSPs](#)

## graceful-restart (Protocols LDP)

Syntax	<pre>graceful-restart {   disable;   helper-disable;   maximum-neighbor-recovery-time value;   reconnect-time seconds;   recovery-time value; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],          [edit protocols ldp],          [edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Configure LDP graceful restart on the LDP master protocol instance or for a specific routing instance.</p>



**NOTE:** When you alter the graceful restart configuration at either the [edit routing-options graceful-restart] or [edit protocols ldp graceful-restart] hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
--------------------------	---------------------------------------------------------------------------------------------------------------------

Related Documentation • [Configuring LDP Graceful Restart on page 4495](#)

## hello-interval (Protocols LDP)

---

<b>Syntax</b>	<code>hello-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp targeted-hello],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit protocols ldp targeted-hello],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for LDP targeted hellos added in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the <b>hello-interval</b> statement.
<b>Options</b>	<b><i>seconds</i></b> —Length of time between transmission of hello packets. <b>Range:</b> 1 through 65,535 seconds <b>Default:</b> 5 seconds for link hello messages, 15 seconds for targeted hello messages
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the LDP Timer for Hello Messages on page 4492</a></li></ul>

---

## helper-disable (LDP)

---

<b>Syntax</b>	helper-disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.
<b>Default</b>	Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Graceful Restart on page 4495</a></li></ul>

## hold-time (Protocols LDP)

---

<b>Syntax</b>	<code>hold-time seconds;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp targeted-hello],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit protocols ldp targeted-hello],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for LDP targeted hellos added in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the <b>hold-time</b> statement.
<b>Options</b>	<b>seconds</b> —Hold-time value. <b>Range:</b> 1 through 65,535 seconds <b>Default:</b> 15 seconds for link hello messages, 45 seconds for targeted hello messages
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 4493</a></li></ul>

## ignore-lsp-metrics

<b>Syntax</b>	ignore-lsp-metrics;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering shortcuts]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Cause OSPF to ignore the RSVP LSP metric.  Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 4502</a></li> </ul>

## igp-synchronization

<b>Syntax</b>	igp-synchronization holddown-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.
<b>Options</b>	<b>holddown-interval <i>seconds</i></b> —Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. <b>Default:</b> 10 seconds <b>Range:</b> 10 through 60 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Synchronization with the IGP on the Router on page 4505</a></li> </ul>

## import (Protocols LDP)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Filtering Inbound LDP Label Bindings on page 4435</a></li></ul>

## interface (Protocols LDP)

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     disable;     hello-interval <i>seconds</i>;     hold-time <i>seconds</i>;     transport-address (interface   loopback); }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Enable LDP on one or more router interfaces.
<b>Default</b>	LDP is disabled on all interfaces.
<b>Options</b>	<p><i>interface-name</i>—Name of an interface. To configure all interfaces, specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling LDP on page 4435</a></li> </ul>

## keepalive-interval

---

<b>Syntax</b>	<code>keepalive-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the keepalive interval value.
<b>Options</b>	<b><i>seconds</i></b> —Keepalive value. <b>Range:</b> 1 through 65,535 <b>Default:</b> 10 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Interval for LDP Keepalive Messages on page 4494</a></li></ul>



## keepalive-timeout

<b>Syntax</b>	<code>keepalive-timeout seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.
<b>Options</b>	<b>seconds</b> —Keepalive timeout value. <b>Range:</b> 1 through 65,535 <b>Default:</b> 30 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the LDP Keepalive Timeout on page 4495</a></li> </ul>

## l2-smart-policy

<b>Syntax</b>	<code>l2-smart-policy;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP IPv4 FEC Filtering</a></li> </ul>

## label-withdrawal-delay

---

<b>Syntax</b>	label-withdrawal-delay <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Delay the withdrawal of labels to reduce router workload during IGP convergence.
<b>Options</b>	<b>seconds</b> —Configure the number of seconds to wait before withdrawing labels for the LDP LSPs. <b>Default:</b> 60 seconds <b>Range:</b> 0 through 300 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Label Withdrawal Timer on page 4505</a></li></ul>

## ldp

```
Syntax  ldp {
    (deaggregate | no-deaggregate);
    egress-policy [ policy-names ];
    explicit-null;
    export [ policy-names ];
    graceful-restart {
        disable;
        helper-disable;
        maximum-neighbor-recovery-time seconds;
        reconnect-time seconds;
        recovery-time seconds;
    }
    import [ policy-names ];
    interface (interface-name | all) {
        disable;
        hello-interval seconds;
        hold-time seconds;
        transport-address (interface | router-id);
    }
    keepalive-interval seconds;
    keepalive-timeout seconds;
    log-updown {
        trap disable;
    }
    no-forwarding;
    oam {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        }
    }
    fec fec-address {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
        }
    }
}
```

```
    holddown-interval milliseconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0 | 1 | automatic);
}
no-bfd-liveness-detection;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}
ingress-policy ingress-policy-name;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}
p2mp;
policing {
    fec fec-address {
        ingress-traffic filter-name;
        transit-traffic filter-name;
    }
}
preference preference;
session address {
    authentication-algorithm algorithm;
    authentication-key authentication-key;
    authentication-key-chain key-chain-name;
}
strict-targeted-hellos;
traceoptions {
    file filename <files number <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
```

```

}
track-igp-metric;
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval interval;
  no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable LDP routing on the router or switch.  You must include the <b>ldp</b> statement in the configuration to enable LDP on the router or switch.
<b>Default</b>	LDP is disabled on the router.
<b>Options</b>	The other statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Minimum LDP Configuration on page 4434</a></li> <li>• <a href="#">Enabling and Disabling LDP on page 4435</a></li> </ul>

## ldp-synchronization

---

<b>Syntax</b>	<pre>ldp-synchronization {     disable;     hold-time seconds; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ospf interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i> ], [edit protocols ospf interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.
<b>Options</b>	The other statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Synchronization with the IGP on LDP Links on page 4504</a></li></ul>

---

## log-updown (Protocols LDP)

---

<b>Syntax</b>	log-updown { trap disable; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Disable LDP traps on the router, logical system, or routing instance.
<b>Options</b>	<b>trap disable</b> —Disable LDP traps. <b>Default:</b> LDP traps are enabled on the router.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Disabling SNMP Traps for LDP on page 4504</a></li></ul>

## maximum-neighbor-recovery-time

---

<b>Syntax</b>	<code>maximum-neighbor-recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement changed from <b>maximum-recovery-time</b> to <b>maximum-neighbor-recovery-time</b> in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
<b>Options</b>	<b>seconds</b> —Configure the maximum recovery time, in seconds. <b>Range:</b> 120 through 1800 seconds <b>Default:</b> 140 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 4497</a></li><li>• <a href="#">Configuring Graceful Restart Options for LDP</a></li><li>• <a href="#">no-strict-lsa-checking on page 2303</a></li><li>• <a href="#">recovery-time</a></li></ul>



## no-forwarding

---

<b>Syntax</b>	no-forwarding;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Do not add ingress routes to the inet.0 routing table even if <b>traffic-engineering bgp-igp</b> (configured at the [edit protocols mpls] hierarchy level) is enabled.
<b>Default</b>	The <b>no-forwarding</b> statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when <b>traffic-engineering bgp-igp</b> is enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 4500</a></li> <li>• <a href="#">Configuring Virtual-Router Routing Instances in VPNs</a></li> </ul>

## policing (Protocols LDP)

---

<b>Syntax</b>	<pre>policing {     fec <i>fec-address</i> {         ingress-traffic <i>filter-name</i>;         transit-traffic <i>filter-name</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable policing of forwarding equivalence classes (FECs) for LDP.
<b>Options</b>	<p><b>fec <i>fec-address</i></b>—Specify the address for the FEC.</p> <p><b>ingress-traffic <i>filter-name</i></b>—Specify the name of the filter for policing ingress FEC traffic.</p> <p><b>transit-traffic <i>filter-name</i></b>—Specify the name of the filter for policing transit FEC traffic.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Policers for LDP FECs</i></li></ul>

## preference (Protocols LDP)

---

<b>Syntax</b>	<code>preference <i>preference</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit protocols ldp interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the route preference level for LDP routes.
<b>Options</b>	<p><i>preference</i>—Preferred value.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 9</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Route Preferences on page 4495</a></li> </ul>

## reconnect-time

---

<b>Syntax</b>	<code>reconnect-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
<b>Options</b>	<b>seconds</b> —Time required for reconnection. <b>Range:</b> 30 through 300 <b>Default:</b> 60 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Graceful Restart on page 4495</a> on <i>LDP Feature Guide for Routing Devices</i></li><li>• <i>Configuring Graceful Restart Options for LDP</i></li></ul>

## recovery-time

---

<b>Syntax</b>	<code>recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the amount of time a router waits for LDP to restart gracefully.
<b>Options</b>	<b>seconds</b> —Configure the recovery time, in seconds. <b>Range:</b> 120 through 1800 seconds <b>Default:</b> 140 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 4497</a></li></ul>

## session (ldp)

---

<b>Syntax</b>	<pre> session address {   authentication-algorithm <i>algorithm</i>;   authentication-key <i>authentication-key</i>;   authentication-key-chain <i>key-chain-name</i>; } </pre>
<b>Hierarchy Level</b>	<pre> [edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>authentication-algorithm</b> statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Specify the address for the remote end of the LDP session.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the TCP MD5 Signature for LDP Sessions on page 4502</a></li> </ul>

## session-protection

---

<b>Syntax</b>	<code>session-protection {     timeout <i>seconds</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Description</b>	Configure when an LDP session is torn down and resigaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.
<b>Options</b>	<b>timeout <i>seconds</i></b> —Time in seconds before the LDP session is torn down and resigaled. <b>Range:</b> 1 through 65,535 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring LDP Session Protection on page 4503</a></li></ul>

## strict-targeted-hellos

---

<b>Syntax</b>	<code>strict-targeted-hellos;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling Strict Targeted Hello Messages for LDP on page 4435</a></li></ul>

## targeted-hello

---

<b>Syntax</b>	targeted-hello { hello-interval <i>seconds</i> ; hold-time <i>seconds</i> ; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the LDP timer and LDP hold time for targeted hellos.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the LDP Timer for Hello Messages on page 4492</a></li> <li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 4493</a></li> </ul>

## traceoptions (Protocols LDP)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>match-on address</b> option for the <b>filter</b> flag modifier added in Junos OS Release 10.4.</p> <p><b>nsr-synchronization</b> and <b>p2mp-nsr-synchronization</b> operations for <b>flag</b> statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Specify LDP protocol-level trace options.
<b>Default</b>	The default LDP protocol-level trace options are inherited from the routing protocols <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>ldp-log</b>. We recommend that you place LDP tracing output in the file <b>ldp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"><li>• <b>address</b>—Operation of address and address withdrawal messages</li><li>• <b>binding</b>—Label-binding operations</li><li>• <b>error</b>—Error conditions</li><li>• <b>event</b>—Protocol events</li></ul>



- **initialization**—Operation of initialization messages
- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **nsr-synchronization**—Nonstop active routing synchronization events
- **p2mp-nsr-synchronization**—Point-to-multipoint nonstop active routing synchronization events
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
  - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
    - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
    - **fec**—Filter based on the FEC associated with the traced object.
    - **policy policy-name**—Specify the filter policy.
  - **receive**—Packets being received.
  - **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent all users from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable any user to read the log file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Tracing LDP Protocol Traffic on page 4442](#)
- *Network Management Administration Guide for Routing Devices*

---

## track-igp-metric

---

**Syntax** track-igp-metric;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],  
[edit protocols ldp],  
[edit routing-instances *routing-instance-name* protocols ldp]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring LDP to Use the IGP Route Metric on page 4500](#)

## traffic-statistics (Protocols LDP)

<b>Syntax</b>	<pre>traffic-statistics {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     interval <i>seconds</i>;     no-penultimate-hop; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>ldp</i>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>],          [edit protocols <i>ldp</i>],          [edit routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the LDP statistics operation.          Enclose the name within quotation marks. All files are placed in the directory <i>/var/log</i>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of LDP statistics files. When a statistics file named <i>ldp-stat</i> reaches its maximum size, it is renamed <i>ldp-stat.0</i>, then <i>ldp-stat.1</i>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.</p> <p><b>Range:</b> 2 through 1000  <b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must include the <b>size</b> statement to specify the maximum file size.</p> <p><b>interval <i>seconds</i></b>—(Optional) Specify the interval at which the statistics are polled and written to the file.  <b>Default:</b> 300 seconds (5 minutes)</p> <p><b>no-penultimate-hop</b>—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p><b>no-world-readable</b>—(Optional) Prevent all users from reading the log file.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <i>ldp-stat</i> reaches this size, it is renamed <i>ldp-stat.0</i>. When <i>ldp-stat</i> again reaches this size, <i>ldp-stat.0</i> is renamed <i>ldp-stat.1</i> and <i>ldp-stat</i> is renamed <i>ldp-stat.0</i>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.</p> <p><b>Syntax:</b> <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB  <b>Range:</b> 10 KB through the maximum file size supported on your system</p>

**Default:** 1 MB

If you specify a maximum file size, you also must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable log file access for all users.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Collecting LDP Statistics on page 4440</a></li></ul>
------------------------------	----------------------------------------------------------------------------------------------------------

## transport-address

<b>Syntax</b>	<code>transport-address (interface   router-id);</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp],          [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],          [edit protocols ldp],          [edit protocols ldp interface <i>interface-name</i>],          [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.</p>
<b>Default</b>	<b>router-id</b>
<b>Options</b>	<p><b>interface</b>—The first IP address on the interface is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the <b>interface</b> option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the <b>router-id</b> option.</p> <p><b>router-id</b>—The router identifier is used as the transport address. Unless otherwise configured, the router identifier is the loopback address.</p>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Specifying the Transport Address Used by LDP on page 4439</a></li> </ul>



# Administration

- [Routine Monitoring on page 4565](#)
- [Operational Mode Commands on page 4567](#)

## Routine Monitoring

---

- [Verifying That MPLS Is Working Correctly on page 4565](#)

### Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

1. [Verifying the Physical Layer on the Switches on page 4565](#)
2. [Verifying the Routing Protocol on page 4566](#)
3. [Verifying the Core Interfaces Being Used for the MPLS Traffic on page 4566](#)
4. [Verifying RSVP on page 4566](#)

### Verifying the Physical Layer on the Switches

---

**Purpose** Verify that the interfaces are up. Perform this verification task on each of the switches.

**Action** user@switch> **show interfaces xe-\* terse**

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up			
xe-0/0/1.0	up	up			
xe-0/0/2.0	up	up			
xe-0/0/3.0	up	up	inet	2.2.2.1/16	
xe-0/0/4.0	up	up			
xe-0/0/5.0	up	up	inet mpls	10.1.5.1/24	
xe-0/0/6.0	up	up	inet mpls	10.1.6.1/24	

**Meaning** The **show interfaces terse** command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and

xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

### Verifying the Routing Protocol

---

**Purpose** Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

**Action** user@switch> **show ospf neighbor**

Address	Interface	State	ID	Pri	Dead
127.1.1.1	xe-0/0/5	Full	10.10.10.10	128	39

**Meaning** The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the [Junos OS Routing Protocols and Policies Command Reference](#).

### Verifying the Core Interfaces Being Used for the MPLS Traffic

---

**Purpose** Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

**Action** user@switch> **show mpls interface**

Interface	State	Administrative groups
ge-0/0/5	Up	<none>
ge-0/0/6	Up	<none>

**Meaning** The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

### Verifying RSVP

---

**Purpose** Verify the state of the RSVP session. You should perform this verification task on each of the switches.



```
user@switch> show ldp session
```

```
Ingress RSVP: 1 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
127.1.1.3   127.1.1.1   Up      0  1 FF      -    300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
127.1.1.1   127.1.1.3   Up      0  1 FF  299968    -  lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** This output confirms that the RSVP sessions are up.

**Related Documentation**

- [Configuring MPLS on Provider Edge Switches on page 4468](#)
- [Configuring MPLS on Provider Switches on page 4471](#)

## Operational Mode Commands

- [clear ldp neighbor](#)
- [clear ldp session](#)
- [clear ldp statistics](#)
- [clear mpls lsp](#)
- [clear rsvp session](#)
- [clear rsvp statistics](#)
- [monitor label-switched-path](#)
- [ping mpls bgp](#)
- [ping mpls l2circuit](#)
- [ping mpls l3vpn](#)
- [ping mpls ldp](#)
- [ping mpls lsp-end-point](#)
- [ping mpls rsvp](#)
- [request mpls lsp adjust-autobandwidth](#)
- [show ldp database](#)
- [show ldp fec-filters](#)
- [show ldp interface](#)
- [show ldp neighbor](#)
- [show ldp path](#)
- [show ldp route](#)

- `show ldp session`
- `show ldp statistics`
- `show ldp traffic-statistics`
- `show security keychain`
- `show link-management`
- `show link-management peer`
- `show link-management routing`
- `show link-management statistics`
- `show link-management te-link`
- `show mpls call-admission-control`
- `show mpls cspf`
- `show mpls diffserv-te`
- `show route forwarding-table`
- `show mpls interface`
- `show mpls lsp`
- `show mpls lsp autobandwidth`
- `show mpls path`
- `show mpls static-lsp`
- `show rsvp interface`
- `show rsvp neighbor`
- `show rsvp session`
- `show rsvp statistics`
- `show rsvp version`
- `show ted database`
- `show ted link`
- `show ted protocol`
- `traceroute mpls ldp`
- `traceroute mpls rsvp`

## clear ldp neighbor

---

<b>Syntax</b>	clear ldp neighbor <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )> < <i>neighbor</i> >
<b>Description</b>	Tear down Label Distribution Protocol (LDP) neighbor connections.
<b>Options</b>	<p><b>none</b>—Tear down connections with all LDP neighbors for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>neighbor</i></b>—(Optional) Clear an LDP session for the specified neighbor (IP address) only.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ldp neighbor on page 4611</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ldp neighbor on page 4569</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ldp neighbor

```
user@host> clear ldp neighbor
```

## clear ldp session

---

<b>Syntax</b>	<code>clear ldp session</code> <code>&lt;destination&gt;</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Clear Label Distribution Protocol (LDP) sessions.
<b>Options</b>	<b>none</b> —Clear LDP sessions for all destinations for all routing instances.  <b><i>destination</i></b> —(Optional) Clear an LDP session for the specified destination (IP address).  <b><i>instance instance-name</i></b> —(Optional) Clear the LDP session for the specified routing instance only.  <b><i>logical-system (all   logical-system-name)</i></b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ldp session on page 4619</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear ldp session on page 4570</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ldp session

```
user@host> clear ldp session
```

## clear ldp statistics

---

<b>Syntax</b>	clear ldp statistics <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set all Label Distribution Protocol (LDP) statistics to zero.
<b>Options</b>	<p><b>none</b>—Set all LDP statistics to zero for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ldp statistics on page 4625</a></li> <li>• <a href="#">show ldp traffic-statistics on page 4629</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ldp statistics on page 4571</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ldp statistics

```
user@host> clear ldp statistics
```

## clear mpls lsp

---

**List of Syntax**    [Syntax on page 4572](#)  
                          [Syntax \(EX and QFX Series Switches\) on page 4572](#)

**Syntax**    clear mpls lsp  
              <autobandwidth>  
              <logical-system (all | *logical-system-name*)>  
              <name *name*>  
              <optimize | optimize-aggressive>  
              <path *regular-expression*>  
              <statistics>

**Syntax (EX and QFX Series Switches)**    clear mpls lsp  
                                                  <autobandwidth>  
                                                  <name *name*>  
                                                  <optimize | optimize-aggressive>  
                                                  <path *regular-expression*>  
                                                  <statistics>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              Command introduced in Junos OS Release 9.5 for EX Series switches.  
                              Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

**Description**    Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.



**CAUTION:** This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.

---

**Options**    **none**—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

**autobandwidth**—(Optional) Clear LSP autobandwidth counters.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**name *name***—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.

**optimize | optimize-aggressive**—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

**path *regular-expression***—(Optional) Clear the specific LSP path matching the specified regular expression.

**statistics**—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (**name** and **path** options) on transit routers.

**Required Privilege Level**

clear

**Related Documentation**

- [show mpls lsp on page 4665](#)
- [show rsvp session on page 4698](#)

**List of Sample Output** [clear mpls lsp on page 4573](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear mpls lsp

```
user@host> clear mpls lsp
```

## clear rsvp session

---

<b>List of Syntax</b>	<a href="#">Syntax on page 4574</a> <a href="#">Syntax (EX and QFX Series Switches) on page 4574</a>
<b>Syntax</b>	<pre>clear rsvp session &lt;connection-destination address&gt; &lt;connection-source address&gt; &lt;gracefully&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;lsp-id identifier&gt; &lt;name name&gt; &lt;optimize-fast-reroute&gt; &lt;tunnel-id identifier&gt;</pre>
<b>Syntax (EX and QFX Series Switches)</b>	<pre>clear rsvp session &lt;connection-destination address&gt; &lt;connection-source address&gt; &lt;gracefully&gt; &lt;lsp-id identifier&gt; &lt;name name&gt; &lt;optimize-fast-reroute&gt; &lt;tunnel-id identifier&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Reset and restart Resource Reservation Protocol (RSVP) sessions.
<b>Options</b>	<p><b>none</b>—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.</p> <p><b>connection-source address</b>—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p><b>connection-destination address</b>—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p><b>gracefully</b>—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>lsp-id identifier</b>—(Optional) LSP identifier (source port) for the RSVP sender template.</p> <p><b>name name</b>—(Optional) Reset and restart the specified RSVP session.</p> <p><b>optimize-fast-reroute</b>—(Optional) Begin fast reroute optimization.</p>



**tunnel-id** *identifier*—(Optional) Tunnel identifier (destination port) for the RSVP session.

**Required Privilege Level**

clear

**Related Documentation**

- [clear mpls lsp on page 4572](#)
- [show rsvp session on page 4698](#)

**List of Sample Output** [clear rsvp session on page 4575](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear rsvp session](#)

```
user@host> clear rsvp session
```

## clear rsvp statistics

---

<b>List of Syntax</b>	<a href="#">Syntax on page 4576</a> <a href="#">Syntax (EX Series Switches) on page 4576</a>
<b>Syntax</b>	clear rsvp statistics <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	clear rsvp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Clear Resource Reservation Protocol (RSVP) packet and error statistics.
<b>Options</b>	<b>none</b> —Clear RSVP packet and error statistics.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show rsvp statistics on page 4707</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear rsvp statistics on page 4576</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear rsvp statistics

```
user@host> clear rsvp statistics
```

## monitor label-switched-path

**Syntax** `monitor label-switched-path lsp-name`  
`<logical-system (logical-system-name)>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Logical system support introduced in Junos OS Release 9.4.  
 Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

**Description** Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

**Options** `logical-system ( logical-system-name )`—(Optional) Perform this operation on all logical systems or on a particular logical system.

*lsp-name*—Name of the LSP.

**Additional Information** You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the `c` key. To control the output of the **monitor label-switched-path** command while it is running, use the keys listed in [Table 350 on page 4577](#). The keys are not case-sensitive.

**Table 350: Output Control Keys for the monitor label-switched-path Command**

Key	Action
c	Clears the screen and refreshes the display for this LSP.
f	Freezes the display, preventing new information from being displayed.
l	Monitors a different LSP. After you type l, you can type the new LSP name.
n	Displays information about the next LSP (whose name is alphabetically higher than the current LSP name) configured on the router.
p	Goes to the previous LSP (whose name is alphabetically lower than the current LSP name) configured on the router.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws, or restarts, the data display for this LSP.

**Required Privilege Level** trace

**List of Sample Output** [monitor label-switched-path on page 4579](#)

**Output Fields** [Table 351 on page 4578](#) describes the output fields for the **monitor label-switched-path** command. Output fields are listed in the approximate order in which they appear.

**Table 351: monitor label-switched-path Output Fields**

Field Name	Field Description
(1)	Displays the following information: <ul style="list-style-type: none"> <li>• <b>hostname</b>—Name of the router.</li> <li>• <b>Seconds</b>—Time elapsed since this display was started.</li> <li>• <b>Time</b>—Current local time.</li> </ul>
(2)	<b>Delay</b> —Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.
(3)	Displays the following: <ul style="list-style-type: none"> <li>• <b>To</b>—Destination address of the LSP.</li> <li>• <b>From</b>—Originating address of the LSP.</li> <li>• <b>State</b>—Current state of the LSP: <b>Up</b> or <b>Down</b>.</li> </ul>
(4)	Displays the following: <ul style="list-style-type: none"> <li>• <b>LSPName</b>—Name of the LSP.</li> <li>• <b>Type</b>—Type of LSP: <b>Ingress</b>, <b>Egress</b>, or <b>Transit</b>.</li> </ul>
(5)	Displays the following: <ul style="list-style-type: none"> <li>• <b>Label in</b>—Incoming label of the LSP.</li> <li>• <b>Label out</b>—Outgoing label of the LSP.</li> </ul>
(6)	<b>Port number</b> —Port number for the sending router, the port number for the receiving router, and the protocol ID. For MPLS traffic engineering applications, the protocol ID is always 0.
(7/8)	<b>Record route</b> —All intermediate and egress router addresses for this LSP.
(9/10/11)	Displays traffic statistics: <ul style="list-style-type: none"> <li>• <b>Output packets</b>—Number of packets that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.</li> <li>• <b>Output bytes</b>—Number of bytes that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.</li> </ul>
(12)	Displays any errors the router encountered while attempting to retrieve information on the LSP.
(13)	Lists the keyboard commands you can use to navigate to other LSPs. For a description of the keyboard commands, see <a href="#">Table 350 on page 4577</a> .

## Sample Output

### monitor label-switched-path

```
user@host> monitor label-switched-path
(1) host                               Seconds: 112           Time: 15:32:22
(2)                                     Delay: 0/0/0
(3) To 10.10.10.16, From 10.10.10.17, state: Up
(4) LSPname: k, type: Ingress
(5) Label in: -, Label out: 126000
(6) Port number: sender 1, receiver 45583, protocol 0
(7) Record Route: <self> 192.168.224.196
(8)   192.168.224.202 192.168.224.179
(9) Traffic statistics:                Current delta
(10)  Output packets:                  0                [0]
(11)  Output bytes:                   0                [0]
(12)
(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',
    LSP='l'
```

## ping mpls bgp

---

**Syntax**    ping mpls bgp *fec*  
             <bottom-label-ttl>  
             <count *count*>  
             <destination *address*>  
             <detail>  
             <exp *forwarding-class*>  
             <instance *routing-instance-name*>  
             <logical-system (all | *logical-system-name*)>  
             <size *bytes*>  
             <source *source-address*>  
             <sweep>

**Release Information**    Command introduced in Junos OS Release 11.1.

**Description**    Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a **ping mpls bgp** command.

**Options**    **bottom-label-ttl**—(Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.

**count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination** *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**fec**—Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

**instance** *routing-instance-name*—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**size** *bytes*—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

**source** *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls bgp fec count on page 4581](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. To display the error codes, use the **detail** option (for example, **ping mpls bgp 10.255.245.222 detail**).

## Sample Output

### ping mpls bgp fec count

```
user@host> ping mpls bgp 10.255.245.222 count 10
!!!xxx...x--- 1sping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

## ping mpls l2circuit

---

**Syntax** ping mpls l2circuit (interface *interface-name* | virtual-circuit *virtual-circuit-id* neighbor *address*)  
<count *count*>  
<destination *address*>  
<detail>  
<exp *forwarding-class*>  
<logical-system (all | *logical-system-name*)>  
reply-mode (application-level-control-channel | ip-udp | no-reply)  
<size *bytes*>  
<source *source-address*>  
<sweep>  
<v1>

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.

**Description** Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a ping mpls l2circuit command. You can also issue this command within logical systems.

**Options** **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination** *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**interface** *interface-name*—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**reply-mode**—(Optional) Reply mode for the ping request. This option has the following suboptions:

**application-level-control-channel**—Reply using an application level control channel.

**ip-udp**—Reply using an IPv4 or IPv6 UDP packet.

**no-reply**—Do not reply to the ping request.



**NOTE:** The **reply-mode** option and its suboptions **application-level-control-channel**, **ip-udp**, and **no-reply** are also available in Junos OS Release 10.2R4 and 10.3R2.

---



**size bytes**—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

**source source-address**—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**vl**—(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

**virtual-circuit virtual-circuit-id neighbor address**—Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

**Additional Information** You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls l2circuit interface on page 4583](#)  
[ping mpls l2circuit virtual-circuit detail on page 4583](#)  
[ping mpls l2circuit interface <interface-name> reply-mode on page 4584](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

### ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
```

Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms

#### ping mpls l2circuit interface <interface-name> reply-mode

```
user@host> ping mpls l2circuit interface lt-1/2/0.21 reply-mode application-level-control-channel
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## ping mpls l3vpn

<b>Syntax</b>	<pre>ping mpls l3vpn prefix <i>prefix-name</i> &lt;l3vpn-name&gt; &lt;bottom-label-ttl&gt; &lt;count <i>count</i>&gt; &lt;destination <i>address</i>&gt; &lt;detail&gt; &lt;exp forwarding-class&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;size <i>bytes</i>&gt; &lt;source <i>source-address</i>&gt; &lt;sweep&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The <b>size</b> and <b>sweep</b> options were introduced in Junos OS Release 9.6.</p>
<b>Description</b>	<p>Check the operability of an MPLS Layer 3 virtual private network (VPN) connection.</p> <p>Press Ctrl+c to interrupt a <b>ping mpls l3vpn</b> command.</p>
<b>Options</b>	<p><b>bottom-label-ttl</b>—(Optional) Display the time-to-live value for the bottom label in the label stack.</p> <p><b>count <i>count</i></b>—(Optional) Number of ping requests to send. If <b>count</b> is not specified, five ping requests are sent. The range of values is 1 through <b>1,000,000</b>. The default value is <b>5</b>.</p> <p><b>destination <i>address</i></b>—(Optional) Specify an address other than the default (<b>127.0.0.1/32</b>) for the ping echo requests. The address can be anything within the <b>127/8</b> subnet.</p> <p><b>detail</b>—(Optional) Display detailed information about the echo requests sent and received.</p> <p><b>exp forwarding-class</b>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p><b>l3vpn-name</b>—(Optional) Layer 3 VPN name.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p><b>prefix <i>prefix-name</i></b>—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.</p> <p><b>size <i>bytes</i></b>—(Optional) Size of the label-switched path (LSP) ping request packet (<b>96</b> through <b>65468</b> bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.</p>

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes. The echo response is dropped because the PAD TLV is included in the echo response, making it too large.

If the Layer 3 VPN traffic transits a route reflector within the network, the **ping mpls l3vpn** command does not work.

**Required Privilege Level** network

**List of Sample Output** [ping mpls l3vpn on page 4586](#)  
[ping mpls l3vpn detail on page 4586](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. When an echo reply is received with an error code, the packets are not counted in the received packets count, and are counted separately.

## Sample Output

### ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!!!
--- 1sping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

### ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
```

```
--- lsping statistics ---  
5 packets transmitted, 5 packets received, 0% packet loss
```

## ping mpls ldp

---

**Syntax**    ping mpls ldp *fec*  
             <count *count*>  
             <destination *address*>  
             <detail>  
             <exp *forwarding-class*>  
             <instance *routing-instance-name*>  
             <logical-system (all | *logical-system-name*)>  
             <p2mp root-addr *ip-address* lsp-id *identifier*>  
             <size *bytes*>  
             <source *source-address*>  
             <sweep>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                             Command introduced in Junos OS Release 9.0 for EX Series switches.  
                             **size** and **sweep** options introduced in Junos OS Release 9.6.  
                             **instance** option introduced in Junos OS Release 10.0.  
                             **p2mp**, **root-address**, and **lsp-id** options introduced in Junos OS Release 11.2.  
                             Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**    Check the operability of MPLS LDP-signaled label-switched path (LSP) connections.  
                             Type Ctrl+c to interrupt a **ping mpls** command.

**Options**    **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through **1,000,000**. The default value is **5**.

**destination** *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.

**exp** *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

**fec**—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

**instance** *routing-instance-name*—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

**logical-system** (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

**p2mp root-addr** *ip-address* **lsp-id** *identifier*—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

**size** *bytes*—(Optional) Size of the LSP ping request packet (**88** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller

than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *Junos OS MPLS Applications Library for Routing Devices*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls ldp fec count on page 4589](#)  
[ping mpls ldp p2mp root-addr lsp-id on page 4589](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

### ping mpls ldp p2mp root-addr lsp-id

```
user@host> ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

```
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
    Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
    Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
    Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
    Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```



## ping mpls lsp-end-point

<b>Syntax</b>	<pre>ping mpls lsp-end-point <i>prefix-name</i> &lt;count <i>count</i>&gt; &lt;destination <i>address</i>&gt; &lt;detail&gt; &lt;exp <i>forwarding-class</i>&gt; &lt;instance <i>routing-instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;size <i>bytes</i>&gt; &lt;source <i>source-address</i>&gt; &lt;sweep&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The <b>size</b> and <b>sweep</b> options were introduced in Junos OS Release 9.6.</p> <p>The <b>instance</b> option was introduced in Junos OS Release 10.0.</p>
<b>Description</b>	<p>Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a <b>ping mpls</b> command.</p>
<b>Options</b>	<p><b>count</b> <i>count</i>—(Optional) Number of ping requests to send. If <b>count</b> is not specified, five ping requests are sent. The range of values is 1 through <b>1,000,000</b>. The default value is <b>5</b>.</p> <p><b>destination</b> <i>address</i>—(Optional) Specify an address other than the default (<b>127.0.0.1/32</b>) for the ping echo requests. The address can be anything within the <b>127/8</b> subnet.</p> <p><b>detail</b>—(Optional) Display detailed information about the echo requests sent and received.</p> <p><b>exp</b> <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p><b>instance</b> <i>routing-instance-name</i>—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.</p> <p><b>logical-system</b> (all   <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p><b>prefix-name</b>—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.</p> <p><b>size</b> <i>bytes</i>—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is <b>88</b> bytes. If the endpoint is RSVP-based, the minimum size of the packet is <b>100</b> bytes. The maximum size in either case is <b>65468</b> bytes.</p> <p><b>source</b> <i>source-address</i>—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (<b>lo.0</b>).</p> <p><b>sweep</b>—(Optional) Automatically determine the size of the maximum transmission unit (MTU).</p>

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls lsp-end-point detail on page 4592](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### [ping mpls lsp-end-point detail](#)

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## ping mpls rsvp

**Syntax** ping mpls rsvp  
 <lsp-name>  
 <count count>  
 <destination address>  
 <detail>  
 <dynamic-bypass>  
 <egress egress-address>  
 <exp forwarding-class>  
 <interface interface-name>  
 <logical-system (all | logical-system-name)>  
 <manual-bypass>  
 <multipoint>  
 <size bytes>  
 <source source-address>  
 <standby standby-path-name>  
 <sweep>

**Release Information** Command introduced before Junos OS Release 7.4.  
 The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.  
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
 The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.

**Description** Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

**Options** **count count**—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination address**—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.



**NOTE:** When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

**dynamic-bypass**—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

**egress *egress-address***—(Optional) Only the specified egress router or switch responds to the ping request.

**exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

**interface**—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.

***lsp-name***—Ping an RSVP-signaled LSP using an LSP name.

**manual-bypass**—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

**multipoint**—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

**size *bytes***—(Optional) Size of the LSP ping request packet (100 through 65468 bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

**standby *standby-path-name***—(Optional) Name of the standby path.

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls rsvp \(Echo Reply Received\) on page 4595](#)  
[ping mpls rsvp \(Echo Reply with Error Code\) on page 4595](#)

[ping mpls rsvp detail on page 4595](#)

[ping mpls rsvp multipoint egress detail count on page 4595](#)

[ping mpls rsvp multipoint detail count on page 4595](#)

[ping mpls rsvp destination detail count size on page 4596](#)

[ping mpls rsvp destination detail sweep size on page 4596](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
loss
```

### ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

### ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

### ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
Local transmit time: 1205310695s 215737us
Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

### ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 ms
Local transmit time: 1205310615s 347317us
Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
```

```
Local transmit time: 1205310615s 347262us
Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
Local transmit time: 1205310615s 347167us
Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

#### ping mpls rsvp destination detail count size

```
user@host> ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
    Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
    Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

#### ping mpls rsvp destination detail sweep size

```
user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
    Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
    Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
    Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
    Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
    Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
    Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
    Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
    Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
    Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
    Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms
```

```
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
    Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
    Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
    Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
    Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
    Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
    Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
    Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
    Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

## request mpls lsp adjust-autobandwidth

---

<b>List of Syntax</b>	<a href="#">Syntax on page 4598</a> <a href="#">Syntax (EX and QFX Series Switches) on page 4598</a>
<b>Syntax</b>	<pre>request mpls lsp adjust-autobandwidth &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;name <i>lsp-name</i>&gt;</pre>
<b>Syntax (EX and QFX Series Switches)</b>	<pre>request mpls lsp adjust-autobandwidth &lt;name <i>lsp-name</i>&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	<p>Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).</p> <p>Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.</p> <p>During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as <i>make-before-break</i>.</p>
<b>Options</b>	<p><b>none</b>—Manually trigger a bandwidth allocation adjustment for all active LSP paths.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>lsp-name</i></b>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.</p>
<b>Additional Information</b>	<p>For this command to work properly, the following conditions must exist:</p> <ul style="list-style-type: none"><li>• Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the <b>request mpls lsp adjust-autobandwidth</b> command.</li><li>• The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.</li></ul>
<b>Required Privilege Level</b>	clear, maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">auto-bandwidth on page 4516</a></li><li>• <a href="#">Configuring Automatic Bandwidth Allocation for LSPs on page 4482</a></li></ul>



List of Sample Output [request mpls lsp adjust-auto-bandwidth on page 4599](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

### Sample Output

[request mpls lsp adjust-auto-bandwidth](#)

```
user@host> request mpls lsp adjust-auto-bandwidth
```

## show ldp database

---

<b>Syntax</b>	<pre>show ldp database &lt;brief   detail   extensive&gt; &lt;inet   l2circuit&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;session <i>session</i>&gt; p2mp</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display entries in the LDP database.
<b>Options</b>	<p><b>none</b>—Display standard information about all entries in the LDP database for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>inet   l2circuit</b>—(Optional) Display only IPv4 or Layer 2 circuit bindings.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>session <i>session</i></b>—(Optional) Display database for the specified session only. <b><i>session</i></b> is the destination address of the LDP session.</p> <p><b>p2mp</b>—(Optional) Display point-to-multipoint binding information.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show ldp database (master) on page 4602</a></p> <p><a href="#">show ldp database (standby) on page 4603</a></p> <p><a href="#">show ldp database l2circuit detail on page 4604</a></p> <p><a href="#">show ldp database l2circuit extensive on page 4604</a></p> <p><a href="#">show ldp database p2mp (master) on page 4604</a></p> <p><a href="#">show ldp database p2mp (standby) on page 4605</a></p> <p><a href="#">show ldp database p2mp (master) on page 4605</a></p> <p><a href="#">show ldp database p2mp (standby) on page 4605</a></p> <p><a href="#">show ldp database session on page 4606</a></p> <p><a href="#">show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4606</a></p> <p><a href="#">show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4607</a></p>
<b>Output Fields</b>	<p><a href="#">Table 352 on page 4601</a> describes the output fields for the <b>show ldp database</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 352: show ldp database Output Fields

Field Name	Field Description	Level of Output
Input label database	Label received from the other router.	All levels
Output label database	Label advertised to the other router.	All levels
<i>session-identifier</i>	Session identifier, which includes the local and remote label space identifiers.	All levels
Label	Label binding to a route prefix.	All levels
Prefix	<p>Route prefix.</p> <p>It can be one of the following values:</p> <ul style="list-style-type: none"> <li>IP prefix.</li> <li>Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.</li> <li>Layer 2 encapsulation type.</li> </ul> <p>Layer 2 encapsulation types are displayed in the format <b>L2CKT control word status encapsulation-type vc-number</b>, for example, <b>L2CKT CtlfWord FRAME RELAY VC 2</b></p> <ul style="list-style-type: none"> <li><b>control-word-status</b>—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> <li>NoCtrlWord</li> <li>CtrlWord</li> </ul> </li> <li><b>encapsulation-type</b>—Encapsulation type: <ul style="list-style-type: none"> <li>FRAME RELAY</li> <li>ATM AAL5</li> <li>ATM CELL</li> <li>VLAN</li> <li>ETHERNET</li> <li>CISCO_HDLC</li> <li>PPP</li> </ul> </li> <li><b>VC number</b>—Virtual circuit number. It can have any numeric value.</li> <li><b>(Stale)</b>—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time.</li> </ul>	All levels
MTU	MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.	detail
VCCV Control Channel types	<p>Virtual Circuit Connection Verification (VCCV) control channel types.</p> <ul style="list-style-type: none"> <li>MPLS router alert label</li> <li>MPLS PW label with TTL=1</li> </ul>	extensive

Table 352: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
VCCV Control Verification types	The only valid VCCV control verification type is <b>LSP ping</b> .	<b>extensive</b>
TDM payload size	Size of the Time Division Multiplex (TDM) payload.	All levels
TDM bitrate	Bit rate for the TDM traffic.	All levels
Requested VLAN ID	(VLANs) VLAN identifier of the Layer 2 circuit.	<b>detail</b>
Cell bundle size	(ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.	<b>detail</b>
State	State of the label binding: <ul style="list-style-type: none"> <li>• <b>Active</b>—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.</li> <li>• <b>New</b>—New label that has not yet been distributed.               <ul style="list-style-type: none"> <li>• <b>MapRcv</b>—Waiting to receive a label mapping message.</li> <li>• <b>MapSend</b>—Waiting to send a label mapping message.</li> <li>• <b>RelRcv</b>—Waiting to receive a label release message.</li> <li>• <b>RelRsnd</b>—Waiting to receive a label release message before resending label mapping message.</li> <li>• <b>RelSend</b>—Waiting to send a label release message.</li> <li>• <b>ReqSend</b>—Waiting to send a label request message.</li> <li>• <b>W/dSend</b>—Waiting to send a label withdrawal message.</li> </ul> </li> </ul>	<b>detail</b>
Age	Time elapsed since the binding was created.	<b>detail</b>

## Sample Output

### show ldp database (master)

```

user@host> show ldp database extensive
Input label database, 10.255.107.232:0--10.255.107.236:0
Label Prefix
299840 10.255.107.232/32
      State: Active
      Age: 9:35
      Entropy Label Capability: No
      3 10.255.107.236/32
      State: Active
      Age: 9:35
      Entropy Label Capability: No
      299776 L2CKT CtrlWord VLAN VC 100
      MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R
      Bit: 1
      State: Active
      Age: 9:35
      Entropy Label Capability: No
      VCCV Control Channel types:

```

```

PWE3 control word
MPLS router alert label
MPLS PW label with TTL=1
VCCV Control Verification types:
LSP ping
BFD with PW-ACH-encapsulation for Fault Detection
BFD with IP/UDP-encapsulation for Fault Detection

```

Output label database, 10.255.107.232:0--10.255.107.236:0

```

Label Prefix
3      10.255.107.232/32
      State: Active
      Age: 9:35
      Entropy Label Capability: No
299776 10.255.107.236/32
      State: Active
      Age: 9:35
      Entropy Label Capability: No

```

### show ldp database (standby)

user@host> show ldp database extensive

Input label database, 10.255.107.236:0--10.255.107.234:0

```

Label Prefix
299808 10.255.107.230/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0
Label Prefix
301136 10.255.107.232/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0
Label Prefix
3      10.255.107.234/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0
Label Prefix
302480 10.255.107.236/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0

```

Output label database, 10.255.107.236:0--10.255.107.234:0

```

Label Prefix
299904 10.255.107.230/32
      State: Active
      Age: 1d 2:46:36
299936 10.255.107.232/32
      State: Active
      Age: 1d 2:46:36

```

```
299872      10.255.107.234/32
            State: Active
            Age: 1d 2:46:36
      3      10.255.107.236/32
            State: Active
            Age: 1d 2:46:36
299952      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 1d 2:46:36
```

#### show ldp database l2circuit detail

```
user@host> show ldp database l2circuit detail
Input label database, 10.255.245.44:0--10.255.245.45:0
  Label      Prefix
  100176     L2CKT CtrlWord ATM CELL (VC Mode) VC 100
            Cell bundle size: 80
            State: Active
            Age: 9:48
  100256     L2CKT CtrlWord FRAME RELAY VC 101
            MTU: 4470
            State: Active
            Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
  Label      Prefix
  100048     L2CKT CtrlWord ATM CELL (VC Mode) VC 100
            Cell bundle size: 80
            State: Active
            Age: 9:48
  100112     L2CKT CtrlWord FRAME RELAY VC 101
            MTU: 4470
            State: Active
            Age: 9:48
```

#### show ldp database l2circuit extensive

```
user@host> show ldp database l2circuit extensive
Input label database, 10.255.245.198:0--10.255.245.194:0
  Label      Prefix
  299872     L2CKT CtrlWord PPP VC 100
            MTU: 4470
            VCCV Control Channel types:
              MPLS router alert label
              MPLS PW label with TTL=1
            VCCV Control Verification types:
              LSP ping
  Label      Prefix
            State: Active
            Age: 19:23:08
```

#### show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive

Input label database, 10.255.107.232:0--10.255.107.236:0
  Label      Prefix
  569649     P2MP root-addr 10.255.107.232, lsp-id 16777217
            State: Active
            Age: 2d 6:41:46
```

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0

```
Label Prefix
299888 P2MP root-addr 10.255.107.230, lsp-id 16777217
State: Active
Age: 2d 6:41:35
```

#### show ldp database p2mp (standby)

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0

```
Label Prefix
299968 P2MP root-addr 10.255.107.230, lsp-id 16777217
State: Active
Age: 4d 22:21:57
Standby binding state:
Map messages: 1
Release messages: 0
```

Output label database, 10.255.107.236:0--10.255.107.232:0

```
Label Prefix
3 P2MP root-addr 10.255.107.232, lsp-id 1
State: Active
Age: 4d 22:21:57
```

#### show ldp database p2mp (master)

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.232:0--10.255.107.236:0

```
Label Prefix
569649 P2MP root-addr 10.255.107.232, lsp-id 16777217
State: Active
Age: 2d 6:41:46
```

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0

```
Label Prefix
299888 P2MP root-addr 10.255.107.230, lsp-id 16777217
State: Active
Age: 2d 6:41:35
```

#### show ldp database p2mp (standby)

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0

```
Label Prefix
299968 P2MP root-addr 10.255.107.230, lsp-id 16777217
State: Active
Age: 4d 22:21:57
Standby binding state:
Map messages: 1
Release messages: 0
```

```
Output label database, 10.255.107.236:0--10.255.107.232:0
Label Prefix
 3      P2MP root-addr 10.255.107.232, lsp-id 1
      State: Active
      Age: 4d 22:21:57
```

#### show ldp database session

```
user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
Label Prefix
100002 10.255.245.197/32
100003 10.255.245.196/32
100004 10.0.0.194/32
 3      10.1.1.195/32
100000 L2CKT NoCtrlWord FRAME RELAY VC 1
100001 L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
Label Prefix
100003 10.255.245.197/32
100004 10.1.1.195/32
100002 10.255.245.196/32
 3      10.0.0.194/32
100000 L2CKT CtrlWord FRAME RELAY VC 2
100001 L2CKT NoCtrlWord FRAME RELAY VC 1
```

#### show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp database
Input label database, 1.1.1.2:0--1.1.1.3:0
Label Prefix
299808 1.1.1.2/32
 3      1.1.1.3/32
299792 1.1.1.6/32
299776 10.255.2.227/32
299840 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299824 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

Output label database, 1.1.1.2:0--1.1.1.3:0
Label Prefix
 3      1.1.1.2/32
299776 1.1.1.3/32
299808 1.1.1.6/32
299792 10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
Label Prefix
299856 1.1.1.2/32
299792 1.1.1.3/32
 3      1.1.1.6/32
299776 10.255.2.227/32
299888 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808 P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840 P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872 P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

Output label database, 1.1.1.2:0--1.1.1.6:0
Label Prefix
 3      1.1.1.2/32
```



```

299776    1.1.1.3/32
299808    1.1.1.6/32
299792    10.255.2.227/32

```

### show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 10.255.2.227:0--1.1.1.3:0
  Label    Prefix
299808     1.1.1.2/32
   3       1.1.1.3/32
299792     1.1.1.6/32
299776     10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.3:0
  Label    Prefix
299856     1.1.1.2/32
299776     1.1.1.3/32
299792     1.1.1.6/32
   3       10.255.2.227/32

Input label database, 10.255.2.227:0--1.1.1.6:0
  Label    Prefix
299856     1.1.1.2/32
299792     1.1.1.3/32
   3       1.1.1.6/32
299776     10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.6:0
  Label    Prefix
299856     1.1.1.2/32
299776     1.1.1.3/32
299792     1.1.1.6/32
   3       10.255.2.227/32
299888     P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808     P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824     P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840     P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872     P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

```

## show ldp fec-filters

<b>Syntax</b>	show ldp fec-filters <fec> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.
<b>Options</b>	<b>fec</b> —(Optional) Display FEC filter information for the specified FEC.  <b>instance <i>instance-name</i></b> —(Optional) Display FEC filter information for the specified instance.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp fec-filters on page 4608</a>
<b>Output Fields</b>	<a href="#">Table 353 on page 4608</a> lists the output fields for the <b>show ldp fec-filters</b> command. Output fields are listed in the approximate order in which they appear.

**Table 353: show ldp fec-filters Output Fields**

Field Name	Field Description
Ingress	Names of the FEC filters on the ingress routers.
Transit	Names of the FEC filters on the transit routers.

## Sample Output

### show ldp fec-filters

```
user@host> show ldp fec-filters 10/8
10.22.1.2/32
  Ingress: f1-10.22.1.2/32 (index: 3)
  Transit: (null) (index: 0)
```

## show ldp interface

<b>Syntax</b>	show ldp interface <brief   detail   extensive> <interface-name> <instance instance-name> <logical-system (all   logical-system-name)>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display the status of Label Distribution Protocol (LDP)-enabled interfaces.
<b>Options</b>	<p><b>none</b>—Display standard status information about all LDP-enabled interface for all routing instances.</p> <p><b>interface-name</b>—(Optional) Display information for the specified interface.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance instance-name</b>—(Optional) Display information for the specified routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp interface extensive on page 4610</a>
<b>Output Fields</b>	<a href="#">Table 354 on page 4609</a> describes the output fields for the <b>show ldp interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 354: show ldp interface Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Interface name.	All levels
<b>Label space ID</b>	Label space identifier that the router is advertising on the interface.	All levels
<b>Nbr count</b>	Number of neighbors on the interface.	All levels
<b>Next hello</b>	How long until the next hello packet is sent on this interface, in seconds.	All levels
<b>Hello interval</b>	One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used.	<b>detail</b> <b>extensive</b>
<b>Hold time</b>	Configured hold time, in seconds.	<b>detail</b> <b>extensive</b>

Table 354: show ldp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	extensive
Local hello interval	Locally configured hello interval.	extensive

## Sample Output

### show ldp interface extensive

```
user@host> show ldp interface extensive
Interface          Label space ID      Nbr count  Next hello
fe-0/0/3.0         10.255.245.6:0      2          0
Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
Local hello interval: 2, Index: 69
```

## show ldp neighbor

<b>Syntax</b>	show ldp neighbor <brief   detail   extensive> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )> <neighbor-address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>neighbor-address</b> option added in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display Label Distribution Protocol (LDP) neighbor information.
<b>Options</b>	<b>none</b> —Display standard information about LDP neighbors for all routing instances.  <b>brief   detail   extensive</b> —(Optional) Display the specified level of output.  <b>instance <i>instance-name</i></b> —(Optional) Display information for the specified routing instance.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>neighbor-address</b> —(Optional) Display information about the specified LDP neighbor.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ldp neighbor on page 4569</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ldp neighbor extensive on page 4612</a>
<b>Output Fields</b>	<a href="#">Table 355 on page 4611</a> describes the output fields for the <b>show ldp neighbor</b> command. Output fields are listed in the approximate order in which they appear.

Table 355: show ldp neighbor Output Fields

Field Name	Field Description	Level of Output
Address	IP address of the neighbor.	All levels
Interface	Interface over which the neighbor was discovered.	All levels
Label space ID	Label space identifier advertised by the neighbor.	All levels
Hold time	Remaining hold time before the neighbor expires, in seconds.	All levels
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	<b>detail</b>

Table 355: show ldp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configuration sequence	Counter that increments whenever the neighbor changes its configuration.	detail
Up for	Length of time the LDP neighbor has been in operation.	detail extensive
Reference count	Reference count for the LDP neighbor.	extensive
Hold time	Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers.	extensive
Proposed local/peer	Hold time value proposed by the local router and the peer router.	extensive

## Sample Output

### show ldp neighbor extensive

```

user@host> show ldp neighbor extensive
Address          Interface      Label space ID      Hold Time
192.168.37.23    so-1/0/0.0    10.255.245.5:0      44
  Transport address: 10.255.245.5, Configuration sequence: 6
  Up for 00:03:37
  Reference count: 1
  Hold time: 45, Proposed local/peer: 15/45

```

## show ldp path

<b>Syntax</b>	show ldp path <brief   detail   extensive> <destination> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display Label Distribution Protocol (LDP) label-switched paths (LSPs).
<b>Options</b>	<p><b>none</b>—Display standard information about all LDP LSPs for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict the output to entries that match the specified destination prefix.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp path extensive on page 4614</a>
<b>Output Fields</b>	<a href="#">Table 356 on page 4613</a> describes the output fields for the <b>show ldp path</b> command. Output fields are listed in the approximate order in which they appear.

**Table 356: show ldp path Output Fields**

Field Name	Field Description
<b>Output Session (label)</b>	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.
<b>Input Session (label)</b>	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.
<b>route</b>	MPLS route.
<b>Attached route</b>	Route corresponding to the LSP.
<b>Ingress route</b>	The router acts as the ingress for the LSP.
<b>Reference count</b>	Reference count for the LDP neighbor.

Table 356: show ldp path Output Fields (*continued*)

Field Name	Field Description
Transit route	Names of the forwarding equivalence class (FEC) filters on the transit routers.
Global label	MPLS label that is used globally.

## Sample Output

### show ldp path extensive

```
user@host> show ldp path extensive
Output Session (label)      Input Session (label)
10.255.14.220:0(3)         ( )
  Attached route: 10.255.14.221/32
  Reference count: 3, Global label: 3
10.255.14.220:0(100000)     10.255.14.220:0(3)
  Attached route: 10.255.14.220/32, Ingress route
  Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)     10.255.14.220:0(100001)
  Attached route: 10.255.14.214/32, Ingress route
  Reference count: 2, Transit route, Global label: 100001
```



## show ldp route

<b>Syntax</b>	<pre>show ldp route &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).
<b>Options</b>	<p><b>none</b>—Display standard information about all entries in the LDP internal topology table for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.</p> <p><b>instance instance-name</b>—(Optional) Display entries for the specified routing instance only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show ldp route detail on page 4617</a></p> <p><a href="#">show ldp route extensive on page 4617</a></p>
<b>Output Fields</b>	<a href="#">Table 357 on page 4615</a> describes the output fields for the <b>show ldp route</b> command. Output fields are listed in the approximate order in which they appear.

**Table 357: show ldp route Output Fields**

Field Name	Field Description
<b>Destination</b>	Destination prefix.
<b>Next-hop intf/lsp/table</b>	Interface that is the next hop to the destination prefix.
<b>Next-hop address</b>	IP address of the next hop.
<b>Session ID</b>	LDP session ID.

Table 357: show ldp route Output Fields (*continued*)

Field Name	Field Description
Route flags	Information about the route. For example, the <b>Ingress TTL propagate</b> flag indicates that the time-to-live (TTL) value is being propagated with the route.
Bound to outgoing label	The route has been bound to LSPs with the label being distributed for that LSP.
Topology entry	The topology that the route is bound to.
Ingress route status	Status of the ingress route. For example, it could be <b>Active</b> or <b>Inactive</b> .
Last modified	The length of time since the ingress route status last changed.

## Sample Output

### show ldap route detail

```

user@host> show ldap route 10.255.8.5 detail
Destination      Next-hop intf/lsp      Next-hop address
10.255.8.5/32     f1
  Session ID 10.255.170.84:0--10.255.170.92:0
                    fe-0/0/0.0      192.168.100.2
  Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/1.0
  Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/2.0
  Session ID 10.255.170.84:0--10.255.8.3:0
  Bound to outgoing label 299776, Topology entry: 0x8c38a80
  BFD dest addr   BFD state LSP-ping Next-hop addr Next-hop intf/lsp
127.0.0.64       up        up        192.168.100.2 fe-0/0/0.0
127.0.1.64       up        up        so-0/2/1.0
127.0.2.64       up        up        so-0/2/2.0
127.0.3.64       up        up        f1
.....

```

### show ldap route extensive

```

user@host> show ldap route extensive

Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.0/30       ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.4/30       ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.8/30       ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.12/30      ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.16/30      ge-1/2/0.18
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.18/32      ge-1/2/0.18
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.20/30      ge-1/2/1.21
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.21/32      ge-1/2/1.21
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
192.168.0.1/32    ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
192.168.0.2/32    ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0

```

```

                                ge-1/2/0.18                10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.3/32   ge-1/2/1.21                10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.4/32   ge-1/2/1.21                10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Bound to outgoing label 299808, Topology entry: 0x92a483c
  Ingress route status: Active, Last modified: 00:01:19 ago
  Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.5/32   ge-1/2/0.18                10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Bound to outgoing label 299792, Topology entry: 0x92a47f8
  Ingress route status: Active, Last modified: 00:01:19 ago
  Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.6/32   lo0.6
  Bound to outgoing label 3, Topology entry: 0x92a4a5c
  Ingress route status: Inactive
  Route type: Egress route
  Route flags: None
```

## show ldp session

<b>Syntax</b>	show ldp session <brief   detail   extensive> <destination> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display information about Label Distribution Protocol (LDP) sessions.
<b>Options</b>	<p><b>none</b>—Display standard information about all LDP sessions for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict LDP session display to the specified address.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <b><i>instance-name</i></b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ldp session on page 4570</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ldp session brief on page 4623</a> <a href="#">show ldp session detail on page 4623</a> <a href="#">show ldp session extensive on page 4623</a>
<b>Output Fields</b>	Table 358 on page 4619 describes the output fields for the <b>show ldp session</b> command. Output fields are listed in the approximate order in which they appear.

Table 358: show ldp session Output Fields

Field Name	Field Description	Level of Output
Address	Transport address of the session.	any
State	State of the session: <b>Nonexistent</b> , <b>Connecting</b> , <b>Initialized</b> , <b>OpenRec</b> , <b>OpenSent</b> , <b>Operational</b> , or <b>Closing</b> . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt.	any
Connection	TCP connection state: <b>Closed</b> , <b>Opening</b> , or <b>Open</b> .	any
Hold time	Time remaining until the session will be closed, in seconds.	any

Table 358: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Session ID</b>	LDP identifiers of the peers of this session.	detail extensive
<b>Next keepalive</b>	Time until next keepalive is sent, in seconds.	detail extensive
<b>Active</b>	Whether the local router is playing the active role in the session and during session establishment.	detail extensive
<b>Passive</b>	Whether the local router is playing the passive role in the session and during session establishment.	detail extensive
<b>Maximum PDU</b>	Maximum protocol data unit (PDU) size (packet size) for the session.	detail extensive
<b>Hold time</b>	Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the <b>keepalive-timeout</b> statement configured at the <b>[edit protocols ldp]</b> hierarchy level.	detail extensive
<b>Neighbor count</b>	Number of neighbors that are contributing to the session.	detail extensive
<b>Keepalive interval</b>	Keepalive interval, in seconds.	detail extensive
<b>Connect retry interval</b>	TCP connection retry interval, in seconds.	detail extensive
<b>Local address</b>	Local transport address.	detail extensive
<b>Remote address</b>	Remote transport address.	detail extensive
<b>Up for</b>	Time that this session has been up.	detail extensive
<b>Last down</b>	Time since the session last went down.	detail extensive

Table 358: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Reason	Reason the session went down: <ul style="list-style-type: none"> <li>Aborted graceful restart</li> <li>Authentication key was changed</li> <li>Bad type length value (TLV)</li> <li>Bad protocol data unit (PDU) packets</li> <li>Command-line interface (CLI) command</li> <li>Connect time expired</li> <li>Connection error</li> <li>Connection reset</li> <li>Error during initialization</li> <li>Hold time expired</li> <li>No adjacency or all adjacencies down</li> <li>Notification received</li> <li>Received notification from peer</li> <li>Unexpected End of File (EOF)</li> <li>Unknown reason</li> </ul>	detail extensive
Number of session flaps	Number of times the session changes from up to down.	detail extensive
Restarting	LDP is in the process of gracefully restarting.	detail extensive
Capabilities advertised	LDP capabilities advertised to a peer.	detail extensive
Capabilities received	LDP capabilities received from a peer.	detail extensive
Protection	Information about the status of MPLS LDP session protection.	detail extensive
restart complete in <i>nnn msec</i>	Amount of time (in milliseconds) remaining until graceful restart is declared complete.	detail extensive
Local	Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent. <ul style="list-style-type: none"> <li><b>Restart</b>—Status of the graceful restart feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li><b>Helper mode</b>—Status of the helper mode feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li><b>Reconnect time</b>—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is <b>60000 msec</b> and is not configurable. (<b>Reconnect timeout</b> refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.)</li> </ul>	detail extensive

Table 358: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Remote</b>	<p>Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.</li> </ul>	<b>detail extensive</b>
<b>Local maximum recovery time</b>	Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).	<b>detail extensive</b>
<b>Next-hop addresses received</b>	Next-hop addresses received on the session.	<b>detail extensive</b>
<b>Queue depth</b>	Number of messages that are queued for sending to the peers in the group.	<b>extensive</b>
<b>Message type</b>	<p>Type of message being sent:</p> <ul style="list-style-type: none"> <li>• <b>Initialization</b>—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established.</li> <li>• <b>Keepalive</b>—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them.</li> <li>• <b>Notification</b>—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer.</li> <li>• <b>Address</b>—Message sent by an LSR to an LDP peer to advertise interface addresses.</li> <li>• <b>Address withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address.</li> <li>• <b>Label mapping</b>—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC).</li> <li>• <b>Label request</b>—Message sent by an LSR to an LDP peer to request a label mapping for an FEC.</li> <li>• <b>Label withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping.</li> <li>• <b>Label release</b>—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released.</li> <li>• <b>Label abort</b>—Message sent by an LSR to an LDP peer to abort a label request message.</li> <li>• <b>Total</b>—Messages sent and received during the lifetime of the session.</li> <li>• <b>Last 5 seconds</b>—Messages sent and received during the current session.</li> </ul>	<b>extensive</b>



## Sample Output

### show ldp session brief

```
user@host> show ldp session brief
  Address      State      Connection  Hold time
10.255.72.160  Operational Open        21
10.255.72.164  Operational Open        20
10.255.72.172  Operational Open        21
```

### show ldp session detail

```
user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
```

### show ldp session extensive

```
user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 2 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:05:37
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
```

Queue depth: 0				
Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	33	33	1	1
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	7	5	0	0
Label request	0	0	0	0
Label withdraw	3	1	0	0
Label release	1	3	0	0
Label abort	0	0	0	0

## show ldp statistics

<b>Syntax</b>	show ldp statistics <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display Label Distribution Protocol (LDP) statistics.
<b>Options</b>	<p><b>none</b>—Display LDP statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear ldp statistics on page 4571</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ldp statistics on page 4628</a>
<b>Output Fields</b>	<a href="#">Table 359 on page 4625</a> lists the output fields for the <b>show ldp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 359: show ldp statistics Output Fields**

Field Name	Field Description
Total Sent, Received	Total number of each message type sent and received.
Last 5 seconds Sent, Received	Number of each message type sent and received in the last 5 seconds.

Table 359: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Message type</b>	<p>LDP message types:</p> <ul style="list-style-type: none"> <li>• <b>Hello</b>—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor.</li> <li>• <b>Initialization</b>—Messages that indicate an LDP session has started.</li> <li>• <b>Keepalive</b>—Messages that ensure that the keepalive timeout is not exceeded.</li> <li>• <b>Notification</b>—Advisory information and signal error information.</li> <li>• <b>Address</b>—Messages with address information.</li> <li>• <b>Address withdrawal</b>—Messages regarding address withdrawal.</li> <li>• <b>Label mapping</b>—Messages with label mapping information.</li> <li>• <b>Label request</b>—Request for a label mapping from a neighboring router.</li> <li>• <b>Label withdrawal</b>—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use.</li> <li>• <b>Label release</b>—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use.</li> <li>• <b>Label abort</b>—Messages about label interruptions.</li> <li>• <b>All UDP</b>—All hello messages sent by LSRs to the well-known UDP port, 646.</li> <li>• <b>All TCP</b>—All LDP session messages.</li> </ul>

Table 359: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Event type</b>	<p>LDP events and errors:</p> <ul style="list-style-type: none"> <li>• <b>Sessions opened</b>—Number of LDP sessions that have been opened.</li> <li>• <b>Sessions closed</b>—Number of LDP sessions that have been closed.</li> <li>• <b>Topology changes</b>—Number of changes to the known LDP topology.</li> <li>• <b>No interface</b>—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages.</li> <li>• <b>No session</b>—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers.</li> <li>• <b>No adjacency</b>—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected.</li> <li>• <b>Unknown version</b>—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment.</li> <li>• <b>Malformed PDU</b>—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session.  An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment.  An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length).</li> <li>• <b>Malformed message</b>—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them.  An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code.  An LDP message is considered to be malformed if the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU.  The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component.  The LDP message is considered to be malformed if the message is missing one or more mandatory parameters.</li> <li>• <b>Unknown message type</b>—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message.</li> <li>• <b>Inappropriate message</b>—The message is not of the type that the receiver expects to receive.</li> <li>• <b>Malformed TLV</b>—The TLV Length is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR.</li> <li>• <b>Bad TLV value</b>—The TLV Length is too large.</li> <li>• <b>Missing TLV</b>—The TLV is missing one or more mandatory parameters.</li> <li>• <b>PDU too large</b>—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.</li> </ul>
<b>Total</b>	Total number of each event or error.
<b>Last 5 seconds</b>	Number of each event or error in the last 5 seconds.

## Sample Output

### show ldp statistics

```
user@host> show ldp statistics
```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	265	263	2	2
Initialization	2	2	0	0
Keepalive	112	111	1	0
Notification	0	0	0	0
Address	2	2	0	0
Address withdraw	0	0	0	0
Label mapping	7	6	0	0
Label request	0	0	0	0
Label withdraw	2	0	0	0
Label release	0	2	0	0
Label abort	0	0	0	0
All UDP	265	263	2	2
All TCP	123	121	1	0

Event type	Total	Last 5 seconds	
		Sent	Received
Sessions opened	2		0
Sessions closed	0		0
Topology changes	11		0
No interface	0		0
No session	0		0
No adjacency	0		0
Unknown version	0		0
Malformed PDU	0		0
Malformed message	0		0
Unknown message type	0		0
Inappropriate message	0		0
Malformed TLV	0		0
Bad TLV value	0		0
Missing TLV	0		0
PDU too large	0		0

## show ldp traffic-statistics


<b>Syntax</b>	<pre>show ldp traffic-statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;p2mp&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>p2mp</b> option added in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>
<b>Description</b>	Display Label Distribution Protocol (LDP) traffic statistics.
<div>  <b>NOTE:</b> If nonstop active routing features is configured, <b>show ldp traffic-statistics</b> command is not supported on backup Routing Engines. </div>	
<b>Options</b>	<p><b>none</b>—Display LDP traffic statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>p2mp</b>—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p>
<b>Additional Information</b>	To collect output from this command on a periodic basis, configure the <a href="#">traffic-statistics</a> statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> .
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ldp statistics on page 4571</a></li> <li>• <a href="#">Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain</a></li> <li>• <a href="#">Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show ldp traffic-statistics on page 4630</a></p> <p><a href="#">show ldp traffic-statistics p2mp on page 4631</a></p> <p><a href="#">show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 4631</a></p> <p><a href="#">show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute) on page 4631</a></p>
<b>Output Fields</b>	<p><a href="#">Table 360 on page 4630</a> lists the output fields for the <b>show ldp traffic-statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 360: show ldp traffic-statistics Output Fields

Field Name	Field Description
<b>Message type</b>	LDP message types.
<b>FEC</b>	Forwarding equivalence class (FEC) for which LDP traffic statistics are collected.  For P2MP LSPs, FEC appears as a combination of root address and the LSP ID ( <b>root_addr:lsp_id</b> ).  For M-LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address ( <b>root_addr:lsp_id/grp,src</b> ).
<b>Type</b>	Type of traffic originating from a router, either <b>Ingress</b> (originating from this router) or <b>Transit</b> (forwarded through this router).
<b>Packets</b>	Number of packets passed by the FEC since its LSP came up.
<b>Bytes</b>	Number of bytes of data passed by the FEC since its LSP came up.
<b>Shared</b>	Whether a label is shared by prefixes: <b>Yes</b> or <b>No</b> . A <b>Yes</b> value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
<b>Nextthop</b>	The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)
<b>Label</b>	For multipoint LDP with multicast-only fast reroute (MoFRR), the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop.  Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
<b>Backup route</b>	For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.

## Sample Output

### show ldp traffic-statistics

```
user@host> show ldp traffic-statistics
```

FEC	Type	Packets	Bytes	Shared
10.35.3.0/30	Transit	0	0	Yes
	Ingress	0	0	No
10.35.10.1/32	Transit	0	0	Yes



	Ingress	0	0	No
10.255.245.214/32	Transit	0	0	No
	Ingress	11	752	No
192.168.37.36/30	Transit	0	0	Yes
	Ingress	0	0	No

FEC(root_addr:lsp_id)	Nexthop	Packets	Bytes	Shared
10.255.72.160:16777217	192.168.8.81	152056	14597376	No
	192.168.8.1	152056	14597376	No
	192.168.8.65	152056	14597376	No

#### show ldp traffic-statistics p2mp

```
user@host> show ldp traffic-statistics p2mp
FEC(root_addr:lsp_id) Nexthop      Packets      Bytes Shared
10.255.72.160:16777217 192.168.8.81 152056      14597376 No
                        192.168.8.1 152056      14597376 No
                        192.168.8.65 152056      14597376 No
```

#### show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp traffic-statistics p2mp
P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)  Nexthop      Packets      Bytes
Shared
11.99.0.73:239.10.0.1,11.98.0.10 11.99.0.117 243408      121217184
No
                        11.99.0.13 236286      117670428
No
11.99.0.73:239.10.0.2,11.98.0.10 11.99.0.117 248800      123902400
No
                        11.99.0.13 240759      119897982
No
11.99.0.73:239.10.0.1,11.98.0.20 11.99.0.117 250286      124642428
No
                        11.99.0.13 243741      121383018
No
11.99.0.73:239.10.0.2,11.98.0.20 11.99.0.117 252970      125979060
No
                        11.99.0.13 245218      122118564
No
```

#### show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show ldp traffic-statistics p2mp

P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)  Nexthop      Packets      Bytes
Shared
```

1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568		
1.3.8.2	0	0
No		
1.3.4.2	0	0
No		
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route		
1.3.4.2	0	0
No		
1.3.8.2	0	0
No		
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600		
1.3.8.2	0	0
No		
1.3.4.2	0	0
No		
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route		
1.3.4.2	0	0
No		
1.3.8.2	0	0
No		

## show security keychain

<b>Syntax</b>	show security keychain <brief   detail>
<b>Release Information</b>	Command introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.
<b>Options</b>	<b>none</b> —Display information about authentication keychains. <b>brief   detail</b> —(Optional) Display the specified level of output.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show security keychain brief on page 4635</a> <a href="#">show security keychain detail on page 4635</a>
<b>Output Fields</b>	<a href="#">Table 312 on page 3990</a> describes the output fields for the <b>show security keychain</b> command. Output fields are listed in the approximate order in which they appear.

**Table 361: show security keychain Output Fields**

Field Name	Field Description	Level of Output
<b>keychain</b>	The name of the keychain in operation.	All levels
<b>Active-ID Send</b>	Number of routing protocols packets sent with the active key.	All levels
<b>Active-ID Receive</b>	Number of routing protocols packets received with the active key.	All levels
<b>Next-ID Send</b>	Number of routing protocols packets sent with the next key.	All levels
<b>Next-ID Receive</b>	Number of routing protocols packets received with the next key.	All levels
<b>Transition</b>	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
<b>Tolerance</b>	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels

Table 361: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Id</b>	Identification number configured for the current key.	<b>detail</b>
<b>Algorithm</b>	Authentication algorithm configured for the current key.	<b>detail</b>
<b>State</b>	<p>State of the current key.</p> <p>The value can be:</p> <ul style="list-style-type: none"> <li>• <b>receive</b></li> <li>• <b>send</b></li> <li>• <b>send-receive</b></li> </ul> <p>For the active key, the <b>State</b> can be <b>send-receive</b>, <b>send</b>, or <b>receive</b>. For keys that have a future start time, the <b>State</b> is <b>inactive</b>. Compare the <b>State</b> field to the <b>Mode</b> field.</p>	<b>detail</b>
<b>Option</b>	<p>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.</p> <p>The values can be:</p> <ul style="list-style-type: none"> <li>• <b>basic</b>—Based on RFC 5304.</li> <li>• <b>isis-enhanced</b>—Based on RFC 5310.</li> </ul> <p>The default value is <b>basic</b>. When you configure the <b>isis-enhanced</b> option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.</p> <p>When you configure <b>basic</b> (or do not include the <b>options</b> statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.</p> <p>Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</p>	<b>detail</b>
<b>Start-time</b>	Time that the current key became active.	<b>detail</b>

Table 361: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Mode</b>	<p>Mode of each key (Informational only.)</p> <p>The value can be</p> <ul style="list-style-type: none"> <li>• <b>receive</b></li> <li>• <b>send</b></li> <li>• <b>send-receive</b></li> </ul> <p>The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the <b>Mode</b> can be <b>send-receive</b>, <b>send</b>, or <b>receive</b>, regardless of the configured start-time. Compare the <b>Mode</b> field to the <b>State</b> field.</p>	<b>detail</b>

## Sample Output

### show security keychain brief

```

user@host> show security keychain brief
keychain          Active-ID      Next-ID      Transition  Tolerance
                  Send  Receive    Send  Receive
hakr              3     3           1     1         1d 23:58    3600

```

### show security keychain detail

```

user@host> show security keychain detail
keychain          Active-ID      Next-ID      Transition  Tolerance
                  Send  Receive    Send  Receive
hakr              3     3           1     1         1d 23:58    3600
Id 3, Algorithm hmac-md5, State send-receive, Option basic
Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
Id 1, Algorithm hmac-md5, State inactive, Option basic
Start-time Fri Aug 20 11:30:57 2010, Mode send-receive

```

## show link-management

<b>Syntax</b>	show link-management
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management peer on page 4640</a></li> <li>• <a href="#">show link-management routing on page 4642</a></li> <li>• <a href="#">show link-management statistics on page 4645</a></li> <li>• <a href="#">show link-management te-link on page 4647</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management on page 4639</a>
<b>Output Fields</b>	<a href="#">Table 362 on page 4636</a> describes the output fields for the <b>show link-management</b> command. Output fields are listed in the approximate order in which they appear.

**Table 362: show link-management Output Fields**

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: <b>Up</b> or <b>Down</b> .
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.
State	State of the control channel: <b>Up</b> or <b>Down</b> .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295.

Table 362: show link-management Output Fields (*continued*)

Field Name	Field Description
<b>Flags</b>	Code that provides information about the control channel. Currently supports only code value <b>R</b> , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
<b>TE links</b>	Traffic-engineered links that are managed by their peer.
<b>TE link name</b>	Name of the traffic-engineered link.
<b>State</b>	State of the traffic-engineered link: <b>Up</b> , <b>Down</b> , or <b>Init</b> .
<b>Local identifier</b>	Identifier of the local side of the link.
<b>Remote identifier</b>	Identifier of the remote side of the link.
<b>Local address</b>	Address of the local side of the link.
<b>Remote address</b>	Address of the remote side of the link.
<b>Encoding</b>	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> .
<b>Switching</b>	Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
<b>Total bandwidth</b>	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
<b>Available bandwidth</b>	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
<b>Name</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Down</b> .
<b>Local ID</b>	Identifier of the local side of the interface.
<b>Remote ID</b>	Identifier of the remote side of the interface.
<b>Bandwidth</b>	Bandwidth, in bps or Mbps, of the member interface.
<b>Used</b>	Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .

Table 362: show link-management Output Fields (*continued*)

Field Name	Field Description
LSP-name	LSP name.



## Sample Output

### show link-management

```

user@host> show link-management
Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    24547      24547 Up          1027      1026
TE links:
  pro4-ba

TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps
  Name          State Local ID Remote ID      Bandwidth Used  LSP-name
  so-1/0/2      Up          21271      0      155.52Mbps    No

```

## show link-management peer

<b>Syntax</b>	<code>show link-management peer</code> <code>&lt;name <i>peer-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) peer link information.
<b>Options</b>	<b>none</b> —Display all peer link information.  <b>name <i>peer-name</i></b> —(Optional) Display information for the specified peer only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 4636</a></li> <li>• <a href="#">show link-management routing on page 4642</a></li> <li>• <a href="#">show link-management statistics on page 4645</a></li> <li>• <a href="#">show link-management te-link on page 4647</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management peer on page 4641</a>
<b>Output Fields</b>	Table 363 on page 4640 describes the output fields for the <b>show link-management peer</b> command. Output fields are listed in the approximate order in which they appear.

**Table 363: show link-management peer Output Fields**

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: <b>Up</b> or <b>Down</b> .
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.

Table 363: show link-management peer Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	State of the control channel: <b>Up</b> or <b>Down</b> .
<b>TxSeqNum</b>	Sequence number of the hello message being sent to the peer. The range of values is <b>1</b> through <b>4,294,967,295</b> .
<b>RcvSeqNum</b>	Sequence number of the last hello message received from the peer. The range of values is <b>0</b> through <b>4,294,967,295</b> .
<b>Flags</b>	Code that provides information about the control channel. Currently supports only code value <b>R</b> , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
<b>TE links</b>	Traffic-engineered links that are managed by their peer.

## Sample Output

### show link-management peer

```

user@host> show link-management peer
Peer name: sonet, System identifier: 41448
State: Up, Control address: 70.70.70.70
Hello interval: 10000, Hello dead interval: 30000
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    3265           0 ConfSnd         1          0 R
TE links:
to-sonet

```

## show link-management routing

<b>Syntax</b>	show link-management routing <peer <name <i>name</i> >   te-link <name <i>name</i> >> <resource <name <i>name</i> >>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.
<b>Options</b>	<p><b>none</b>—Display all peer and traffic-engineered link information.</p> <p><b>peer &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all peers or for the specified peer only.</p> <p><b>resource &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all resources or for the specified resource only.</p> <p><b>te-link &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 4636</a></li> <li>• <a href="#">show link-management peer on page 4640</a></li> <li>• <a href="#">show link-management statistics on page 4645</a></li> <li>• <a href="#">show link-management te-link on page 4647</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management routing on page 4644</a>
<b>Output Fields</b>	Table 364 on page 4642 describes the output fields for the <b>show link-management routing</b> command. Output fields are listed in the approximate order in which they appear.

**Table 364: show link-management routing Output Fields**

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down.
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.

Table 364: show link-management routing Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	State of the control channel.
<b>TE link name</b>	Traffic-engineered link name.
<b>State</b>	State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .
<b>Local identifier</b>	Identifier of the local side of the link.
<b>Remote identifier</b>	Identifier of the remote side of the link.
<b>Local address</b>	Address of the local side of the link.
<b>Remote address</b>	Address of the remote side of the link.
<b>Encoding</b>	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
<b>Total bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
<b>Available bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
<b>Resource</b>	Forwarding adjacency LSP information.
<b>Type</b>	Type of resource. The type is always a forwarding adjacency LSP.
<b>State</b>	State of the LSP: <b>Up</b> or <b>Down</b> .
<b>System Identifier</b>	Internal identifier for the peer. The range of values is <b>0</b> through <b>64,000</b> .
<b>Total bandwidth</b>	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
<b>Traffic parameters</b>	<ul style="list-style-type: none"> <li>• <b>Encoding</b>—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b>, <b>Ethernet</b>, and <b>Packet</b>.</li> <li>• <b>Switching</b>—Type of switching that can be performed on the traffic-engineered link: <b>PSC-1</b> and <b>Packet</b>.</li> <li>• <b>Granularity</b>—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always <b>unknown</b>.</li> </ul>

## Sample Output

### show link-management routing

```
user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
Control-channel          State
fe-0/1/0.0               Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
Control-channel          State
fe-0/1/2.0               Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
Control-channel          State
so-0/2/0.0               State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
Control-channel          State
so-0/2/1.0               State

...

TE link name: __rpd:fe-0/1/0.0, State: Up
Local identifier: 2147483649, Remote identifier: 0,
Local address: 192.168.37.66, Remote address: 192.168.37.66,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
Local identifier: 2147483650, Remote identifier: 0,
Local address: 192.168.37.73, Remote address: 192.168.37.73,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
Local identifier: 2147483651, Remote identifier: 0,
Local address: 192.168.37.82, Remote address: 192.168.37.95,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: 0bps, Traffic parameters: Encoding: Packet, Switching: Packet,
Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown
```

## show link-management statistics

<b>Syntax</b>	show link-management statistics <peer <name <i>name</i> >>
<b>Release Information</b>	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display statistical information for Link Management Protocol (LMP) packets.
<b>Options</b>	<b>none</b> —Display information for all peers.  <b>peer &lt;name <i>name</i>&gt;</b> —(Optional) Display information for all peers or for the specified peer only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 4636</a></li> <li>• <a href="#">show link-management peer on page 4640</a></li> <li>• <a href="#">show link-management routing on page 4642</a></li> <li>• <a href="#">show link-management te-link on page 4647</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management statistics on page 4646</a>
<b>Output Fields</b>	Table 365 on page 4645 describes the output fields for the <b>show link-management statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 365: show link-management statistics Output Fields**

Field Name	Field Description
<b>Received packets</b>	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Received bad packets</b>	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Small packets</b>	Number of packets that are too small.
<b>Wrong protocol version</b>	Number of packets specifying the wrong LMP version.
<b>Messages for unknown peer</b>	Number of packets destined for an unknown peer.
<b>Messages for bad state</b>	Number of packets indicating a state that does not match the recipient.
<b>Stale acknowledgments</b>	Number of <b>configAck</b> and <b>LinkSummaryAck</b> packets received that have a stale message ID.

Table 365: show link-management statistics Output Fields (*continued*)

Field Name	Field Description
<b>Stale negative acknowledgments</b>	Number of <b>configNack</b> and <b>LinkSummaryNack</b> packets received that have a stale message ID.
<b>Sent packets</b>	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Retransmitted packets</b>	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Dropped packets</b>	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

## Sample Output

### show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
  Small packets: 0
  Wrong protocol version: 0
  Messages for unknown peer: 0
  Messages for bad state: 0
  Stale acknowledgments: 0
  Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1

```



## show link-management te-link

<b>Syntax</b>	show link-management te-link <brief   detail> <name <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.
<b>Options</b>	<b>none</b> —Display information for all traffic-engineered links.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>name <i>name</i></b> —(Optional) Display information for the specified traffic-engineered link only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 4636</a></li> <li>• <a href="#">show link-management peer on page 4640</a></li> <li>• <a href="#">show link-management routing on page 4642</a></li> <li>• <a href="#">show link-management statistics on page 4645</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management te-link on page 4648</a>
<b>Output Fields</b>	<a href="#">Table 366 on page 4647</a> describes the output fields for the <b>show link-management te-link</b> command. Output fields are listed in the approximate order in which they appear.

**Table 366: show link-management te-link Output Fields**

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> .

Table 366: show link-management te-link Output Fields (*continued*)

Field Name	Field Description
<b>Switching</b>	Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
<b>Total bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
<b>Available Bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
<b>Name</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Down</b> .
<b>Local ID</b>	Identifier of the local side of the interface.
<b>Remote ID</b>	Identifier of the remote side of the interface.
<b>Bandwidth</b>	Bandwidth, in bps or Mbps, of the member interface.
<b>Used</b>	Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .
<b>LSP-name</b>	LSP name.

## Sample Output

### show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-bd  Dn      43077      0             0bps No
TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-be  Up      43076      0            10Mbps Yes  e2elasp-bf

```

## show mpls call-admission-control

<b>List of Syntax</b>	<a href="#">Syntax on page 4649</a> <a href="#">Syntax (EX Series Switches) on page 4649</a>
<b>Syntax</b>	<pre>show mpls call-admission-control &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-name&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show mpls call-admission-control &lt;lsp-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.
<b>Options</b>	<p><b>none</b>—Display CAC information for all LSPs.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>lsp-name</i></b>—(Optional) Display CAC information for the specified LSP only.</p>
<b>Additional Information</b>	The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls call-admission-control on page 4650</a>
<b>Output Fields</b>	<a href="#">Table 367 on page 4649</a> describes the output fields for the <b>show mpls call-admission-control</b> command. Output fields are listed in the approximate order in which they appear.

**Table 367: show mpls call-admission-control Output Fields**

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at <b>ct0</b> ) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.
LSP name	LSP pathname.

Table 367: show mpls call-admission-control Output Fields (*continued*)

Field Name	Field Description
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

## Sample Output

### show mpls call-admission-control

```

user@host# show mpls call-admission-control

LSP name: pro1-be
*Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary  prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby  sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

## show mpls cspf

<b>List of Syntax</b>	<a href="#">Syntax on page 4651</a> <a href="#">Syntax (EX Series Switches) on page 4651</a>
<b>Syntax</b>	<pre>show mpls cspf &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls cspf
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.
<b>Options</b>	<p><b>none</b>—Display MPLS CSFP statistics.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls cspf on page 4652</a>
<b>Output Fields</b>	<a href="#">Table 368 on page 4651</a> describes the output fields for the <b>show mpls cspf</b> command. Output fields are listed in the approximate order in which they appear.

**Table 368: show mpls cspf Output Fields**

Field Name	Field Description
<b>Queue length</b>	Number of LSPs queued for automatic path computation.
<b>current</b>	Current queue length.
<b>maximum</b>	Maximum queue length (high-water mark).
<b>dequeued</b>	Number of aborted computation attempts.
<b>Paths</b>	Counters for label-switched path computations.
<b>total</b>	Sum of the next four fields.
<b>successful</b>	Number of path computations that were successfully completed.
<b>no route</b>	Number of path computations that failed because the destination is unreachable.

Table 368: show mpls cspf Output Fields (*continued*)

Field Name	Field Description
<b>Sys Error</b>	Number of path computations that failed because of lack of memory.
<b>CSPFs</b>	Total number of CSPF computations. A single path might require multiple CSPF computations.
<b>Time</b>	Time, in seconds, required to perform the label-switched path computation.
<b>Total</b>	Total amount of time consumed by the CSPF path computation algorithm.
<b>CSPFs</b>	Total number of CSPF computations.
<b>Avg per CSPF</b>	Average amount of time required for each CSPF computation.
<b>% of rpd</b>	Percentage of routing process CPU used in the CSPF computation.

## Sample Output

show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length  current      maximum    dequeued
              0           0           0
Paths         total      successful    no route   sys error   CSPFs
              0           0           0           0           0
Time (secs)   total      CSPFs    avg per CSPF    % of rpd
              0.000000    0.000000    0.000000    0.0000

```

## show mpls diffserv-te

<b>List of Syntax</b>	<a href="#">Syntax on page 4653</a> <a href="#">Syntax (EX Series Switches) on page 4653</a>
<b>Syntax</b>	show mpls diffserve-te <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show mpls diffserve-te
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.
<b>Options</b>	<b>none</b> —Display DiffServ classes and priorities used by MPLS LSPs.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls diffserv-te on page 4654</a>
<b>Output Fields</b>	<a href="#">Table 369 on page 4653</a> describes the output fields for the <b>show mpls diffserv-te</b> command. Output fields are listed in the approximate order in which they appear.

**Table 369: show mpls diffserv-te Output Fields**

Field Name	Field Description
<b>Bandwidth model</b>	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
<b>TE class</b>	DiffServ traffic engineering class.
<b>Traffic class</b>	MPLS class type that corresponds to the DiffServ traffic engineering class: <ul style="list-style-type: none"> <li>• <b>ct0</b>—Best effort</li> <li>• <b>ct1</b>—Assured forwarding</li> <li>• <b>ct2</b>—Expedited forwarding</li> <li>• <b>ct3</b>—Network control</li> </ul>
<b>Priority</b>	MPLS preemption priority for this class type, a value from 0 through 7. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

## Sample Output

`show mpls diffserv-te`

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class    Traffic class    Priority
te0         ct0              3
te1         ct1              2
```



## show route forwarding-table

<b>Syntax</b>	<pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;ccc ccc-interface-name&gt; &lt;destination&gt; &lt;family family-name&gt; &lt;label label&gt; &lt;matching ip_prefix&gt; &lt;multicast&gt; &lt;vpn vpn&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.
<b>Options</b>	<p><b>none</b>—Display the routes in the forwarding table.</p> <p><b>detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>ccc</b>—(Optional) Display the specified circuit cross-connect interface name for entries to match.</p> <p><b>destination</b>—(Optional) Display the destination prefix.</p> <p><b>family family-name</b>—(Optional) Display routing table entries for the specified family: <b>ethernet-switching, inet, inet6, iso, mpls, vlan classification</b>.</p> <p><b>label label</b>—(Optional) Display route entries for the specified label name.</p> <p><b>matching ip_prefix</b>—(Optional) Display route entries for the specified IP prefix.</p> <p><b>multicast</b>—(Optional) Display route entries for multicast routes.</p> <p><b>vpn vpn</b>—(Optional) Display route entries for the specified VPN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Example: Configuring MPLS on EX Series Switches</i></li> <li><i>Configuring MPLS on Provider Switches (CLI Procedure)</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show route forwarding-table on page 4657</a></p> <p><a href="#">show route forwarding-table summary on page 4658</a></p> <p><a href="#">show route forwarding-table extensive on page 4658</a></p> <p><a href="#">show route forwarding-table ccc on page 4660</a></p> <p><a href="#">show route forwarding-table family (MPLS) on page 4660</a></p>

[show route forwarding-table family \(IPv6\) on page 4660](#)  
[show route forwarding-table label on page 4661](#)  
[show route forwarding-table matching on page 4661](#)  
[show route forwarding-table multicast on page 4661](#)

**Output Fields** Table 288 on page 3147 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used instead of the **extensive** keyword.

Table 370: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
<b>Routing table</b>	Name of the routing table (for example, <b>inet</b> , <b>inet6</b> , <b>mpls</b> ).	All levels
<b>Address family</b>	Address family (for example, <b>IP</b> , <b>IPv6</b> , <b>ISO</b> , <b>MPLS</b> ).	All levels
<b>Destination</b>	Destination of the route.	<b>detail</b> , <b>extensive</b>
<b>Route Type (Type)</b>	How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li><b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li><b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li><b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li><b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li><b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li><b>ignore (ignr)</b>—Ignore this route.</li> <li><b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li><b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li><b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul>	All levels
<b>Route reference (RtRef)</b>	Number of routes to reference.	<b>detail</b> , <b>extensive</b>
<b>Flags</b>	Route type flags: <ul style="list-style-type: none"> <li><b>none</b>—No flags are enabled.</li> <li><b>accounting</b>—Route has accounting enabled.</li> <li><b>cached</b>—Cache route.</li> <li><b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li><b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li><b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li><b>static</b>—Static route.</li> </ul>	<b>extensive</b>
<b>Nexthop</b>	IP address of the next hop to the destination.	<b>detail</b> , <b>extensive</b>

Table 370: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Next hop type (Type)</b>	<p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcr)</b>—Regular multicast next hop</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b> —Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul>	<b>detail, extensive</b>
<b>Index</b>	Software index of the next hop that is used to route the traffic for a given prefix.	<b>detail, extensive none</b>
<b>Route interface-index</b>	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	<b>extensive</b>
<b>Reference (NhRef)</b>	Number of routes that refer to this next hop.	<b>none detail, extensive</b>
<b>Next-hop interface (Netif)</b>	Interface used to reach the next hop.	<b>none detail, extensive</b>
<b>Alternate forward nh index</b>	Index number of the alternate next hop interface. Seen with <b>multicast</b> option only.	<b>extensive</b>
<b>Next-hop L3 Interface</b>	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the <b>multicast</b> option.	<b>extensive</b>
<b>Next-hop L2 Interfaces</b>	The next hop layer 2 interfaces. Seen with <b>multicast</b> option only.	<b>extensive</b>

## Sample Output

### show route forwarding-table

```

user@switch> show route forwarding-table

Routing table: default.inet

```

Internet:							
Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	user	2	0:12:f2:21:cf:0	ucst	333	5	me0.0
default	perm	0		rjct	36	2	
0.0.0.0/32	perm	0		dscd	34	1	
2.2.2.0/24	intf	0		rslv	1309	1	ae0.0
2.2.2.0/32	dest	0	2.2.2.0	recv	1307	1	ae0.0
2.2.2.1/32	dest	0	0:21:59:cc:89:c0	ucst	1320	1	ae0.0
2.2.2.2/32	intf	0	2.2.2.2	loc1	1308	2	
2.2.2.2/32	dest	0	2.2.2.2	loc1	1308	2	
2.2.2.255/32	dest	0	2.2.2.255	bcst	1306	1	ae0.0
3.3.3.0/24	intf	0		rslv	1313	1	ae1.0
3.3.3.0/32	dest	0	3.3.3.0	recv	1311	1	ae1.0
3.3.3.1/32	intf	0	3.3.3.1	loc1	1312	2	
3.3.3.1/32	dest	0	3.3.3.1	loc1	1312	2	
3.3.3.2/32	dest	0	0:21:59:cc:89:c1	ucst	1321	24	ae1.0
3.3.3.255/32	dest	0	3.3.3.255	bcst	1310	1	ae1.0
4.4.4.0/24	user	0	3.3.3.2	ucst	1321	24	ae1.0
8.8.8.8/32	user	0	3.3.3.2	ucst	1321	24	ae1.0
9.9.9.9/32	intf	0	9.9.9.9	loc1	1280	1	
10.10.10.10/32	user	0	3.3.3.2	ucst	1321	24	ae1.0
10.93.8.0/21	intf	0		rslv	323	1	me0.0
10.93.8.0/32	dest	0	10.93.8.0	recv	321	1	me0.0
10.93.13.238/32	intf	0	10.93.13.238	loc1	322	2	
10.93.13.238/32	dest	0	10.93.13.238	loc1	322	2	
10.93.15.254/32	dest	0	0:12:f2:21:cf:0	ucst	333	5	me0.0
10.93.15.255/32	dest	0	10.93.15.255	bcst	320	1	me0.0
14.14.14.0/24	ifdn	0		rslv	1319	1	ge-0/0/25.0
14.14.14.0/32	iddn	0	14.14.14.0	recv	1317	1	ge-0/0/25.0
14.14.14.2/32	user	0		rjct	36	2	
14.14.14.2/32	intf	0	14.14.14.2	loc1	1318	2	
14.14.14.2/32	iddn	0	14.14.14.2	loc1	1318	2	
14.14.14.255/32	iddn	0	14.14.14.255	bcst	1316	1	ge-0/0/25.0
224.0.0.0/4	perm	1		mdsc	35	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	31	3	
224.0.0.5/32	user	1	224.0.0.5	mcst	31	3	
255.255.255.255/32	perm	0		bcst	32	1	

### show route forwarding-table summary

```
user@switch> show route forwarding-table summary
```

```
Routing table: default.inet
```

```
Internet:
```

```

user:          6 routes
perm:          5 routes
intf:          8 routes
dest:         12 routes
ifdn:          1 routes
iddn:          3 routes
```

### show route forwarding-table extensive

```
user@switch> show route forwarding-table summary
```

```
Routing table: default.inet [Index 0]
```

```
Internet:
```

```
Destination: default
```

```
Route type: user
```

```
Route reference: 2
```

```
Route interface-index: 0
```

```

Flags: sent to PFE, rt nh decoupled
Nexthop: 0:12:f2:21:cf:0
Next-hop type: unicast          Index: 333      Reference: 5
Next-hop interface: me0.0

Destination: default
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: none
Next-hop type: reject          Index: 36       Reference: 2

Destination: 0.0.0.0/32
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Next-hop type: discard         Index: 34       Reference: 1

Destination: 2.2.2.0/24
Route type: interface
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Next-hop type: resolve         Index: 1309     Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.0/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.0
Next-hop type: receive         Index: 1307     Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.1/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 0:21:59:cc:89:c0
Next-hop type: unicast         Index: 1320     Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.2/32
Route type: interface
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308     Reference: 2

Destination: 2.2.2.2/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: none
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308     Reference: 2

Destination: 2.2.2.255/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.255
Next-hop type: broadcast       Index: 1306     Reference: 1
Next-hop interface: ae0.0

```

**show route forwarding-table ccc**

```

user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
ge-0/0/0.10      (CCC) user    0 3.3.3.2          Push 300112 1343    2 ae1.0

```

**show route forwarding-table family (MPLS)**

```

user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm    0
0                user    0                recv  49    3
1                user    0                recv  49    3
2                user    0                recv  49    3
299776           user    0                Pop   1334   2 ge-0/0/0.10
299792           user    0                Pop   1339   2 ge-0/0/0.14
299808           user    0                Pop   1341   2 ge-0/0/0.2
299824           user    0                Pop   1344   2 ge-0/0/0.11
299840           user    0                Pop   1345   2 ge-0/0/0.13
299856           user    0                Pop   1346   2 ge-0/0/0.18
299872           user    0                Pop   1347   2 ge-0/0/0.16
299888           user    0                Pop   1348   2 ge-0/0/0.7
299904           user    0                Pop   1349   2 ge-0/0/0.20
299920           user    0                Pop   1350   2 ge-0/0/0.19
299936           user    0                Pop   1351   2 ge-0/0/0.17
299952           user    0                Pop   1352   2 ge-0/0/0.9
299968           user    0                Pop   1353   2 ge-0/0/0.1
299984           user    0                Pop   1354   2 ge-0/0/0.12
300000           user    0                Pop   1355   2 ge-0/0/0.8
300016           user    0                Pop   1356   2 ge-0/0/0.4
300032           user    0                Pop   1357   2 ge-0/0/0.5
300048           user    0                Pop   1358   2 ge-0/0/0.3
300064           user    0                Pop   1359   2 ge-0/0/0.15
ge-0/0/0.1       (CCC) user    0 3.3.3.2          Push 300064 1340    2 ae1.0
ge-0/0/0.2       (CCC) user    0 3.3.3.2          Push 299872 1328    2 ae1.0
ge-0/0/0.3       (CCC) user    0 3.3.3.2          Push 299792 1323    2 ae1.0
ge-0/0/0.4       (CCC) user    0 3.3.3.2          Push 300016 1337    2 ae1.0
ge-0/0/0.5       (CCC) user    0 3.3.3.2          Push 299824 1325    2 ae1.0
ge-0/0/0.7       (CCC) user    0 3.3.3.2          Push 299920 1331    2 ae1.0
ge-0/0/0.8       (CCC) user    0 3.3.3.2          Push 299840 1326    2 ae1.0
ge-0/0/0.9       (CCC) user    0 3.3.3.2          Push 299888 1329    2 ae1.0
ge-0/0/0.10      (CCC) user    0 3.3.3.2          Push 300112 1343    2 ae1.0
ge-0/0/0.11      (CCC) user    0 3.3.3.2          Push 299776 1322    2 ae1.0
ge-0/0/0.12      (CCC) user    0 3.3.3.2          Push 299952 1333    2 ae1.0
ge-0/0/0.13      (CCC) user    0 3.3.3.2          Push 300096 1342    2 ae1.0
ge-0/0/0.14      (CCC) user    0 3.3.3.2          Push 299984 1335    2 ae1.0
ge-0/0/0.15      (CCC) user    0 3.3.3.2          Push 299936 1332    2 ae1.0
ge-0/0/0.16      (CCC) user    0 3.3.3.2          Push 299808 1324    2 ae1.0
ge-0/0/0.17      (CCC) user    0 3.3.3.2          Push 300000 1336    2 ae1.0
ge-0/0/0.18      (CCC) user    0 3.3.3.2          Push 300032 1338    2 ae1.0
ge-0/0/0.19      (CCC) user    0 3.3.3.2          Push 299904 1330    2 ae1.0
ge-0/0/0.20      (CCC) user    0 3.3.3.2          Push 299856 1327    2 ae1.0

```

**show route forwarding-table family (IPv6)**

```

user@switch> show route forwarding-table family inet6

```

```

Routing table: default.inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  44    1
::/128           perm  0                dscd  42    1
ff00::/8         perm  0                mdsc  43    1
ff02::1/128      perm  0 ff02::1          mcst  39    1

```

```

Routing table: default-switch.inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  530   1
::/128           perm  0                dscd  528   1
2:1::3a00/312    user  0                indr  131070 2
                  comp  572   1
2:1::3a82/320     user  0                indr  131071 3
                  comp  573   1
2:1::3af0/320     user  0                indr  131071 3
                  comp  573   1
2:1:0:ff00::/56   user  0                mdsc  529   2
ff00::/8         perm  0                mdsc  529   2
ff02::1/128      perm  0 ff02::1          mcst  526   1

```

```

Routing table: __master.anon__.inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  554   1
::/128           perm  0                dscd  552   1
ff00::/8         perm  0                mdsc  553   1
ff02::1/128      perm  0 ff02::1          mcst  550   1

```

### show route forwarding-table label

```
user@switch> show route forwarding-table label 29976
```

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
299776           user  0                Pop   1334   2 ge-0/0/0.10

```

### show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3
```

```

Routing table: default.inet
Internet:

```

### show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast
```

```

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
224.0.0.0/4       perm  1                mdsc  35    1
224.0.0.1/32      perm  0 224.0.0.1          mcst  31    3
224.0.0.5/32      user  1 224.0.0.5          mcst  31    3

```

```

Routing table: __master.anon__.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
224.0.0.0/4       perm  0                mdsc  1289   1

```

```
224.0.0.1/32      perm      0 224.0.0.1      mcst 1285      1
```

Routing table: default.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	



## show mpls interface

<b>List of Syntax</b>	<a href="#">Syntax on page 4663</a> <a href="#">Syntax (EX Series Switches) on page 4663</a>
<b>Syntax</b>	show mpls interface <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show mpls interface
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.
<b>Options</b>	<b>none</b> —Display information about MPLS-enabled interfaces.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Additional Information</b>	MPLS is enabled on an interface when the interface is configured with both the <b>set protocol mpls interface <i>interface-name</i></b> and <b>set interface <i>interface-name</i> unit 0 family mpls</b> statements.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls interface on page 4664</a>
<b>Output Fields</b>	<a href="#">Table 371 on page 4663</a> describes the output fields for the <b>show mpls interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 371: show mpls interface Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Dn</b> (down).
<b>Administrative groups</b>	Administratively assigned colors of the link.
<b>Maximum labels</b>	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the <b>maximum-labels</b> statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.

Table 371: show mpls interface Output Fields (*continued*)

Field Name	Field Description
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the <b>protection-revert-time</b> statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.
Always mark connection protection tlv	Enabled or Disabled: Enabled indicates that the <b>always-mark-connection-protection-tlv</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the <b>switch-away-lsps</b> statement must be configured.
Switch away lsps	Enabled or Disabled: Enabled indicates that the <b>switch-away-lsps</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

## Sample Output

### show mpls interface

```

user@host> show mpls interface

Interface: ge-0/2/1.57
  State: Up
  Administrative group: <none>
  Maximum labels: 5
  Static protection revert time: 5 seconds
  Always mark connection protection tlv: Disabled
  Switch away lsps : Disabled

```

## show mpls lsp

**List of Syntax**    [Syntax on page 4665](#)  
                          [Syntax \(EX Series Switches\) on page 4665](#)

**Syntax**    show mpls lsp  
                  <brief | detail | extensive | terse>  
                  <autobandwidth>  
                  <bidirectional | unidirectional>  
                  <bypass>  
                  <count-active-routes>  
                  <defaults>  
                  <descriptions>  
                  <down | up>  
                  <externally-controlled>  
                  <externally-provisioned>  
                  <logical-system (all | *logical-system-name*)>  
                  <lsp-type>  
                  <name *name*>  
                  <p2mp>  
                  <statistics>  
                  <transit>

**Syntax (EX Series Switches)**    show mpls lsp  
                  <brief | detail | extensive | terse>  
                  <bidirectional | unidirectional>  
                  <bypass>  
                  <descriptions>  
                  <down | up>  
                  <externally-controlled>  
                  <externally-provisioned>  
                  <lsp-type>  
                  <name *name*>  
                  <p2mp>  
                  <statistics>  
                  <transit>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                  **defaults** option added in Junos OS Release 8.5.  
                  Command introduced in Junos OS Release 9.5 for EX Series switches.  
                  **autobandwidth** option added in Junos OS Release 11.4.  
                  **externally-controlled** option added in Junos OS Release 12.3.  
                  **externally-provisioned** option added in Junos OS Release 13.3.  
                  Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.

**Description**    Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options**    **none**—Display standard information about all configured and active dynamic MPLS LSPs.  
                  **brief | detail | extensive | terse**—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

**autobandwidth**—(Optional) Display automatic bandwidth information. This option is explained separately (see [show mpls lsp autobandwidth](#)).

**bidirectional | unidirectional**—(Optional) Display bidirectional or unidirectional LSP information, respectively.

**bypass**—(Optional) Display LSPs used for protecting other LSPs.

**count-active-routes**—(Optional) Display active routes for LSPs.

**defaults**—(Optional) Display the MPLS LSP default settings.

**descriptions**—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

**down | up**—(Optional) Display only LSPs that are inactive or active, respectively.

**externally-controlled**—(Optional) Display the LSPs that are under the control of an external Path Computation Element (PCE).

**externally-provisioned**—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

**name *name***—(Optional) Display information about the specified LSP or group of LSPs.

**p2mp**—(Optional) Display information about point-to-multipoint LSPs.

**statistics**—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.



**NOTE:** If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored. (Bypass LSPs are not supported on QFX Series switches.)

When used with the `bypass` option (`show mpls lsp bypass statistics`), display statistics for the traffic that flows only through the bypass LSP.

**transit**—(Optional) Display LSPs transiting this routing device.

**Required Privilege Level** view

**Related Documentation**

- [clear mpls lsp on page 4572](#)
- [show mpls lsp autobandwidth on page 4681](#)

**List of Sample Output**

- [show mpls lsp defaults on page 4674](#)
- [show mpls lsp descriptions on page 4674](#)
- [show mpls lsp detail on page 4674](#)
- [show mpls lsp extensive on page 4675](#)
- [show mpls lsp ingress extensive on page 4676](#)
- [show mpls lsp extensive \(automatic bandwidth adjustment enabled\) on page 4677](#)
- [show mpls lsp p2mp on page 4678](#)
- [show mpls lsp p2mp detail on page 4678](#)
- [show mpls lsp detail count-active-routes on page 4679](#)
- [show mpls lsp statistics extensive on page 4679](#)

**Output Fields** [Table 372 on page 4667](#) describes the output fields for the `show mpls lsp` command. Output fields are listed in the approximate order in which they appear.

**Table 372: show mpls lsp Output Fields**

Field Name	Field Description	Level of Output
<b>Ingress LSP</b>	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
<b>Egress LSP</b>	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
<b>Transit LSP</b>	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
<b>P2MP name</b>	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format <code>identifier:vpls:router-id:routing-instance-name</code> . The <i>identifier</i> is automatically generated by Junos OS.	All levels

Table 372: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>P2MP branch count</b>	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
<b>P</b>	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
<b>address</b>	( <b>detail</b> and <b>extensive</b> ) Destination (egress routing device) of the LSP.	<b>detail extensive</b>
<b>To</b>	Destination (egress routing device) of the session.	<b>brief</b>
<b>From</b>	Source (ingress routing device) of the session.	<b>brief detail</b>
<b>State</b>	State of the LSP handled by this RSVP session: <b>Up</b> , <b>Dn</b> (down), or <b>Restart</b> .	<b>brief detail</b>
<b>Active Route</b>	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table ( <b>mpls.0</b> ).	<b>detail extensive</b>
<b>Rt</b>	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the routing table is the primary MPLS table ( <b>mpls.0</b> ).	<b>brief</b>
<b>P</b>	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	<b>brief</b>
<b>ActivePath</b>	(Ingress LSP) Name of the active path: <b>Primary</b> or <b>Secondary</b> .	<b>detail extensive</b>
<b>LSPname</b>	Name of the LSP.	<b>brief detail</b>
<b>Statistics</b>	Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).	<b>extensive</b>
<b>Aggregate statistics</b>	Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the <b>clear mpls lsp statistics</b> command.	<b>extensive</b>
<b>Packets</b>	Displays the number of packets transmitted over the LSP.	<b>brief extensive</b>
<b>Bytes</b>	Displays the number of bytes transmitted over the LSP.	<b>brief extensive</b>
<b>DiffServInfo</b>	Type of LSP: multiclass LSP ( <b>multiclass diffServ-TE LSP</b> ) or Differentiated-Services-aware traffic engineering LSP ( <b>diffServ-TE LSP</b> ).	<b>detail</b>
<b>LSPtype</b>	Type of LSP: static <b>Static configured</b> or dynamic <b>Dynamic configured</b> . Also indicates if the LSP is a <b>Penultimate hop popping</b> LSP or an <b>Ultimate hop popping</b> LSP.	<b>detail extensive</b>
<b>Bypass</b>	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels

Table 372: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>LSPpath</b>	Indicates whether the RSVP session is for the primary or secondary LSP path. <b>LSPpath</b> can be either <b>primary</b> or <b>secondary</b> and can be displayed on the ingress, egress, and transit routing devices.	<b>detail</b>
<b>Bidir</b>	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
<b>Bidirectional</b>	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
<b>FastReroute desired</b>	Fast reroute has been requested by the ingress routing device.	<b>detail</b>
<b>Link protection desired</b>	Link protection has been requested by the ingress routing device.	<b>detail</b>
<b>Node/Link protection desired</b>	Link protection has been requested by the ingress routing device.	<b>detail extensive</b>
<b>LoadBalance</b>	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: <b>Most-fill</b> , <b>Least-fill</b> , or <b>Random</b> .	<b>detail extensive</b>
<b>Signal type</b>	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: <b>DS0</b> , <b>DS3</b> , <b>STS-1</b> , <b>STM-1</b> , or <b>STM-4</b> .	All levels
<b>Encoding type</b>	LSP encoding type: <b>Packet</b> , <b>Ethernet</b> , <b>PDH</b> , <b>SDH/SONET</b> , <b>Lambda</b> , or <b>Fiber</b> .	All levels
<b>Switching type</b>	Type of switching on the links needed for the LSP: <b>Fiber</b> , <b>Lambda</b> , <b>Packet</b> , <b>TDM</b> , or <b>PSC-1</b> .	All levels
<b>GPID</b>	Generalized Payload Identifier (identifier of the payload carried by an LSP): <b>HDLC</b> , <b>Ethernet</b> , <b>IPv4</b> , <b>PPP</b> , or <b>Unknown</b> .	All levels
<b>Protection</b>	Configured protection capability desired for the LSP: <b>Extra</b> , <b>Enhanced</b> , <b>none</b> , <b>One plus one</b> , <b>One to one</b> , or <b>Shared</b> .	All levels
<b>Upstream label in</b>	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
<b>Upstream label out</b>	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels
<b>Suggested label received</b>	(Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent.	All levels
<b>Suggested label sent</b>	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
<b>Autobandwidth</b>	(Ingress LSP) The LSP is performing autobandwidth allocation.	<b>detail extensive</b>
<b>MinBW</b>	(Ingress LSP) Configured minimum value of the LSP, in bps.	<b>detail extensive</b>

Table 372: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>MaxBW</b>	(Ingress LSP) Configured maximum value of the LSP, in bps.	<b>detail extensive</b>
<b>Dynamic MinBW</b>	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	<b>detail extensive</b>
<b>Adjustment Timer</b>	(Ingress LSP) Configured value for the <b>adjust-timer</b> statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	<b>detail extensive</b>
<b>Adjustment Threshold</b>	(Ingress LSP) Configured value for the <b>adjust-threshold</b> statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	<b>detail extensive</b>
<b>Time for Next Adjustment</b>	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	<b>detail extensive</b>
<b>Time of Last Adjustment</b>	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	<b>detail extensive</b>
<b>Max AvgBW util</b>	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	<b>detail extensive</b>
<b>Overflow limit</b>	(Ingress LSP) Configured value of the threshold overflow limit.	<b>detail extensive</b>
<b>Overflow sample count</b>	(Ingress LSP) Current value for the overflow sample count.	<b>detail extensive</b>
<b>Bandwidth Adjustment in <i>nnn</i> second(s)</b>	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	<b>detail extensive</b>
<b>Underflow limit</b>	(Ingress LSP) Configured value of the threshold underflow limit.	<b>detail extensive</b>
<b>Underflow sample count</b>	(Ingress LSP) Current value for the underflow sample count.	<b>detail extensive</b>
<b>Underflow Max AvgBW</b>	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	<b>detail extensive</b>
<b>Active path indicator</b>	(Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path.  *Primary long Standby short	<b>detail extensive</b>
<b>Primary</b>	(Ingress LSP) Name of the primary path.	<b>detail extensive</b>
<b>Secondary</b>	(Ingress LSP) Name of the secondary path.	<b>detail extensive</b>



Table 372: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Standby</b>	(Ingress LSP) Name of the path in standby mode.	<b>detail extensive</b>
<b>State</b>	(Ingress LSP) State of the path: <b>Up</b> or <b>Dn</b> (down).	<b>detail extensive</b>
<b>COS</b>	(Ingress LSP) Class-of-service value.	<b>detail extensive</b>
<b>Bandwidth per class</b>	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	<b>detail extensive</b>
<b>Priorities</b>	(Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	<b>detail extensive</b>
<b>OptimizeTimer</b>	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	<b>detail extensive</b>
<b>SmartOptimizeTimer</b>	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	<b>detail extensive</b>
<b>Reoptimization in xxx seconds</b>	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	<b>detail extensive</b>
<b>Computed ERO (S [L] denotes strict [loose] hops)</b>	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict ( <b>S</b> ) or loose ( <b>L</b> ).	<b>detail extensive</b>
<b>CSPF metric</b>	(Ingress LSP) Constrained Shortest Path First metric for this path.	<b>detail extensive</b>

Table 372: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Received RRO</b>	<p>(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If <b>Received RRO</b> is different from <b>Computed ERO</b>, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> <li>• <b>0x01</b>—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the <b>SESSION_ATTRIBUTE</b> object of the corresponding Path message.</li> <li>• <b>0x02</b>—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).</li> <li>• <b>0x03</b>—Combination of <b>0x01</b> and <b>0x02</b>.</li> <li>• <b>0x04</b>—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.</li> <li>• <b>0x08</b>—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the <b>Local protection available</b> bit is set but the <b>Node protection</b> bit is cleared.</li> <li>• <b>0x09</b>—Detour is established. Combination of <b>0x01</b> and <b>0x08</b>.</li> <li>• <b>0x10</b>—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted.</li> <li>• <b>0x20</b>—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently.</li> <li>• <b>0xb</b>—Detour is in use. Combination of <b>0x01</b>, <b>0x02</b>, and <b>0x08</b>.</li> </ul>	<b>detail extensive</b>
<b>Index number</b>	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	<b>extensive</b>
<b>Date</b>	(Ingress LSP) Date of the LSP event.	<b>extensive</b>
<b>Time</b>	(Ingress LSP) Time of the LSP event.	<b>extensive</b>
<b>Event</b>	(Ingress LSP) Description of the LSP event.	<b>extensive</b>
<b>Created</b>	(Ingress LSP) Date and time the LSP was created.	<b>extensive</b>
<b>Resv style</b>	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).	<b>brief detail extensive</b>
<b>Labelin</b>	Incoming label for this LSP.	<b>brief detail</b>
<b>Labelout</b>	Outgoing label for this LSP.	<b>brief detail</b>

Table 372: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>LSPname</b>	Name of the LSP.	<b>brief detail</b>
<b>Time left</b>	Number of seconds remaining in the lifetime of the reservation.	<b>detail</b>
<b>Since</b>	Date and time when the RSVP session was initiated.	<b>detail</b>
<b>Tspec</b>	Sender's traffic specification, which describes the sender's traffic parameters.	<b>detail</b>
<b>Port number</b>	Protocol ID and sender or receiver port used in this RSVP session.	<b>detail</b>
<b>PATH rcvfrom</b>	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	<b>detail</b>
<b>PATH sentto</b>	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	<b>detail</b>
<b>RESV rcvfrom</b>	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the <b>PATH rcvfrom</b> field, indicates that the RSVP negotiation is complete.	<b>detail</b>
<b>Record route</b>	Recorded route for the session, taken from the record route object.	<b>detail</b>
<b>Soft preempt</b>	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	<b>detail</b>
<b>Soft preemption pending</b>	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	<b>detail</b>
<b>MPLS-TE LSP Defaults</b>	Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> <li>• <b>LSP Holding Priority</b>—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully.</li> <li>• <b>LSP Setup Priority</b>—Determines whether a new LSP that preempts an existing LSP can be established.</li> <li>• <b>Hop Limit</b>—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress).</li> <li>• <b>Bandwidth</b>—Specifies the bandwidth in bits per second for the LSP.</li> <li>• <b>LSP Retry Timer</b>—Length of time in seconds that the ingress router waits between attempts to establish the primary path.</li> </ul>	<b>defaults</b>

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

## Sample Output

### show mpls lsp defaults

```
user@host> show mpls lsp defaults
MPLS-TE LSP Defaults
  LSP Holding Priority      0
  LSP Setup Priority       7
  Hop Limit                255
  Bandwidth                 0
  LSP Retry Timer          30 seconds
```

### show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To          LSP name          Description
10.0.0.195  to-sanjose                 to-sanjose-desc
10.0.0.195  to-sanjose-other-desc      other-desc
Total 2 displayed, Up 2, Down 0
```

### show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
    10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

Transit LSP: 0 sessions  
Total 0 displayed, Up 0, Down 0

### show mpls lsp extensive

user@host> show mpls lsp extensive  
Ingress LSP: 4 sessions

```

1.1.1.1
  From: 3.3.3.3, State: Up, ActiveRoute: 0, LSPname: m120b-to-mx960
  ActivePath: DEFAULT (primary)
  FastReroute desired
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary  DEFAULT      State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 310)
10.0.35.5 S 10.0.15.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.34.4(flag=1) 10.0.14.1
50 Sep 13 16:08:19.712 Record Route: 10.0.35.5(flag=1) 10.0.15.1
49 Sep 13 16:08:16.720 Record Route: 10.0.34.4(flag=1) 10.0.14.1
48 Sep 13 16:08:16.699 Fast-reroute Detour Up
47 Sep 13 16:08:13.702 Record Route: 10.0.34.4 10.0.14.1
46 Sep 13 16:08:13.702 Up
45 Sep 13 16:08:13.672 Originate make-before-break call
44 Sep 13 16:08:13.672 CSPF: computation result accepted 10.0.34.4 10.0.14.1

43 Sep 13 16:08:13.672 Selected as active path
42 Sep 13 16:08:13.672 Make-before-break: Switched to new instance
41 Sep 13 16:08:01.685 Pending path switchover, skip CSPF run[3 times]
40 Sep 13 16:06:33.910 Deselected as active
39 Sep 13 16:06:33.910 Pending path switchover, skip CSPF run

38 Sep 13 16:06:19.521 Record Route: 10.0.35.5 10.0.15.1
37 Sep 13 16:06:19.518 ResvTear received
36 Sep 13 16:06:19.518 Fast-reroute Detour Down
35 Sep 13 16:06:16.676 Record Route: 10.0.35.5(flag=1) 10.0.15.1
34 Sep 13 16:06:13.670 Record Route: 10.0.35.5 10.0.15.1
33 Sep 13 16:06:13.670 Up
32 Sep 13 16:06:13.569 Pending path switchover, skip CSPF run

31 Sep 13 16:06:13.569 CSPF: link down/deleted:
10.0.34.3(3.3.3.3:79)(m120-b-re1.00/3.3.3.3)->0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)

30 Sep 13 16:06:13.552 Pending path switchover, skip CSPF run

29 Sep 13 16:06:13.552 CSPF: link down/deleted:
0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)->0.0.0.0(4.4.4.4:0)(m10i-a-re0.00/4.4.4.4)

28 Sep 13 16:06:13.549 Originate make-before-break call
27 Sep 13 16:06:13.549 CSPF: computation result accepted 10.0.35.5 10.0.15.1

26 Sep 13 16:06:13.548 Tunnel local repaired
25 Sep 13 16:06:13.546 Record Route: 10.0.23.2 10.0.12.1
24 Sep 13 16:06:13.546 10.0.34.3: Tunnel local repaired
23 Sep 13 16:06:13.546 10.0.34.3: Down
22 Sep 13 16:03:46.842 Fast-reroute Detour Up

```

```

21 Sep 13 16:03:42.730 Record Route: 10.0.34.4(flag=1) 10.0.14.1
20 Sep 13 16:03:39.836 Selected as active path
19 Sep 13 16:03:39.834 Record Route: 10.0.34.4 10.0.14.1
18 Sep 13 16:03:39.834 Up
17 Sep 13 16:03:39.698 Originate Call
16 Sep 13 16:03:39.698 CSPF: computation result accepted 10.0.34.4 10.0.14.1

15 Sep 13 16:03:39.697 Clear Call
14 Sep 13 16:03:39.696 Deselected as active
13 Sep 13 16:03:37.837 Record Route: 10.0.34.4 10.0.14.1
12 Sep 13 16:03:32.829 Fast-reroute Detour Down
11 Sep 13 16:02:15.493 Record Route: 10.0.34.4(flag=1) 10.0.14.1
10 Sep 13 16:02:15.486 Fast-reroute Detour Up
9 Sep 13 16:02:12.468 Record Route: 10.0.34.4 10.0.14.1
8 Sep 13 16:02:07.460 Fast-reroute Detour Down
7 Sep 13 15:57:46.741 Fast-reroute Detour Up
6 Sep 13 15:57:40.768 Record Route: 10.0.34.4(flag=1) 10.0.14.1
5 Sep 13 15:57:37.761 Selected as active path
4 Sep 13 15:57:37.760 Record Route: 10.0.34.4 10.0.14.1
3 Sep 13 15:57:37.760 Up
2 Sep 13 15:57:37.733 Originate Call
1 Sep 13 15:57:37.733 CSPF: computation result accepted 10.0.34.4 10.0.14.1

```

Created: Fri Sep 13 15:57:38 2013

Total 1 displayed, Up 1, Down 0

Egress LSP: 4 sessions, 6 detours

Total 0 displayed, Up 0, Down 0

Transit LSP: 6 sessions, 1 detours

#### 1.1.1.1

```

From: 3.3.3.3, LSPstate: Up, ActiveRoute: 0
LSPname: m120b-to-mx960, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 302288
Resv style: 1 FF, Label in: 300416, Label out: 302288
Time left: 147, Since: Fri Sep 13 16:08:16 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 4 receiver 13955 protocol 0
Detour branch from 10.0.34.4, to skip 1.1.1.1, Up
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0
PATH rcvfrom: 10.0.34.4 (ge-4/3/7.0) 7 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.35.5 (ge-3/1/0.0) 7 pkts
RESV rcvfrom: 10.0.35.5 (ge-3/1/0.0) 7 pkts
Explicit route: 10.0.35.5 10.0.15.1
Record route: 10.0.34.3 10.0.34.4 <self>10.0.35.5 10.0.15.1
Label in: 300416, Label out: 302288
Total 1 displayed, Up 1, Down 0

```

### show mpls lsp ingress extensive

```
user@host> show mpls lsp ingress extensive
```

Ingress LSP: 1 sessions

#### 50.0.0.1

```

From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
ActivePath: (primary)

```

```

LSPtype: Static Configured
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  OptimizeTimer: 300
  SmartOptimizeTimer: 180
  Reoptimization in 240 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    1.1.1.2 4.4.4.1 5.5.5.2
  17 Aug 3 13:17:33.601 CSPF: computation result ignored, new path less avail
bw[3 times]
  16 Aug 3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
times]
  15 Aug 3 12:54:36.678 Selected as active path
  14 Aug 3 12:54:36.676 Record Route: 1.1.1.2 4.4.4.1 5.5.5.2
  13 Aug 3 12:54:36.676 Up
  12 Aug 3 12:54:33.924 Deselected as active
  11 Aug 3 12:54:33.924 Originate Call
  10 Aug 3 12:54:33.923 Clear Call
  9 Aug 3 12:54:33.923 CSPF: computation result accepted 1.1.1.2 4.4.4.1
5.5.5.2
  8 Aug 3 12:54:33.922 2.2.2.2: No Route toward dest
  7 Aug 3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
times]
  6 Aug 3 12:35:03.830 Selected as active path
  5 Aug 3 12:35:03.828 Record Route: 2.2.2.2 3.3.3.2
  4 Aug 3 12:35:03.827 Up
  3 Aug 3 12:35:03.814 Originate Call
  2 Aug 3 12:35:03.814 CSPF: computation result accepted 2.2.2.2 3.3.3.2
  1 Aug 3 12:34:34.921 CSPF failed: no route toward 50.0.0.1
Created: Tue Aug 3 12:34:35 2010
Total 1 displayed, Up 1, Down 0

```

#### show mpls lsp extensive (automatic bandwidth adjustment enabled)

```

user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  Node/Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 300bps, MaxBW: 1000bps, Dynamic MinBW: 1000bps
  Adjustment Timer: 300 secs AdjustThreshold: 25%
  Max AvgBW util: 963.739bps, Bandwidth Adjustment in 0 second(s).
  Min BW Adjust Interval: 1000, MinBW Adjust Threshold (in %): 50
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 9, Underflow Max AvgBW: 614.421bps

  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180

```

```

    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    192.168.0.6(flag=0x20) 10.0.0.18(Label=299792) 192.168.0.4(flag=0x20)
10.0.0.22(Label=3)
    12 Apr 30 10:25:17.024 Make-before-break: Switched to new instance
    11 Apr 30 10:25:16.023 Record Route: 192.168.0.6(flag=0x20)
10.0.0.18(Label=299792) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    10 Apr 30 10:25:16.023 Up
    9 Apr 30 10:25:16.023 Automatic Autobw adjustment succeeded: BW changes from
300 bps to 1000 bps
    8 Apr 30 10:25:15.946 Originate make-before-break call
    7 Apr 30 10:25:15.946 CSPF: computation result accepted 10.0.0.18 10.0.0.22

    6 Apr 30 10:16:42.891 Selected as active path
    5 Apr 30 10:16:42.891 Record Route: 192.168.0.6(flag=0x20)
10.0.0.18(Label=299776) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    4 Apr 30 10:16:42.890 Up
    3 Apr 30 10:16:42.828 Originate Call
    2 Apr 30 10:16:42.828 CSPF: computation result accepted 10.0.0.18 10.0.0.22

    1 Apr 30 10:16:14.064 CSPF: could not determine self[2 times]
Created: Tue Apr 30 10:15:16 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show mpls lsp p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1        p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1        p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show mpls lsp p2mp detail

```

user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1
  LoadBalance: Random

```



```

    Encoding type: Packet, Switching type: Packet, GPID: IPv4
    *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
    192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
        192.168.208.17
    P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
    From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
    ActivePath: path1 (primary)
    P2MP name: p2mp-lsp2
    LoadBalance: Random
    Encoding type: Packet, Switching type: Packet, GPID: IPv4
    *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
    192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
        192.168.208.17
Total 2 displayed, Up 2, Down 0

```

#### show mpls lsp detail count-active-routes

```

user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
    From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
    ActivePath: (primary)
    LSPtype: Static Configured
    LoadBalance: Random
    Autobandwidth
    MinBW: 5Mbps MaxBW: 250Mbps
    Adjustment Timer: 300 secs
    Max AvgBW util: 60.2599Mbps, Bandwidth Adjustment in 0 second(s).
    Overflow limit: 0, Overflow sample count: 0
    Encoding type: Packet, Switching type: Packet, GPID: IPv4
    *Primary State: Up
        Priorities: 7 0
        Bandwidth: 5Mbps
        SmartOptimizeTimer: 180
        Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
    10.252.0.177 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
        10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

#### show mpls lsp statistics extensive

```

user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

```

```
192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPName: E-D
  Statistics: Packets 302, Bytes 28992
  Aggregate statistics: Packets 302, Bytes 28992
  ActivePath: (primary)
  LSPType: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
  6 Oct  3 11:18:28.281 Selected as active path
  5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
  4 Oct  3 11:18:28.280 Up
  3 Oct  3 11:18:27.995 Originate Call
  2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22

  1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0
```

## show mpls lsp autobandwidth

<b>Syntax</b>	<code>show mpls lsp autobandwidth</code> <brief   detail   extensive> <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Display automatic bandwidth information for the LSP(s).
<b>Options</b>	<p><b>brief   detail   extensive</b> — (Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b> — (Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mpls lsp on page 4665</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mpls lsp autobandwidth on page 4682</a>
<b>Output Fields</b>	<a href="#">Table 373 on page 4681</a> describes the output fields for the <code>show mpls lsp autobandwidth</code> command. Output fields are listed in the approximate order in which they appear.

**Table 373: show mpls lsp autobandwidth Output Fields**

Field Name	Field Description	Level of Output
<b>To</b>	Destination (egress routing device) of the session.	All Levels
<b>From</b>	Source (ingress routing device) of the session.	All Levels
<b>LSPname</b>	Name of the LSP.	All Levels
<b>Min BW</b>	(Ingress LSP) Configured minimum value of the LSP, in bps.	<b>detail extensive</b>
<b>Max BW</b>	(Ingress LSP) Configured maximum value of the LSP, in bps.	<b>detail extensive</b>
<b>Max AvgBW util</b>	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	<b>detail extensive</b>
<b>Overflow limit</b>	(Ingress LSP) Configured value of the threshold overflow limit.	<b>detail extensive</b>
<b>Overflow sample count</b>	(Ingress LSP) Current value for the overflow sample count.	<b>detail extensive</b>
<b>Underflow limit</b>	(Ingress LSP) Configured value of the threshold underflow limit.	<b>detail extensive</b>

Table 373: show mpls lsp autobandwidth Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Underflow sample count</b>	(Ingress LSP) Current value for the underflow sample count.	<b>detail extensive</b>
<b>Adjustment Timer</b>	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	<b>detail extensive</b>
<b>Adjustment Threshold</b>	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	<b>detail extensive</b>
<b>Time for Next Adjustment</b>	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	<b>detail extensive</b>
<b>Time of Last Adjustment</b>	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	<b>detail extensive</b>
<b>Last BW</b>	Previous active bandwidth of the LSP.	<b>detail extensive</b>
<b>Last Requested BW</b>	Bandwidth requested in the previous automatic bandwidth adjustment.	<b>detail extensive</b>
<b>Last Signaled BW</b>	Bandwidth signaled in the previous automatic bandwidth adjustment.	<b>detail extensive</b>
<b>Highest Watermark BW</b>	Maximum bandwidth used by the LSP.	<b>detail extensive</b>
<b>Total AutoBw Adjustments</b>	Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.	<b>detail extensive</b>
<b>Successful Adjustments</b>	Number of successful automatic bandwidth adjustments.	<b>detail extensive</b>
<b>Failed Adjustments</b>	Number of failed automatic bandwidth adjustments.	<b>detail extensive</b>

## Sample Output

### show mpls lsp autobandwidth

```

user@host> show mpls lsp autobandwidth extensive
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun 3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0

```



## show mpls path

<b>List of Syntax</b>	<a href="#">Syntax on page 4684</a> <a href="#">Syntax (EX Series Switches) on page 4684</a>
<b>Syntax</b>	show mpls path <logical-system (all   <i>logical-system-name</i> )> <path-name>
<b>Syntax (EX Series Switches)</b>	show mpls path <path-name>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).
<b>Options</b>	<b>none</b> —Display standard information about all MPLS LSPs.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>path-name</b> —(Optional) Display information about the specified LSP only.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls path on page 4684</a>
<b>Output Fields</b>	<a href="#">Table 374 on page 4684</a> describes the output fields for the <b>show mpls path</b> command. Output fields are listed in the approximate order in which they appear.

**Table 374: show mpls path Output Fields**

Field Name	Field Description
<b>Path name</b>	Information about ingress LSPs. Each path has one line of output.
<b>Address</b>	Addresses of the routing devices that form the LSP.
<b>Strict/loose address</b>	Whether the address is configured as a strict or loose address.

## Sample Output

### show mpls path

```

user@host> show mpls path
Path name      Address          Strict/loose address
p1             123.456.55.6    Strict
               123.456.1.6     Loose
p2             191.456.1.4     Strict
  
```

## show mpls static-lsp

**Syntax** show mpls static-lsp  
 <brief | detail | extensive | terse>  
 <bypass>  
 <descriptions>  
 <down | up>  
 <ingress>  
 <logical-system (all | *logical-system-name*)>  
 <lsp-type>  
 <name *name*>  
 <statistics>  
 <transit>

**Release Information** Command introduced in Junos OS Release 10.1.

**Description** Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options** **none**—Display standard information about all configured and active static MPLS LSPs.

**brief | detail | extensive | terse**—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

**bypass**—(Optional) Display LSPs used for protecting other static LSPs.

**descriptions**—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the **[edit protocols mpls static-label-switched-path *path-name* bypass]**, **[edit protocols mpls static-label-switched-path *path-name* ingress]**, or **[edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*]** hierarchy levels. Only static LSPs with a description are displayed.

**down | up**—(Optional) Display only static LSPs that are inactive or active, respectively.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

**name *name***—(Optional) Display information about the specified static LSP or group of LSPs.

**statistics**—(Optional) Display accounting information about static LSPs.

**transit**—(Optional) Display static LSPs transiting this routing device.

**Required Privilege Level** view

**List of Sample Output** [show mpls static-lsp extensive on page 4687](#)  
[show mpls static-lsp statistics ingress on page 4687](#)

**Output Fields** [Table 375 on page 4686](#) describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

**Table 375: show mpls static-lsp Output Fields**

Field Name	Field Description	Level of Output
<b>Ingress LSPs</b>	Information about the static LSPs on the ingress routing device. Each session has one line of output.	All levels
<b>Transit LSPs</b>	Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
<b>Bypass LSPs</b>	Information about the bypass LSPs configured on the routing device. Each session has one line of output.	All levels
<b>LSPname</b>	Name of the static LSP.	All levels
<b>To</b>	Destination (egress routing device) of the session.	All levels
<b>State</b>	State of the static LSP handled by this RSVP session: <b>Up</b> , <b>Dn</b> (down), or <b>Restart</b> .	All levels
<b>Packets</b>	Number of packet transiting the static LSP ( <b>statistics</b> option only).	All levels
<b>Bytes</b>	Number of bytes transiting the static LSP ( <b>statistics</b> option only).	All levels
<b>Nexthop</b>	IP address for the next-hop router for the static LSP.	<b>detail, extensive</b>
<b>Bypass</b>	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
<b>Link protection desired</b>	Link protection has been requested by the ingress routing device.	<b>detail, extensive</b>
<b>LabelOperation</b>	Label operation to perform: <b>Push</b> , <b>Pop</b> , <b>Swap</b> .	<b>detail, extensive</b>
<b>Outgoing-label</b>	Outgoing label to use for the MPLS packet in either push or swap label operations.	<b>detail, extensive</b>
<b>Created</b>	(Ingress LSP) Date and time the static LSP was created.	<b>extensive</b>
<b>Bandwidth</b>	Bandwidth configured for the static LSP.	<b>detail, extensive</b>
<b>Resv style</b>	(Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).	All levels



## Sample Output

### show mpls static-lsp extensive

```
user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPname: alpha-to-beta, To: 192.168.14.1
State: Dn
Nexthop: 192.168.10.1
LabelOperation: Push, Outgoing-label: 1000001
Created: Thu Jan 14 16:44:43 2010
Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0
```

### show mpls static-lsp statistics ingress

```
user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPname           To           State      Packets      Bytes
alpha-to-beta     192.168.14.1 Dn         NA           NA
Total 1, displayed 1, Up 0, Down 1
```

## show rsvp interface

<b>List of Syntax</b>	<a href="#">Syntax on page 4688</a> <a href="#">Syntax (EX Series Switches) on page 4688</a>
<b>Syntax</b>	<pre>show rsvp interface &lt;brief   detail   extensive&gt; &lt;link-management&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show rsvp interface &lt;brief   detail   extensive&gt; &lt;link-management&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p>
<b>Description</b>	Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.
<b>Options</b>	<p><b>none</b>—Display standard information about the status of RSVP-enabled interfaces and packet statistics.</p> <p><b>brief   detail   extensive   link-management</b>—(Optional) Display the specified level of output.</p> <p><b>link-management</b>—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show rsvp interface brief on page 4691</a> <a href="#">show rsvp interface detail on page 4691</a> <a href="#">show rsvp interface extensive on page 4691</a> <a href="#">show rsvp interface link-management on page 4692</a>
<b>Output Fields</b>	<a href="#">Table 376 on page 4688</a> lists the output fields for the <b>show rsvp interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 376: show rsvp interface Output Fields**

Field Name	Field Description	Level of Output
<b>RSVP interface</b>	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
<b>Interface</b>	Name of the interface.	All levels

Table 376: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Index</b>	Index of the interface.	<b>detail</b>
<b>State</b>	State of the interface. <ul style="list-style-type: none"> <li>• <b>Disabled</b>—No traffic engineering information is displayed.</li> <li>• <b>Down</b>—Interface is not operational.</li> <li>• <b>Enabled</b>—Displays traffic engineering information.</li> <li>• <b>Up</b>—Interface is operational.</li> </ul>	All levels
<b>NoAuthentication</b>	Interface does not support RSVP authentication.	<b>detail</b>
<b>NoAggregate</b>	Interface does not support refresh reduction.	<b>detail</b>
<b>NoReliable</b>	Interface does not support refresh reduction message ID extension.	<b>detail</b>
<b>NoLinkProtection</b>	Interface does not support link protection.	<b>detail</b>
<b>HelloInterval</b>	Frequency at which RSVP hellos are sent on this interface (in seconds).	<b>detail</b>
<b>Address</b>	IP address of the local interface.	<b>detail</b>
<b>Active control channel</b>	Next-hop link address to transmit messages.	None specified
<b>TElink</b>	Traffic-engineered links that are managed by the peer they are associated with.	None specified
<b>Active resv</b>	Number of reservations that are actively reserving bandwidth on the interface.	All levels
<b>PreemptionCnt</b>	Number of times an RSVP session was preempted on this interface.	<b>detail</b>
<b>Update threshold</b>	Percentage change in reserved bandwidth to trigger an IGP update.	<b>detail</b>
<b>Subscription</b>	User-configured subscription factor.	All levels
<b>bc number</b>	Bandwidth allocated for the specified bandwidth constraint.	<b>extensive</b>
<b>ct number</b>	Bandwidth allocated for the specified class type.	<b>extensive</b>
<b>Static BW</b>	Total interface bandwidth, in bps.	All levels
<b>Available BW</b>	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	all levels
<b>Reserved BW</b>	Currently reserved bandwidth, in bps.	All levels
<b>SoftPreemptionCnt</b>	Number of times a soft preemption occurred on this interface. This number is not included in the <b>PreemptionCnt</b> value.	<b>detail</b>

Table 376: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Overbooked BW</b>	Currently overbooked bandwidth, in bps, by class type (ct0 through ct3).	<b>detail</b>
<b>Highwater mark</b>	Highest bandwidth that has ever been reserved on this interface, in bps.	<b>brief</b>
<b>PacketType</b>	Type of RSVP packet.	<b>detail</b>
<b>Total Sent</b>	Total number of packets sent.	<b>detail</b>
<b>Total Received</b>	Total number of packets received since RSVP was enabled.	<b>detail</b>
<b>Last 5 seconds Sent</b>	Number of packets sent in the last 5 seconds.	<b>detail</b>
<b>Last 5 seconds Received</b>	Number of packets received in the last 5 seconds.	<b>detail</b>
<b>Path</b>	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	<b>detail</b>
<b>PathErr</b>	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	<b>detail</b>
<b>PathTear</b>	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	<b>detail</b>
<b>Resv</b>	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	<b>detail</b>
<b>ResvErr</b>	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	<b>detail</b>
<b>ResvTear</b>	Statistics about ResvTear messages, which remove reservation states along a path.	<b>detail</b>
<b>Hello</b>	Number of RSVP hello packets that have been sent to and received from the neighbor.	<b>detail</b>
<b>Ack</b>	Acknowledge message for refresh reductions.	<b>detail</b>
<b>Srefresh</b>	Summary refresh messages.	<b>detail</b>
<b>EndtoEnd RSVP</b>	Statistics for the number of end-to-end RSVP messages sent.	<b>detail</b>
<b>Queue</b>	CoS transmit queue number and its associated forwarding class designation.	<b>extensive</b>
<b>TxRate</b>	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	<b>extensive</b>
<b>Priority</b>	Weight of the queue relative to other configured queues, in percentage.	<b>extensive</b>

Table 376: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>queue-priority-value</i>	Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only.	extensive

## Sample Output

### show rsvp interface brief

```

user@host> show rsvp interface brief
RSVP interface: 1 active
      Active Subscr- Static   Available   Reserved   Highwater
Interface State resv  iption  BW         BW         mark
de0.0      Up      1      23%     10Mbps     989.992kbps 1.31Mbps     1.31Mbps

```

### show rsvp interface detail

```

user@host> show rsvp interface detail
so-0/1/1.0 Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  SoftPreemptionCnt1
  OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
  PacketType          Total          Last 5 seconds
      Sent      Received      Sent      Received
  Path                16           0           1           0
  PathErr              0           0           0           0
  PathTear             1           0           0           0
  Resv                  0          11           0           1
  ResvErr               0           0           0           0
  ResvTear              0           0           0           0
  Hello                 66          67           1           1
  Ack                   0           0           0           0
  Srefresh              0           0           0           0
  EndtoEnd RSVP        0           0           0           0
...

```

### show rsvp interface extensive

```

user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps
  bc2 = (ct2+ct3), StaticBW 311.04Mbps
  bc3 = ct3, StaticBW 155.52Mbps
  ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
  ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

```

```
ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Queue      TxRate      Priority Exact
0          155.52Mbps      25%     Low
1          155.52Mbps      25%     Low
2          155.52Mbps      25%     Low
3          155.52Mbps      25%     Low
```

#### show rsvp interface link-management

```
user@host> show rsvp interface link-management
```

```
RSVP interface: 2 active
```

```
PEER-C State: Up
```

```
Active Control Channel: so-0/1/0.0
```

```
TElink: TElnk1, Link ID: 37811
```

```
ActiveResv 0, PreemptionCnt 0
```

```
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps
```

```
TElink: TElnk2, Link ID: 37808
```

```
ActiveResv 1, PreemptionCnt 0
```

```
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps
```

```
PEER-B State: Up
```

```
Active Control Channel: so-1/0/0.0
```

```
TElink: TElnkAB1, Link ID: 1598
```

```
ActiveResv 0, PreemptionCnt 0
```

```
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```

```
TElink: TElnkAB2, Link ID: 1597
```

```
ActiveResv 0, PreemptionCnt 0
```

```
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```

## show rsvp neighbor

<b>List of Syntax</b>	<a href="#">Syntax on page 4693</a> <a href="#">Syntax (EX Series Switches) on page 4693</a>
<b>Syntax</b>	show rsvp neighbor <brief   detail> <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show rsvp neighbor <brief   detail>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.
<b>Options</b>	<b>none</b> —Display standard information about RSVP neighbors.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show rsvp neighbor on page 4697</a> <a href="#">show rsvp neighbor detail on page 4697</a>
<b>Output Fields</b>	<a href="#">Table 377 on page 4693</a> lists the output fields for the <b>show rsvp neighbor</b> command. Output fields are listed in the approximate order in which they appear.

**Table 377: show rsvp neighbor Output Fields**

Field Name	Field Description	Level of Output
<b>RSVP neighbor</b>	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
<b>via</b>	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	<b>detail</b>
<b>Address</b>	Address of a learned neighbor.	All levels
<b>Idle</b>	Length of time the neighbor has been idle, in seconds.	All levels

Table 377: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Up/Dn</b>	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	All levels
<b>Up cnt and Down cnt</b>	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	<b>detail</b>
<b>status</b>	<p>State of the RSVP neighbor:</p> <ul style="list-style-type: none"> <li>• <b>Up</b>—Routing device can detect RSVP Hello messages from the neighbor.</li> <li>• <b>Down</b>—Routing device has received one of the following indications: <ul style="list-style-type: none"> <li>• Communication failure from the neighbor.</li> <li>• Communication from IGP that the neighbor is unavailable.</li> <li>• Change in the sequence numbers in the RSVP Hello messages sent by the neighbor.</li> </ul> </li> <li>• <b>Restarting</b>—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled.</li> <li>• <b>Restarted</b>—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures.</li> <li>• <b>Dead</b>—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down.</li> </ul>	<b>detail</b>
<b>LastChange</b>	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <b>hh:mm:ss</b> .	All levels
<b>Last changed time</b>	Time elapsed since the neighbor state changed either from up to down or from down to up.	<b>detail</b>
<b>HelloInt</b>	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
<b>HelloTx/Rx</b>	Number of hello packets sent to and received from the neighbor.	All levels
<b>Hello</b>	Number of RSVP hello packets that have been sent to and received from the neighbor.	<b>detail</b>
<b>Message received</b>	Number of Path and Resv messages that this routing device has received from the neighbor.	<b>detail</b>
<b>Remote Instance</b>	Identification provided by the remote routing device during Hello message exchange.	<b>detail</b>
<b>Local Instance</b>	Identification sent to the remote routing device during Hello message exchange.	<b>detail</b>



Table 377: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Refresh reduction</b>	<p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. <b>Refresh reduction</b> can have the following values:</p> <ul style="list-style-type: none"> <li>• <b>operational</b>—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961.</li> <li>• <b>incomplete</b>—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices.</li> <li>• <b>no operational</b>—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions.</li> </ul>	<b>detail</b>
<b>Remote end</b>	<p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Remote routing device has requested refresh reduction during RSVP message exchanges.</li> <li>• <b>disabled</b>—Remote routing device does not require refresh reduction.</li> </ul>	<b>detail</b>
<b>Ack-extension</b>	<p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Both local and remote routing devices support the ack-extension (RFC 2961).</li> <li>• <b>disabled</b>—Remote routing device does not support the ack-extension.</li> </ul>	<b>detail</b>
<b>Link protection</b>	<p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Link protection feature has been turned on, protecting the neighbor with a bypass LSP.</li> <li>• <b>disabled</b>—No link protection feature has been enabled for this neighbor.</li> </ul>	<b>detail</b>
<b>LSP name</b>	Name of the bypass LSP.	<b>detail</b>
<b>Bypass LSP</b>	<p>Status of the bypass LSP. It can have the following values:</p> <ul style="list-style-type: none"> <li>• <b>does not exist</b>—Bypass LSP is not available.</li> <li>• <b>connecting</b>—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment.</li> <li>• <b>operational</b>—Bypass LSP is up and running.</li> <li>• <b>down</b>—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.</li> </ul>	<b>detail</b>
<b>Backup routes</b>	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	<b>detail</b>
<b>Backup LSPs</b>	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	<b>detail</b>
<b>Bypass explicit route</b>	Explicit route object's (ERO) path that is taken by the bypass LSP.	<b>detail</b>

Table 377: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Restart time</b>	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	<b>detail</b>
<b>Recovery time</b>	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	<b>detail</b>

## Sample Output

### show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203   0 3/2    13:01      3   366/349
192.168.207.207   0 1/0    22:49      3   448/448
```

### show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasy1 status: Up
  Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
  Message received: 632
  Hello: sent 673, received 656, interval 3 sec
  Remote instance: 0x6432838a, Local instance: 0x74b72e36
  Refresh reduction: operational
    Remote end: enabled, Ack-extension: enabled
  Link protection: enabled
    LSP name: Bypass_to_192.168.207.203
    Bypass LSP: operational, Backup routes: 1, Backup LSPs: 0
    Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

## show rsvp session

---

<b>List of Syntax</b>	<a href="#">Syntax on page 4698</a> <a href="#">Syntax (EX and QFX Series Switches) on page 4698</a>
<b>Syntax</b>	<pre>show rsvp session &lt;brief   detail   extensive   terse&gt; &lt;bidirectional   unidirectional&gt; &lt;bypass&gt; &lt;down   up&gt; &lt;externally-provisioned&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-type&gt; &lt;name <i>session-name</i>&gt; &lt;p2mp&gt; &lt;session-type&gt; &lt;statistics&gt; &lt;te-link <i>te-link</i>&gt;</pre>
<b>Syntax (EX and QFX Series Switches)</b>	<pre>show rsvp session &lt;brief   detail   extensive   terse&gt; &lt;bidirectional   unidirectional&gt; &lt;bypass&gt; &lt;down   up&gt; &lt;externally-provisioned&gt; &lt;interface <i>interface-name</i>&gt; &lt;lsp-type&gt; &lt;name <i>session-name</i>&gt; &lt;p2mp&gt; &lt;session-type&gt; &lt;statistics&gt; &lt;te-link <i>te-link</i>&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. <b>externally-provisioned</b> option added in Junos OS Release 13.3. Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.
<b>Description</b>	Display information about Resource Reservation Protocol (RSVP) sessions.
<b>Options</b>	<p><b>none</b>—Display standard information about all RSVP sessions.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>bidirectional   unidirectional</b>—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.</p> <p><b>bypass</b>—(Optional) Display RSVP sessions for bypass LSPs.</p> <p><b>down   up</b>—(Optional) Display only LSPs that are inactive or active, respectively.</p>

**externally-provisioned**—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

**interface *interface-name***—(Optional) Display RSVP sessions for the specified interface only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.
- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

**name *session-name***—(Optional) Display information about the named session.

**p2mp**—(Optional) Display point-to-multipoint information.

***session-type***—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

**statistics**—(Optional) Display packet statistics.

**te-link *te-link***—(Optional) Display sessions with reservations on the specified TE link.

**Required Privilege Level**

view

**Related Documentation**

- [clear rsvp session on page 4574](#)

**List of Sample Output**

[show rsvp session on page 4703](#)  
[show rsvp session statistics on page 4703](#)  
[show rsvp session detail on page 4704](#)  
[show rsvp session detail \(Path MTU Output Field\) on page 4704](#)  
[show rsvp session detail \(GMPLS\) on page 4704](#)  
[show rsvp session extensive on page 4705](#)  
[show rsvp session p2mp \(Ingress Router\) on page 4705](#)  
[show rsvp session p2mp \(Transit Router\) on page 4706](#)

**Output Fields**

[Table 378 on page 4700](#) describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 378: show rsvp session Output Fields

Field Name	Field Description	Level of Output
<b>Ingress RSVP</b>	Information about ingress RSVP sessions.	<b>detail</b>
<b>Ingress RSVP</b>	Information about ingress RSVP sessions. Each session has one line of output.	All levels
<b>Egress RSVP</b>	Information about egress RSVP sessions.	All levels
<b>Transit RSVP</b>	Information about the transit RSVP sessions.	All levels
<b>P2MP name</b>	(Appears only when the <b>p2mp</b> option is specified). Name of the point-to-multipoint LSP path.	All levels
<b>P2MP branch count</b>	(Appears only when the <b>p2mp</b> option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
<b>To</b>	Destination (egress routing device) of the session.	All levels
<b>From</b>	Source (ingress routing device) of the session.	All levels
<b>State</b>	State of the path: <b>Up</b> , <b>Down</b> , or <b>AdminDn</b> . <b>AdminDn</b> indicates that the LSP is being taken down gracefully.	All levels
<b>Address</b>	Destination (egress routing device) of the LSP.	<b>detail</b>
<b>From</b>	Source (ingress routing device) of the session.	<b>detail</b>
<b>LSPstate</b>	State of the LSP that is being handled by this RSVP session. It can be either <b>Up</b> , <b>Dn</b> (down), or <b>AdminDn</b> . <b>AdminDn</b> indicates that the LSP is being taken down gracefully.	<b>brief detail</b>
<b>Rt</b>	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the routing table is the primary MPLS table ( <b>mpls.0</b> ).	<b>brief</b>
<b>Active Route</b>	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table ( <b>mpls.0</b> ).	<b>detail</b>
<b>LSPname</b>	Name of the LSP.	<b>brief detail</b>
<b>LSPpath</b>	Indicates whether the RSVP session is for the primary or secondary LSP path. <b>LSPpath</b> can be either <b>primary</b> or <b>secondary</b> and can be displayed on the ingress, egress, and transit routing devices. <b>LSPpath</b> can also indicate when a graceful LSP deletion has been triggered.	<b>detail</b>
<b>Bypass</b>	(Egress routing device) Destination address for the bypass LSP.	<b>detail</b>

Table 378: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bidir</b>	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	<b>detail</b>
<b>Bidirectional</b>	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	<b>detail</b>
<b>Upstream label in</b>	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	<b>detail</b>
<b>Upstream label out</b>	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	<b>detail</b>
<b>Recovery label received</b>	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	<b>detail</b>
<b>Recovery label sent</b>	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	<b>detail</b>
<b>Suggested label received</b>	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	<b>detail</b>
<b>Suggested label sent</b>	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	<b>detail</b>
<b>Resv style or Style</b>	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).	<b>brief detail</b>
<b>Label in</b>	Incoming label for this LSP.	<b>brief detail</b>
<b>Label out</b>	Outgoing label for this LSP.	<b>brief detail</b>
<b>Time left</b>	Number of seconds remaining in the lifetime of the reservation.	<b>brief detail</b>
<b>Since</b>	Date and time when the RSVP session was initiated.	<b>detail</b>
<b>Tspec</b>	Sender's traffic specification, which describes the sender's traffic parameters.	<b>detail</b>
<b>DiffServ info</b>	Indicates whether the LSP is a multiclass LSP ( <b>multiclass diffServ-TE LSP</b> ) or a Differentiated-Services-aware traffic engineering LSP ( <b>diffServ-TE LSP</b> ).	<b>detail</b>
<b>bandwidth</b>	Bandwidth for each class type ( <b>ct0</b> , <b>ct1</b> , <b>ct2</b> , or <b>ct3</b> ).	<b>detail</b>
<b>Port number</b>	Protocol ID and sender/receiver port used in this RSVP session.	<b>detail</b>
<b>Attrib flags</b>	<b>Non-PHP</b> indicates that ultimate hop popping has been requested by the LSP using this RSVP session	<b>extensive</b>

Table 378: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>FastReroute desired</b>	Fast reroute has been requested by the ingress routing device.	<b>detail</b>
<b>Soft preemption desired</b>	Soft preemption has been requested by the ingress routing device.	<b>detail</b>
<b>FastReroute desired</b>	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	<b>detail extensive</b>
<b>Link protection desired</b>	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	<b>detail extensive</b>
<b>Node/Link protection desired</b>	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	<b>detail extensive</b>
<b>Type</b>	<p>LSP type:</p> <ul style="list-style-type: none"> <li>• <b>Link protected LSP</b>—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (<b>extensive</b>).</li> <li>• <b>Node/Link protected LSP</b>—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (<b>extensive</b>).</li> <li>• <b>Protection down</b>—LSP is not currently protected.</li> <li>• <b>Bypass LSP</b>—LSP that is used to protect one or more user LSPs in case of link failure.</li> <li>• <b>Backup LSP at Point-of-Local-Repair (PLR)</b>—LSP that has been temporarily established to protect a user LSP at the ingress of a failed link.</li> <li>• <b>Backup LSP at Merge Point (MP)</b>—LSP that has been temporarily established to protect a user LSP at the egress of a failed link.</li> </ul>	<b>detail extensive</b>
<b>New bypass</b>	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	<b>extensive</b>
<b>Link protection up, using <i>bypass-name</i></b>	Link protection (the bypass name is also displayed) has been activated for the LSP.	<b>extensive</b>
<b>Creating backup LSP, link down</b>	A <b>link down</b> event occurred, and traffic is being switched over to the bypass LSP.	<b>extensive</b>
<b>Deleting backup LSP, protected LSP restored</b>	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	<b>extensive</b>
<b>Path mtu</b>	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the <b>allow-fragmentation</b> statement configured at the <b>[edit protocols mpls path-mtu]</b> hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	<b>detail</b>



Table 378: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>PATH rcvfrom</b>	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	<b>detail</b>
<b>Adspec</b>	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	<b>detail</b>
<b>PATH sentto</b>	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	<b>detail</b>
<b>Explct route</b>	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	<b>detail</b>
<b>Record route</b>	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	<b>detail</b>

## Sample Output

### show rsvp session

```

user@host> show rsvp session
Ingress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.245.214 10.255.245.212 AdminDn 0 1 FF - 22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.245.194 10.255.245.195 Up 0 1 FF 39811 - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up 0 1 FF 3 - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.245.198 10.255.245.197 Up 0 1 SE 100000 3 pro3-de
Total 1 displayed, Up 1, Down 0

```

### show rsvp session statistics

```

user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To          From          State Packets Bytes LSPname
10.255.245.24 10.255.245.22 Up 0 0 pro3-bd
10.255.245.24 10.255.245.22 Up 44868 2333136 pro3-bd-2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
To          From          State Packets Bytes LSPname
10.255.245.22 10.255.245.24 Up 0 0 pro3-db
10.255.245.22 10.255.245.24 Up 0 0 pro3-db-2
Total 2 displayed, Up 2, Down 0

```

Transit RSVP: 0 sessions  
Total 0 displayed, Up 0, Down 0

#### show rsvp session detail

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

#### show rsvp session detail (Path MTU Output Field)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left: -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding
```

#### show rsvp session detail (GMPLS)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left: -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
```

```

Adspec: sent MTU 1500
PATH MTU: received 0
PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
Explct route: 100.100.100.100 93.93.93.93
Record route: <self> 100.100.100.100 93.93.93.93
Total 1 displayed, Up 0, Down 1
Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show rsvp session extensive

```

user@host> show rsvp session extensive
Ingress RSVP: 1 sessions

192.168.0.4
  From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299808
  Resv style: 1 FF, Label in: -, Label out: 299808
  Time left: -, Since: Thu Sep 20 15:54:20 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 61576 protocol 0
  Attrib flags: Non-PHP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
  RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
  Explct route: 10.0.0.18 10.0.0.22
  Record route: <self> 10.0.0.18 10.0.0.22
  Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 140, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
  Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show rsvp session p2mp (Ingress Router)

```

user@host> show rsvp session p2mp

```

```
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
To          From          State   Rt Style Labelin Labelout LSPName
10.255.10.95 10.255.10.2 Up      0  1 SE  -        3 to-pe1
P2MP name: test2, P2MP branch count: 2
To          From          State   Rt Style Labelin Labelout LSPName
10.255.10.23 10.255.10.2 Up      0  1 SE  -        299776 to-pe3
10.255.10.16 10.255.10.2 Up      0  1 SE  -        299776 to-pe4
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

#### show rsvp session p2mp (Transit Router)

```
user@host> show rsvp session p2mp
Ingress RSVP: 1 sessions
P2MP name: test, P2MP branch count: 1
To          From          State   Rt Style Labelin Labelout LSPName
10.255.10.23 10.255.10.95 Up      0  1 SE  -        299792 to-pe2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
P2MP name: test, P2MP branch count: 1
To          From          State   Rt Style Labelin Labelout LSPName
10.255.10.95 10.255.10.2 Up      0  1 SE  3        -        to-pe1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions
P2MP name: test2, P2MP branch count: 2
To          From          State   Rt Style Labelin Labelout LSPName
10.255.10.23 10.255.10.2 Up      0  1 SE  299776   299808 to-pe3
10.255.10.16 10.255.10.2 Up      0  1 SE  299776   299856 to-pe4
Total 2 displayed, Up 2, Down 0
```

## show rsvp statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 4707</a> <a href="#">Syntax (EX Series Switches) on page 4707</a>
<b>Syntax</b>	show rsvp statistics <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show rsvp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Resource Reservation Protocol (RSVP) packet and error statistics.
<b>Options</b>	<b>none</b> —Display RSVP packet and error statistics.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear rsvp statistics on page 4576</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show rsvp statistics on page 4710</a>
<b>Output Fields</b>	<a href="#">Table 379 on page 4707</a> describes the output fields for the <b>show rsvp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 379: show rsvp statistics Output Fields**

Field Name	Field Description
<b>Packet Type</b>	Statistics about different RSVP messages.
<b>Total Sent</b>	Total number of packets sent since RSVP was enabled.
<b>Total Received</b>	Total number of packets received since RSVP was enabled.
<b>Last 5 seconds Sent</b>	Total number of packets sent in the last 5 seconds.
<b>Last 5 seconds Received</b>	Number of packets received in the last 5 seconds.
<b>Path</b>	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.
<b>PathErr</b>	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.

Table 379: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
<b>PathTear</b>	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.
<b>Resv FF</b>	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
<b>Resv WF</b>	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
<b>Res SE</b>	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
<b>ResvErr</b>	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
<b>ResvTear</b>	Statistics about ResvTear messages, which remove reservation states along a path.
<b>ResvConf</b>	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
<b>Ack</b>	Acknowledge message for refresh reductions.
<b>SRefresh</b>	Summary refresh messages.
<b>Hello</b>	Number of RSVP hello packets that have been sent to and received from the neighbor.
<b>EndtoEnd RSVP</b>	Statistics for the number of End-to-end RSVP messages.
<b>Errors</b>	Statistics about errored RSVP packets.
<b>Rcv pkt bad length</b>	The packet was not processed because its length is inappropriate.
<b>Rcv pkt unknown type</b>	The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> .
<b>Rcv pkt bad version</b>	The packet is not an RSVP version 1 packet.
<b>Rcv pkt auth fail</b>	The packet failed authentication checks.
<b>Rcv pkt bad checksum</b>	The RSVP checksum check failed.
<b>Rcv pkt bad format</b>	General packet processing failed because the packet was badly formed.
<b>Memory allocation fail</b>	An internal resource failure occurred.
<b>No path information</b>	A reservation was received, but no sender is active.
<b>Resv style conflict</b>	The same session contains inconsistent reservation styles.

Table 379: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Port conflict</b>	There were inconsistent port numbers for the same session.
<b>Resv no interface</b>	An interface for the receive reservation packets cannot be located.
<b>PathErr to client</b>	Number of PathErr packets delivered to the local client.
<b>ResvErr to client</b>	Number of ResvErr packets delivered to the local client.
<b>Path timeout</b>	Number of times the sender timed out because the path was removed.
<b>Resv timeout</b>	Number of times the receiver timed out because the reservation was removed.
<b>Message out-of-order</b>	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
<b>Unknown ack msg</b>	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.
<b>Recv nack</b>	If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again.
<b>Recv duplicated msg-id</b>	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
<b>No TE-link to rcv Hop</b>	Counter of packets discarded because a TE link was not found.
<b>Rcv pkt disabled interface</b>	Number of RSVP packets received on an interface that is not enabled for RSVP.
<b>Transmit buffer full</b>	Number of times the buffer for assembling an outgoing RSVP message was not large enough.
<b>Transmit failure</b>	Number of times the RSVP task failed to send out a packet.
<b>Receive failure</b>	Number of times the RSVP task failed to read an incoming packet.
<b>P2MP RESV discarded by appl</b>	Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.
<b>Rate limit</b>	Number of RSVP packets dropped due to rate limiting.
<b>Err msg loop detected</b>	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

## Sample Output

### show rsvp statistics

```

user@host> show rsvp statistics

```

PacketType	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Path	355	408	0	0
PathErr	2	13	0	0
PathTear	101	139	0	0
Resv FF	0	0	0	0
Resv WF	0	0	0	0
Resv SE	419	225	0	0
ResvErr	0	0	0	0
ResvTear	0	13	0	0
ResvConf	0	0	0	0
Ack	682	1414	0	0
SRefresh	395198	236030	5	2
Hello	578809	578221	4	4
EndtoEnd RSVP	0	0	0	0

	Total	Last 5 seconds
Errors		
Rcv pkt bad length	0	0
Rcv pkt unknown type	0	0
Rcv pkt bad version	0	0
Rcv pkt auth fail	0	0
Rcv pkt bad checksum	0	0
Rcv pkt bad format	0	0
Memory allocation fail	0	0
No path information	10	0
Resv style conflict	0	0
Port conflict	0	0
Resv no interface	0	0
PathErr to client	38	0
ResvErr to client	0	0
Path timeout	8	0
Resv timeout	57	0
Message out-of-order	0	0
Unknown ack msg	2978	0
Recv nack	86	0
Recv duplicated msg-id	5	0
No TE-link to recv Hop	0	0
Rcv pkt disabled interface	0	0
Transmit buffer full	0	0
Transmit failure	0	0
Receive failure	0	0
P2MP RESV discarded by appl	0	0
Rate limit	306	0
Err msg loop detected	0	0



## show rsvp version

<b>List of Syntax</b>	<a href="#">Syntax on page 4711</a> <a href="#">Syntax (EX Series Switches) on page 4711</a>
<b>Syntax</b>	show rsvp version <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show rsvp version
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.
<b>Options</b>	<b>none</b> —Display RSVP protocol settings.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show rsvp version on page 4712</a>
<b>Output Fields</b>	<a href="#">Table 380 on page 4711</a> describes the output fields for the <b>show rsvp version</b> command. Output fields are listed in the approximate order in which they appear.

**Table 380: show rsvp version Output Fields**

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: <b>Enabled</b> or <b>Disabled</b> .
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: <b>Aggressive</b> , <b>Disabled</b> , or <b>Normal</b> . The default is <b>Normal</b> .
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.
Graceful deleting timeout	Currently configured value for the <b>graceful-deletion-timeout</b> statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

Table 380: show rsvp version Output Fields (*continued*)

Field Name	Field Description
<b>NSR Mode</b>	Status of the nonstop active routing feature for RSVP on the restarting device: <b>Disabled</b> , <b>Enabled/Master</b> , or <b>Enabled/Standby</b> .
<b>NSR State</b>	<p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> <li>• <b>Idle</b></li> <li>• <b>TE-link sync complete</b></li> <li>• <b>Neighbor sync complete</b></li> <li>• <b>Path state sync complete</b></li> <li>• <b>Resv state sync complete</b></li> <li>• <b>Bypass sync complete</b></li> <li>• <b>Init sync complete</b></li> </ul>
<b>Setup protection</b>	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: <b>Enabled</b> or <b>Disabled</b>
<b>Graceful restart</b>	Status of the graceful restart feature for RSVP on the restarting routing device: <b>Enabled</b> or <b>Disabled</b> .
<b>Restart helper mode</b>	Status of the helper mode feature: <b>Enabled</b> or <b>Disabled</b> . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
<b>Maximum helper restart time</b>	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
<b>Maximum helper recovery time</b>	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
<b>Restart time</b>	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
<b>Recovery time</b>	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
<b>P2p transit LSP nexthop mode</b>	Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are <b>Chained</b> or <b>Unchained</b>
<b>P2mp transit LSP nexthop mode</b>	Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are <b>Chained</b> or <b>Unchained</b>

## Sample Output

### show rsvp version

```
user@host> show rsvp version
```

```
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol:           Enabled
  R(refresh timer):        30 seconds
  K(keep multiplier):      3
  Preemption:              Normal
  Soft-preemption cleanup:  30 seconds
  Graceful deletion timeout: 30 seconds
  NSR mode:                Enabled/Master
  NSR state:                Init sync complete
  Setup protection:        Disabled
  Graceful restart:        Disabled
  Restart helper mode:      Enabled
  Maximum helper restart time: 20000 msec
  Maximum helper recovery time: 180000 msec
  Restart time:            0 msec
  P2p transit LSP nexthop mode: Unchained
  P2mp transit LSP nexthop mode: Unchained
```

## show ted database

<b>List of Syntax</b>	<a href="#">Syntax on page 4714</a> <a href="#">Syntax (EX Series Switches) on page 4714</a>
<b>Syntax</b>	<pre>show ted database &lt;brief   detail   extensive&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>system-name</i>&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show ted database &lt;brief   detail   extensive&gt; &lt;<i>system-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p>
<b>Description</b>	Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.
<b>Options</b>	<p><b>none</b>—Display standard information about all entries in the traffic engineering database.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>system-name</i></b>—(Optional) Display traffic engineering database information for a particular system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ted database brief on page 4716</a> <a href="#">show ted database detail on page 4717</a> <a href="#">show ted database extensive on page 4718</a>
<b>Output Fields</b>	<p><a href="#">Table 381 on page 4714</a> describes the output fields for the <b>show ted database</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 381: show ted database Output Fields**

Field Name	Field Description	Level of Output
TED database	Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.	All levels
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	<b>brief</b>

Table 381: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
NodeID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	extensive
Type	Type of node. It can be either <b>Rtr</b> (router) or <b>Net</b> (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: <ul style="list-style-type: none"> <li>IS-IS(1)—IS-IS Level 1.</li> <li>IS-IS(2)—IS-IS Level 2.</li> <li>OSPF (area-number)—OSPF from the specified area.</li> </ul>	All levels
To	Address on the far end of a link.	detail extensive
Local	Address of the local interface being used to reach the remote node.	detail extensive
Remote	Address of the interface on the remote node.	detail extensive
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail extensive
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail extensive
Metric	Configured traffic engineering metric.	extensive
Static BW	Total interface bandwidth in bps.	extensive
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the <b>subscription</b> statement when configuring RSVP.	extensive
Available BW [priority]	(Must include <b>diffserv-te</b> statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	extensive
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	extensive
Available BW [TE-class]	(Must include the <b>diffserv-te</b> statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	extensive

Table 381: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Static BW [CT-class]</b>	Total interface bandwidth used by an MPLS traffic class, in bps.	<b>extensive</b>
Interface Switching Capability Descriptor ( <i>n</i> )	<p>Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. <i>n</i> is the index number.</p> <ul style="list-style-type: none"> <li>• <b>Switching type</b>—Type of switching to be performed on a particular link: <ul style="list-style-type: none"> <li>• <b>PSC-1</b>—Packet switch-capable 1</li> <li>• <b>PSC-2</b>—Packet switch-capable 2</li> <li>• <b>PSC-3</b>—Packet switch-capable 3</li> <li>• <b>PSC-4</b>—Packet switch-capable 4</li> <li>• <b>L2SC</b>—Layer-2-switch-capable</li> <li>• <b>TDM</b>—Time-division-multiplexing-capable</li> <li>• <b>LSC</b>—Lambda switch-capable</li> <li>• <b>FSC</b>—Fiber switch-capable</li> </ul> </li> <li>• <b>Encoding type</b>—Encoding of the LSP being requested: <ul style="list-style-type: none"> <li>• <b>Packet</b></li> <li>• <b>Ethernet</b></li> <li>• <b>ANSI/ETSI PDH</b></li> <li>• <b>Reserved</b></li> <li>• <b>SDH /SONET</b></li> <li>• <b>Digital Wrapper</b></li> <li>• <b>Lambda (photonic)</b></li> <li>• <b>Fiber</b></li> <li>• <b>FiberSDH/SONET</b></li> </ul> </li> <li>• <b>Maximum LSP BW [priority] bps</b>—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <ul style="list-style-type: none"> <li>• <b>[<i>n</i>]</b>—Priority level. The range is from <b>0</b> (high) through <b>7</b> (low).</li> <li>• <b><i>n</i> Mbps</b>—Amount of the maximum bandwidth.</li> </ul> </li> <li>• <b>Minimum LSP BW</b>—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <b>Minimum LSP BW</b> is displayed only when <b>switching type</b> is <b>PSC-1</b> or <b>TDM</b>.</li> <li>• <b>Interface MTU</b>—Displayed only when <b>switching type</b> is <b>TDM</b>.</li> <li>• <b>Interface supports standard SONET/SDH</b>—Displayed only when <b>switching type</b> is <b>TDM</b>.</li> </ul>	<b>extensive</b>

## Sample Output

### show ted database brief

```

user@host> show ted database brief
TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---   3178    2     0
Router-B.00                      ---   3152    2     0

```

```

Router-B.02          Net      802      0      2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-C.00          ---      3126      2      0
Router-C.02          Net       38      0      2 IS-IS(2)
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-D.00          ---      3144      2      0
Router-D.02          Net       723      0      2 IS-IS(2)
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-D.03          Net       607      0      2 IS-IS(2)
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-E.00          ---      3178      2      0
Router-E.02          Net       131      0      2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-F.00          ---      3153      2      0
Router-F.02          Net       769      0      2 IS-IS(2)
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```

#### show ted database detail

```

TED database: 12 ISIS nodes 0 INET nodes
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-A.00          ---      2913      2      0
Router-B.00          ---      2887      2      0
Router-B.02          Net       537      0      2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-C.00          ---      2861      2      0
Router-C.02          Net       597      0      2 IS-IS(2)
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                    Type Age(s) LnkIn LnkOut Protocol
Router-D.00          ---      2879      2      0
Router-D.02          Net       458      0      2 IS-IS(2)

```

```

To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
ID                                     Type Age(s) LnkIn LnkOut Protocol
Router-D.03                           Net   342    0      2 IS-IS(2)
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
ID                                     Type Age(s) LnkIn LnkOut Protocol
Router-E.00                           ---  2913    2      0
Router-E.02                           Net   640    0      2 IS-IS(2)
To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
ID                                     Type Age(s) LnkIn LnkOut Protocol
Router-F.00                           ---  2888    2      0
Router-F.02                           Net   504    0      2 IS-IS(2)
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0

```

### show ted database extensive

```

user@host> show ted database extensive
TED database: 12 ISIS nodes 0 INET nodes
NodeID: Router-A.00
Type: ---, Age: 3067 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.00
Type: ---, Age: 3041 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.02
Type: Net, Age: 691 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-C.00
Type: ---, Age: 3015 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-C.02
Type: Net, Age: 751 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0

```



```

Metric: 0
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-D.00
Type: ---, Age: 3034 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-D.02
Type: Net, Age: 613 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-D.03
Type: Net, Age: 497 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-E.00

```

```
Type: ---, Age: 3068 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-E.02
Type: Net, Age: 21 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-F.00
Type: ---, Age: 3043 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-F.02
Type: Net, Age: 659 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
```

## show ted link

<b>List of Syntax</b>	<a href="#">Syntax on page 4721</a> <a href="#">Syntax (EX Series Switches) on page 4721</a>
<b>Syntax</b>	show ted link <brief   detail> <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show ted link <brief   detail>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.
<b>Options</b>	<b>none</b> —Display standard information about traffic engineering database link information.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ted link brief on page 4722</a> <a href="#">show ted link detail on page 4722</a>
<b>Output Fields</b>	<a href="#">Table 382 on page 4721</a> describes the output fields for the <b>show ted link</b> command. Output fields are listed in the approximate order in which they appear.

**Table 382: show ted link Output Fields**

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>brief</b>
-->ID	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>brief</b>
<i>hostname</i>	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>detail</b>
<i>hostname</i>	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>detail</b>

Table 382: show ted link Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels
Local	Address of the local interface being used to reach the remote node.	<b>detail extensive</b>
Remote	Address of the interface on the remote node.	<b>detail extensive</b>
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	<b>detail</b>
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	<b>detail</b>

## Sample Output

### show ted link brief

```

user@host> show ted link brief
ID                               ->ID                               LocalPath LocalBW
Router-B.02                      Router-A.00                        0 0bps
Router-B.02                      Router-B.00                        0 0bps
Router-C.02                      Router-B.00                        0 0bps
Router-C.02                      Router-C.00                        0 0bps
Router-D.02                      Router-F.00                        0 0bps
Router-D.02                      Router-D.00                        0 0bps
Router-D.03                      Router-D.00                        0 0bps
Router-D.03                      Router-C.00                        0 0bps
Router-E.02                      Router-A.00                        0 0bps
Router-E.02                      Router-E.00                        0 0bps
Router-F.02                      Router-E.00                        0 0bps
Router-F.02                      Router-F.00                        0 0bps

```

### show ted link detail

```

user@host> show ted link detail
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-B.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-C.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-C.02->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0

```

```

LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-D.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-D.02->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-D.03->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-D.03->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-E.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-E.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-F.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
Router-F.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
LocalPath: 0, Metric: 0, AvailBW: 0bps
localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps

```

## show ted protocol

<b>List of Syntax</b>	<a href="#">Syntax on page 4724</a> <a href="#">Syntax (EX Series Switches) on page 4724</a>
<b>Syntax</b>	show ted protocol <brief   detail> <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show ted protocol <brief   detail>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.
<b>Options</b>	<p><b>none</b>—Display standard information about the protocols from which the traffic engineering database learned about its nodes.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ted protocol on page 4725</a>
<b>Output Fields</b>	<a href="#">Table 383 on page 4724</a> describes the output fields for the <b>show ted protocol</b> command. Output fields are listed in the approximate order in which they appear.

**Table 383: show ted protocol Output Fields**

Field Name	Field Description
<b>Protocol name</b>	Protocol that reported the node information: <ul style="list-style-type: none"> <li><b>IS-IS(1)</b>—IS-IS Level 1.</li> <li><b>IS-IS(2)</b>—IS-IS Level 2.</li> <li><b>OSPF (<i>area-number</i>)</b>—OSPF from the specified area.</li> </ul>
<b>Credibility</b>	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
<b>Self node</b>	Address the protocol uses as the local address.

## Sample Output

show ted protocol

```
user@host> show ted protocol
Protocol name      Credibility Self node
IS-IS(2)           2 (highest) corriedale.00(123.456.1.11)
IS-IS(1)           1          corriedale.00(123.456.1.11)
```

## traceroute mpls ldp

---

**Syntax** `traceroute mpls <ldp> fec`  
`<destination>`  
`<detail>`  
`<exp>`  
`<fanout>`  
`<logical-system>`  
`<no-resolve>`  
`<paths>`  
`<retries>`  
`<routing-instance>`  
`<source>`  
`<ttl>`  
`<update>`  
`<wait>`

**Release Information** Command introduced in Junos OS Release 8.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Trace route to a remote host for an MPLS label-switched path signaled by the LDP. Use **traceroute mpls ldp** as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

**Options** *fec*—Specify the IP address and optional prefix of the forwarding equivalence class (FEC).  
*destination*—(Optional) Specify the destination address to use when sending probes.  
*detail*—(Optional) Display detailed output.  
*exp*—(Optional) Specify the class-of-service to use when sending probes. The range of values is 0 through 7. The default value is 7.  
*fanout*—(Optional) Specify the maximum number of nexthops to search per node. The range of values is 1 through 16. The default value is 16.  
*logical-system*—(Optional) Specify the name of the logical system for the traceroute attempt.  
*no-resolve*—(Optional) Specify not to resolve the hostname that corresponds to the IP address.  
*paths*—(Optional) Specify the number of paths to search. The range of values is 1 through 255. The default value is 16.  
*retries*—(Optional) Specify the number of times to resend probe. values. The range of values is 1 through 9. The default value is 3.  
*routing-instance routing-instance-name*—(Optional) Specify the name of the routing instance for the traceroute attempt.  
*source source-address*—(Optional) Specify the source address of the outgoing traceroute packets.



**ttl value**—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. The range of values is **1** through **125** and the default value is **64**.

**wait seconds**—(Optional) Specify the number of seconds to wait before resending a probe. The range of values is **5** through **15** and the default value is **10** seconds.

**Required Privilege Level** network

**List of Sample Output** [traceroute mpls ldp on page 4728](#)  
[traceroute mpls ldp detail on page 4728](#)

**Output Fields** [Table 384 on page 4727](#) describes the output fields for the **traceroute mpls ldp fec** command and the **traceroute mpls ldp fec detail** commands. Output fields are listed in the approximate order in which they appear.

**Table 384: traceroute mpls ldp Output Fields**

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the <b>traceroute mpls ldp fec</b> command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is <b>null</b> .	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	<b>detail</b>
Parent	Address of the previous hop. Parent value for the first hop is <b>null</b> .	<b>detail</b>
Return Code	Return code for reporting the result of processing the echo request by the receiver.	<b>detail</b>
Response time	Time for the echo request to reach the receiver.	<b>detail</b>
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is <b>none</b> .	<b>detail</b>

Table 384: traceroute mpls ldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Label Stack	Label stack used to forward the packet.	<b>detail</b>

## Sample Output

### traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4
```

```
Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
ttl  Label Protocol Address Previous Hop Probe Status
 1  100016 LDP      24.24.24.1 (null) Success
 2  100000 LDP      20.20.20.2 24.24.24.1 Success
 3      3 LDP      22.22.22.4 20.20.20.2 Egress
```

```
Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

### traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail
```

```
Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
Hop 24.24.24.1 Depth 1
  Parent (null)
  Return code: Label switched at stack-depth 1
  Response time 165.93 msec
  Multipath type: IP bitmask
  Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100032 Protocol LDP

Hop 20.20.20.2 Depth 2
  Parent 24.24.24.1
  Return code: Upstream interface index unknown label-switched at stack-depth
1
  Response time 19.05 msec
  Multipath type: IP bitmask
  Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100000 Protocol LDP

Hop 22.22.22.4 Depth 3
  Parent 20.20.20.2
  Return code: Egress-ok at stack-depth 1
  Response time 0.79 msec
  Multipath type: None
  Label Stack:
    Label 1 Value 3 Protocol LDP
```

## traceroute mpls rsvp

<b>Syntax</b>	<pre>traceroute mpls &lt;rsvp&gt; <i>lsp-name</i> &lt;detail&gt; &lt;egress&gt; &lt;exp&gt; &lt;logical-system&gt; &lt;multipoint&gt; &lt;no-resolve&gt; &lt;retries&gt; &lt;source <i>source-address</i>&gt; &lt;ttl&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.2.</p> <p><b>egress</b>, <b>multipoint</b>, and <b>ttl</b> options added in Junos OS Release 11.2.</p>
<b>Description</b>	<p>Trace route to a remote host for an MPLS LSP signaled by RSVP. Use <b>traceroute mpls rsvp</b> as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)</p>
<b>Options</b>	<p><b><i>lsp-name</i></b>—Specify the name of the LSP to be traced.</p> <p><b>detail</b>—(Optional) Display detailed output.</p> <p><b>egress</b>—(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.</p> <p><b>exp</b>—(Optional) Specify the class of service to use when sending probes. The range of values is 0 through 7. The default value is 7.</p> <p><b>logical-system</b>—(Optional) Specify the name of the logical system for the traceroute attempt.</p> <p><b>multipoint</b>—(Optional) Perform a trace route on a point-to-multipoint LSP.</p> <p><b>no-resolve</b>—(Optional) Specify not to resolve the hostname that corresponds to the IP address.</p> <p><b>retries</b>—(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.</p> <p><b>source <i>source-address</i></b>—(Optional) Specify the source address of the outgoing traceroute packets.</p> <p><b>ttl</b>—(Optional) Specify the number of hops to follow before forcing the trace route to quit.</p>
<b>Required Privilege Level</b>	network
<b>List of Sample Output</b>	<p><a href="#">traceroute mpls rsvp on page 4731</a></p> <p><a href="#">traceroute mpls rsvp detail on page 4731</a></p>

[traceroute mpls rsvp multipoint \(branch node for sub-LSPs\) on page 4732](#)  
[traceroute mpls rsvp multipoint \(single-hop sub-LSPs\) on page 4732](#)

**Output Fields** Table 385 on page 4730 describes the output fields for the **traceroute mpls rsvp *lsp-name*** and **traceroute mpls rsvp *lsp-name* detail** commands. Output fields are listed in the approximate order in which they appear.

**Table 385: traceroute mpls rsvp Output Fields**

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the <b>traceroute mpls rsvp <i>lsp-name</i></b> command.	all levels
ttl	Time-to-live value of the labeled packet.	none specified
Label	MPLS label used to forward the packets along the LSP.	none specified
Protocol	Signaling protocol used. For this command, it is RSVP-TE.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays <b>Success</b> if the trace to a hop is successful or <b>Egress</b> if the trace has reached the last router on the path.	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	<b>detail</b>
Parent	Address of the previous hop. Parent value for the first hop is null.	<b>detail</b>
Return Code	Return code for reporting the result of processing the echo request by the receiver.	<b>detail</b>
Sender timestamp	Displays the timestamp when the MPLS echo request is sent to the next hop.	<b>detail</b>
Receiver timestamp	Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.	<b>detail</b>
Response time	Time for the echo request to reach the receiver.	<b>detail</b>
MTU	Size of the largest packet that includes the label stack forwarded to the next hop.	<b>detail</b>

Table 385: traceroute mpls rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none.	<b>detail</b>
Label stack	Label stack used to forward the packet.	<b>detail</b>
Path	Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.	all levels

## Sample Output

### traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta
```

```
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	299792	RSVP-TE	192.168.1.2	(null)	Success
2	299803	RSVP-TE	192.168.2.3	192.168.1.2	Success
3	3	RSVP-TE	192.168.3.4	192.168.2.3	Egress

```
Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

### traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
```

```
Probe options: retries 3, exp 7
```

```
Hop 192.168.1.2 Depth 1
```

```
Probe status: Success
```

```
Parent: (null)
```

```
Return code: Label-switched at stack-depth 1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
```

```
Response time: 26.99 msec
```

```
MTU: Unknown
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299792 Protocol RSVP-TE
```

```
Hop 192.168.2.3 Depth 2
```

```
Probe status: Success
```

```
Parent: 192.168.1.2
```

```
Return code: Upstream interface index unknown label-switched at stack-depth
```

```
1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
```

```
Response time: 26.55 msec
```

```
MTU: 1518
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299803 Protocol RSVP-TE
```

### traceroute mpls rsvp multipoint (branch node for sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	300000	RSVP-TE	81.1.2.2	(null)	Success
2	299968	RSVP-TE	81.2.3.3	81.1.2.2	Success
3	299952	RSVP-TE	81.3.4.4	81.2.3.3	Success
4	299920	RSVP-TE	81.4.6.6	81.3.4.4	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
4	299920	RSVP-TE	81.4.5.5	81.3.4.4	Egress

Path 2 via lt-1/2/0.102 destination 127.0.0.64

### traceroute mpls rsvp multipoint (single-hop sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.2.2	(null)	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.8.8	(null)	Egress

Path 2 via lt-1/2/0.108 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.9.9	(null)	Egress

Path 3 via lt-1/2/0.109 destination 127.0.0.64

## CHAPTER 56

# Troubleshooting

- [Troubleshooting Procedures on page 4733](#)

## Troubleshooting Procedures

---

- [Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch on page 4733](#)

### Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch

The following issues exist in the operation of MPLS features on QFX Series devices and on the EX4600 switch. In each case, the described behavior is the expected behavior.

- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.

#### Related Documentation

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)





## PART 16

# Multicast

- [Overview on page 4737](#)
- [Configuration on page 4777](#)
- [Administration on page 5047](#)



## CHAPTER 57

# Overview

- [Introduction to PIM Basics on page 4737](#)
- [Introduction to PIM Sparse Mode on page 4741](#)
- [Introduction to Static RP on page 4744](#)
- [Introduction to Anycast RP on page 4745](#)
- [Introduction to PIM Bootstrap Router on page 4745](#)
- [Introduction to PIM Filtering on page 4746](#)
- [Introduction to PIM RPT and SPT Cutover on page 4748](#)
- [Introduction to IGMP on page 4757](#)
- [Introduction to IGMP Snooping on page 4761](#)
- [Introduction to MLD on page 4764](#)
- [Introduction to MSDP on page 4767](#)
- [Introduction to Source-Specific Multicast on page 4769](#)
- [Introduction to Multicast VLAN Registration on page 4773](#)

## Introduction to PIM Basics

---

- [PIM Overview on page 4737](#)
- [PIM on Aggregated Interfaces on page 4740](#)

## PIM Overview

The predominant multicast routing protocol in use on the Internet today is Protocol Independent Multicast, or PIM. The type of PIM used on the Internet is PIM sparse mode. PIM sparse mode is so accepted that when the simple term “PIM” is used in an Internet context, some form of sparse mode operation is assumed.

PIM emerged as an algorithm to overcome the limitations of dense-mode protocols such as the Distance Vector Multicast Routing Protocol (DVMRP), which was efficient for dense clusters of multicast receivers, but did not scale well for the larger, sparser, groups encountered on the Internet. The Core Based Trees (CBT) Protocol was intended to support sparse mode as well, but CBT, with its all-powerful core approach, made placement of the core critical, and large conference-type applications (many-to-many)

resulted in bottlenecks in the core. PIM was designed to avoid the dense-mode scaling issues of DVMRP and the potential performance issues of CBT at the same time.

PIM is one of the most rapidly evolving specifications on the Internet today. Since its introduction in 1995, PIM has already seen two major revisions to its packet structure (PIM version 1 [PIMv1] and PIM version 2 [PIMv2]), two major RFCs (RFC 2362 obsoleted RFC 2117), and numerous drafts describing major components of PIM, such as many-to-many trees and source-specific multicast (SSM). Long-lasting RFCs are not a feature of PIM, and virtually all of PIM must be researched, understood, and implemented directly from Internet drafts. In fact, no current RFC describes PIMv1 at all. The drafts have all expired, and PIMv1 was never issued as an official RFC.

PIM itself is not nonstandard or unstable, however. PIM has been a promising multicast routing protocol since its inception, especially PIM sparse mode, the first real sparse-mode multicast routing protocol. Work continues on PIM in a number of areas, from bidirectional trees to network management, and the rapid pace of development makes drafts essential for PIM.

PIMv1 and PIMv2 can coexist on the same routing device and even on the same interface. The main difference between PIMv1 and PIMv2 is the packet format. PIMv1 messages use Internet Group Management Protocol (IGMP) packets, whereas PIMv2 has its own IP protocol number (103) and packet structure. All routing devices connecting to an IP subnet such as a LAN must use the same PIM version. Some PIM implementations can recognize PIMv1 packets and automatically switch the routing device interface to PIMv1. Because the difference between PIMv1 and PIMv2 involves the message format, but not the meaning of the message or how the routing device processes the PIM message, a routing device can easily mix PIMv1 and PIMv2 interfaces.

PIM is used for efficient routing to multicast groups that might span wide-area and interdomain internetworks. It is called “protocol independent” because it does not depend on a particular unicast routing protocol. Junos OS supports bidirectional mode, sparse mode, dense mode, and sparse-dense mode.

PIM operates in several modes: bidirectional mode, sparse mode, dense mode, and sparse-dense mode. In sparse-dense mode, some multicast groups are configured as dense mode (flood-and-prune, [S,G] state) and others are configured as sparse mode (explicit join to rendezvous point [RP], [\*G] state).

PIM drafts also establish a mode known as PIM source-specific mode, or PIM SSM. In PIM SSM there is only one specific source for the content of a multicast group within a given domain.

Because the PIM mode you choose determines the PIM configuration properties, you first must decide whether PIM operates in bidirectional, sparse, dense, or sparse-dense mode in your network. Each mode has distinct operating advantages in different network environments.

- In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to a downstream routing device unless the downstream routing device has sent an explicit request (by means of a join message) to the rendezvous point (RP) routing device to receive this traffic. The RP serves as the root of the shared multicast delivery tree and is responsible for forwarding multicast data from different sources to the receivers.

Sparse mode is well suited to the Internet, where frequent interdomain join messages and prune messages are common.



**NOTE:** On all the EX series switches (except EX4300 and EX9200) and QFX5100 series switches, the rate limit is set to 1pps per SG to avoid overwhelming the rendezvous point (RP), First hop router (FHR) with PIM-sparse mode (PIM-SM) register messages and cause CPU hogs. This rate limit helps in improving scaling and convergence times by avoiding duplicate packets being trapped, and tunneled to RP in software.

- Bidirectional PIM is similar to sparse mode, and is especially suited to applications that must scale to support a large number of dispersed sources and receivers. In bidirectional PIM, routing devices build shared bidirectional trees and do not switch to a source-based tree. Bidirectional PIM scales well because it needs no source-specific (S,G) state. Instead, it builds only group-specific (\*,G) state.
- Unlike sparse mode and bidirectional mode, in which data is forwarded only to routing devices sending an explicit PIM join request, dense mode implements a *flood-and-prune* mechanism, similar to the Distance Vector Multicast Routing Protocol (DVMRP). In dense mode, a routing device receives the multicast data on the incoming interface, then forwards the traffic to the outgoing interface list. Flooding occurs periodically and is used to refresh state information, such as the source IP address and multicast group pair. If the routing device has no interested receivers for the data, and the outgoing interface list becomes empty, the routing device sends a PIM prune message upstream.

Dense mode works best in networks where few or no prunes occur. In such instances, dense mode is actually more efficient than sparse mode.

- Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as “dense” is not mapped to an RP. Instead, data packets destined for that group are forwarded by means of PIM dense mode rules. A group specified as “sparse” is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. Sparse-dense mode is useful in networks implementing auto-RP for PIM sparse mode.

## Basic PIM Network Components

---

PIM dense mode requires only a multicast source and series of multicast-enabled routing devices running PIM dense mode to allow receivers to obtain multicast content. Dense mode makes sure that all multicast traffic gets everywhere by periodically flooding the network with multicast traffic, and relies on prune messages to make sure that subnets where all receivers are uninterested in that particular multicast group stop receiving packets.

PIM sparse mode is more complicated and requires the establishment of special routing devices called *rendezvous points (RPs)* in the network core. These routing devices are where upstream join messages from interested receivers meet downstream traffic from the source of the multicast group content. A network can have many RPs, but PIM sparse mode allows only one RP to be active for any multicast group.

If there is only one RP in a routing domain, the RP and adjacent links might become congested and form a single point of failure for all multicast traffic. Thus, multiple RPs are the rule, but the issue then becomes how other multicast routing devices find the RP that is the source of the multicast group the receiver is trying to join. This RP-to-group mapping is controlled by a special *bootstrap router (BSR)* running the PIM BSR mechanism. There can be more than one bootstrap router as well, also for single-point-of-failure reasons.

The bootstrap router does not have to be an RP itself, although this is a common implementation. The bootstrap router's main function is to manage the collection of RPs and allow interested receivers to find the source of their group's multicast traffic. PIM bootstrap messages are sourced from the loopback address, which is always up. The loopback address must be routable. If it is not routable, then the bootstrap router is unable to send bootstrap messages to update the RP domain members. The **show pim bootstrap** command displays only those bootstrap routers that have routable loopback addresses.

PIM SSM can be seen as a subset of a special case of PIM sparse mode and requires no specialized equipment other than that used for PIM sparse mode (and IGMP version 3).

Bidirectional PIM RPs, unlike RPs for PIM sparse mode, do not need to perform PIM Register tunneling or other specific protocol action. Bidirectional PIM RPs implement no specific functionality. RP addresses are simply a location in the network to rendezvous toward. In fact, for bidirectional PIM, RP addresses need not be loopback interface addresses or even be addresses configured on any routing device, as long as they are covered by a subnet that is connected to a bidirectional PIM-capable routing device and advertised to the network.

**Related Documentation**

- *Supported IP Multicast Protocol Standards* in the *Multicast Protocols Feature Guide for Routing Devices*

## PIM on Aggregated Interfaces

If you configure PIM on an aggregated (**ae-** or **as-**) interface, each of the interfaces in the aggregate is included in the multicast output interface list and carries the single stream

of replicated packets in a load-sharing fashion. The multicast aggregate interface is “expanded” into its constituent interfaces in the next-hop database.

- Related Documentation**
- [PIM Overview on page 4737](#)
  - [interface on page 4931](#)

## Introduction to PIM Sparse Mode

- [Understanding PIM Sparse Mode on page 4741](#)
- [Designated Router on page 4744](#)

### Understanding PIM Sparse Mode

A Protocol Independent Multicast (PIM) sparse-mode domain uses reverse-path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A (\*,G) PIM join message is sent toward the RP from the receiver's designated router (DR). (By definition, this message is actually called a join/prune message, but for clarity in this description, it is called either join or prune, depending on its context.) The join message is multicast hop by hop upstream to the ALL-PIM-ROUTERS group (224.0.0.13) by means of each router's RPF interface until it reaches the RP. The RP router receives the (\*,G) PIM join message and adds the interface on which it was received to the outgoing interface list (OIL) of the rendezvous-point tree (RPT) forwarding state entry. This builds the RPT connecting the receiver with the RP. The RPT remains in effect, even if no active sources generate traffic.



**NOTE:** State—the (\*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. S is the source IP address, G is the multicast group address, and \* represents any source sending to group G. Routers keep track of the multicast forwarding state for the incoming and outgoing interfaces for each group.

When a source becomes active, the source DR encapsulates multicast data packets into a PIM register message and sends them by means of unicast to the RP router.

If the RP router has interested receivers in the PIM sparse-mode domain, it sends a PIM join message toward the source to build a shortest-path tree (SPT) back to the source. The source sends multicast packets out on the LAN, and the source DR encapsulates the packets in a PIM register message and forwards the message toward the RP router by means of unicast. The RP router receives PIM register messages back from the source, and thus adds a new source to the distribution tree, keeping track of sources in a PIM table. Once an RP router receives packets natively (with S,G), it sends a register stop message to stop receiving the register messages by means of unicast.

In actual application, many receivers with multiple SPTs are involved in a multicast traffic flow. To illustrate the process, we track the multicast traffic from the RP router to one receiver. In such a case, the RP router begins sending multicast packets down the RPT

toward the receiver's DR for delivery to the interested receivers. When the receiver's DR receives the first packet from the RPT, the DR sends a PIM join message toward the source DR to start building an SPT back to the source. When the source DR receives the PIM join message from the receiver's DR, it starts sending traffic down all SPTs. When the first multicast packet is received by the receiver's DR, the receiver's DR sends a PIM prune message to the RP router to stop duplicate packets from being sent through the RPT. In turn, the RP router stops sending multicast packets to the receiver's DR, and sends a PIM prune message for this source over the RPT toward the source DR to halt multicast packet delivery to the RP router from that particular source.

If the RP router receives a PIM register message from an active source but has no interested receivers in the PIM sparse-mode domain, it still adds the active source into the PIM table. However, after adding the active source into the PIM table, the RP router sends a register stop message. The RP router discovers the active source's existence and no longer needs to receive advertisement of the source (which utilizes resources).



**NOTE:** If the number of PIM join messages exceeds the configured MTU, the messages are fragmented in IPv6 PIM sparse mode. To avoid the fragmentation of PIM join messages, the multicast traffic receives the interface MTU instead of the path MTU.

---

The major characteristics of PIM sparse mode are as follows:

- Routers with downstream receivers join a PIM sparse-mode tree through an explicit join message.
- PIM sparse-mode RPs are the routers where receivers meet sources.
- Senders announce their existence to one or more RPs, and receivers query RPs to find multicast sessions.
- Once receivers get content from sources through the RP, the last-hop router (the router closest to the receiver) can optionally remove the RP from the shared distribution tree (\*;G) if the new source-based tree (S,G) is shorter. Receivers can then get content directly from the source.

The transitional aspect of PIM sparse mode from shared to source-based tree is one of the major features of PIM, because it prevents overloading the RP or surrounding core links.

There are related issues regarding source, RPs, and receivers when sparse mode multicast is used:

- Sources must be able to send to all RPs.
- RPs must all know one another.
- Receivers must send explicit join messages to a known RP.
- Receivers initially need to know only one RP (they later learn about others).



- Receivers can explicitly prune themselves from a tree.
- Receivers that never transition to a source-based tree are effectively running Core Based Trees (CBT).


PIM sparse mode has standard features for all of these issues.

### Rendezvous Point

The RP router serves as the information exchange point for the other routers. All routers in a PIM domain must provide mapping to an RP router. It is the only router that needs to know the active sources for a domain—the other routers just need to know how to reach the RP. In this way, the RP matches receivers with sources.

The RP router is downstream from the source and forms one end of the shortest-path tree. As shown in [Figure 14-6 on page 4743](#), the RP router is upstream from the receiver and thus forms one end of the rendezvous-point tree.

**Figure 14-6: Rendezvous Point as Part of the RPT and SPT**



The benefit of using the RP as the information exchange point is that it reduces the amount of state in non-RP routers. No network flooding is required to provide non-RP routers information about active sources.

### RP Mapping Options

RPs can be learned by one of the following mechanisms:

- Static configuration
- Anycast RP
- Auto-RP
- Bootstrap router

We recommend a static RP mapping with anycast RP and a bootstrap router (BSR) with auto-RP configuration, because static mapping provides all the benefits of a bootstrap router and auto-RP without the complexity of the full BSR and auto-RP mechanisms.

#### Related Documentation

- [Understanding Static RP on page 4744](#)
- [Understanding RP Mapping with Anycast RP on page 4745](#)
- [Understanding the PIM Bootstrap Router on page 4745](#)
- [Understanding PIM Auto-RP](#)

## Designated Router

In a PIM sparse mode (PIM-SM) domain, there are two types of designated routers to consider:

- The receiver DR sends PIM join and PIM prune messages from the receiver network toward the RP.
- The source DR sends PIM register messages from the source network to the RP.

Neighboring PIM routers multicast periodic PIM hello messages to each other every 30 seconds (the default). The PIM hello message usually includes a holdtime value for the neighbor to use, but this is not a requirement. If the PIM hello message does not include a holdtime value, a default timeout value (in Junos OS, 105 seconds) is used. On receipt of a PIM hello message, a router stores the IP address and priority for that neighbor. If the DR priorities match, the router with the highest IP address is selected as the DR.

If a DR fails, a new one is selected using the same process of comparing IP addresses.



**NOTE:** In PIM dense mode (PIM-DM), a DR is elected by the same process that PIM-SM uses. However, the only time that a DR has any effect in PIM-DM is when IGMPv1 is used on the interface. (IGMPv2 is the default.) In this case, the DR also functions as the IGMP Query Router because IGMPv1 does not have a Query Router election mechanism.

---

## Introduction to Static RP

- [Understanding Static RP on page 4744](#)

### Understanding Static RP

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the device. However, because PIM must not be configured on the network management interface, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive multicast traffic from the groups served by the RP, the device must determine the IP address of the RP for the source.

You can configure a static rendezvous point (RP) configuration that is similar to static routes. A static configuration has the benefit of operating in PIM version 1 or version 2. When you configure the static RP, the RP address that you select for a particular group must be consistent across all routers in a multicast domain.

One common way for the device to locate RPs is by static configuration of the IP address of the RP. A static configuration is simple and convenient. However, if the statically

defined RP router becomes unreachable, there is no automatic failover to another RP router. To remedy this problem, you can use anycast RP.

- Related Documentation**
- [Configuring Local PIM RPs on page 4795](#)
  - [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4797](#)

## Introduction to Anycast RP

- [Understanding RP Mapping with Anycast RP on page 4745](#)

### Understanding RP Mapping with Anycast RP

Having a single active rendezvous point (RP) per multicast group is much the same as having a single server providing any service. All traffic converges on this single point, although other servers are sitting idle, and convergence is slow when the resource fails. In multicast specifically, there might be closer RPs on the shared tree, so the use of a single RP is suboptimal.

For the purposes of load balancing and redundancy, you can configure anycast RP. You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP fails, sources and receivers are taken to a new RP by means of unicast routing. When you configure anycast RP, you bypass the restriction of having one active RP per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

Anycast means that multiple RP routers share the same unicast IP address. Anycast addresses are advertised by the routing protocols. Packets sent to the anycast address are sent to the nearest RP with this address. Anycast addressing is a generic concept and is used in PIM sparse mode to add load balancing and service reliability to RPs.

Anycast RP is defined in Internet draft draft-ietf-mboned-anycast-rp-08.txt, *Anycast RP Mechanism Using PIM and MSDP*. To access Internet RFCs and drafts, go to the IETF website at <http://www.ietf.org>.

- Related Documentation**
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4797](#)
  - [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 4805](#)
  - [Example: Configuring PIM Anycast With or Without MSDP on page 4799](#)

## Introduction to PIM Bootstrap Router

- [Understanding the PIM Bootstrap Router on page 4745](#)

### Understanding the PIM Bootstrap Router

To determine which router is the rendezvous point (RP), all routers within a PIM sparse-mode domain collect bootstrap messages. A PIM sparse-mode domain is a group

of routers that all share the same RP router. The domain bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routers use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

**Related  
Documentation**

- [Configuring PIM Bootstrap Properties for IPv4 or IPv6](#)

---

## Introduction to PIM Filtering

- [Understanding Multicast Message Filters on page 4746](#)
- [Filtering MAC Addresses on page 4747](#)
- [Filtering RP and DR Register Messages on page 4747](#)

### Understanding Multicast Message Filters

Multicast sources and routers generate a considerable number of control messages, especially when using PIM sparse mode. These messages form distribution trees, locate rendezvous points (RPs) and designated routers (DRs), and transition from one type of tree to another. In most cases, this multicast messaging system operates transparently and efficiently. However, in some configurations, more control over the sending and receiving of multicast control messages is necessary.

You can configure multicast filtering to control the sending and receiving of multicast control messages.

To prevent unauthorized groups and sources from registering with an RP router, you can define a routing policy to reject PIM register messages from specific groups and sources and configure the policy on the designated router or the RP router.

- If you configure the reject policy on an RP router, it rejects incoming PIM register messages from the specified groups and sources. The RP router also sends a register stop message by means of unicast to the designated router. On receiving the register stop message, the designated router sends periodic null register messages for the specified groups and sources to the RP router.
- If you configure the reject policy on a designated router, it stops sending PIM register messages for the specified groups and sources to the RP router.



**NOTE:** If you have configured the reject policy on an RP router, we recommend that you configure the same policy on all the RP routers in your multicast network.

---



**NOTE:** If you delete a group and source address from the reject policy configured on an RP router and commit the configuration, the RP router will register the group and source only when the designated router sends a null register message.

---

- Related Documentation**
- [Filtering MAC Addresses on page 4747](#)
  - [Filtering RP and DR Register Messages on page 4747](#)
  - [Filtering MSDP SA Messages on page 4769](#)

## Filtering MAC Addresses

When a router is exclusively configured with multicast protocols on an interface, multicast sets the interface media access control (MAC) filter to multicast promiscuous mode, and the number of multicast groups is unlimited. However, when the router is not exclusively used for multicasting and other protocols such as OSPF, Routing Information Protocol version 2 (RIPv2), or Network Time Protocol (NTP) are configured on an interface, each of these protocols individually requests that the interface program the MAC filter to pick up its respective multicast group only. In this case, without multicast configured on the interface, the maximum number of multicast MAC filters is limited to 20. For example, the maximum number of interface MAC filters for protocols such as OSPF (multicast group 224.0.0.5) is 20, unless a multicast protocol is also configured on the interface.

No configuration is necessary for MAC filters.

## Filtering RP and DR Register Messages

You can filter Protocol Independent Multicast (PIM) register messages sent from the designated router (DR) or to the rendezvous point (RP). The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, more control over which sources an RP discovers, or which sources a DR notifies other RPs about, is desired. A high degree of control over PIM register messages is provided by RP and DR register message filtering. Message filtering also prevents unauthorized groups and sources from registering with an RP router.

Register messages that are filtered at a DR are not sent to the RP, but the sources are available to local users. Register messages that are filtered at an RP arrive from source DRs, but are ignored by the router. Sources on multicast group traffic can be limited or directed by using RP or DR register message filtering alone or together.

If the action of the register filter policy is to discard the register message, the router needs to send a register-stop message to the DR. Register-stop messages are throttled to prevent malicious users from triggering them on purpose to disrupt the routing process.

Multicast group and source information is encapsulated inside unicast IP packets. This feature allows the router to inspect the multicast group and source information before sending or accepting the PIM register message.

Incoming register messages to an RP are passed through the configured register message filtering policy before any further processing. If the register message is rejected, the RP router sends a register-stop message to the DR. When the DR receives the register-stop message, the DR stops sending register messages for the filtered groups and sources to the RP. Two fields are used for register message filtering:

- Group multicast address

- Source address

The syntax of the existing policy statements is used to configure the filtering on these two fields. The **route-filter** statement is useful for multicast group address filtering, and the **source-address-filter** statement is useful for source address filtering. In most cases, the action is to **reject** the register messages, but more complex filtering policies are possible.

Filtering cannot be performed on other header fields, such as DR address, protocol, or port. In some configurations, an RP might not send register-stop messages when the policy action is to discard the register messages. This has no effect on the operation of the feature, but the router will continue to receive register messages.

When anycast RP is configured, register messages can be sent or received by the RP. All the RPs in the anycast RP set need to be configured with the same RP register message filtering policies. Otherwise, it might be possible to circumvent the filtering policy.

#### Related Documentation

- [Understanding RP Mapping with Anycast RP on page 4745](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 4813](#)

---

## Introduction to PIM RPT and SPT Cutover

- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 4748](#)
- [Building an RPT Between the RP and Receivers on page 4749](#)
- [PIM Sparse Mode Source Registration on page 4750](#)
- [Multicast Shortest-Path Tree on page 4753](#)
- [SPT Cutover on page 4754](#)
- [SPT Cutover Control on page 4757](#)

## Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees

In a shared tree, the root of the distribution tree is a router, not a host, and is located somewhere in the core of the network. In the primary sparse mode multicast routing protocol, Protocol Independent Multicast sparse mode (PIM SM), the core router at the root of the shared tree is the rendezvous point (RP). Packets from the upstream source and join messages from the downstream routers “rendezvous” at this core router.

In the RP model, other routers do not need to know the addresses of the sources for every multicast group. All they need to know is the IP address of the RP router. The RP router discovers the sources for all multicast groups.

The RP model shifts the burden of finding sources of multicast content from each router (the (S,G) notation) to the network (the (\*,G) notation knows only the RP). Exactly how the RP finds the unicast IP address of the source varies, but there must be some method to determine the proper source for multicast content for a particular group.

Consider a set of multicast routers without any active multicast traffic for a certain group. When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the distribution tree for that group back to the RP, not to the actual source of the content.

To join the shared tree, or *rendezvous-point tree (RPT)* as it is called in PIM sparse mode, the router must do the following:

- Determine the IP address of the RP for that group. Determining the address can be as simple as static configuration in the router, or as complex as a set of nested protocols.
- Build the shared tree for that group. The router executes an RPF check on the RP address in its routing table, which produces the interface closest to the RP. The router now detects that multicast packets from this RP for this group need to flow into the router on this RPF interface.
- Send a join message out on this interface using the proper multicast protocol (probably PIM sparse mode) to inform the upstream router that it wants to join the shared tree for that group. This message is a (\*G) join message because S is not known. Only the RP is known, and the RP is not actually the source of the multicast packets. The router receiving the (\*G) join message adds the interface on which the message was received to its outgoing interface list (OIL) for the group and also performs an RPF check on the RP address. The upstream router then sends a (\*G) join message out from the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating join messages from the RPF interface, building the shared tree as it goes. The process stops when the join message reaches one of the following:

- The RP for the group that is being joined
- A router along the RPT that already has a multicast forwarding state for the group that is being joined

In either case, the branch is created, and packets can flow from the source to the RP and from the RP to the receiver. Note that there is no guarantee that the shared tree (RPT) is the shortest path tree to the source. Most likely it is not. However, there are ways to “migrate” a shared tree to an SPT once the flow of packets begins. In other words, the forwarding state can transition from (\*G) to (S,G). The formation of both types of tree depends heavily on the operation of the RPF check and the RPF table.

#### Related Documentation

- *Understanding Multicast Reverse Path Forwarding*

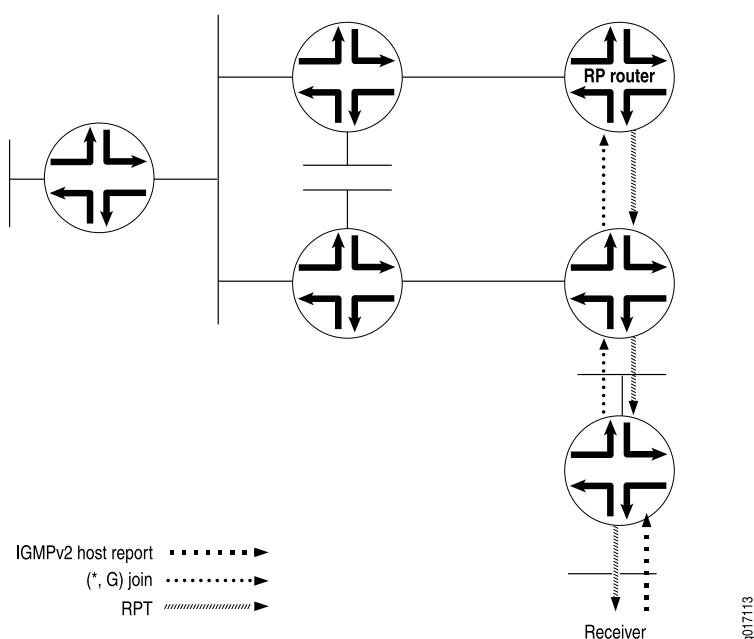
### Building an RPT Between the RP and Receivers

The RPT is the path between the RP and receivers (hosts) in a multicast group (see [Figure 147 on page 4750](#)). The RPT is built by means of a PIM join message from a receiver's DR:

1. A receiver sends a request to join group (G) in an Internet Group Management Protocol (IGMP) host membership report. A PIM sparse-mode router, the receiver's DR, receives

- the report on a directly attached subnet and creates an RPT branch for the multicast group of interest.
2. The receiver's DR sends a PIM join message to its RPF neighbor, the next-hop address in the RPF table, or the unicast routing table.
3. The PIM join message travels up the tree and is multicast to the ALL-PIM-ROUTERS group (224.0.0.13). Each router in the tree finds its RPF neighbor by using either the RPF table or the unicast routing table. This is done until the message reaches the RP and forms the RPT. Routers along the path set up the multicast forwarding state to forward requested multicast traffic back down the RPT to the receiver.

**Figure 147: Building an RPT Between the RP and the Receiver**



## PIM Sparse Mode Source Registration


The RPT is a unidirectional tree, permitting traffic to flow down from the RP to the receiver in one direction. For multicast traffic to reach the receiver from the source, another branch of the distribution tree, called the shortest-path tree, needs to be built from the source's DR to the RP.

The shortest-path tree is created in the following way:

1. The source becomes active, sending out multicast packets on the LAN to which it is attached. The source's DR receives the packets and encapsulates them in a PIM register message, which it sends to the RP router (see [Figure 14-8 on page 4751](#)).
2. When the RP router receives the PIM register message from the source, it sends a PIM join message back to the source.




Figure 148: PIM Register Message and PIM Join Message Exchanged




3. The source's DR receives the PIM join message and begins sending traffic down the SPT toward the RP router (see [Figure 149 on page 4752](#)).
4. Once traffic is received by the RP router, it sends a register stop message to the source's DR to stop the register process.

Figure 149: Traffic Sent from the Source to the RP Router



- The RP router sends the multicast traffic down the RPT toward the receiver (see [Figure 150 on page 4752](#)).

Figure 150: Traffic Sent from the RP Router Toward the Receiver



## Multicast Shortest-Path Tree

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT) as well. Consider a set of multicast routers without any active multicast traffic for a certain group (that is, they have no multicast forwarding state for that group). When a router learns that an interested receiver for that group is on one of its directly connected subnets, the router attempts to join the tree for that group.

To join the distribution tree, the router determines the unicast IP address of the source for that group. This address can be a simple static configuration on the router, or as complex as a set of protocols.

To build the SPT for that group, the router executes an a reverse path forwarding (RPF) check on the source address in its routing table. The RPF check produces the interface closest to the source, which is where multicast packets from this source for this group need to flow into the router.

The router next sends a join message out on this interface using the proper multicast protocol to inform the upstream router that it wants to join the distribution tree for that group. This message is an (S,G) join message because both S and G are known. The router receiving the (S,G) join message adds the interface on which the message was received to its output interface list (OIL) for the group and also performs an RPF check on the source address. The upstream router then sends an (S,G) join message out on the RPF interface toward the source, informing the upstream router that it also wants to join the group.

Each upstream router repeats this process, propagating joins out on the RPF interface, building the SPT as it goes. The process stops when the join message does one of two things:

- Reaches the router directly connected to the host that is the source.
- Reaches a router that already has multicast forwarding state for this source-group pair.

In either case, the branch is created, each of the routers has multicast forwarding state for the source-group pair, and packets can flow down the distribution tree from source to receiver. The RPF check at each router makes sure that the tree is an SPT.

SPTs are always the shortest path, but they are not necessarily short. That is, sources and receivers tend to be on the periphery of a router network, not on the backbone, and multicast distribution trees have a tendency to sprawl across almost every router in the network. Because multicast traffic can overwhelm a slow interface, and one packet can easily become a hundred or a thousand on the opposite side of the backbone, it makes sense to provide a shared tree as a distribution tree so that the multicast source can be located more centrally in the network, on the backbone. This sharing of distribution trees with roots in the core network is accomplished by a multicast rendezvous point.

### Related Documentation


- [Understanding Multicast Rendezvous Points, Shared Trees, and Rendezvous-Point Trees on page 4748](#)

## SPT Cutover

Instead of continuing to use the SPT to the RP and the RPT toward the receiver, a direct SPT is created between the source and the receiver in the following way:


1. Once the receiver's DR receives the first multicast packet from the source, the DR sends a PIM join message to its RPF neighbor (see [Figure 151 on page 4754](#)).
2. The source's DR receives the PIM join message, and an additional (S,G) state is created to form the SPT.
3. Multicast packets from that particular source begin coming from the source's DR and flowing down the new SPT to the receiver's DR. The receiver's DR is now receiving two copies of each multicast packet sent by the source—one from the RPT and one from the new SPT.

**Figure 151: Receiver DR Sends a PIM Join Message to the Source**




4. To stop duplicate multicast packets, the receiver's DR sends a PIM prune message toward the RP router, letting it know that the multicast packets from this particular source coming in from the RPT are no longer needed (see [Figure 152 on page 4755](#)).

Figure 152: PIM Prune Message Is Sent from the Receiver's DR Toward the RP Router




5. The PIM prune message is received by the RP router, and it stops sending multicast packets down to the receiver's DR. The receiver's DR is getting multicast packets only for this particular source over the new SPT. However, multicast packets from the source are still arriving from the source's DR toward the RP router (see [Figure 153 on page 4755](#)).

Figure 153: RP Router Receives PIM Prune Message




6. To stop the unneeded multicast packets from this particular source, the RP router sends a PIM prune message to the source's DR (see [Figure 154 on page 4756](#)).

**Figure 154: RP Router Sends a PIM Prune Message to the Source DR**



7. The receiver's DR now receives multicast packets only for the particular source from the SPT (see [Figure 155 on page 4756](#)).

**Figure 155: Source's DR Stops Sending Duplicate Multicast Packets Toward the RP Router**



## SPT Cutover Control

In some cases, the last-hop router needs to stay on the shared tree to the RP and not transition to a direct SPT to the source. You might not want the last-hop router to transition when, for example, a low-bandwidth multicast stream is forwarded from the RP to a last-hop router. All routers between last hop and source must maintain and refresh the SPT state. This can become a resource-intensive activity that does not add much to the network efficiency for a particular pair of source and multicast group addresses.

In these cases, you configure an SPT threshold policy on the last-hop router to control the transition to a direct SPT. An SPT cutover threshold of infinity applied to a source-group address pair means the last-hop router will never transition to a direct SPT. For all other source-group address pairs, the last-hop router transitions immediately to a direct SPT rooted at the source DR.

---

## Introduction to IGMP

- [Understanding Group Membership Protocols on page 4757](#)
- [Understanding IGMP on page 4759](#)

## Understanding Group Membership Protocols

There is a big difference between the multicast protocols used between host and router and between the multicast routers themselves. Hosts on a given subnetwork need to inform their router only whether or not they are interested in receiving packets from a certain multicast group. The source host needs to inform its routers only that it is the source of traffic for a particular multicast group. In other words, no detailed knowledge of the distribution tree is needed by any hosts; only a group membership protocol is needed to inform routers of their participation in a multicast group. Between adjacent routers, on the other hand, the multicast routing protocols must avoid loops as they build a detailed sense of the network topology and distribution tree from source to leaf. So, different multicast protocols are used for the host-router portion and the router-router portion of the multicast network.

Multicast group membership protocols enable a router to detect when a host on a directly attached subnet, typically a LAN, wants to receive traffic from a certain multicast group. Even if more than one host on the LAN wants to receive traffic for that multicast group, the router sends only one copy of each packet for that multicast group out on that interface, because of the inherent broadcast nature of LANs. When the multicast group membership protocol informs the router that there are no interested hosts on the subnet, the packets are withheld and that leaf is pruned from the distribution tree.

The Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) Protocol are the standard IP multicast group membership protocols: IGMP and MLD have several versions that are supported by hosts and routers:

- IGMPv1—The original protocol defined in RFC 1112. An explicit join message is sent to the router, but a timeout is used to determine when hosts leave a group. This process wastes processing cycles on the router, especially on older or smaller routers.
- IGMPv2—Defined in RFC 2236. Among other features, IGMPv2 adds an explicit leave message to the join message so that routers can more easily determine when a group has no interested listeners on a LAN.
- IGMPv3—Defined in RFC 3376. Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or *source-specific multicast (SSM)*.
- MLDv1—Defined in RFC 2710. MLDv1 is similar to IGMPv2.
- MLDv2—Defined in RFC 3810. MLDv2 similar to IGMPv3.

The various versions of IGMP and MLD are backward compatible. It is common for a router to run multiple versions of IGMP and MLD on LAN interfaces. Backward compatibility is achieved by dropping back to the most basic of all versions run on a LAN. For example, if one host is running IGMPv1, any router attached to the LAN running IGMPv2 can drop back to IGMPv1 operation, effectively eliminating the IGMPv2 advantages. Running multiple IGMP versions ensures that both IGMPv1 and IGMPv2 hosts find peers for their versions on the router.



**CAUTION:** On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

---

**Related  
Documentation**

- [Examples: Configuring MLD on page 4852](#)



## Understanding IGMP

The IPv4 address scheme assigns class D addresses for IP multicast. IGMP is the protocol that uses these addresses, which can be in the range 224.0.0.0 to 239.255.255.255. The following addresses have specific functions or are unavailable:

- 224.0.0.0 is reserved—you cannot assign it to a group.
- 224.0.0.1 is the all-hosts address—a packet sent to this address reaches all hosts on a subnet.
- 224.0.0.2 is the all-routers address—a packet sent to this address reaches all routers on a subnet.

This implementation of IGMP complies with IGMP versions 1, 2, and 3. IGMPv3 supports source-specific join and leave messages and is backward compatible with IGMPv1 and IGMPv2.

IGMPv2 mode interfaces exchange the following types of messages between routers and hosts:

- Group membership queries
- Group membership reports
- Leave group membership messages

IGMPv3 mode interfaces exchange the following types of messages with IGMPv3 hosts:

- Group membership queries
- IGMPv3 group membership reports

IGMP manages the membership of hosts and routers in multicast groups. IP hosts use IGMP to report their multicast group memberships to any immediately neighboring multicast routers. Multicast routers use IGMP to learn, for each of their attached physical networks, which groups have members.

A router receives explicit join and prune messages from those neighboring routers that have downstream group members. When PIM is the multicast protocol in use, IGMP begins the process as follows:

1. To join a multicast group, G, a host conveys its membership information through IGMP.
2. The router then forwards data packets addressed to a multicast group G to only those interfaces on which explicit join messages have been received.
3. A designated router (DR) sends periodic join and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. One or more routers are automatically or statically designated as the RP, and all routers must explicitly join through the RP.
4. Each router along the path toward the RP builds a wild card (any-source) state for the group and sends join and prune messages toward the RP.

The term *route entry* is used to refer to the state maintained in a router to represent the distribution tree.

A route entry can include such fields as:

- source address
- group address
- incoming interface from which packets are accepted
- list of outgoing interfaces to which packets are sent
- timers
- flag bits

The wild card route entry's incoming interface points toward the RP.

The outgoing interfaces point to the neighboring downstream routers that have sent join and prune messages toward the RP as well as the directly connected hosts that have requested membership to group G.

5. This state creates a shared, RP-centered, distribution tree that reaches all group members.

IGMP is also used as the transport for several related multicast protocols (for example, Distance Vector Multicast Routing Protocol [DVMRP] and Protocol Independent Multicast version 1 [PIMv1]).

IGMP is an integral part of IP and must be enabled on all routers and hosts that need to receive IP multicast traffic.

For each attached network, a multicast router can be either a querier or a nonquerier. The querier router periodically sends general query messages to solicit group membership information. Hosts on the network that are members of a multicast group send report messages. When a host leaves a group, it sends a leave group message.

IGMP version 3 (IGMPv3) supports inclusion and exclusion lists. Inclusion lists enable you to specify which sources can send to a multicast group. This type of multicast group is called a source-specific multicast (SSM) group, and its multicast address is 232/8.

IGMPv3 provides support for source filtering. For example, a router can specify particular routers from which it accepts or rejects traffic. With IGMPv3, a multicast router can learn which sources are of interest to neighboring routers.

Exclusion mode works the opposite of an inclusion list. It allows any source but the ones listed to send to the SSM group.

IGMPv3 interoperates with versions 1 and 2 of the protocol. However, to remain compatible with older IGMP hosts and routers, IGMPv3 routers must also implement versions 1 and 2 of the protocol. IGMPv3 supports the following membership-report record types: mode is allowed, allow new sources, and block old sources.

#### **Related Documentation**

- *Supported IP Multicast Protocol Standards*

---

## Introduction to IGMP Snooping

---

- [IGMP Snooping Overview on page 4761](#)

### IGMP Snooping Overview

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This IGMP snooping topic includes:

- [How IGMP Snooping Works on page 4761](#)
- [How IGMP Snooping Works with Routed VLAN Interfaces on page 4762](#)
- [How Hosts Join and Leave Multicast Groups on page 4762](#)
- [IGMP Snooping and Forwarding Interfaces on page 4762](#)
- [General Forwarding Rules on page 4763](#)
- [Using a Switch as an IGMP Querier on page 4763](#)

---

### How IGMP Snooping Works

---

A switch usually learns unicast MAC addresses by checking the source address field of the frames it receives and then sends any traffic for that unicast address only to the appropriate interface. However, a multicast MAC address can never be the source address for a packet. As a result, when a switch receives traffic for a multicast destination address, it floods the traffic on the relevant VLAN, which can cause a significant amount of traffic to be sent unnecessarily.

IGMP snooping prevents this flooding. When you enable IGMP snooping, the switch monitors IGMP packets between receivers and multicast routers and uses the content of the packets to build a multicast cache table—a database of multicast groups and the interfaces that are connected to members of the groups. When the switch receives multicast packets, it uses the cache table to selectively forward the traffic to only the interfaces that are connected to members of the appropriate multicast groups.



**NOTE:** IGMP snooping is enabled by default on the default VLAN only. With versions of Junos OS for the QFX Series previous to 13.2, IGMP snooping is enabled by default on all VLANs.



**NOTE:** You cannot configure IGMP snooping on a secondary (private) VLAN.

### How IGMP Snooping Works with Routed VLAN Interfaces

---

A switch can use a routed VLAN interface (RVI) to forward traffic between VLANs that connect to it. IGMP snooping works with Layer 2 interfaces and RVIs to forward multicast traffic in a switched network.

When a switch receives a multicast packet, its Packet Forwarding Engines perform a multicast lookup on the packet to determine how to forward the packet to its local interfaces. From the results of the lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces that have ports local to the Packet Forwarding Engine. If the list includes an RVI, the switch provides a bridge multicast group ID for the RVI to the Packet Forwarding Engine.

For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID, which identifies the Layer 2 interfaces in the VLAN that are interested in receiving the multicast stream. The Packet Forwarding Engine then forwards multicast traffic to bridge multicast IDs that have multicast receivers for a given multicast group.

### How Hosts Join and Leave Multicast Groups

---

Hosts can join multicast groups in two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

To leave a multicast group, either a host cannot respond to the periodic general IGMP queries, which results in a “silent leave” (the only leave option for IGMPv1), or a host can send a group-specific IGMPv2 leave message.

### IGMP Snooping and Forwarding Interfaces

---

To determine how to forward multicast traffic, a switch with IGMP snooping enabled maintains information about the following interfaces in its multicast forwarding table:

- Multicast-router interfaces—These interfaces lead toward multicast routers or IGMP queriers.
- Group-member interfaces—These interfaces lead toward hosts that are members of multicast groups.

The switch learns about these interfaces by monitoring IGMP traffic. If an interface receives IGMP queries or Protocol Independent Multicast (PIM) updates, the switch adds the interface to its multicast forwarding table as a multicast-router interface. If an interface receives membership reports for a multicast group, the switch adds the interface to its multicast forwarding table as a group-member interface.

Table entries for interfaces that the switch learns about are subject to aging. For example, if a learned multicast-router interface does not receive IGMP queries or PIM hellos within a certain interval, the switch removes the entry for that interface from its multicast forwarding table.



**NOTE:** For a switch to learn multicast-router interfaces and group-member interfaces, an IGMP querier must exist in the network. This is often a multicast router, but if there is no multicast router on the local network, you can configure the switch itself to be an IGMP querier.

You can statically configure an interface to be a multicast-router interface or a group-member interface. The switch adds a static interface to its multicast forwarding table without having to learn about the interface, and the entry in the table is not subject to aging. You can have a mix of statically configured and dynamically learned interfaces on a switch.

### General Forwarding Rules

Multicast traffic received on a switch interface in a VLAN on which IGMP snooping is enabled is forwarded according to the following rules.

IGMP traffic is forwarded as follows:

- IGMP general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- IGMP group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- IGMP reports received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to the other host interfaces in the VLAN.

Multicast traffic that is not IGMP traffic is forwarded as follows:

- A multicast packet with a destination address of 224.0.0.0/24 is flooded to all other interfaces on the VLAN.
- An unregistered multicast packet—that is, a packet for a group that has no current members—is forwarded to all multicast-router interfaces in the VLAN.
- A registered multicast packet is forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.

### Using a Switch as an IGMP Querier

If IGMP snooping is enabled on a pure Layer 2 local network (that is, Layer 3 is not enabled on the network), and there is not multicast router in the network, multicast traffic might not be properly forwarded through the network. This problem occurs if the local network is configured such that multicast traffic must be forwarded between switches in order to reach a multicast receiver. In this case, an upstream switch does not forward multicast traffic to a downstream switch (and therefore to the multicast receivers

attached to the downstream switch) because the downstream switch does not forward IGMP reports to the upstream switch. You can solve this problem by configuring one of the switches to be an IGMP querier. This switch sends periodic general query packets to all the switches in the network, which ensures that the snooping membership tables are updated and prevents any multicast traffic loss.

If you configure multiple switches to be IGMP queriers, the switch with the highest (greatest) IGMP querier source address takes precedence and acts as the querier. Switches with lower IGMP querier source addresses stop sending IGMP queries unless they do not receive IGMP queries for 255 seconds. If a switch with a lower IGMP querier source address does not receive any IGMP queries during that period, it starts sending queries again.

To configure a standalone switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

To configure a QFabric Node device switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

#### Related Documentation

- [Example: Configuring IGMP Snooping on page 4849](#)
- [Configuring IGMP Snooping on page 4848](#)
- [Changing the IGMP Snooping Group Timeout Value](#)
- [Monitoring IGMP Snooping on page 5047](#)
- [Configuring IGMP on page 4826](#)
- RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments*
- IGMPv1—See RFC 1112, *Host extensions for IP multicasting*.
- IGMPv2—See RFC 2236, *Internet Group Management Protocol, Version 2*.
- IGMPv3—See RFC 3376, *Internet Group Management Protocol, Version 3*.

## Introduction to MLD

---

- [Understanding MLD on page 4764](#)

### Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners.

In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 156 on page 4765](#)). The querier routing device on the right is the receiver's DR.

**Figure 156: Routing Devices Start Up on a Subnet**




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In [Figure 157 on page 4766](#), the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




**NOTE:** In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

Figure 157: Querier Routing Device Is Determined




The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address `FF02::1` at short intervals to all attached subnets to solicit group membership information (see Figure 158 on page 4766). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.

Figure 158: General Query Message Is Issued



If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see Figure 159 on page 4766). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.


Figure 159: Reports Are Received by the Querier Routing Device






If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 160 on page 4767](#)).

**Figure 160: Host Has No Interested Receivers and Sends a Done Message to Routing Device**



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 161 on page 4767](#)).

**Figure 161: Host Address Timer Expires and Address Is Removed from Multicast Address List**



- Related Documentation**
- [Enabling MLD on page 4856](#)
  - [Example: Recording MLD Join and Leave Events on page 4870](#)
  - [Example: Modifying the MLD Robustness Variable on page 4861](#)

## Introduction to MSDP

- [Understanding MSDP on page 4767](#)
- [Filtering MSDP SA Messages on page 4769](#)

## Understanding MSDP

The Multicast Source Discovery Protocol (MSDP) is used to connect multicast routing domains. It typically runs on the same router as the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP). Each MSDP router establishes adjacencies with internal and external MSDP peers similar to the way BGP establishes peers. These peer routers inform each other about active sources within the domain. When they detect

active sources, the routers can send PIM sparse-mode explicit join messages to the active source.

The peer with the higher IP address passively listens to a well-known port number and waits for the side with the lower IP address to establish a Transmission Control Protocol (TCP) connection. When a PIM sparse-mode RP that is running MSDP becomes aware of a new local source, it sends source-active type, length, and values (TLVs) to its MSDP peers. When a source-active TLV is received, a peer-reverse-path-forwarding (peer-RPF) check (not the same as a multicast RPF check) is done to make sure that this peer is in the path that leads back to the originating RP. If not, the source-active TLV is dropped. This TLV is counted as a “rejected” source-active message.

The MSDP peer-RPF check is different from the normal RPF checks done by non-MSDP multicast routers. The goal of the peer-RPF check is to stop source-active messages from looping. Router R accepts source-active messages originated by Router S only from neighbor Router N or an MSDP mesh group member. For more information about configuring MSDP mesh groups, see [“Example: Configuring MSDP with Active Source Limits and Mesh Groups” on page 4880](#).

Router R locates its MSDP peer-RPF neighbor (Router N) deterministically. A series of rules is applied in a particular order to received source-active messages, and the first rule that applies determines the peer-RPF neighbor. All source-active messages from other routers are rejected.

The six rules applied to source-active messages originating at Router S received at Router R from Router X are as follows:

1. If Router X originated the source-active message (Router X is Router S), then Router X is also the peer-RPF neighbor, and its source-active messages are accepted.
2. If Router X is a member of the Router R mesh group, or is the configured peer, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
3. If Router X is the BGP next hop of the active multicast RPF route toward Router S (Router X installed the route on Router R), then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
4. If Router X is an external BGP (EBGP) or internal BGP (IBGP) peer of Router R, and the last autonomous system (AS) number in the BGP AS-path to Router S is the same as Router X's AS number, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
5. If Router X uses the same next hop as the next hop to Router S, then Router X is the peer-RPF neighbor, and its source-active messages are accepted.
6. If Router X fits none of these criteria, then Router X is not an MSDP peer-RPF neighbor, and its source-active messages are rejected.

The MSDP peers that receive source-active TLVs can be constrained by BGP reachability information. If the AS path of the network layer reachability information (NLRI) contains the receiving peer's AS number prepended second to last, the sending peer is using the receiving peer as a next hop for this source. If the split horizon information is not being received, the peer can be pruned from the source-active TLV distribution list.

- Related Documentation**
- [Configuring MSDP on page 4876](#)

## Filtering MSDP SA Messages

Along with applying MSDP source active (SA) filters on all external MSDP sessions (in and out) to prevent SAs for groups and sources from leaking in and out of the network, you need to apply bootstrap router (BSR) filters. Applying a BSR filter to the boundary of a network prevents foreign BSR messages (which announce RP addresses) from leaking into your network. Since the routers in a PIM sparse-mode domain need to know the address of only one RP router, having more than one in the network can create issues.

If you did not use multicast scoping to create boundary filters for all customer-facing interfaces, you might want to use PIM join filters. Multicast scopes prevent the actual multicast data packets from flowing in or out of an interface. PIM join filters prevent PIM sparse-mode state from being created in the first place. Since PIM join filters apply only to the PIM sparse-mode state, it might be more beneficial to use multicast scoping to filter the actual data.



**NOTE:** When you apply firewall filters, firewall action modifiers, such as **log**, **sample**, and **count**, work only when you apply the filter on an inbound interface. The modifiers do not work on an outbound interface.

- Related Documentation**
- [Understanding Multicast Administrative Scoping](#)
  - [Filtering Incoming PIM Join Messages on page 4812](#)
  - [Example: Configuring PIM BSR Filters on page 4809](#)

## Introduction to Source-Specific Multicast

- [Source-Specific Multicast Groups Overview on page 4769](#)
- [Understanding PIM Source-Specific Mode on page 4770](#)
- [PIM SSM on page 4771](#)

## Source-Specific Multicast Groups Overview

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM implemented in Junos OS has the efficient explicit join procedures of Protocol Independent Multicast (PIM) sparse mode but eliminates the immediate shared tree and rendezvous point (RP) procedures using (\*,G) pairs. The (\*) is a wildcard referring to any source sending to group G, and "G" refers to the IP multicast group. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The "S" refers to the source's unicast IP address, and the "G" refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. Although ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click a link in a browser, the receiver is notified about the group

information, but not the source information. With SSM, the client receives both source and group information.

SSM is ideal for one-to-many multicast services such as network entertainment channels. However, many-to-many multicast services might require ASM.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an Internet Group Management Protocol version 3 (IGMPv3) or Multicast Listener Discovery version 2 (MLDv2) stack, or you need to configure SSM mapping from IGMPv1 or IGMPv2 to IGMPv3. An IGMPv3 stack provides the capability of a host operating system to use the IGMPv3 protocol. IGMPv3 is available for Windows XP, Windows Vista, and most UNIX operating systems.

SSM mapping allows operators to support an SSM network without requiring all hosts to support IGMPv3. This support exists in static (S,G) configurations, but SSM mapping also supports dynamic per-source group state information, which changes as hosts join and leave the group using IGMP.

SSM is typically supported with a subset of IGMPv3 and PIM sparse mode known as *PIM SSM*. Using SSM, a client can receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the client and the source, but builds the SPT without the help of an RP.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through the Multicast Source Discovery Protocol (MSDP).

## Understanding PIM Source-Specific Mode

RFC 1112, the original multicast RFC, supported both many-to-many and one-to-many models. These came to be known collectively as any-source multicast (ASM) because ASM allowed one or many sources for a multicast group's traffic. However, an ASM network must be able to determine the locations of all sources for a particular multicast group whenever there are interested listeners, no matter where the sources might be located in the network. In ASM, the key function of *source discovery* is a required function of the network itself.

Multicast source discovery appears to be an easy process, but in sparse mode it is not. In dense mode, it is simple enough to flood traffic to every router in the whole network so that every router learns the source address of the content for that multicast group. However, the flooding presents scalability and network resource use issues and is not a viable option in sparse mode.

PIM sparse mode (like any sparse mode protocol) achieves the required source discovery functionality without flooding at the cost of a considerable amount of complexity. The RP routers must be added and must know all multicast sources, and complicated shared distribution trees must be built to the RPs.

In an environment where many sources come and go, such as for a videoconferencing service, ASM is appropriate. However, by ignoring the many-to-many model and focusing attention on the one-to-many source-specific multicast (SSM) model, several commercially promising multicast applications, such as television channel distribution

over the Internet, might be brought to the Internet much more quickly and efficiently than if full ASM functionality were required of the network.

PIM SSM is simpler than PIM sparse mode because only the one-to-many model is supported. Initial commercial multicast Internet applications are likely to be available to *subscribers* (that is, receivers that issue join messages) from only a single source (a special case of SSM covers the need for a backup source). PIM SSM therefore forms a subset of PIM sparse mode. PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

PIM SSM introduces new terms for many of the concepts in PIM sparse mode. PIM SSM can technically be used in the entire 224/4 multicast address range, although PIM SSM operation is guaranteed only in the 232/8 range (232.0.0/24 is reserved). The new SSM terms are appropriate for Internet video applications and are summarized in

[Table 386 on page 4771](#).

**Table 386: ASM and SSM Terminology**

Term	Any-Source Multicast	Source-Specific Multicast
Address identifier	G	S,G
Address designation	group	channel
Receiver operations	join, leave	subscribe, unsubscribe
Group address range	224/4 excluding 232/8	224/4 (guaranteed only for 232/8)

Although PIM SSM describes receiver operations as *subscribe* and *unsubscribe*, the same PIM sparse mode join and leave messages are used by both forms of the protocol. The terminology change distinguishes ASM from SSM even though the receiver messages are identical.

## PIM SSM

PIM source-specific multicast (SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3) to allow a client to receive multicast traffic directly from the source. PIM SSM uses the PIM sparse-mode functionality to create an SPT between the receiver and the source, but builds the SPT without the help of an RP.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.


You can also configure the Junos OS to accept any-source multicast (ASM) join messages (\*G) for group addresses that are within the default or configured range of source-specific multicast (SSM) groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

An SSM-configured network has distinct advantages over a traditionally configured PIM sparse-mode network. There is no need for shared trees or RP mapping (no RP is required), or for RP-to-RP source discovery through MSDP.

Deploying SSM is easy. You need to configure PIM sparse mode on all router interfaces and issue the necessary SSM commands, including specifying IGMPv3 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group member interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3, are used in PIM SSM. As sources become active and start sending multicast packets, interested receivers in the SSM group receive the multicast packets.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3), announcing a desire to join group G and source S (see [Figure 162 on page 4772](#)). The directly connected PIM sparse-mode router, the receiver's DR, sends an (S,G) join message to its RPF neighbor for the source. Notice in [Figure 162 on page 4772](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

**Figure 162: Receiver Announces Desire to Join Group G and Source S**




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 163 on page 4772](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

**Figure 163: Router 3 (Last-Hop Router) Joins the Source Tree**



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 164 on page 4773](#)).

Figure 164: (S,G) State Is Built Between the Source and the Receiver



To configure additional SSM groups, include the **ssm-groups** statement at the **[edit routing-options multicast]** hierarchy level.

#### Related Documentation

- [Source-Specific Multicast Groups Overview on page 4769](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895](#)

## Introduction to Multicast VLAN Registration

- [Understanding Multicast VLAN Registration on page 4773](#)

### Understanding Multicast VLAN Registration

Multicast VLAN registration (MVR) enables you to efficiently distribute IPTV multicast streams across an Ethernet ring-based Layer 2 network and reduce the amount of bandwidth consumed by this multicast traffic.

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which IPTV multicast traffic flows throughout the Layer 2 network. A Juniper Networks EX Series switch or QFX Series switch that is enabled for MVR selectively forwards IPTV multicast traffic from interfaces on the MVLAN (source interfaces) to hosts that are connected to interfaces that are not part of the MVLAN. These interfaces are known as *MVR receiver ports*. The MVR receiver ports can receive traffic from a port on the MVLAN but cannot send traffic onto the MVLAN, and they remain in their own VLANs for bandwidth and security reasons.

This topic includes:

- [How MVR Works on page 4773](#)

#### How MVR Works

In many ways, MVR is similar to IGMP snooping. Both MVR and IGMP snooping monitor IGMP join and leave messages and build forwarding tables based on the media access control (MAC) addresses of the hosts sending those IGMP messages. Whereas IGMP snooping operates within a given VLAN to regulate multicast traffic, MVR can operate

with hosts on different VLANs in a Layer 2 network to selectively deliver IPTV multicast traffic to requesting hosts, thereby reducing the amount of bandwidth needed to forward multicast traffic.

When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs. Interfaces that are on the MVLAN itself cannot be MVR receiver ports for that MVLAN.



**NOTE:** MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

### ***MVR Modes***

MVR operates in two modes: MVR transparent mode and MVR proxy mode. Both modes enable MVR to forward only one copy of a multicast stream to the Layer 2 network.

- [MVR Transparent Mode on page 4774](#)
- [MVR Proxy Mode on page 4774](#)

### ***MVR Transparent Mode***

In MVR transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. Transparent mode is the default mode.

The switch handles IGMP packets destined for both the multicast source VLAN and multicast receiver VLANs in the same way that it handles them when MVR is not being used. That is, when a host on a VLAN sends IGMP join and leave messages, the switch floods the messages to all router interfaces in the VLAN. Similarly, when a VLAN receives IGMP queries from its router interfaces, it floods the queries to all interfaces in the VLAN.

If a host on a multicast receiver port joins an MVR group on the multicast receiver VLAN, the appropriate bridging entry is added and the MVLAN forwards that group's IPTV multicast traffic on that port (even though that port is not in the MVLAN). Likewise, if a host on a multicast receiver port leaves an MVR group on the multicast receiver VLAN, the appropriate bridging entry is deleted, and the MVLAN stops forwarding that group's IPTV multicast traffic on that port. In addition, you can configure the switch to statically install the bridging entries on the multicast receiver VLAN.

### ***MVR Proxy Mode***

When you use MVR in proxy mode, the switch acts as a proxy for any MVR group in both the upstream and downstream directions. In the downstream direction, the switch acts as the querier for the groups in the MVR receiver VLANs. In the upstream direction, the switch originates the IGMP reports and leaves and answers IGMP queries from multicast routers. When the MVR receiver VLANs receive IGMP joins and leaves, the switch creates bridging entries on the MVLAN as needed, as it does in MVR transparent mode. In addition, the switch sends out IGMP joins and leaves on the MVLAN based on these bridging entries.



Configuring MVR proxy mode on the MVLAN automatically enables IGMP snooping proxy mode on all MVR receiver VLANs as well as on the MVLAN.

**Related  
Documentation**

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5546](#)
- *Example: Configuring Multicast VLAN Registration*
- *Configuring Multicast VLAN Registration (CLI Procedure)*



## CHAPTER 58

# Configuration

- [PIM Basics on page 4777](#)
- [PIM Designated Router on page 4784](#)
- [PIM Sparse Mode on page 4786](#)
- [Static RP on page 4795](#)
- [Anycast RP on page 4798](#)
- [PIM Bootstrap Router on page 4807](#)
- [PIM Filtering on page 4810](#)
- [PIM RPT and SPT Cutover on page 4815](#)
- [PIM and the BFD Protocol on page 4821](#)
- [IGMP on page 4826](#)
- [IGMP Snooping on page 4847](#)
- [MLD on page 4852](#)
- [MSDP on page 4875](#)
- [Source-Specific Multicast on page 4891](#)
- [PIM Configuration Statements on page 4902](#)
- [IGMP Configuration Statements on page 4972](#)
- [IGMP Snooping Configuration Statements on page 4995](#)
- [MSDP Configuration Statements on page 5019](#)
- [Source-Specific Multicast Configuration Statements on page 5040](#)

## PIM Basics

---

- [Changing the PIM Version on page 4778](#)
- [Modifying the PIM Hello Interval on page 4778](#)
- [Preserving Multicast Performance by Disabling Response to the ping Utility on page 4779](#)
- [Configuring PIM Trace Options on page 4780](#)
- [Disabling PIM on page 4782](#)

## Changing the PIM Version

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults.

To configure the PIM version, include the **version** statement:

```
version (1 | 2);
```

## Modifying the PIM Hello Interval

Routing devices send hello messages at a fixed interval on all PIM-enabled interfaces. By using hello messages, routing devices advertise their existence as PIM routing devices on the subnet. With all PIM-enabled routing devices advertised, a single designated router for the subnet is established.

When a routing device is configured for PIM, it sends a hello message at a 30-second default interval. The interval range is from 0 through 255. When the interval counts down to 0, the routing device sends another hello message, and the timer is reset. A routing device that receives no response from a neighbor in 3.5 times the interval value drops the neighbor. In the case of a 30-second interval, the amount of time a routing device waits for a response is 105 seconds.

If a PIM hello message contains the hold-time option, the neighbor timeout is set to the hold-time sent in the message. If a PIM hello message does not contain the hold-time option, the neighbor timeout is set to the default hello hold time.

To modify how often the routing device sends hello messages out of an interface:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface fe-3/0/2.0]  
user@host# set hello-interval 255
```

2. Verify the configuration by checking the **Hello Option Holdtime** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail  
Instance: PIM.master  
Interface: fe-3/0/2.0  
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse  
Hello Option Holdtime: 255 seconds  
Hello Option DR Priority: 1  
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms  
Join Suppression supported  
Rx Join: Group Source Timeout  
225.1.1.1 192.168.195.78 0  
225.1.1.1 0
```

```
Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

```
Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 255 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

**Related Documentation** • [show pim neighbors on page 5175](#) in the [CLI Explorer](#)

## Preserving Multicast Performance by Disabling Response to the ping Utility

The ping utility uses ICMP Echo messages to verify connectivity to any device with an IP address. However, in the case of multicast applications, a single ping sent to a multicast address can degrade the performance of routers because the stream of packets is replicated multiple times.

You can disable the router's response to ping (ICMP Echo) packets sent to multicast addresses. The system responds normally to unicast ping packets.

To disable the router's response to ping packets sent to multicast addresses:

1. Include the **no-multicast-echo** statement:

```
[edit system]
user@host# set no-multicast-echo
```

2. Verify the configuration by checking the **echo drops with broadcast or multicast destination address** field in the output of the **show system statistics icmp** command.

```
user@host> show system statistics icmp

icmp:
0 drops due to rate limit
0 calls to icmp_error
0 errors not generated because old message was icmp
Output histogram:
echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
100 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
echo: 21
21 message responses generated
```

**Related Documentation** • [Configuring Junos OS to Disable the Routing Engine Response to Multicast Ping Packets](#) in the [Junos OS Administration Library for Routing Devices](#)

- *show system statistics icmp* in the [CLI Explorer](#)

## Configuring PIM Trace Options

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
<b>all</b>	Trace all operations.
<b>assert</b>	Trace assert messages, which are used to resolve which of the parallel routers connected to a multiaccess LAN is responsible for forwarding packets to the LAN.
<b>autorp</b>	Trace bootstrap, RP, and auto-RP messages.
<b>bidirectional-df-election</b>	Trace bidirectional PIM designated-forwarder (DF) election events.
<b>bootstrap</b>	Trace bootstrap messages, which are sent periodically by the PIM domain's bootstrap router and are forwarded, hop by hop, to all routers in that domain.
<b>general</b>	Trace general events.
<b>graft</b>	Trace graft and graft acknowledgment messages.
<b>hello</b>	Trace hello packets, which are sent so that neighboring routers can discover one another.
<b>join</b>	Trace join messages, which are sent to join a branch onto the multicast distribution tree.
<b>mdt</b>	Trace messages related to multicast data tunnels.
<b>normal</b>	Trace normal events.
<b>nsr-synchronization</b>	Trace nonstop routing synchronization events
<b>packets</b>	Trace all PIM packets.
<b>policy</b>	Trace poison-route-reverse packets.
<b>prune</b>	Trace prune messages, which are sent to prune a branch off the multicast distribution tree.

Flag	Description
<b>register</b>	Trace register and register-stop messages. Register messages are sent to the RP when a multicast source first starts sending to a group.
<b>route</b>	Trace routing information.
<b>rp</b>	Trace candidate RP advertisements.
<b>state</b>	Trace state transitions.
<b>task</b>	Trace task processing.
<b>timer</b>	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on PIM packets of a particular type.

To configure tracing operations for PIM:

1. (Optional) Configure tracing at the [**routing-options** hierarchy level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the PIM trace file.

```
[edit protocols pim traceoptions]
user@host# set file pim-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols pim traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols pim traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols pim traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags.

Suppose you are troubleshooting issues with PIM version 1 control packets that are received on an interface configured for PIM version 2. The following example shows how to trace messages associated with this problem.

```
[edit protocols pim traceoptions]
user@host# set flag packets | match "Rx V1 Require V2"
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/pim-trace
```

**Related  
Documentation**

- [PIM Overview on page 4737](#)
- *Tracing and Logging Junos OS Operations* in the *Junos OS Administration Library for Routing Devices*

## Disabling PIM

By default, when configured, the PIM protocol is enabled on all interfaces for all families. If desired, you can disable PIM at the protocol, interface, or family hierarchy levels.

The hierarchy in which you configure PIM is critical. In general, the most specific configuration takes precedence. However, if PIM is disabled at the protocol level, then any disable statements with respect to an interface or family are ignored.

For example, the order of precedence for disabling PIM on a particular interface family is:

1. If PIM is disabled at the **[edit protocols pim interface *interface-name* family]** hierarchy level, then PIM is disabled for that interface family.
2. If PIM is not configured at the **[edit protocols pim interface *interface-name* family]** hierarchy level, but is disabled at the **[edit protocols pim interface *interface-name*]** hierarchy level, then PIM is disabled for all families on the specified interface.
3. If PIM is not configured at either the **[edit protocols pim interface *interface-name* family]** hierarchy level or the **[edit protocols pim interface *interface-name*]** hierarchy level, but is disabled at the **[edit protocols pim]** hierarchy level, then the PIM protocol is disabled globally for all interfaces and all families.

The following sections describe how to disable PIM at the various hierarchy levels.

- [Disabling the PIM Protocol on page 4782](#)
- [Disabling PIM On an Interface on page 4783](#)
- [Disabling PIM for a Family on page 4783](#)
- [Disabling PIM for a Rendezvous Point on page 4784](#)

### Disabling the PIM Protocol

---

You can explicitly disable the PIM protocol. Disabling the PIM protocol disables the protocol for all interfaces and all families. This is accomplished at the **[edit protocols pim]** hierarchy level:

```
[edit protocols]
pim {
  disable;
}
```

To disable the PIM protocol:

1. Include the **disable** statement.



```
user@host# set protocols pim disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

### Disabling PIM On an Interface

You can disable the PIM protocol on a per-interface basis. This is accomplished at the **[edit protocols pim interface *interface-name*]** hierarchy level:

```
[edit protocols]
pim {
  interface interface-name {
    disable;
  }
}
```

To disable PIM on an interface:

1. Include the **disable** statement.

```
user@host# set protocols pim interface fe-0/1/0 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

### Disabling PIM for a Family

You can disable the PIM protocol on a per-family basis. This is accomplished at the **[edit protocols pim family]** hierarchy level:

```
[edit protocols]
pim {
  family inet {
    disable;
  }
  family inet6 {
    disable;
  }
}
```

To disable PIM for a family:

1. Include the **disable** statement.

```
user@host# set protocols pim family inet disable
```

```
user@host# set protocols pim family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

### Disabling PIM for a Rendezvous Point

---

You can disable the PIM protocol for a rendezvous point (RP) on a per-family basis. This is accomplished at the **[edit protocols pim rp local family]** hierarchy level:

```
[edit protocols]
pim {
  rp {
    local {
      family inet {
        disable;
      }
      family inet6 {
        disable;
      }
    }
  }
}
```

To disable PIM for an RP family:

1. Use the **disable** statement.

```
user@host# set protocols pim rp local family inet disable
user@host# set protocols pim rp local family inet6 disable
```

2. (Optional) Verify your configuration settings before committing them by using the **show protocols pim** command.

```
user@host# run show protocols pim
```

## PIM Designated Router

---

- [Configuring Interface Priority for PIM Designated Router Selection on page 4784](#)
- [Configuring PIM Designated Router Election on Point-to-Point Links on page 4785](#)

### Configuring Interface Priority for PIM Designated Router Selection

A designated router (DR) sends periodic join messages and prune messages toward a group-specific rendezvous point (RP) for each group for which it has active members. When a Protocol Independent Multicast (PIM) router learns about a source, it originates a Multicast Source Discovery Protocol (MSDP) source-address message if it is the DR on the upstream interface.

By default, every PIM interface has an equal probability (priority 1) of being selected as the DR. Configuring the interface DR priority helps ensure that changing an IP address does not alter your forwarding model.

To configure the interface designated router priority:

1. This example shows the configuration for the routing instance. Configure the interface globally or in the routing instance.

```
[edit routing-instances PIM.master protocols pim interface ge-0/0/0.0 family inet]
user@host# set priority 5
```

2. Verify the configuration by checking the **Hello Option DR Priority** field in the output of the **show pim neighbors detail** command.

```
user@host> show pim neighbors detail
```

```
Instance: PIM.master
Interface: ge-0/0/0.0
Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 5
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
Rx Join: Group Source Timeout
225.1.1.1 192.168.195.78 0
225.1.1.1 0
```

```
Interface: lo0.0
Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

```
Interface: pd-6/0/0.32768
Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 0
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

#### Related Documentation

- [Configuring PIM Designated Router Election on Point-to-Point Links on page 4785](#)
- [Understanding PIM Sparse Mode on page 4741](#)
- [show pim neighbors on page 5175](#) in the CLI Explorer

## Configuring PIM Designated Router Election on Point-to-Point Links

To comply with the latest PIM drafts, enable designated router (DR) election on all PIM interfaces, including point-to-point (P2P) interfaces. (DR election is enabled by default on all other interfaces.) One of the two routers might join a multicast group on its P2P link interface. The DR on that link is responsible for initiating the relevant join messages.

To enable DR election on point-to-point interfaces:

1. On both point-to-point link routers, configure the router globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set dr-election-on-p2p
```

2. Verify the configuration by checking the **State** field in the output of the **show pim interfaces** command. The possible values for the **State** field are DR, NotDR, and P2P. When a point-to-point link interface is elected to be the DR, the interface state becomes DR instead of P2P.

3. If the **show pim interfaces** command continues to report the P2P state, consider running the **restart routing** command on both routers on the point-to-point link. Then recheck the state.



**CAUTION:** Do not restart a software process unless specifically asked to do so by your Juniper Networks customer support representative. Restarting a software process during normal operation of a routing platform could cause interruption of packet forwarding and loss of data.

[edit]  
user@host# run restart routing

#### Related Documentation

- [Understanding PIM Sparse Mode on page 4741](#)
- [Configuring Interface Priority for PIM Designated Router Selection on page 4784](#)
- [show pim interfaces on page 5150](#) in the CLI Explorer

## PIM Sparse Mode

---

- [Enabling PIM Sparse Mode on page 4786](#)
- [Configuring PIM Join Load Balancing on page 4787](#)
- [Modifying the Join State Timeout on page 4790](#)
- [Example: Enabling Join Suppression on page 4791](#)

### Enabling PIM Sparse Mode

In PIM sparse mode (PIM-SM), the assumption is that very few of the possible receivers want packets from a source, so the network establishes and sends packets only on branches that have at least one leaf indicating (by message) a desire for the traffic. WANs are appropriate networks for sparse-mode operation.

By default, PIM is disabled. When you enable PIM, it operates in sparse mode by default. You do not need to configure Internet Group Management Protocol (IGMP) version 2 for a sparse mode configuration. After you enable PIM, by default, IGMP version 2 is also enabled.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default for rendezvous point (RP) mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. The following example explicitly configures PIMv2 on the interfaces.

You can configure PIM sparse mode globally or for a routing instance. This example shows how to configure PIM sparse mode globally on all interfaces. It also shows how to configure a static RP router and how to configure the non-RP routers.

To configure the router properties for PIM sparse mode:

1. Configure the static RP router.

```
[edit protocols pim]
user@host# set rp local family inet address 192.168.3.253
```

2. Configure the RP router interfaces. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
[edit protocols pim]
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

3. Configure the non-RP routers. Include the following configuration on all of the non-RP routers.

```
[edit protocols pim]
user@host# set rp static address 192.168.3.253 version 2
user@host# set interface all mode sparse
user@host# set interface all version 2
user@host# set interface fxp0.0 disable
```

4. Monitor the operation of PIM sparse mode.

- [show pim interfaces](#)
- [show pim join](#)
- [show pim neighbors](#)
- [show pim rps](#)

**Related Documentation**

- [Understanding PIM Sparse Mode on page 4741](#)

## Configuring PIM Join Load Balancing

By default, PIM join messages are sent toward a source based on the RPF routing table check. If there is more than one equal-cost path toward the source, then one upstream interface is chosen to send the join message. This interface is also used for all downstream traffic, so even though there are alternative interfaces available, the multicast load is concentrated on one upstream interface and routing device.

For PIM sparse mode, you can configure PIM join load balancing to spread join messages and traffic across equal-cost upstream paths (interfaces and routing devices) provided by unicast routing toward a source. PIM join load balancing is only supported for PIM sparse mode configurations.

PIM join load balancing is supported on draft-rosen multicast VPNs (also referred to as dual PIM multicast VPNs). PIM join load balancing is not supported on multiprotocol BGP-based multicast VPNs (also referred to as next-generation Layer 3 VPN multicast). When PIM join load balancing is enabled in a draft-rosen Layer 3 VPN scenario, the load balancing is achieved based on the join counts for the far-end PE routing devices, not for any intermediate P routing devices.

If an internal BGP (IBGP) multipath forwarding VPN route is available, the Junos OS uses the multipath forwarding VPN route to send join messages to the remote PE routers to achieve load balancing over the VPN.

By default, when multiple PIM joins are received for different groups, all joins are sent to the same upstream gateway chosen by the unicast routing protocol. Even if there are multiple equal-cost paths available, these alternative paths are not utilized to distribute multicast traffic from the source to the various groups.

When PIM join load balancing is configured, the PIM joins are distributed equally among all equal-cost upstream interfaces and neighbors. Every new join triggers the selection of the least-loaded upstream interface and neighbor. If there are multiple neighbors on the same interface (for example, on a LAN), join load balancing maintains a value for each of the neighbors and distributes multicast joins (and downstream traffic) among these as well.

Join counts for interfaces and neighbors are maintained globally, not on a per-source basis. Therefore, there is no guarantee that joins for a particular source are load-balanced. However, the joins for all sources and all groups known to the routing device are load-balanced. There is also no way to administratively give preference to one neighbor over another: all equal-cost paths are treated the same way.

You can configure message filtering globally or for a routing instance. This example shows the global configuration.

You configure PIM join load balancing on the non-RP routers in the PIM domain.

1. Determine if there are multiple paths available for a source (for example, an RP) with the output of the **show pim join extensive** or **show pim source** commands.

```
user@host> show pim join extensive
Instance: PIM.master Family: INET

Group: 224.1.1.1
Source: *
RP: 10.255.245.6
Flags: sparse,rptree,wildcard
Upstream interface: t1-0/2/3.0
Upstream neighbor: 192.168.38.57
Upstream state: Join to RP
Downstream neighbors:
    Interface: t1-0/2/1.0
```

```

192.168.38.16 State: JOIN Flags; SRW Timeout: 164
Group: 224.2.127.254
Source: *
RP: 10.255.245.6
Flags: sparse,rptree,wildcard
Upstream interface: so-0/3/0.0
Upstream neighbor: 192.168.38.47
Upstream state: Join to RP
Downstream neighbors:
Interface: t1-0/2/3.0
192.168.38.16 State: JOIN Flags; SRW Timeout: 164

```

Note that for this router, the RP at IP address 10.255.245.6 is the source for two multicast groups: 224.1.1.1 and 224.2.127.254. This router has two equal-cost paths through two different upstream interfaces (**t1-0/2/3.0** and **so-0/3/0.0**) with two different neighbors (192.168.38.57 and 192.168.38.47). This router is a good candidate for PIM join load balancing.

2. On the non-RP router, configure PIM join load balancing.

```

[edit protocols pim rp]
user@host# set static address 10.10.10.1
user@host# set interface all mode sparse version 2
user@host# set join-load-balance

```

The static address is the address of the RP.

3. Monitor the operation.

If load balancing is enabled for this router, the number of PIM joins sent on each interface is shown in the output for the **show pim interfaces** command.

```

user@host> show pim interfaces
Instance: PIM.master

```

Name	Stat	Mode	IP V	State	NbrCnt	JoinCnt	DR address
lo0.0	Up	Sparse	4 2	DR	0	0	10.255.168.58
pe-1/2/0.32769	Up	Sparse	4 2	P2P	0	0	
so-0/3/0.0	Up	Sparse	4 2	P2P	1	1	
t1-0/2/1.0	Up	Sparse	4 2	P2P	1	0	
t1-0/2/3.0	Up	Sparse	4 2	P2P	1	1	
lo0.0	Up	Sparse	6 2	DR	0	0	fe80::2a0:a5ff:4b7

Note that the two equal-cost paths shown by the **show pim interfaces** command now have nonzero join counts. If the counts differ by more than one and were zero (0) when load balancing commenced, an error occurs (joins before load balancing are not redistributed). The join count also appears in the **show pim neighbors detail** output:

```

user@host> show pim neighbors detail
Interface: so-0/3/0.0

```

```

Address: 192.168.38.46, IPv4, PIM v2, Mode: Sparse, Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 1689116164
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

```

```

Address: 192.168.38.47, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 102 remaining
Hello Option DR Priority: 1

```

```
Hello Option Generation ID: 792890329
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Interface: t1-0/2/3.0
```

```
Address: 192.168.38.56, IPv4, PIM v2, Mode: Sparse, Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 678582286
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.38.57, IPv4, PIM v2, Join Count: 1
BFD: Disabled
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1854475503
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Note that the join count is nonzero on the two load-balanced interfaces toward the upstream neighbors.

PIM join load balancing only takes effect when the feature is configured. Prior joins are not redistributed to achieve perfect load balancing. In addition, if an interface or neighbor fails, the new joins are redistributed among remaining active interfaces and neighbors. However, when the interface or neighbor is restored, prior joins are not redistributed. The **clear pim join-distribution** command redistributes the existing flows to new or restored upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you perform PIM join redistribution during a maintenance window.

#### Related Documentation

- *clear pim join-distribution* in the [CLI Explorer](#)
- [show pim interfaces on page 5150](#) in the [CLI Explorer](#)
- [show pim neighbors on page 5175](#) in the [CLI Explorer](#)
- [show pim source on page 5186](#) in the [CLI Explorer](#)

## Modifying the Join State Timeout

This section describes how to configure the join state timeout.

A downstream router periodically sends join messages to refresh the join state on the upstream router. If the join state is not refreshed before the timeout expires, the join state is removed.

By default, the join state timeout is 210 seconds. You can change this timeout to allow additional time to receive the join messages. Because the messages are called join-prune messages, the name used is the **join-prune-timeout** statement.

To modify the timeout, include the **join-prune-timeout** statement:

```
user@host# set protocols pim join-prune-timeout 230
```

The join timeout value can be from 210 through 240 seconds.



Related Documentation • [join-prune-timeout on page 4933](#)

## Example: Enabling Join Suppression

This example describes how to enable PIM join suppression.

- [Requirements on page 4791](#)
- [Overview on page 4791](#)
- [Configuration on page 4793](#)
- [Verification on page 4795](#)

### Requirements

---

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 4786](#).

### Overview

---

PIM join suppression enables a router on a multiaccess network to defer sending join messages to an upstream router when it sees identical join messages on the same network. Eventually, only one router sends these join messages, and the other routers suppress identical messages. Limiting the number of join messages improves scalability and efficiency by reducing the number of messages sent to the same router.


This example includes the following statements:

- **override-interval**—Sets the maximum time in milliseconds to delay sending override join messages. When a router sees a prune message for a join it is currently suppressing, it waits before it sends an override join message. Waiting helps avoid multiple downstream routers sending override join messages at the same time. The override interval is a random timer with a value of 0 through the maximum override value.
- **propagation-delay**—Sets a value in milliseconds for a prune pending timer, which specifies how long to wait before executing a prune on an upstream router. During this period, the router waits for any prune override join messages that might be currently suppressed. The period for the prune pending timer is the sum of the **override-interval** value and the value specified for **propagation-delay**.
- **reset-tracking-bit**—Enables PIM join suppression on each multiaccess downstream interface. This statement resets a tracking bit field (T-bit) on the LAN prune delay hello option from the default of 1 (join suppression disabled) to 0 (join suppression enabled).

When multiple identical join messages are received, a random join suppression timer is activated, with a range of 66 through 84 milliseconds. The timer is reset each time join suppression is triggered.

Figure 165 on page 4792 shows the topology used in this example.

**Figure 165: Join Suppression**



The items in the figure represent the following functions:

- Host 0 is the multicast source.
- Host 1, Host 2, Host 3, and Host 4 are receivers.
- Router R0 is the first-hop router and the RP.
- Router R1 is an upstream router.
- Routers R2, R3, R4, and R5 are downstream routers in the multicast LAN.

This example shows the configuration of the downstream devices: Routers R2, R3, R4, and R5.

### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set protocols pim traceoptions file pim.log
set protocols pim traceoptions file size 5m
set protocols pim traceoptions file world-readable
set protocols pim traceoptions flag join detail
set protocols pim traceoptions flag prune detail
set protocols pim traceoptions flag normal detail
set protocols pim traceoptions flag register detail
set protocols pim rp static address 10.255.112.160
set protocols pim interface all mode sparse
set protocols pim interface all version 2
set protocols pim interface fxp0.0 disable
set protocols pim reset-tracking-bit
set protocols pim propagation-delay 500
set protocols pim override-interval 4000
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PIM join suppression on a non-RP downstream router in the multicast LAN:

1. Configure PIM sparse mode on the interfaces.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set rp static address 10.255.112.160
[edit protocols pim]
user@host# set interface all mode sparse version 2
[edit protocols pim]
user@host# set interface all version 2
[edit protocols pim]
user@host# set interface fxp0.0 disable
```

2. Enable the join suppression timer.

```
[edit protocols pim]
user@host# set reset-tracking-bit
```

3. Configure the prune override interval value.

```
[edit protocols pim]
user@host# set override-interval 4000
```

4. Configure the propagation delay of the link.

```
[edit protocols pim]
user@host# set propagation-delay 500
```

5. (Optional) Configure PIM tracing operations.

```
[edit protocols pim]
user@host# set traceoptions file pim.log size 5m world-readable
[edit protocols pim]
user@host# set traceoptions flag join detail
[edit protocols pim]
user@host# set traceoptions flag normal detail
[edit protocols pim]
user@host# set traceoptions flag register detail
```

6. If you are done configuring the device, commit the configuration.

```
[edit protocols pim]
user@host# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show protocols
pim {
  traceoptions {
    file pim.log size 5m world-readable;
    flag join detail;
    flag prune detail;
    flag normal detail;
    flag register detail;
  }
  rp {
    static {
      address 10.255.112.160;
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
  reset-tracking-bit;
  propagation-delay 500;
  override-interval 4000;
}
```

## Verification

To verify the configuration, run the following commands on the upstream and downstream routers:

- `show pim join extensive`
- `show multicast route extensive`

## Related Documentation

- [Example: Configuring the PIM Assert Timeout on page 4815](#)
- [Example: Configuring PIM RPF Selection](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 4818](#)
- [Enabling PIM Sparse Mode on page 4786](#)
- [PIM Overview on page 4737](#)

## Static RP

- [Configuring Local PIM RPs on page 4795](#)
- [Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4797](#)

## Configuring Local PIM RPs

Local RP configuration makes the routing device a statically defined RP. Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You can configure a local RP globally or for a routing instance. This example shows how to configure a local RP in a routing instance for IPv4 or IPv6.

To configure the routing device's RP properties:

1. Configure the routing instance as the local RP.

```
[routing-instances VPN-A protocols pim]
user@host# set rp local
```

2. Configure the IP protocol family and IP address.

IPv6 PIM hello messages are sent to every interface on which you configure **family inet6**, whether at the PIM level of the hierarchy or not. As a result, if you configure an interface with both **family inet** at the **[edit interface *interface-name*]** hierarchy level and **family inet6** at the **[edit protocols pim interface *interface-name*]** hierarchy level, PIM sends both IPv4 and IPv6 hellos to that interface.

By default, PIM operates in sparse mode on an interface. If you explicitly configure sparse mode, PIM uses this setting for all IPv6 multicast groups. However, if you configure sparse-dense mode, PIM does not accept IPv6 multicast groups as dense groups and operates in sparse mode over them.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set family inet6 address 2001:db8:85a3::8a2e:370:7334
user@host# set family inet address 10.1.2.254
```

3. (IPv4 only) Configure the routing device's RP priority.



**NOTE:** The priority statement is not supported for IPv6, but is included here for informational purposes. The routing device's priority value for becoming the RP is included in the bootstrap messages that the routing device sends. Use a smaller number to increase the likelihood that the routing device becomes the RP for local multicast groups. Each PIM routing device uses the priority value and other factors to determine the candidate RPs for a particular group range. After the set of candidate RPs is distributed, each routing device determines algorithmically the RP from the candidate RP set using a hash function. By default, the priority value is set to 1. If this value is set to 0, the bootstrap router can override the group range being advertised by the candidate RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set priority 5
```

4. Configure the groups for which the routing device is the RP.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which this routing device can be the RP.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set group-ranges fec0::/10
user@host# set group-ranges 10.1.2.0/24
```

5. (IPv4 only) Modify the local RP hold time.

If the local routing device is configured as an RP, it is considered a candidate RP for its local multicast groups. For candidate RPs, the hold time is used by the bootstrap router to time out RPs, and applies to the bootstrap RP-set mechanism. The RP hold time is part of the candidate RP advertisement message sent by the local routing device to the bootstrap router. If the bootstrap router does not receive a candidate RP advertisement from an RP within the hold time, it removes that routing device from its list of candidate RPs. The default hold time is 150 seconds.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set hold-time 200
```

6. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp local]
user@host# set override
```

7. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

**Related  
Documentation**

- [PIM Overview on page 4737](#)
- [Understanding MLD on page 4764](#)

## Configuring the Static PIM RP Address on the Non-RP Routing Device

Consider statically defining an RP if the network does not have many different RPs defined or if the RP assignment does not change very often. The Junos IPv6 PIM implementation supports only static RP configuration. Automatic RP announcement and bootstrap routers are not available with IPv6.

You configure a static RP address on the non-RP routing device. This enables the non-RP routing device to recognize the local statically defined RP. For example, if R0 is a non-RP router and R1 is the local RP router, you configure R0 with the static RP address of R1. The static IP address is the routable address assigned to the loopback interface on R1. In the following example, the loopback address of the RP is 2001:db8:85a3::8a2e:370:7334.

You can configure a static RP address globally or for a routing instance. This example shows how to configure a static RP address in a routing instance for IPv6.

To configure the static RP address:

1. On a non-RP routing device, configure the routing instance to point to the routable address assigned to the loopback interface of the RP.

```
[routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334
```



**NOTE:** Logical systems are also supported. You can configure a static RP address in a logical system only if the logical system is not directly connected to a source.

2. (Optional) Set the PIM sparse mode version.

For each static RP address, you can optionally specify the PIM version. The default PIM version is version 1.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 version 2
```



**NOTE:** The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIM version 1 is the default for RP mode ([edit pim rp static address *address*]). PIM version 2 is the default for interface mode ([edit pim interface *interface-name*]). Explicitly configured versions override the defaults.

---

3. (Optional) Set the group address range.

By default, a routing device running PIM is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12). The following example limits the groups for which the 2001:db8:85a3::8a2e:370:7334 address can be the RP.

```
[edit routing-instances VPN-A protocols pim rp]
user@host# set static address 2001:db8:85a3::8a2e:370:7334 group-ranges fec0::/10
```

The RP that you select for a particular group must be consistent across all routers in a multicast domain.

4. (Optional) Override dynamic RP for the specified group address range.

If you configure both static RP mapping and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for the given static RP group range, and allow dynamic RP mapping for all other groups.

If you exclude this statement from the configuration and you use both static and dynamic RP mechanisms for different group ranges within the same routing instance, the dynamic RP mapping takes precedence over the static RP mapping, even if static RP is defined for a specific group range.

```
[edit routing-instances VPN-A protocols pim rp static address
  2001:db8:85a3::8a2e:370:7334]
user@host# set override
```

5. Monitor the operation of PIM by running the **show pim** commands. Run **show pim ?** to display the supported commands.

- Related Documentation**
- [PIM Overview on page 4737](#)
  - [Understanding MLD on page 4764](#)

---

## Anycast RP

- [Example: Configuring PIM Anycast With or Without MSDP on page 4799](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 4802](#)
- [Configuring a PIM Anycast RP Router Using Only PIM on page 4803](#)
- [Configuring All PIM Anycast Non-RP Routers on page 4804](#)
- [Example: Configuring Multiple RPs in a Domain with Anycast RP on page 4805](#)



## Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32;
        primary;
        address 198.58.3.253/32;
      }
    }
  }
}
```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
  msdp {
    peer 198.58.3.250 {
      local-address 198.58.3.254;
    }
  }
}
```



**NOTE:** If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

---

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
```

```

        primary;
    }
    address 198.58.3.253/32;
}
}
}
}

```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```

protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, use lo0 primary
          }
        }
      }
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}
}

```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
  }
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

## Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}
```

```

    }
  }

```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```

protocols {
  msdp {
    peer 198.58.3.250 {
      local-address 198.58.3.254;
    }
  }
}

```

## Configuring a PIM Anycast RP Router Using Only PIM

In this example, configure an RP using the **lo0** loopback interface, which is always up. Use the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this case, the router ID is 198.58.3.254/32 and the shared RP address is 198.58.3.253/32. Add the flag statement **primary** to the first address. Using this flag selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}

```

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse**, and include the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

Use the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, use lo0 primary
          }
        }
      }
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}
```

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

## Configuring All PIM Anycast Non-RP Routers

Use the **mode** statement at the **[edit protocols pim rp interface all]** hierarchy level to specify sparse mode on all interfaces. Then add the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
```

## Example: Configuring Multiple RPs in a Domain with Anycast RP

This example shows how to configure anycast RP on each RP router in the PIM-SM domain. With this configuration you can deploy more than one RP for a single group range. This enables load balancing and redundancy.

- [Requirements on page 4805](#)
- [Overview on page 4805](#)
- [Configuration on page 4805](#)
- [Verification on page 4807](#)

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See “[Enabling PIM Sparse Mode](#)” on [page 4786](#).

### Overview

When you configure anycast RP, the RP routers in the PIM-SM domain use a shared address. In this example, the shared address is 10.1.1.2/32. Anycast RP uses Multicast Source Discovery Protocol (MSDP) to discover and maintain a consistent view of the active sources. Anycast RP also requires an RP selection method, such as static, auto-RP, or bootstrap RP. This example uses static RP and shows only one RP router configuration.

### Configuration

<b>CLI Quick Configuration</b>	To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the <b>[edit]</b> hierarchy level.
<b>RP Routers</b>	<pre> set interfaces lo0 unit 0 family inet address 192.168.132.1/32 primary set interfaces lo0 unit 0 family inet address 10.1.1.2/32 set protocols msdp local-address 192.168.132.1 set protocols msdp peer 192.168.12.1 set protocols pim rp local address 10.1.1.2 set routing-options router-id 192.168.132.1 </pre>
<b>Non-RP Routers</b>	<pre> set protocols pim rp static address 10.1.1.2 </pre>

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure anycast RP:

1. On each RP router in the domain, configure the shared anycast address on the router's loopback address.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.1.1.2/32
```

2. On each RP router in the domain, make sure that the router's regular loopback address is the primary address for the interface, and set the router ID.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 192.168.132.1/32 primary
```

```
[edit routing-options]
user@host# set router-id 192.168.132.1
```

3. On each RP router in the domain, configure the local RP address, using the shared address.

```
[edit protocols pim]
user@host# set rp local address 10.1.1.2
```

4. On each RP router in the domain, create MSDP sessions to the other RPs in the domain.

```
[edit protocols msdp]
user@host# set local-address 192.168.132.1
user@host# set peer 192.168.12.1
```

5. On each non-RP router in the domain, configure a static RP address using the shared address.

```
[edit protocols pim]
user@host# set rp static address 10.1.1.2
```

6. If you are done configuring the devices, commit the configuration.

```
user@host# commit
```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
lo0 {
  unit 0 {
    family inet {
      address 192.168.132.1/32 {
        primary;
      }
      address 10.1.1.2/32;
    }
  }
}
```



```

    }
  }
}

```

*On the RP routers:*

```

user@host# show protocols
msdp {
  local-address 192.168.132.1;
  peer 192.168.12.1;
}
pim {
  rp {
    local {
      address 10.1.1.2;
    }
  }
}
}

```

*On the non-RP routers:*

```

user@host# show protocols
pim {
  rp {
    static {
      address 10.1.1.2;
    }
  }
}

user@host# show routing-options
router-id 192.168.132.1;

```

### Verification

To verify the configuration, run the `show pim rps extensive inet` command.

#### Related Documentation

- [Example: Configuring PIM Anycast With or Without MSDP on page 4799](#)
- [Understanding PIM Sparse Mode on page 4741](#)
- [Understanding RP Mapping with Anycast RP on page 4745](#)

## PIM Bootstrap Router

- [Configuring PIM Bootstrap Properties for IPv4 or IPv6 on page 4807](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 4809](#)
- [Example: Configuring PIM BSR Filters on page 4809](#)

### Configuring PIM Bootstrap Properties for IPv4 or IPv6

For correct operation, every multicast router within a PIM domain must be able to map a particular multicast group address to the same rendezvous point (RP). The bootstrap

router mechanism is one way that a multicast router can learn the set of group-to-RP mappings. Bootstrap routers are supported in IPv4 and IPv6.

To determine which routing device is the RP, all routing devices within a PIM domain collect bootstrap messages. A PIM domain is a contiguous set of routing devices that implement PIM. All devices are configured to operate within a common boundary. The domain's bootstrap router initiates bootstrap messages, which are sent hop by hop within the domain. The routing devices use bootstrap messages to distribute RP information dynamically and to elect a bootstrap router when necessary.

You can configure bootstrap properties globally or for a routing instance. This example shows the global configuration.

To configure the bootstrap router properties:

1. Configure the bootstrap priority.

By default, each routing device has a bootstrap priority of 0, which means the routing device can never be the bootstrap router. The routing device with the highest priority value is elected to be the bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router. A simple bootstrap configuration assigns a bootstrap priority value to a routing device.



**NOTE:** In the IPv4-only configuration, specifying a bootstrap priority of 0 disables the bootstrap function and does not cause the routing device to send BSR packets with a 0 in the priority field. In the combined IPv4 and IPv6 configuration, specifying a bootstrap priority of 0 does not disable the function, but causes the routing device to send BSR packets with a 0 in the priority field. To disable the bootstrap function in the IPv4 and IPv6 configuration, delete the `bootstrap` statement.

```
user@host# edit protocols pim rp
user@host# set bootstrap family inet priority 3
```

2. (Optional) Create import and export policies to control the flow of bootstrap messages to and from the RP, and apply the policies to PIM. Import and export policies are useful when some of the routers in your PIM domain have interfaces that connect to other PIM domains. Configuring a policy prevents bootstrap messages from crossing domain boundaries. The **import** statement prevents messages from being imported into the RP. The **export** statement prevents messages from being exported from the RP.

```
[edit protocols pim rp]
user@host# set bootstrap family inet import pim-bootstrap-import
user@host# set bootstrap family inet export pim-bootstrap-export
user@host# exit
```

3. Configure the policies.

```
user@host# edit policy-options policy-statement pim-bootstrap-import
[edit policy-options policy-statement pim-bootstrap-import]
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit
```

```

user@host# edit policy-options policy-statement pim-bootstrap-export
user@host# set from interface se-0/0/0
user@host# set then reject
user@host# exit

```

4. Monitor the operation of PIM bootstrap routers by running the `show pim bootstrap` command.

#### Related Documentation

- [Understanding PIM Sparse Mode on page 4741](#)
- [Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain on page 4809](#)
- [show pim bootstrap on page 5148](#) in the CLI Explorer

### Example: Rejecting PIM Bootstrap Messages at the Boundary of a PIM Domain

In this example, the `from interface so-0-1/0 then reject` policy statement rejects bootstrap messages from the specified interface (the example is configured for both IPv4 and IPv6 operation):

```

protocols {
  pim {
    rp {
      bootstrap {
        family inet {
          priority 1;
          import pim-import;
          export pim-export;
        }
        family inet6 {
          priority 1;
          import pim-import;
          export pim-export;
        }
      }
    }
  }
}
policy-options {
  policy-statement pim-import {
    from interface so-0/1/0;
    then reject;
  }
  policy-statement pim-export {
    to interface so-0/1/0;
    then reject;
  }
}

```

### Example: Configuring PIM BSR Filters

Configure a filter to prevent BSR messages from entering or leaving your network. Add this configuration to all routers:

```
protocols {
  pim {
    rp {
      bootstrap-import no-bsr;
      bootstrap-export no-bsr;
    }
  }
}
policy-options {
  policy-statement no-bsr {
    then reject;
  }
}
```

---

## PIM Filtering

- [Configuring Interface-Level PIM Neighbor Policies on page 4810](#)
- [Filtering Outgoing PIM Join Messages on page 4811](#)
- [Filtering Incoming PIM Join Messages on page 4812](#)
- [Configuring Register Message Filters on a PIM RP and DR on page 4813](#)

### Configuring Interface-Level PIM Neighbor Policies

You can configure a policy to filter unwanted PIM neighbors. In the following example, the PIM interface compares neighbor IP addresses with the IP address in the policy statement before any hello processing takes place. If any of the neighbor IP addresses (primary or secondary) match the IP address specified in the prefix list, PIM drops the hello packet and rejects the neighbor.

If you configure a PIM neighbor policy after PIM has already established a neighbor adjacency to an unwanted PIM neighbor, the adjacency remains intact until the neighbor hold time expires. When the unwanted neighbor sends another hello message to update its adjacency, the router recognizes the unwanted address and rejects the neighbor.

To configure a policy to filter unwanted PIM neighbors:

1. Configure the policy. The neighbor policy must be a properly structured policy statement that uses a prefix list (or a route filter) containing the neighbor primary address (or any secondary IP addresses) in a prefix list, and the **reject** option to reject the unwanted address.

```
[edit policy-options]
user@host# set prefix-list nbrGroup 1 20.20.20.1/32
user@host# set policy-statement nbr-policy from prefix-list nbrGroup1
user@host# set policy-statement nbr-policy then reject
```

2. Configure the interface globally or in the routing instance. This example shows the configuration for the routing instance.

```
[edit routing-instances PIM.master protocols pim]
user@host# set neighbor-policy nbr-policy
```

3. Verify the configuration by checking the **Hello dropped on neighbor policy** field in the output of the **show pim statistics** command.

- Related Documentation**
- [Understanding PIM Sparse Mode on page 4741](#)
  - *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
  - [show pim statistics on page 5189](#) in the [CLI Explorer](#)

## Filtering Outgoing PIM Join Messages

When the core of your network is using MPLS, PIM join and prune messages stop at the customer edge (CE) routers and are not forwarded toward the core, because these routers do not have PIM neighbors on the core-facing interfaces. When the core of your network is using IP, PIM join and prune messages are forwarded to the upstream PIM neighbors in the core of the network.

When the core of your network is using a mix of IP and MPLS, you might want to filter certain PIM join and prune messages at the upstream egress interface of the CE routers.

You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. If the filter is removed, the join is sent after the PIM periodic join timer expires.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, and then apply the policy.

The following example filters PIM join and prune messages for group addresses 224.0.1.2 and 225.1.1.1.

1. In configuration mode, create the policy.

```
user@host# set policy-options policy-statement block-groups term t1 from route-filter
224.0.1.2/32 exact
user@host# set policy-options policy-statement block-groups term t1 from route-filter
225.1.1.1/32 exact
user@host# set policy-options policy-statement block-groups term t1 then reject
user@host# set policy-options policy-statement block-groups term last then accept
```

2. Verify the policy configuration by running the **show policy-options** command.

```
user@host# show policy-options
policy-statement block-groups {
  term t1 {
    from {
      route-filter 224.0.1.2/32 exact;
      route-filter 225.1.1.1/32 exact;
      then reject;
    }
    term last {
      then accept;
    }
  }
}
```

3. Apply the PIM join and prune message filter.

```
user@host> set protocols pim export block-groups
```

4. After the configuration is committed, use the **show pim statistics** command to verify that outgoing PIM join and prune messages are being filtered.

```
user@host> show pim statistics | grep filtered
RP Filtered Source                0
Rx Joins/Prunes filtered          0
Tx Joins/Prunes filtered          254
```

The egress filter count is shown on the **Tx Joins/Prunes filtered** line.

**Related Documentation**

- [Filtering Incoming PIM Join Messages on page 4812](#)

## Filtering Incoming PIM Join Messages

Multicast scoping controls the propagation of multicast messages. Whereas multicast scoping prevents the actual multicast data packets from flowing in or out of an interface, PIM join filters prevent a state from being created in a router. A state—the (\*,G) or (S,G) entries—is the information used for forwarding unicast or multicast packets. Using PIM join filters prevents the transport of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. Also, PIM join filters reduce the potential for denial-of-service (DoS) attacks and PIM state explosion—large numbers of PIM join messages forwarded to each router on the rendezvous-point tree (RPT), resulting in memory consumption.

To use PIM join filters to efficiently restrict multicast traffic from certain source addresses, create and apply the routing policy across all routers in the network.

See [Table 387 on page 4812](#) for a list of match conditions.

**Table 387: PIM Join Filter Match Conditions**

Match Condition	Matches On
<b>interface</b>	Router interface or interfaces specified by name or IP address
<b>neighbor</b>	Neighbor address (the source address in the IP header of the join and prune message)
<b>route-filter</b>	Multicast group address embedded in the join and prune message
<b>source-address-filter</b>	Multicast source address embedded in the join and prune message

The following example shows how to create a PIM join filter. The filter is composed of a route filter and a source address filter—**bad-groups** and **bad-sources**, respectively. the **bad-groups** filter prevents (\*,G) or (S,G) join messages from being received for all groups listed. The **bad-sources** filter prevents (S,G) join messages from being received for all sources listed. The **bad-groups** filter and **bad-sources** filter are in two different terms. If route filters and source address filters are in the same term, they are logically ANDed.

To filter incoming PIM join messages:

1. Configure the policy.

```
[edit policy-statement pim-join-filter term bad-groups]
user@host# set from route-filter 224.0.1.2/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term bad-sources]
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject

[edit policy-statement pim-join-filter term last]
user@host# set then accept
```

2. Apply one or more policies to routes being imported into the routing table from PIM.

```
[edit protocols pim]
user@host# set import pim-join-filter
```

3. Verify the configuration by checking the output of the **show pim join** and **show policy** commands.

#### Related Documentation

- [Understanding Multicast Administrative Scoping](#)
- [Filtering Outgoing PIM Join Messages on page 4811](#)
- [show pim join on page 5153](#) in the [CLI Explorer](#)
- [show policy](#) in the [CLI Explorer](#)

## Configuring Register Message Filters on a PIM RP and DR

PIM register messages are sent to the rendezvous point (RP) by a designated router (DR). When a source for a group starts transmitting, the DR sends unicast PIM register packets to the RP.

Register messages have the following purposes:

- Notify the RP that a source is sending to a group.
- Deliver the initial multicast packets sent by the source to the RP for delivery down the shortest-path tree (SPT).

The PIM RP keeps track of all active sources in a single PIM sparse mode domain. In some cases, you want more control over which sources an RP discovers, or which sources a DR notifies other RPs about. A high degree of control over PIM register messages is provided by RP or DR register message filtering. Message filtering prevents unauthorized groups and sources from registering with an RP router.

You configure RP or DR register message filtering to control the number and location of multicast sources that an RP discovers. You can apply register message filters on a DR to control outgoing register messages, or apply them on an RP to control incoming register messages.

When anycast RP is configured, all RPs in the anycast RP set need to be configured with the same register message filtering policy.

You can configure message filtering globally or for a routing instance. These examples show the global configuration.

To configure an RP filter to drop the register packets for multicast group range 224.1.1.0/24 from source address 10.10.94.2:

1. On the RP, configure the policy.

```
[edit policy-options policy-statement incoming-policy-for-rp from]
user@host# set route-filter 224.1.1.0/24 orlonger
user@host# set source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the RP.

```
[edit protocols pim rp]
user@host# set rp-register-policy incoming-policy-for-rp
user@host# set local address 10.10.10.5
user@host# exit
```

To configure a DR filter to prevent sending register packets for group range 224.1.1.0/24 and source address 10.10.10.1/32:

1. On the DR, configure the policy.

```
[edit policy-options policy-statement outgoing-policy-for-rp]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.10.1/32 exact
user@host# set then reject
user@host# exit
```

2. Apply the policy to the DR.

The static address is the address of the RP to which you do not want the DR to send the filtered register messages.

```
[edit protocols pim rp]
user@host# set dr-register-policy outgoing-policy-for-dr
user@host# set static 10.10.10.3
user@host# exit
```

To configure a policy expression to accept register messages for multicast group 224.1.1.5 but reject those for 224.1.1.1:

1. On the RP, configure the policies.

```
[edit policy-options policy-statement reject_224_1_1_1]
user@host# set from route-filter 224.1.1.0/24 orlonger
user@host# set from source-address-filter 10.10.94.2/32 exact
user@host# set then reject
user@host# exit
```

```
[edit policy-options policy-statement accept_224_1_1_5]
user@host# set term one from route-filter 224.1.1.5/32 exact
```



```

user@host# set term one from source-address-filter 10.10.94.2/32 exact
user@host# set term one then accept
user@host# set term two then reject
user@host# exit

```

2. Apply the policies to the RP.

```

[edit protocols pim rp]
user@host# set rp-register-policy [ reject_224_1_1_1 | accept_224_1_1_5 ]
user@host# set local address 10.10.10.5

```

To monitor the operation of the filters, run the **show pim statistics** command. The command output contains the following fields related to filtering:

- RP Filtered Source
- Rx Joins/Prunes filtered
- Tx Joins/Prunes filtered
- Rx Register msgs filtering drop
- Tx Register msgs filtering drop

#### Related Documentation

- [PIM Sparse Mode Source Registration on page 4750](#)
- [Filtering RP and DR Register Messages on page 4747](#)
- [show pim statistics on page 5189](#) in the CLI Explorer

## PIM RPT and SPT Cutover

- [Example: Configuring the PIM Assert Timeout on page 4815](#)
- [Example: Configuring the PIM SPT Threshold Policy on page 4818](#)

### Example: Configuring the PIM Assert Timeout

This example shows how to configure the timeout period for a PIM assert forwarder.

- [Requirements on page 4815](#)
- [Overview on page 4816](#)
- [Configuration on page 4817](#)

#### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 4786](#).

## Overview

---

The role of PIM assert messages is to determine the forwarder on a network with multiple routers. The forwarder is the router that forwards multicast packets to a network with multicast group members. The forwarder is generally the same as the PIM DR.

A router sends an assert message when it receives a multicast packet on an interface that is listed in the outgoing interface list of the matching routing entry. Receiving a message on an outgoing interface is an indication that more than one router forwards the same multicast packets to a network.

In [Figure 166 on page 4817](#), both routing devices R1 and R2 forward multicast packets for the same (S,G) entry on a network. Both devices detect this situation and both devices send assert messages on the Ethernet network. An assert message contains, in addition to a source address and group address, a unicast cost metric for sending packets to the source, and a preference metric for the unicast cost. The preference metric expresses a preference between unicast routing protocols. The routing device with the smallest preference metric becomes the forwarder (also called the assert winner). If the preference metrics are equal, the device that sent the lowest unicast cost metric becomes the forwarder. If the unicast metrics are also equal, the routing device with the highest IP address becomes the forwarder. After the transmission of assert messages, only the forwarder continues to forward messages on the network.


When an assert message is received and the RPF neighbor is changed to the assert winner, the assert timer is set to an assert timeout period. The assert timeout period is restarted every time a subsequent assert message for the route entry is received on the incoming interface. When the assert timer expires, the routing device resets its RPF neighbor according to its unicast routing table. Then, if multiple forwarders still exist, the forwarders reenter the assert message cycle. In effect, the assert timeout period determines how often multicast routing devices enter a PIM assert message cycle.

The range is from 5 through 210 seconds. The default is 180 seconds.

Assert messages are useful for LANs that connect multiple routing devices and no hosts.

[Figure 166 on page 4817](#) shows the topology for this example.

Figure 166: PIM Assert Topology



### Configuration

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an assert timeout:

1. Configure the timeout period, in seconds.  

```
[edit protocols pim]
user@host# set assert-timeout 60
```
2. (Optional) Trace assert messages.  

```
[edit protocols pim]
user@host# set traceoptions file PIM.log
user@host# set traceoptions flag assert detail
```
3. If you are done configuring the device, commit the configuration.  

```
user@host# commit
```
4. To verify the configuration, run the following commands:
  - `show pim join`
  - `show pim statistics`

#### Related Documentation

- [Configuring PIM Trace Options on page 4780](#)
- [SPT Cutover on page 4754](#)
- [SPT Cutover Control on page 4757](#)

## Example: Configuring the PIM SPT Threshold Policy

This example shows how to apply a policy that suppresses the transition from the rendezvous-point tree (RPT) rooted at the RP to the shortest-path tree (SPT) rooted at the source.

- [Requirements on page 4818](#)
- [Overview on page 4818](#)
- [Configuration on page 4819](#)
- [Verification on page 4821](#)

---

### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Configure PIM Sparse Mode on the interfaces. See [“Enabling PIM Sparse Mode” on page 4786](#).

---

### Overview

Multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or through an SPT rooted at the source. In some cases, the last-hop routing device needs to stay on the shared RPT to the RP and not transition to a direct SPT to the source. Receiving the multicast data traffic on SPT is optimal but introduces more state in the network, which might not be desirable in some multicast deployments. Ideally, low-bandwidth multicast streams can be forwarded on the SPT, and high-bandwidth streams can use the SPT. This example shows how to configure such a policy.

This example includes the following settings:

- **spt-threshold**—Enables you to configure an SPT threshold policy on the last-hop routing device to control the transition to a direct SPT. When you include this statement in the main PIM instance, the PE router stays on the RPT for control traffic.
- **infinity**—Applies an SPT cutover threshold of infinity to a source-group address pair, so that the last-hop routing device never transitions to a direct SPT. For all other source-group address pairs, the last-hop routing device transitions immediately to a direct SPT rooted at the source DR. This statement must reference a properly configured policy to set the SPT cutover threshold for a particular source-group pair to infinity. The use of values other than infinity for the SPT threshold is not supported. You can configure more than one policy.
- **policy-statement**—Configures the policy. The simplest type of SPT threshold policy uses a route filter and source address filter to specify the multicast group and source

addresses and to set the SPT threshold for that pair of addresses to infinity. The policy is applied to the main PIM instance.

This example sets the SPT transition value for the source-group pair 10.10.10.1 and 224.1.1.1 to infinity. When the policy is applied to the last-hop router, multicast traffic from this source-group pair never transitions to a direct SPT to the source. Traffic will continue to arrive through the RP. However, traffic for any other source-group address combination at this router transitions to a direct SPT to the source.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.

Note these points when configuring the SPT threshold policy:

- Configuration changes to the SPT threshold policy affect how the routing device handles the SPT transition.
- When the policy is configured for the first time, the routing device continues to transition to the direct SPT for the source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- If you do not clear the PIM-join state when you apply the infinity policy configuration for the first time, you must apply it before the PE router is brought up.
- When the policy is deleted for a source-group address pair for the first time, the routing device does not transition to the direct SPT for that source-group address pair until the PIM-join state is cleared with the **clear pim join** command.
- When the policy is changed for a source-group address pair for the first time, the routing device does not use the new policy until the PIM-join state is cleared with the **clear pim join** command.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set policy-options policy-statement spt-infinity-policy term one from route-filter
  224.1.1.1/32 exact
set policy-options policy-statement spt-infinity-policy term one from source-address-filter
  10.10.10.1/32 exact
set policy-options policy-statement spt-infinity-policy term one then accept
set policy-options policy-statement spt-infinity-policy term two then reject
set protocols pim spt-threshold infinity spt-infinity-policy
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure an SPT threshold policy:

1. Apply the policy.

```
[edit]
user@host# edit protocols pim
[edit protocols pim]
user@host# set spt-threshold infinity spt-infinity-policy
[edit protocols pim]
user@host# exit
```

2. Configure the policy.

```
[edit]
user@host# edit policy-options policy-statement spt-infinity-policy
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from route-filter 224.1.1.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one from source-address-filter 10.10.10.1/32 exact
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term one then accept
[edit policy-options policy-statement spt-infinity-policy]
user@host# set term two then reject
[edit policy-options policy-statement spt-infinity-policy]
user@host# exit
policy-statement {
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

4. Clear the PIM join cache to force the configuration to take effect.

```
[edit]
user@host# run clear pim join
```

### Results

Confirm your configuration by entering the **show policy-options** command and the **show protocols** command from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show policy-options
policy-statement spt-infinity-policy {
  term one {
    from {
      route-filter 224.1.1.1/32 exact;
      source-address-filter 10.10.10.1/32 exact;
    }
    then accept;
  }
  term two {
    then reject;
  }
}
```

```

    }
  }

user@host# show protocols
pim {
  spt-threshold {
    infinity spt-infinity-policy;
  }
}

```

### Verification

To verify the configuration, run the `show pim join` command.

**Related Documentation**

- [SPT Cutover Control on page 4757](#)

## PIM and the BFD Protocol

- [Configuring BFD for PIM on page 4821](#)
- [Configuring BFD Authentication for PIM on page 4823](#)

### Configuring BFD for PIM

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the Protocol Independent Multicast (PIM) hello hold time, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the `clear bfd adaptation` command to return BFD interval timers to their configured values. The `clear bfd adaptation` command is hitless, meaning that the command does not affect traffic flow on the routing device.

You must specify the minimum transmit and minimum receive intervals to enable BFD on PIM.

To enable failure detection:

1. Configure the interface globally or in a routing instance.

This example shows the global configuration.

```
[edit protocols pim]
user@host# edit interface fe-1/0/0.0 family inet bfd-liveness-detection
```

2. Configure the minimum transmit interval.

This is the minimum interval after which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set transmit-interval 350
```

3. Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-receive-interval 350
```

4. (Optional) Configure other BFD settings.

As an alternative to setting the receive and transmit intervals separately, configure one interval for both.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set minimum-interval 350
```

5. Configure the threshold for the adaptation of the BFD session detection time.

When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set detection-time threshold 800
```

6. Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set multiplier 50
```

7. Configure the BFD version.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set version 1
```

8. Specify that BFD sessions should not adapt to changing network conditions.

We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

```
[edit protocols pim interface fe-1/0/0.0 family inet bfd-liveness-detection]
user@host# set no-adaptation
```

9. Verify the configuration by checking the output of the **show bfd session** command.

**Related Documentation**

- *show bfd session* in the [CLI Explorer](#)



## Configuring BFD Authentication for PIM

Beginning with Junos OS Release 9.6, you can configure authentication for Bidirectional Forwarding Detection (BFD) sessions running over Protocol Independent Multicast (PIM). Routing instances are also supported. The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the PIM protocol.
2. Associate the authentication keychain with the PIM protocol.
3. Configure the related security authentication keychain.

The following sections provide instructions for configuring and viewing BFD authentication on PIM:

- [Configuring BFD Authentication Parameters on page 4823](#)
- [Viewing Authentication Information for BFD Sessions on page 4824](#)

### Configuring BFD Authentication Parameters

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

To configure BFD authentication:

1. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on a PIM route or routing instance.

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
algorithm keyed-sha-1
```



**NOTE:** Nonstop active routing (NSR) is not supported with the **meticulous-keyed-md5** and **meticulous-keyed-sha-1** authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

2. Specify the keychain to be used to associate BFD sessions on the specified PIM route or routing instance with the unique security authentication keychain attributes.

The keychain you specify must match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

```
[edit protocols pim]
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication
keychain bfd-pim
```



**NOTE:** The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

3. Specify the unique security authentication information for BFD sessions:

- The matching keychain name as specified in Step 2.
- At least one key, a unique integer between 0 and 63. Creating multiple keys allows multiple clients to use the BFD session.
- The secret data used to allow access to the session.
- The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security]
```

```
user@host# set authentication-key-chains key-chain bfd-pim key 53 secret  
$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHM start-time 2009-06-14.10:00:00
```

4. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols pim]
```

```
user@host# set interface ge-0/1/5 family inet bfd-liveness-detection authentication  
loose-check
```

5. (Optional) View your configuration by using the **show bfd session detail** or **show bfd session extensive** command.

6. Repeat these steps to configure the other end of the BFD session.

---

### Viewing Authentication Information for BFD Sessions

You can view the existing BFD authentication configuration by using the **show bfd session detail** and **show bfd session extensive** commands.

The following example shows BFD authentication configured for the **ge-0/1/5** interface. It specifies the keyed SHA-1 authentication algorithm and a keychain name of **bfd-pim**. The authentication keychain is configured with two keys. Key 1 contains the secret data “\$9\$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHM” and a start time of June 1, 2009, at 9:46:02 AM PST. Key 2 contains the secret data “\$9\$a5jiKW9L.reP38ny.TszF2/9” and a start time of June 1, 2009, at 3:29:20 PM PST.

```
[edit protocols pim]
```

```
interface ge-0/1/5 {  
  family inet {  
    bfd-liveness-detection {  
      authentication {  
        key-chain bfd-pim;  
        algorithm keyed-sha-1;  
      }  
    }  
  }  
}
```

```

}
[edit security]
authentication key-chains {
  key-chain bfd-pim {
    key 1 {
      secret "$9$ggaJDmPQ6/tJgF/AtREVsyPsnCtUHm";
      start-time "2009-6-1.09:46:02 -0700";
    }
    key 2 {
      secret "$9$a5jiKW9l.reP38ny.TszF2/9";
      start-time "2009-6-1.15:29:20 -0700";
    }
  }
}
}

```

If you commit these updates to your configuration, you see output similar to the following example. In the output for the **show bfd session detail** command, **Authenticate** is displayed to indicate that BFD authentication is configured. For more information about the configuration, use the **show bfd session extensive** command. The output for this command provides the keychain name, the authentication algorithm and mode for each client in the session, and the overall BFD authentication configuration status, keychain name, and authentication algorithm and mode.

#### show bfd session detail

```
user@host# show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
50.0.0.2	Up	ge-0/1/5.0	0.900	0.300	3
Client PIM, TX interval 0.300, RX interval 0.300, <b>Authenticate</b>					
Session up time 3d 00:34					
Local diagnostic None, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					

#### show bfd session extensive

```
user@host# show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
50.0.0.2	Up	ge-0/1/5.0	0.900	0.300	3
Client PIM, TX interval 0.300, RX interval 0.300, <b>Authenticate</b>					
keychain bfd-pim, algo keyed-sha-1, mode strict					
Session up time 00:04:42					
Local diagnostic None, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					
Min async interval 0.300, min slow interval 1.000					
Adaptive async TX interval 0.300, RX interval 0.300					
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3					
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3					
Local discriminator 2, remote discriminator 2					
Echo mode disabled/inactive					
<b>Authentication enabled/active, keychain bfd-pim, algo keyed-sha-1, mode strict</b>					

**Related Documentation**

- *Understanding Bidirectional Forwarding Detection Authentication for PIM*

- [Configuring BFD for PIM on page 4821](#)
- [authentication-key-chains on page 5466](#)
- [bfd-liveness-detection on page 4911](#)
- *show bfd session* in the [CLI Explorer](#)

## IGMP

---

- [Configuring IGMP on page 4826](#)
- [Enabling IGMP on page 4828](#)
- [Changing the IGMP Version on page 4829](#)
- [Modifying the IGMP Host-Query Message Interval on page 4830](#)
- [Modifying the IGMP Last-Member Query Interval on page 4831](#)
- [Specifying Immediate-Leave Host Removal for IGMP on page 4831](#)
- [Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 4832](#)
- [Accepting IGMP Messages from Remote Subnetworks on page 4833](#)
- [Modifying the IGMP Query Response Interval on page 4834](#)
- [Modifying the IGMP Robustness Variable on page 4835](#)
- [Limiting the Maximum IGMP Message Rate on page 4836](#)
- [Enabling IGMP Static Group Membership on page 4836](#)
- [Recording IGMP Join and Leave Events on page 4843](#)
- [Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 4844](#)
- [Tracing IGMP Protocol Traffic on page 4845](#)
- [Disabling IGMP on page 4847](#)

## Configuring IGMP

Before you begin:

1. Determine whether the router is directly attached to any multicast sources. Receivers must be able to locate these sources.
2. Determine whether the router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
3. Determine whether to configure multicast to use sparse, dense, or sparse-dense mode. Each mode has different configuration considerations.
4. Determine the address of the RP if sparse or sparse-dense mode is used.
5. Determine whether to locate the RP with the static configuration, BSR, or auto-RP method.

6. Determine whether to configure multicast to use its own RPF routing table when configuring PIM in sparse, dense, or sparse-dense mode.
7. Configure the SAP and SDP protocols to listen for multicast session announcements. See *Configuring the Session Announcement Protocol*.

To configure the Internet Group Management Protocol (IGMP), include the **igmp** statement:

```
igmp {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-policy [ policy-names ];
    immediate-leave;
    oif-map map-name;
    promiscuous-mode;
    ssm-map ssm-map-name;
    static {
      group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
          source-count number;
          source-increment increment;
        }
      }
    }
  }
  version version;
}
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, IGMP is enabled on all interfaces on which you configure Protocol Independent Multicast (PIM), and on all broadcast interfaces on which you configure the Distance Vector Multicast Routing Protocol (DVMRP).



**NOTE:** You can configure IGMP on an interface without configuring PIM. PIM is generally not needed on IGMP downstream interfaces. Therefore, only one “pseudo PIM interface” is created to represent all IGMP downstream (IGMP-only) interfaces on the router. This reduces the amount of router resources, such as memory, that are consumed. You must configure PIM on upstream IGMP interfaces to enable multicast routing, perform reverse-path forwarding for multicast data packets, populate the multicast forwarding table for upstream interfaces, and in the case of bidirectional PIM and PIM sparse mode, to distribute IGMP group memberships into the multicast routing domain.

## Enabling IGMP

The Internet Group Management Protocol (IGMP) manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use IGMP to learn which groups have members on each of their attached physical networks. IGMP must be enabled for the router to receive IPv4 multicast packets. IGMP is only needed for IPv4 networks, because multicast is handled differently in IPv6 networks. IGMP is automatically enabled on all IPv4 interfaces on which you configure PIM and on all IPv4 broadcast interfaces when you configure DVMRP.

If IGMP is not running on an interface—either because PIM and DVMRP are not configured on the interface or because IGMP is explicitly disabled on the interface—you can explicitly enable IGMP.

To explicitly enable IGMP:

1. If PIM and DVMRP are not running on the interface, explicitly enable IGMP by including the interface name.

```
[edit protocols igmp]
user@host# set interface fe-0/0/0.0
```

2. See if IGMP is disabled on any interfaces. In the following example, IGMP is disabled on a Gigabit Ethernet interface.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
interface ge-1/0/0.0 {
  disable;
}
```

3. Enable IGMP on the interface by deleting the **disable** statement.

```
[edit protocols igmp]
delete interface ge-1/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols igmp]
user@host# show
interface fe-0/0/0.0;
```

```
interface ge-1/0/0.0;
```

5. Verify the operation of IGMP on the interfaces by checking the output of the **show igmp interface** command.

#### Related Documentation

- [Understanding IGMP on page 4759](#)
- [Disabling IGMP on page 4847](#)
- [show igmp interface on page 5089](#)

## Changing the IGMP Version

By default, the routing device runs IGMPv2. Routing devices running different versions of IGMP determine the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version.

To enable source-specific multicast (SSM) functionality, you must configure version 3 on the host and the host's directly connected routing device. If a source address is specified in a multicast group that is statically configured, the version must be set to IGMPv3.

If a static multicast group is configured with the source address defined, and the IGMP version is configured to be version 2, the source is ignored and only the group is added. In this case, the join is treated as an IGMPv2 group join.

If you configure the IGMP version setting at the individual interface hierarchy level, it overrides the **interface all** statement.

If you have already configured the routing device to use IGMP version 1 (IGMPv1) and then configure it to use IGMPv2, the routing device continues to use IGMPv1 for up to 6 minutes and then uses IGMPv2.

To change to IGMPv3 for SSM functionality:

1. Configure the IGMP interface.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0 version 3
```

2. Verify the configuration by checking the version field in the output of the **show igmp interfaces** command. The **show igmp statistics** command has version-specific output fields, such as V1 Membership Report, V2 Membership Report, and V3 Membership Report.



**CAUTION:** On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

#### Related Documentation

- [Understanding IGMP on page 4759](#)
- [show pim interfaces on page 5150](#)

- [show igmp statistics on page 5093](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

## Modifying the IGMP Host-Query Message Interval

The objective of IGMP is to keep routers up to date with group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The IGMP querier router periodically sends general host-query messages on each attached network to solicit membership information. The messages are sent to the all-systems multicast group address, 224.0.0.1.

The query interval, the response interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of IGMP messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols igmp]  
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the IGMP Query Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

### Related Documentation

- [Understanding IGMP on page 4759](#)
- [Modifying the IGMP Query Response Interval on page 4834](#)
- [Modifying the IGMP Robustness Variable on page 4835](#)
- [show igmp interface on page 5089](#)
- [show igmp statistics on page 5093](#)



## Modifying the IGMP Last-Member Query Interval

The last-member query interval is the maximum amount of time between group-specific query messages, including those sent in response to leave-group messages. You can configure this interval to change the amount of time it takes a routing device to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group message from a host, the routing device sends multiple group-specific queries to the group being left. The querier sends a specific number of these queries at a specific interval. The number of queries sent is called the last-member query count. The interval at which the queries are sent is called the last-member query interval. Because both settings are configurable, you can adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-member query count x (times) the last-member query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-member query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols igmp]
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the IGMP Last Member Query Interval field in the output of the **show igmp interfaces** command.



**NOTE:** You can configure the last-member query count by configuring the robustness variable. The two are always equal.

### Related Documentation

- [Modifying the IGMP Robustness Variable on page 4835](#)
- [show pim interfaces on page 5150](#)

## Specifying Immediate-Leave Host Removal for IGMP

The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave on an interface:

1. Configure immediate leave on the IGMP interface.

```
[edit protocols IGMP]
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the Immediate Leave field in the output of the `show igmp interface` command.

**Related  
Documentation**

- [Understanding IGMP on page 4759](#)
- [show igmp interface on page 5089](#)

## Filtering Unwanted IGMP Reports at the IGMP Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted IGMP reports at the interface level. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only IGMP group addresses (for IGMPv2) by using the policy's **route-filter** statement to match the group address. You define the policy to match IGMP (source, group) addresses (for IGMPv3) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.



**CAUTION:** On MX Series platforms, IGMPv2 and IGMPv3 cannot be configured together on the same interface. Configuring both together causes unexpected behavior in multicast traffic forwarding.

To filter unwanted IGMP reports:

1. Configure an IGMPv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set then reject
```

2. Configure an IGMPv3 policy.

```
[edit policy-statement reject_policy_v3]
user@host# set from route-filter 224.1.1.1/32 exact
user@host# set from route-filter 239.0.0.0/8 orlonger
user@host# set from source-address-filter 10.0.0.0/8 orlonger
user@host# set from source-address-filter 127.0.0.0/8 orlonger
user@host# set then reject
```

3. Apply the policies to the IGMP interfaces on which you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running IGMPv2, and **ge-0/1/1.0** is running IGMPv3.

```
[edit protocols igmp]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v2
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v3
```

4. Verify the operation of the filter by checking the Rejected Report field in the output of the **show igmp statistics** command.

#### Related Documentation

- [Understanding IGMP on page 4759](#)
- [Example: Configuring Policy Chains and Route Filters](#)
- [show igmp statistics on page 5093](#)

## Accepting IGMP Messages from Remote Subnetworks

By default, IGMP interfaces accept IGMP messages only from the same subnet. Including the **promiscuous-mode** statement enables the routing device to accept IGMP messages from indirectly connected subnets.



**NOTE:** When you enable IGMP on an unnumbered Ethernet interface that uses a /32 loopback address as a donor address, you must configure IGMP promiscuous mode to accept the IGMP packets received on this interface.



**NOTE:** When enabling promiscuous-mode, all routers on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.

To enable IGMP promiscuous mode on an interface:

1. Configure the IGMP interface.

```
[edit protocols igmp]
```

```
user@host# set interface ge-0/1/1.0 promiscuous-mode
```

2. Verify the configuration by checking the Promiscuous Mode field in the output of the **show igmp interface** command.
3. Verify the operation of the filter by checking the Rx non-local field in the output of the **show igmp statistics** command.

**Related  
Documentation**

- [Understanding IGMP on page 4759](#)
- [Configuring the Loopback Interface in the Junos OS Network Interfaces Library for Routing Devices](#)
- [show igmp interface on page 5089](#)
- [show igmp statistics on page 5093](#)

## Modifying the IGMP Query Response Interval

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. Configuring this interval allows you to adjust the burst peaks of IGMP messages on the subnet. Set a larger interval to make the traffic less bursty. Bursty traffic refers to an uneven pattern of data transmission: sometimes a very high data transmission rate, whereas at other times a very low data transmission rate.

The query response interval, the host-query interval, and the robustness variable are related in that they are all variables that are used to calculate the group membership timeout. The group membership timeout is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The group membership timeout is calculated as the (robustness variable x query-interval) + (query-response-interval). If no reports are received for a particular group before the group membership timeout has expired, the routing device stops forwarding remotely originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

```
[edit protocols igmp]
user@host# set query-response-interval 0.4
```

2. Verify the configuration by checking the IGMP Query Response Interval field in the output of the **show igmp interface** command.
3. Verify the operation of the query interval by checking the Membership Query field in the output of the **show igmp statistics** command.

#### Related Documentation

- [Understanding IGMP on page 4759](#)
- [Modifying the IGMP Host-Query Message Interval on page 4830](#)
- [Modifying the IGMP Robustness Variable on page 4835](#)
- [show igmp interface on page 5089](#)
- [show igmp statistics on page 5093](#)

## Modifying the IGMP Robustness Variable

Fine-tune the IGMP robustness variable to allow for expected packet loss on a subnet. The robust count automatically changes certain IGMP message intervals for IGMPv2 and IGMPv3. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

When the query router receives an IGMP leave message on a shared network running IGMPv2, the query router must send an IGMP group query message a specified number of times. The number of IGMP group query messages sent is determined by the robust count.

The value of the robustness variable is also used in calculating the following IGMP message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—The robust count is used to calculate the amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The number of queries is equal to the value of the robustness variable.

In IGMPv3, a change of interface state causes the system to immediately transmit a state-change report from that interface. In case the state-change report is missed by one or more multicast routers, it is retransmitted. The number of times it is retransmitted

is the robust count minus one. In IGMPv3, the robust count is also a factor in determining the group membership interval, the older version querier interval, and the other querier present interval.

By default, the robustness variable is set to 2. You might want to increase this value if you expect a subnet to lose packets.

The number can be from 2 through 10.

To change the value of the robustness variable:

1. Configure the robust count.

When you set the robust count, you are in effect configuring the number of times the querier retries queries on the connected subnets.

```
[edit protocols igmp]
user@host# set robust-count 5
```

2. Verify the configuration by checking the IGMP Robustness Count field in the output of the **show igmp interfaces** command.

#### Related Documentation

- [Modifying the IGMP Host-Query Message Interval on page 4830](#)
- [Modifying the IGMP Query Response Interval on page 4834](#)
- [Modifying the IGMP Last-Member Query Interval on page 4831](#)
- [show pim interfaces on page 5150](#)
- RFC 2236, *Internet Group Management Protocol, Version 2*
- RFC 3376, *Internet Group Management Protocol, Version 3*

## Limiting the Maximum IGMP Message Rate

This section describes how to change the limit for the maximum number of IGMP packets transmitted in 1 second by the router.

Increasing the maximum number of IGMP packets transmitted per second might be useful on a router with a large number of interfaces participating in IGMP.

To change the limit for the maximum number of IGMP packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

#### Related Documentation

- [maximum-transmit-rate \(Protocols IGMP\) on page 4984](#)

## Enabling IGMP Static Group Membership

You can create IGMP static group membership to test multicast forwarding without a receiver host. When you enable IGMP static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts.

The router on which you enable static IGMP group membership must be the designated router (DR) for the subnet. Otherwise, traffic does not flow downstream.

When enabling IGMP static group membership, you cannot configure multiple groups using the **group-count**, **group-increment**, **source-count**, and **source-increment** statements if the **all** option is specified as the IGMP interface.

Class-of-service (CoS) adjustment is not supported with IGMP static group membership.

In this example, you create static group 225.1.1.1.

1. On the DR, configure the static groups to be created by including the **static** statement and **group** statement and specifying which IP multicast address of the group to be created. When creating groups individually, you must specify a unique address for each group.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  static {
    group 225.1.1.1;
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created.

```
user@host> show igmp group
Interface: fe-0/1/2
Group: 225.1.1.1
Source: 10.0.0.2
Last reported by: Local
Timeout: 0 Type: Static
```



**NOTE:** When you configure static IGMP group entries on point-to-point links that connect routing devices to a rendezvous point (RP), the static IGMP group entries do not generate join messages toward the RP.

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. On the DR, configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  static {
    group 225.1.1.1 {
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.2, and 225.1.1.3 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
  Group: 225.1.1.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.2
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.3
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can also configure the group address to be automatically incremented for each group created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and when you do not want the group addresses to be sequential.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. On the DR, configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 group-count 3 group-increment
0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp

interface fe-0/1/2.0 {
  version 3;
```



```
static {
  group 225.1.1.1 {
    group-increment 0.0.0.2;
    group-count 3;
  }
}
```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static groups 225.1.1.1, 225.1.1.3, and 225.1.1.5 have been created.

```
user@host> show igmp group
Interface: fe-0/1/2
  Group: 225.1.1.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.3
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.5
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, and your network is operating in source-specific multicast (SSM) mode, you can also specify that the multicast source address be accepted. This is useful when you want to test forwarding to multicast receivers from a specific multicast source.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you create group 225.1.1.1 and accept IP address 10.0.0.2 as the only source.

1. On the DR, configure the source address by including the **source** statement and specifying the IPv4 address of the source host.

```
[edit protocols igmp]
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2;
    }
  }
}
```

```
    }  
  }
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that source 10.0.0.2 has been accepted.

```
user@host> show igmp group  
Interface: fe-0/1/2  
  Group: 225.1.1.1  
    Source: 10.0.0.2  
    Last reported by: Local  
    Timeout: 0 Type: Static
```

When you create IGMP static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of multicast sources be automatically accepted. This is useful when you want to test forwarding to multicast receivers from more than one specified multicast source.

In this example, you create group 255.1.1.1 and accept addresses 10.0.0.2, 10.0.0.3, and 10.0.0.4 as the sources.

1. On the DR, configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols igmp]  
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count  
3
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp  
  
interface fe-0/1/2.0 {  
  version 3;  
  static {  
    group 225.1.1.1 {  
      source 10.0.0.2 {  
        source-count 3;  
      }  
    }  
  }  
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.3, and 10.0.0.4 have been accepted.

```
user@host> show igmp group  
Interface: fe-0/1/2  
  Group: 225.1.1.1  
    Source: 10.0.0.2  
    Last reported by: Local  
    Timeout: 0 Type: Static  
  Group: 225.1.1.1  
    Source: 10.0.0.3  
    Last reported by: Local
```

```

        Timeout: 0 Type: Static
Group: 225.1.1.1
Source: 10.0.0.4
Last reported by: Local
Timeout: 0 Type: Static

```

When you configure static groups on an interface on which you want to receive multicast traffic, and specify that a number of multicast sources be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately and you do not want the source addresses to be sequential.

In this example, you create group 225.1.1.1 and accept addresses 10.0.0.2, 10.0.0.4, and 10.0.0.6 as the sources.

1. Configure the multicast source address increment by including the **source-increment** statement and specifying the number by which the address should be incremented for each source. The increment is specified in dotted decimal notation similar to an IPv4 address.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 source 10.0.0.2 source-count
3 source-increment 0.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```

interface fe-0/1/2.0 {
  version 3;
  static {
    group 225.1.1.1 {
      source 10.0.0.2 {
        source-count 3;
        source-increment 0.0.0.2;
      }
    }
  }
}

```

3. After you have committed the configuration and after the source is sending traffic, use the **show igmp group** command to verify that static group 225.1.1.1 has been created and that sources 10.0.0.2, 10.0.0.4, and 10.0.0.6 have been accepted.

```

user@host> show igmp group
Interface: fe-0/1/2
  Group: 225.1.1.1
    Source: 10.0.0.2
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.1
    Source: 10.0.0.4
    Last reported by: Local
    Timeout: 0 Type: Static
  Group: 225.1.1.1
    Source: 10.0.0.6

```

Last reported by: Local  
Timeout: 0 Type: Static

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the source address configured. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the source address configured.

If a source address is specified in a multicast group that is statically configured, the IGMP version on the interface must be set to IGMPv3. IGMPv2 is the default value.

In this example, you exclude address 10.0.0.2 as a source for group 225.1.1.1.

1. On the DR, configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv4 source address to exclude.

```
[edit protocols igmp]
```

```
user@host# set interface fe-0/1/2 static group 225.1.1.1 exclude source 10.0.0.2
```

2. After you commit the configuration, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@host> show configuration protocol igmp
```

```
interface fe-0/1/2.0 {  
  version 3;  
  static {  
    group 225.1.1.1 {  
      exclude;  
      source 10.0.0.2;  
    }  
  }  
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show igmp group detail** command to verify that static group 225.1.1.1 has been created and that the static group is operating in exclude mode.

```
user@host> show igmp group detail
```

```
Interface: fe-0/1/2  
Group: 225.1.1.1  
Group mode: Exclude  
Source: 10.0.0.2  
Last reported by: Local  
Timeout: 0 Type: Static
```

#### Related Documentation

- [Enabling MLD Static Group Membership on page 4863](#)
- [group \(Protocols IGMP\) on page 4976](#)
- [group-count \(Protocols IGMP\) on page 4977](#)

- [group-increment \(Protocols IGMP\) on page 4977](#)
- [source-count \(Protocols IGMP\) on page 4991](#)
- [source-increment \(Protocols IGMP\) on page 4991](#)
- [static \(Protocols IGMP\) on page 4992](#)

## Recording IGMP Join and Leave Events

To determine whether IGMP tuning is needed in a network, you can configure the routing device to record IGMP join and leave events. You can record events globally for the routing device or for individual interfaces.

[Table 388 on page 4843](#) describes the recordable IGMP events.

**Table 388: IGMP Event Messages**

ERRMSG Tag	Definition
RPD_IGMP_JOIN	Records IGMP join events.
RPD_IGMP_LEAVE	Records IGMP leave events.
RPD_IGMP_ACCOUNTING_ON	Records when IGMP accounting is enabled on an IGMP interface.
RPD_IGMP_ACCOUNTING_OFF	Records when IGMP accounting is disabled on an IGMP interface.
RPD_IGMP_MEMBERSHIP_TIMEOUT	Records IGMP membership timeout events.

To enable IGMP accounting:

1. Enable accounting globally or on an IGMP interface. This example shows both options.

```
[edit protocols igmp]
user@host# set accounting
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded and filter the events to a system log file with a descriptive filename, such as **igmp-events**.

```
[edit system syslog file igmp-events]
user@host# set any info
user@host# set match ".*RPD_IGMP_JOIN.* | .*RPD_IGMP_LEAVE.* |
.*RPD_IGMP_ACCOUNTING.* | .*RPD_IGMP_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file size when it reaches 100 KB and keeps three files.

```
[edit system syslog file igmp-events]
user@host# set archive size 100000
user@host# set archive files 3
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
```

```
user@host# set archive transfer-interval 24
user@host# set archive start-time 2011-01-07:12:30
```

4. You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start igmp-events

*** igmp-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command
'run monitor start igmp-events '
monitor
```

**Related  
Documentation**

- [Understanding IGMP on page 4759](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 6629](#)

## Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces

The **group-limit** statement enables you to limit the number of IGMP multicast group joins for logical interfaces. When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for IGMP multicast groups, keep the following in mind:

- Each any-source group (\*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in IGMPv3 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (\*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on IGMP logical interfaces using dynamic profiles.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for IGMP multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of IGMP multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs the warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for IGMP multicast group joins.

To limit multicast group joins on an IGMP logical interface:

1. Access the logical interface at the IGMP protocol hierarchy level.

```
[edit]
user@host# edit protocols igmp interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols igmp interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols igmp interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols igmp interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols igmp** command. To verify the operation of IGMP on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show igmp interface** command.

#### Related Documentation

- [Enabling IGMP Static Group Membership on page 4836](#)

## Tracing IGMP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
all	Trace all operations.

Flag	Description
<b>client-notification</b>	Trace notifications.
<b>general</b>	Trace general flow.
<b>group</b>	Trace group operations.
<b>host-notification</b>	Trace host notifications.
<b>leave</b>	Trace leave group messages (IGMPv2 only).
<b>mtrace</b>	Trace mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.
<b>normal</b>	Trace normal events.
<b>packets</b>	Trace all IGMP packets.
<b>policy</b>	Trace policy processing.
<b>query</b>	Trace IGMP membership query messages, including general and group-specific queries.
<b>report</b>	Trace membership report messages.
<b>route</b>	Trace routing information.
<b>state</b>	Trace state transitions.
<b>task</b>	Trace task processing.
<b>timer</b>	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on IGMP packets of a particular type. To configure tracing operations for IGMP:

1. (Optional) Configure tracing at the routing options level to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the IGMP trace file.

```
[edit protocols igmp traceoptions]
user@host# set file igmp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols igmp traceoptions]
user@host# set file files 5
```



4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols igmp traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols igmp traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular multicast group. The following example shows how to flag all events for packets associated with the group IP address.

```
[edit protocols igmp traceoptions]
user@host# set flag group | match 232.1.1.2
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/igmp-trace
```

#### Related Documentation

- [Understanding IGMP on page 4759](#)
- *Tracing and Logging Junos OS Operations* in the *Junos OS Administration Library for Routing Devices*
- [mtrace on page 5072](#) in the [CLI Explorer](#)

## Disabling IGMP

To disable IGMP on an interface, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols igmp interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols igmp interface *interface-name*]**

#### Related Documentation

- [Enabling IGMP on page 4828](#)

## IGMP Snooping

- [Configuring IGMP Snooping on page 4848](#)
- [Configuring VLAN-Specific IGMP Snooping Parameters on page 4849](#)
- [Example: Configuring IGMP Snooping on page 4849](#)
- [Using a Switch as an IGMP Querier on page 4851](#)

## Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).



**NOTE:** You cannot configure IGMP snooping on a secondary VLAN.

To enable IGMP snooping and configure individual options as needed for your network by using the CLI:

1. Enable IGMP snooping on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

2. Configure the switch to immediately remove group membership from interfaces on a VLAN when it receives a leave message through that VLAN, and have it not forward any membership queries for the multicast group to the VLAN (IGMPv2 only):

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name immediate-leave
```

3. Configure an interface to belong to a multicast group:

```
[edit protocols]
user@switch# set igmp-snooping vlan-name interface interface-name static group
group-address
```

4. Configure an interface to forward IGMP queries received from multicast routers.

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name interface interface-name
multicast-router-interface
```

5. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name robust-count 4
```

6. If you want a standalone switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

The switch uses the address that you configure as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.

7. If you want a QFabric Node device to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

- Related Documentation**
- [IGMP Snooping Overview on page 4761](#)
  - [Example: Configuring IGMP Snooping on page 4849](#)
  - [Changing the IGMP Snooping Group Timeout Value](#)
  - [Monitoring IGMP Snooping on page 5047](#)

## Configuring VLAN-Specific IGMP Snooping Parameters

All of the IGMP snooping statements configured with the **igmp-snooping** statement, with the exception of the **traceoptions** statement, can be qualified with the same statement at the VLAN level. To configure IGMP snooping parameters at the VLAN level, include the **vlan** statement:

```
vlan vlan-id;
  immediate-leave;
  interface interface-name {
    group-limit limit;
    host-only-interface;
    multicast-router-interface;
    static {
      group ip-address {
        source ip-address;
      }
    }
  }
  proxy {
    source-address ip-address;
  }
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
}
```

You can include this statement at the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* protocols igmp-snooping]
- [edit routing-instances *routing-instance-name* bridge-domains *bridge-domain-name* protocols igmp-snooping]

- Related Documentation**
- [Multicast Overview](#)
  - [Understanding Multicast Snooping](#)

## Example: Configuring IGMP Snooping

With IGMP snooping enabled, a switch monitors the IGMP (Internet Group Management Protocol) traffic between hosts and multicast routers and uses what it learns to forward multicast traffic to only the downstream interfaces that are connected to interested receivers. This conserves bandwidth by allowing the switch to send multicast traffic to

only those interfaces that are connected to devices that want to receive the traffic (instead of flooding the traffic to all the downstream VLAN interfaces).

This example describes how to configure IGMP snooping:

- [Requirements on page 4850](#)
- [Overview and Topology on page 4850](#)
- [Configuration on page 4850](#)

## Requirements

This example requires Junos OS Release 11.1 or later on a QFX Series product.

Before you configure IGMP snooping, be sure you have:

- Configured the **employee-vlan** VLAN
- Assigned interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3** to **employee-vlan**

## Overview and Topology

In this example you configure an interface to receive multicast traffic from a source and configure some multicast-related behavior for downstream interfaces. The example assumes that IGMP snooping was previously disabled for the VLAN.

[Table 389 on page 4850](#) shows the components of the topology for this example.

**Table 389: Components of the IGMP Snooping Topology**

Components	Settings
VLAN name	<b>employee-vlan</b> , tag 20
Interfaces in <b>employee-vlan</b>	<b>ge-0/0/1</b> , <b>ge-0/0/2</b> , <b>ge-0/0/3</b>
Multicast IP address for <b>employee-vlan</b>	225.100.100.100

## Configuration

To configure basic IGMP snooping on a switch:

### CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into a terminal window:

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

### Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN **employee-vlan**:  

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```
2. Configure a interface to belong to a multicast group:  

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
```
3. Configure an interface to forward IGMP queries received from multicast routers.  

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
```
4. Configure the switch to wait for four timeout intervals before timing out a multicast group on a VLAN:  

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

**Results** Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
vlan employee-vlan {
  robust-count 4;
}
interface ge-0/0/2 {
  multicast-router-interface;
}
interface ge-0/0/3 {
  static {
    group 255.100.100.100;
  }
}
```

### Related Documentation

- [IGMP Snooping Overview on page 4761](#)
- [Configuring IGMP Snooping on page 4848](#)
- [Changing the IGMP Snooping Group Timeout Value](#)
- [Monitoring IGMP Snooping on page 5047](#)
- [Example: Setting Up Bridging with Multiple VLANs.](#)

## Using a Switch as an IGMP Querier

If IGMP snooping is enabled on a pure Layer 2 a local network (that is, Layer 3 is not enabled on the network), and there is not multicast router in the network, multicast traffic might not be properly forwarded through the network. This problem occurs if the local network is configured such that multicast traffic must be forwarded between switches in order to reach a multicast receiver. In this case, an upstream switch does not forward multicast traffic to a downstream switch (and therefore to the multicast receivers attached to the downstream switch) because the downstream switch does not forward

IGMP reports to the upstream switch. You can solve this problem by configuring one of the switches to be an IGMP querier. This switch sends periodic general query packets to all the switches in the network, which ensures that the snooping membership tables are updated and prevents any multicast traffic loss.

If you configure multiple switches to be IGMP queriers, the switch with the highest (greatest) IGMP querier source address takes precedence and acts as the querier. Switches with lower IGMP querier source addresses stop sending IGMP queries unless they do not receive IGMP queries for 255 seconds. If a switch with a lower IGMP querier source address does not receive any IGMP queries during that period, it starts sending queries again.

To configure a standalone switch to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name l2-querier source-address source address
```

To configure a QFabric Node device to act as an IGMP querier, enter the following:

```
[edit protocols]
user@switch# set igmp-snooping vlan vlan-name igmp-querier source-address source address
```

#### Related Documentation

- [IGMP Snooping Overview on page 4761](#)
- [Example: Configuring IGMP Snooping on page 4849](#)
- [Configuring IGMP Snooping on page 4848](#)
- [Changing the IGMP Snooping Group Timeout Value](#)
- [Monitoring IGMP Snooping on page 5047](#)

## MLD

---

- [Examples: Configuring MLD on page 4852](#)

### Examples: Configuring MLD

- [Understanding MLD on page 4853](#)
- [Configuring MLD on page 4855](#)
- [Enabling MLD on page 4856](#)
- [Modifying the MLD Version on page 4857](#)
- [Modifying the MLD Host-Query Message Interval on page 4858](#)
- [Modifying the MLD Query Response Interval on page 4858](#)
- [Modifying the MLD Last-Member Query Interval on page 4859](#)
- [Specifying Immediate-Leave Host Removal for MLD on page 4860](#)
- [Filtering Unwanted MLD Reports at the MLD Interface Level on page 4861](#)
- [Example: Modifying the MLD Robustness Variable on page 4861](#)
- [Limiting the Maximum MLD Message Rate on page 4863](#)
- [Enabling MLD Static Group Membership on page 4863](#)

- [Example: Recording MLD Join and Leave Events on page 4870](#)
- [Configuring the Number of MLD Multicast Group Joins on Logical Interfaces on page 4872](#)
- [Tracing MLD Protocol Traffic on page 4873](#)
- [Disabling MLD on page 4875](#)

### Understanding MLD

The Multicast Listener Discovery (MLD) Protocol manages the membership of hosts and routers in multicast groups. IP version 6 (IPv6) multicast routers use MLD to learn, for each of their attached physical networks, which groups have interested listeners. Each routing device maintains a list of host multicast addresses that have listeners for each subnetwork, as well as a timer for each address. However, the routing device does not need to know the address of each listener—just the address of each host. The routing device provides addresses to the multicast routing protocol it uses, which ensures that multicast packets are delivered to all subnetworks where there are interested listeners. In this way, MLD is used as the transport for the Protocol Independent Multicast (PIM) Protocol.


MLD is an integral part of IPv6 and must be enabled on all IPv6 routing devices and hosts that need to receive IP multicast traffic. The Junos OS supports MLD versions 1 and 2. Version 2 is supported for source-specific multicast (SSM) include and exclude modes.

In include mode, the receiver specifies the source or sources it is interested in receiving the multicast group traffic from. Exclude mode works the opposite of include mode. It allows the receiver to specify the source or sources it is not interested in receiving the multicast group traffic from.

For each attached network, a multicast routing device can be either a querier or a nonquerier. A querier routing device, usually one per subnet, solicits group membership information by transmitting MLD queries. When a host reports to the querier routing device that it has interested listeners, the querier routing device forwards the membership information to the rendezvous point (RP) routing device by means of the receiver's (host's) designated router (DR). This builds the rendezvous-point tree (RPT) connecting the host with interested listeners to the RP routing device. The RPT is the initial path used by the sender to transmit information to the interested listeners. Nonquerier routing devices do not transmit MLD queries on a subnet but can do so if the querier routing device fails.

All MLD-configured routing devices start as querier routing devices on each attached subnet (see [Figure 156 on page 4765](#)). The querier routing device on the right is the receiver's DR.

**Figure 167: Routing Devices Start Up on a Subnet**




To elect the querier routing device, the routing devices exchange query messages containing their IPv6 source addresses. If a routing device hears a query message whose IPv6 source address is numerically lower than its own selected address, it becomes a nonquerier. In [Figure 157 on page 4766](#), the routing device on the left has a source address numerically lower than the one on the right and therefore becomes the querier routing device.




**NOTE:** In the practical application of MLD, several routing devices on a subnet are nonqueriers. If the elected querier routing device fails, query messages are exchanged among the remaining routing devices. The routing device with the lowest IPv6 source address becomes the new querier routing device. The IPv6 Neighbor Discovery Protocol (NDP) implementation drops incoming Neighbor Announcement (NA) messages that have a broadcast or multicast address in the target link-layer address option. This behavior is recommended by RFC 2461.

**Figure 168: Querier Routing Device Is Determined**



The querier routing device sends general MLD queries on the **link-scope all-nodes** multicast address `FF02::1` at short intervals to all attached subnets to solicit group membership information (see [Figure 158 on page 4766](#)). Within the query message is the *maximum response delay* value, specifying the maximum allowed delay for the host to respond with a report message.


**Figure 169: General Query Message Is Issued**



If interested listeners are attached to the host receiving the query, the host sends a report containing the host's IPv6 address to the routing device (see [Figure 159 on page 4766](#)). If the reported address is not yet in the routing device's list of multicast addresses with interested listeners, the address is added to the list and a timer is set for the address. If the address is already on the list, the timer is reset. The host's address is transmitted to the RP in the PIM domain.




Figure 170: Reports Are Received by the Querier Routing Device




If the host has no interested multicast listeners, it sends a done message to the querier routing device. On receipt, the querier routing device issues a multicast address-specific query containing the last **listener query interval** value to the multicast address of the host. If the routing device does not receive a report from the multicast address, it removes the multicast address from the list and notifies the RP in the PIM domain of its removal (see [Figure 160 on page 4767](#)).

Figure 171: Host Has No Interested Receivers and Sends a Done Message to Routing Device



If a done message is not received by the querier routing device, the querier routing device continues to send multicast address-specific queries. If the timer set for the address on receipt of the last report expires, the querier routing device assumes there are no longer interested listeners on that subnet, removes the multicast address from the list, and notifies the RP in the PIM domain of its removal (see [Figure 161 on page 4767](#)).

Figure 172: Host Address Timer Expires and Address Is Removed from Multicast Address List



### Configuring MLD

To configure the Multicast Listener Discovery (MLD) Protocol, include the **mld** statement:

```
mld {
  accounting;
  interface interface-name {
    disable;
    (accounting | no-accounting);
    group-policy [ policy-names ];
    immediate-leave;
  }
}
```

```
oif-map [ map-names ];
passive;
ssm-map ssm-map-name;
static {
    group multicast-group-address {
        exclude;
        group-count number;
        group-increment increment;
        source ip-address {
            source-count number;
            source-increment increment;
        }
    }
}
version version;
}
maximum-transmit-rate packets-per-second;
query-interval seconds;
query-last-member-interval seconds;
query-response-interval seconds;
robust-count number;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

By default, MLD is enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or the Distance Vector Multicast Routing Protocol (DVMRP).

---

### Enabling MLD

The Multicast Listener Discovery (MLD) Protocol manages multicast groups by establishing, maintaining, and removing groups on a subnet. Multicast routing devices use MLD to learn which groups have members on each of their attached physical networks. MLD must be enabled for the router to receive IPv6 multicast packets. MLD is only needed for IPv6 networks, because multicast is handled differently in IPv4 networks. MLD is enabled on all IPv6 interfaces on which you configure PIM and on all IPv6 broadcast interfaces when you configure DVMRP.

MLD specifies different behaviors for multicast listeners and for routers. When a router is also a listener, the router responds to its own messages. If a router has more than one interface to the same link, it needs to perform the router behavior over only one of those interfaces. Listeners, on the other hand, must perform the listener behavior on all interfaces connected to potential receivers of multicast traffic.

If MLD is not running on an interface—either because PIM and DVMRP are not configured on the interface or because MLD is explicitly disabled on the interface—you can explicitly enable MLD.

To explicitly enable MLD:

1. If PIM and DVMRP are not running on the interface, explicitly enable MLD by including the interface name.

```
[edit protocols mld]
user@host# set interface fe-0/0/0.0
```

2. Check to see if MLD is disabled on any interfaces. In the following example, MLD is disabled on a Gigabit Ethernet interface.

```
[edit protocols mld]
user@host# show

interface fe-0/0/0.0;
interface ge-0/0/0.0 {
    disable;
}
```

3. Enable MLD on the interface by deleting the **disable** statement.

```
[edit protocols mld]
delete interface ge-0/0/0.0 disable
```

4. Verify the configuration.

```
[edit protocols mld]
user@host# show

interface fe-0/0/0.0;
interface ge-0/0/0.0;
```

5. Verify the operation of MLD by checking the output of the **show mld interface** command.

### Modifying the MLD Version

By default, the router supports MLD version 1 (MLDv1). To enable the router to use MLD version 2 (MLDv2) for source-specific multicast (SSM) only, include the **version 2** statement.

If you configure the MLD version setting at the individual interface hierarchy level, it overrides configuring the IGMP version using the **interface all** statement.

If a source address is specified in a multicast group that is statically configured, the version must be set to MLDv2.

To change an MLD interface to version 2:

1. Configure the MLD interface.

```
[edit protocols mld]
user@host# set interface fe-0/0/0.0 version 2
```

2. Verify the configuration by checking the **version** field in the output of the **show mld interface** command. The **show mld statistics** command has version-specific output fields, such as the counters in the **MLD Message type** field.

### Modifying the MLD Host-Query Message Interval

---

The objective of MLD is to keep routers up to date with IPv6 group membership of the entire subnet. Routers need not know who all the members are, only that members exist. Each host keeps track of which multicast groups are subscribed to. On each link, one router is elected the querier. The MLD querier router periodically sends general host-query messages on each attached network to solicit membership information. These messages solicit group membership information and are sent to the **link-scope all-nodes** address **FF02::1**. A general host-query message has a maximum response time that you can set by configuring the query response interval.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

By default, host-query messages are sent every 125 seconds. You can change this interval to change the number of MLD messages sent on the subnet.

To modify the query interval:

1. Configure the interval.

```
[edit protocols mld]  
user@host# set query-interval 200
```

The value can be from 1 through 1024 seconds.

2. Verify the configuration by checking the **MLD Query Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

### Modifying the MLD Query Response Interval

---

The query response interval is the maximum amount of time that can elapse between when the querier router sends a host-query message and when it receives a response from a host. You can change this interval to adjust the burst peaks of MLD messages on the subnet. Set a larger interval to make the traffic less bursty.

The query response timeout, the query interval, and the robustness variable are related in that they are all variables that are used to calculate the multicast listener interval. The multicast listener interval is the number of seconds that must pass before a multicast router determines that no more members of a host group exist on a subnet. The multicast listener interval is calculated as the (robustness variable x query-interval) + (1 x query-response-interval). If no reports are received for a particular group before the

multicast listener interval has expired, the routing device stops forwarding remotely-originated multicast packets for that group onto the attached network.

The default query response interval is 10 seconds. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify the query response interval:

1. Configure the interval.

```
[edit protocols mld]
user@host# set query-response-interval 0.5
```

2. Verify the configuration by checking the **MLD Query Response Interval** field in the output of the **show mld interface** command.
3. Verify the operation of the query interval by checking the **Listener Query** field in the output of the **show mld statistics** command.

### Modifying the MLD Last-Member Query Interval

The last-member query interval (also called the last-listener query interval) is the maximum amount of time between group-specific query messages, including those sent in response to done messages sent on the **link-scope-all-routers** address FF02::2. You can lower this interval to reduce the amount of time it takes a router to detect the loss of the last member of a group.

When the routing device that is serving as the querier receives a leave-group (done) message from a host, the routing device sends multiple group-specific queries to the group. The querier sends a specific number of these queries, and it sends them at a specific interval. The number of queries sent is called the last-listener query count. The interval at which the queries are sent is called the last-listener query interval. Both settings are configurable, thus allowing you to adjust the leave latency. The IGMP leave latency is the time between a request to leave a multicast group and the receipt of the last byte of data for the multicast group.

The last-listener query count x (times) the last-listener query interval = (equals) the amount of time it takes a routing device to determine that the last member of a group has left the group and to stop forwarding group traffic.

The default last-listener query interval is 1 second. You can configure a subsecond interval up to one digit to the right of the decimal point. The configurable range is 0.1 through 0.9, then in 1-second intervals 1 through 999,999.

To modify this interval:

1. Configure the time (in seconds) that the routing device waits for a report in response to a group-specific query.

```
[edit protocols mld]
user@host# set query-last-member-interval 0.1
```

2. Verify the configuration by checking the **MLD Last Member Query Interval** field in the output of the **show igmp interfaces** command.



**NOTE:** You can configure the last-member query count by configuring the robustness variable. The two are always equal.

---

### Specifying Immediate-Leave Host Removal for MLD

---

The immediate leave setting is useful for minimizing the leave latency of MLD memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.

The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows MLD to determine when the last host sends a leave message for the multicast group.

When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending MLD group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the MLD leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.

When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both MLD version 1 and MLD version 2.



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the router only knows about the one interested host and does not know about the others.

To enable immediate leave:

1. Configure immediate leave on the MLD interface.

```
[edit protocols mld]  
user@host# set interface ge-0/0/0.1 immediate-leave
```

2. Verify the configuration by checking the **Immediate Leave** field in the output of the **show mld interface** command.

### Filtering Unwanted MLD Reports at the MLD Interface Level

Suppose you need to limit the subnets that can join a certain multicast group. The **group-policy** statement enables you to filter unwanted MLD reports at the interface level.

When the **group-policy** statement is enabled on a router, after the router receives an MLD report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report if the policy matches the defined address or network).

You define the policy to match only MLD group addresses (for MLDv1) by using the policy's **route-filter** statement to match the group address. You define the policy to match MLD (source, group) addresses (for MLDv2) by using the policy's **route-filter** statement to match the group address and the policy's **source-address-filter** statement to match the source address.

To filter unwanted MLD reports:

1. Configure an MLDv1 policy.

```
[edit policy-statement reject_policy_v1]
user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set then reject
```

2. Configure an MLDv2 policy.

```
[edit policy-statement reject_policy_v2]
user@host# set from route-filter fec0:1:1:4::/64 exact
user@host# set from source-address-filter fe80::2e0:81ff:fe05:1a8d/32 orlonger
user@host# set then reject
```

3. Apply the policies to the MLD interfaces where you prefer not to receive specific group or (source, group) reports. In this example, **ge-0/0/0.1** is running MLDv1 and **ge-0/1/1.0** is running MLDv2.

```
[edit protocols mld]
user@host# set interface ge-0/0/0.1 group-policy reject_policy_v1
user@host# set interface ge-0/1/1.0 group-policy reject_policy_v2
```

4. Verify the operation of the filter by checking the **Rejected Report** field in the output of the **show mld statistics** command.

### Example: Modifying the MLD Robustness Variable

This example shows how to configure and verify the MLD robustness variable in a multicast domain.

- [Requirements on page 4861](#)
- [Overview on page 4862](#)
- [Configuration on page 4862](#)
- [Verification on page 4863](#)

#### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM. See [“PIM Overview” on page 4737](#).

### Overview

The MLD robustness variable can be fine-tuned to allow for expected packet loss on a subnet. Increasing the robust count allows for more packet loss but increases the leave latency of the subnetwork.

The value of the robustness variable is used in calculating the following MLD message intervals:

- Group member interval—Amount of time that must pass before a multicast router determines that there are no more members of a group on a network. This interval is calculated as follows: (robustness variable x query-interval) + (1 x query-response-interval).
- Other querier present interval—Amount of time that must pass before a multicast router determines that there is no longer another multicast router that is the querier. This interval is calculated as follows: (robustness variable x query-interval) + (0.5 x query-response-interval).
- Last-member query count—Number of group-specific queries sent before the router assumes there are no local members of a group. The default number is the value of the robustness variable.

By default, the robustness variable is set to 2. The number can be from 2 through 10. You might want to increase this value if you expect a subnet to lose packets.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols mld robust-count 5
```

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To change the value of the robustness variable:

1. Configure the robust count.

```
[edit protocols mld]
```



```
user@host# set robust-count 5
```

- If you are done configuring the device, commit the configuration.

```
[edit protocols mld]
user@host# commit
```

### Verification

To verify the configuration is working properly, check the **MLD Robustness Count** field in the output of the **show mld interfaces** command.

### Limiting the Maximum MLD Message Rate

You can change the limit for the maximum number of MLD packets transmitted in 1 second by the router.

Increasing the maximum number of MLD packets transmitted per second might be useful on a router with a large number of interfaces participating in MLD.

To change the limit for the maximum number of MLD packets the router can transmit in 1 second, include the **maximum-transmit-rate** statement and specify the maximum number of packets per second to be transmitted.

### Enabling MLD Static Group Membership

You can create MLD static group membership to test multicast forwarding without a receiver host. When you enable MLD static group membership, data is forwarded to an interface without that interface receiving membership reports from downstream hosts.

Class-of-service (CoS) adjustment is not supported with MLD static group membership.

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify the number of static groups to be automatically created.

In this example, you create static group ff0e::1:ff05:1a8d.

- Configure the static groups to be created by including the **static** statement and **group** statement and specifying which IPv6 multicast address of the group to be created.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d
```

- After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d;
  }
}
```

- After you have committed the configuration and after the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created.

```
user@host> show mld group
Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```



**NOTE:** You must specify a unique address for each group.

### Automatically create static groups

When you create MLD static group membership to test multicast forwarding on an interface on which you want to receive multicast traffic, you can specify that a number of static groups be automatically created. This is useful when you want to test forwarding to multiple receivers without having to configure each receiver separately.

In this example, you create three groups.

1. Configure the number of static groups to be created by including the **group-count** statement and specifying the number of groups to be created.

```
[edit protocols mld]
```

```
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d group-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld
```

```
interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff0e::1:ff05:1a8d, ff0e::1:ff05:1a8e, and ff0e::1:ff05:1a8f have been created.

```
user@host> show mld group
```

```
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8e
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8f
```

```
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```

### Automatically increment group addresses

When you configure static groups on an interface on which you want to receive multicast traffic and you specify the number of static groups to be automatically created, you can also configure the group address to be automatically incremented by some number of addresses.

In this example, you create three groups and increase the group address by an increment of two for each group.

1. Configure the group address increment by including the **group-increment** statement and specifying the number by which the address should be incremented for each group. The increment is specified in a format similar to an IPv6 address.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d group-count 3
group-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      group-increment ::2;
      group-count 3;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static groups ff0e::1:ff05:1a8d, ff0e::1:ff05:1a8f, and ff0e::1:ff05:1a91 have been created.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8f
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a91
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
```

### Specify multicast source address (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify the multicast source address to be accepted.

If you specify a group address in the SSM range, you must also specify a source.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you create group ff0e::1:ff05:1a8d and accept IPv6 address fe80::2e0:81ff:fe05:1a8d as the only source.

1. Configure the source address by including the **source** statement and specifying the IPv6 address of the source host.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d source
fe80::2e0:81ff:fe05:1a8d
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d;
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that source fe80::2e0:81ff:fe05:1a8d has been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
Group: ff0e::1:ff05:1a8d
Source: fe80::2e0:81ff:fe05:1a8d
Last reported by: Local
Timeout: 0 Type: Static
```

### Automatically specify multicast sources

When you configure static groups on an interface on which you want to receive multicast traffic, you can specify a number of multicast sources to be automatically accepted.

In this example, you create static group ff0e::1:ff05:1a8d and accept fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f as the source addresses.

1. Configure the number of multicast source addresses to be accepted by including the **source-count** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d source
fe80::2e0:81ff:fe05:1a8d source-count 3
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group ff0e::1:ff05:1a8d has been created and that sources fe80::2e0:81ff:fe05:1a8d, fe80::2e0:81ff:fe05:1a8e, and fe80::2e0:81ff:fe05:1a8f have been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8e
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
```

**Automatically  
increment source  
addresses**

When you configure static groups on an interface on which you want to receive multicast traffic, and specify a number of multicast sources to be automatically accepted, you can also specify the number by which the address should be incremented for each source accepted.

In this example, you create static group `ff0e::1:ff05:1a8d` and accept `fe80::2e0:81ff:fe05:1a8d`, `fe80::2e0:81ff:fe05:1a8f`, and `fe80::2e0:81ff:fe05:1a91` as the sources.

1. Configure the number of multicast source addresses to be accepted by including the **source-increment** statement and specifying the number of sources to be accepted.

```
[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d source
fe80::2e0:81ff:fe05:1a8d source-count 3 source-increment ::2
```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```
user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      source fe80::2e0:81ff:fe05:1a8d {
        source-count 3;
        source-increment ::2;
      }
    }
  }
}
```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group** command to verify that static group `ff0e::1:ff05:1a8d` has been created and that sources `fe80::2e0:81ff:fe05:1a8d`, `fe80::2e0:81ff:fe05:1a8f`, and `fe80::2e0:81ff:fe05:1a91` have been accepted.

```
user@host> show mld group

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a8f
    Last reported by: Local
    Timeout: 0 Type: Static
Interface: fe-0/1/2
  Group: ff0e2::1:ff05:1a8d
    Source: fe80::2e0:81ff:fe05:1a91
    Last reported by: Local
    Timeout: 0 Type: Static

Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
  Group mode: Include
  Source: fe80::2e0:81ff:fe05:1a8d
```

```

Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a8f
Last reported by: Local
Timeout: 0 Type: Static
Group: ff0e::1:ff05:1a8d
Group mode: Include
Source: fe80::2e0:81ff:fe05:1a91
Last reported by: Local
Timeout: 0 Type: Static

```

### Exclude multicast source addresses (in SSM mode)

When you configure static groups on an interface on which you want to receive multicast traffic and your network is operating in source-specific multicast (SSM) mode, you can specify that certain multicast source addresses be excluded.

By default the multicast source address configured in a static group operates in include mode. In include mode the multicast traffic for the group is accepted from the configured source address. You can also configure the static group to operate in exclude mode. In exclude mode the multicast traffic for the group is accepted from any address other than the configured source address.

If a source address is specified in a multicast group that is statically configured, the MLD version must be set to MLDv2 on the interface. MLDv1 is the default value.

In this example, you exclude address fe80::2e0:81ff:fe05:1a8d as a source for group ff0e::1:ff05:1a8d.

1. Configure a multicast static group to operate in exclude mode by including the **exclude** statement and specifying which IPv6 source address to be excluded.

```

[edit protocols mld]
user@host# set interface fe-0/1/2 static group ff0e::1:ff05:1a8d exclude source
fe80::2e0:81ff:fe05:1a8d

```

2. After you commit the configuration, use the **show configuration protocol mld** command to verify the MLD protocol configuration.

```

user@host> show configuration protocol mld

interface fe-0/1/2.0 {
  static {
    group ff0e::1:ff05:1a8d {
      exclude;
      source fe80::2e0:81ff:fe05:1a8d;
    }
  }
}

```

3. After you have committed the configuration and the source is sending traffic, use the **show mld group detail** command to verify that static group ff0e::1:ff05:1a8d has been created and that the static group is operating in exclude mode.

```

user@host> show mld group detail
Interface: fe-0/1/2
  Group: ff0e::1:ff05:1a8d
    Group mode: Exclude
    Source: fe80::2e0:81ff:fe05:1a8d
    Last reported by: Local
    Timeout: 0 Type: Static

```

Similar configuration is available for IPv4 multicast traffic using the IGMP protocol.

### Example: Recording MLD Join and Leave Events

This example shows how to determine whether MLD tuning is needed in a network by configuring the routing device to record MLD join and leave events.

- [Requirements on page 4870](#)
- [Overview on page 4870](#)
- [Configuration on page 4871](#)
- [Verification on page 4872](#)

#### Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable IPv6 unicast routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM. See [“PIM Overview” on page 4737](#).

#### Overview

[Table 390 on page 4870](#) describes the recordable MLD join and leave events.

**Table 390: MLD Event Messages**

ERRMSG Tag	Definition
RPD_MLD_JOIN	Records MLD join events.
RPD_MLD_LEAVE	Records MLD leave events.
RPD_MLD_ACCOUNTING_ON	Records when MLD accounting is enabled on an MLD interface.
RPD_MLD_ACCOUNTING_OFF	Records when MLD accounting is disabled on an MLD interface.
RPD_MLD_MEMBERSHIP_TIMEOUT	Records MLD membership timeout events.



### Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols mld interface fe-0/1/0.2 accounting
set system syslog file mld-events any info
set system syslog file mld-events match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
.*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
set system syslog file mld-events archive size 100000
set system syslog file mld-events archive files 3
set system syslog file mld-events archive transfer-interval 1440
set system syslog file mld-events archive archive-sites "ftp://user@host1//var/tmp"
password "anonymous"
set system syslog file mld-events archive archive-sites "ftp://user@host2//var/tmp"
password "test"
```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure recording of MLD join and leave events:

1. Enable accounting globally or on an MLD interface. This example shows the interface configuration.

```
[edit protocols mld]
user@host# set interface fe-0/1/0.2 accounting
```

2. Configure the events to be recorded, and filter the events to a system log file with a descriptive filename, such as **mld-events**.

```
[edit system syslog file mld-events]
user@host# set any info
[edit system syslog file mld-events]
user@host# set match ".*RPD_MLD_JOIN.* | .*RPD_MLD_LEAVE.* |
.*RPD_MLD_ACCOUNTING.* | .*RPD_MLD_MEMBERSHIP_TIMEOUT.*"
```

3. Periodically archive the log file.

This example rotates the file every 24 hours (1440 minutes) when it reaches 100 KB and keeps three files.

```
[edit system syslog file mld-events]
user@host# set archive size 100000
[edit system syslog file mld-events]
user@host# set archive files 3
[edit system syslog file mld-events]
user@host# set archive archive-sites "ftp://user@host1//var/tmp" password
"anonymous"
[edit system syslog file mld-events]
user@host# set archive archive-sites "ftp://user@host2//var/tmp" password "test"
[edit system syslog file mld-events]
user@host# set archive transfer-interval 1440
```

```
[edit system syslog file mld-events]
user@host# set archive start-time 2011-01-07:12:30
```

4. If you are done configuring the device, commit the configuration.

```
[edit system syslog file mld-events]]
user@host# commit
```

### Verification

You can view the system log file by running the **file show** command.

```
user@host> file show mld-events
```

You can monitor the system log file as entries are added to the file by running the **monitor start** and **monitor stop** commands.

```
user@host> monitor start mld-events
```

```
*** mld-events ***
Apr 16 13:08:23 host mgd[16416]: UI_CMDLINE_READ_LINE: User 'user', command 'run
monitor start mld-events '
monitor
```

### Configuring the Number of MLD Multicast Group Joins on Logical Interfaces

---

The **group-limit** statement enables you to limit the number of MLD multicast group joins for logical interfaces. When this statement is enabled on a router running MLD version 2, the limit is applied upon receipt of the group report. Once the group limit is reached, subsequent join requests are rejected.

When configuring limits for MLD multicast groups, keep the following in mind:

- Each any-source group (\*G) counts as one group toward the limit.
- Each source-specific group (S,G) counts as one group toward the limit.
- Groups in MLDv2 exclude mode are counted toward the limit.
- Multiple source-specific groups count individually toward the group limit, even if they are for the same group. For example, (S1, G1) and (S2, G1) would count as two groups toward the configured limit.
- Combinations of any-source groups and source-specific groups count individually toward the group limit, even if they are for the same group. For example, (\*, G1) and (S, G1) would count as two groups toward the configured limit.
- Configuring and committing a group limit on a network that is lower than what already exists on the network results in the removal of all groups from the configuration. The groups must then request to rejoin the network (up to the newly configured group limit).
- You can dynamically limit multicast groups on MLD logical interfaces by using dynamic profiles. For detailed information about creating dynamic profiles, see the *Junos OS Subscriber Management and Services Library*.

Beginning with Junos OS 12.2, you can optionally configure a system log warning threshold for MLD multicast group joins received on the logical interface. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of MLD multicast group joins have been received on the interface. These log messages convey when the configured group limit has been exceeded, when the configured threshold has been exceeded, and when the number of groups drop below the configured threshold.

The **group-threshold** statement enables you to configure the threshold at which a warning message is logged. The range is 1 through 100 percent. The warning threshold is a percentage of the group limit, so you must configure the **group-limit** statement to configure a warning threshold. For instance, when the number of groups exceed the configured warning threshold, but remain below the configured group limit, multicast groups continue to be accepted, and the device logs a warning message. In addition, the device logs a warning message after the number of groups drop below the configured warning threshold. You can further specify the amount of time (in seconds) between the log messages by configuring the **log-interval** statement. The range is 6 through 32,767 seconds.

You might consider throttling log messages because every entry added after the configured threshold and every entry rejected after the configured limit causes a warning message to be logged. By configuring a log interval, you can throttle the amount of system log warning messages generated for MLD multicast group joins.

To limit multicast group joins on an MLD logical interface:

1. Access the logical interface at the MLD protocol hierarchy level.

```
[edit]
user@host# edit protocols mld interface interface-name
```

2. Specify the group limit for the interface.

```
[edit protocols mld interface interface-name]
user@host# set group-limit limit
```

3. (Optional) Configure the threshold at which a warning message is logged.

```
[edit protocols mld interface interface-name]
user@host# set group-threshold value
```

4. (Optional) Configure the amount of time between log messages.

```
[edit protocols mld interface interface-name]
user@host# set log-interval seconds
```

To confirm your configuration, use the **show protocols mld** command. To verify the operation of MLD on the interface, including the configured group limit and the optional warning threshold and interval between log messages, use the **show mld interface** command.

### Tracing MLD Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy

actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
<b>all</b>	Trace all operations.
<b>client-notification</b>	Trace notifications.
<b>general</b>	Trace general flow.
<b>group</b>	Trace group operations.
<b>host-notification</b>	Trace host notifications.
<b>leave</b>	Trace leave group messages.
<b>mtrace</b>	Trace mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.
<b>normal</b>	Trace normal events.
<b>packets</b>	Trace all MLD packets.
<b>policy</b>	Trace policy processing.
<b>query</b>	Trace MLD membership query messages, including general and group-specific queries.
<b>report</b>	Trace membership report messages.
<b>route</b>	Trace routing information.
<b>state</b>	Trace state transitions.
<b>task</b>	Trace task processing.
<b>timer</b>	Trace timer processing.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MLD packets of a particular type. To configure tracing operations for MLD:

- (Optional) Configure tracing at the routing options level to trace all protocol packets.
 

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```
- Configure the filename for the MLD trace file.
 

```
[edit protocols mld traceoptions]
```

```
user@host# set file mld-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols mld traceoptions]
```

```
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols mld traceoptions]
```

```
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols mld traceoptions]
```

```
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with a particular interface. The following example shows how to flag all events for packets associated with the interface name.

```
[edit protocols mld traceoptions]
```

```
user@host# set flag all | match fe-1/0/1.0
```

7. View the trace file.

```
user@host> file list /var/log
```

```
user@host> file show /var/log/mld-trace
```

### Disabling MLD

To disable MLD on an interface, include the **disable** statement:

```
interface interface-name {
  disable;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mld]
- [edit logical-systems *logical-system-name* protocols mld]

#### Related Documentation

- [Configuring IGMP](#)

### MSDP

- [Configuring MSDP on page 4876](#)
- [Tracing MSDP Protocol Traffic on page 4877](#)
- [Configuring the Interface to Accept Traffic from a Remote Source on page 4879](#)
- [Example: Configuring MSDP on page 4879](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880](#)
- [Example: Configuring PIM Anycast With or Without MSDP on page 4886](#)
- [Configuring a PIM Anycast RP Router with MSDP on page 4890](#)

## Configuring MSDP

To configure the Multicast Source Discovery Protocol (MSDP), include the **msdp** statement:

```
msdp {
  disable;
  active-source-limit {
    maximum number;
    threshold number;
  }
  data-encapsulation (disable | enable);
  export [ policy-names ];
  group group-name {
    ... group-configuration ...
  }
  hold-time seconds;
  import [ policy-names ];
  local-address address;
  keep-alive seconds;
  peer address {
    ... peer-configuration ...
  }
  rib-group group-name;
  source ip-prefix </prefix-length> {
    active-source-limit {
      maximum number;
      threshold number;
    }
  }
  sa-hold-time seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  group group-name {
    disable;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    mode (mesh-group | standard);
    peer address {
      ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
      just following ...
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
  peer address {
    disable;
    active-source-limit {
      maximum number;
      threshold number;
    }
  }
}
```

```

    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols]
- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

By default, MSDP is disabled.

#### Related Documentation

- [Example: Configuring MSDP in a Routing Instance](#)
- [Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880](#)

## Tracing MSDP Protocol Traffic

Tracing operations record detailed messages about the operation of routing protocols, such as the various types of routing protocol packets sent and received, and routing policy actions. You can specify which trace operations are logged by including specific tracing flags. The following table describes the flags that you can include.

Flag	Description
<b>all</b>	Trace all operations.
<b>general</b>	Trace general events.
<b>keepalive</b>	Trace keepalive messages.
<b>normal</b>	Trace normal events.
<b>packets</b>	Trace all MSDP packets.
<b>policy</b>	Trace policy processing.
<b>route</b>	Trace MSDP changes to the routing table.
<b>source-active</b>	Trace source-active packets.

Flag	Description
<b>source-active-request</b>	Trace source-active request packets.
<b>source-active-response</b>	Trace source-active response packets.
<b>state</b>	Trace state transitions.
<b>task</b>	Trace task processing.
<b>timer</b>	Trace timer processing.

You can configure MSDP tracing for all peers, for all peers in a particular group, or for a particular peer.

In the following example, tracing is enabled for all routing protocol packets. Then tracing is narrowed to focus only on MSDP peers in a particular group. To configure tracing operations for MSDP:

1. (Optional) Configure tracing by including the **traceoptions** statement at the **[edit routing-options]** hierarchy level and set the **all-packets-trace** and **all** flags to trace all protocol packets.

```
[edit routing-options traceoptions]
user@host# set file all-packets-trace
user@host# set flag all
```

2. Configure the filename for the MSDP trace file.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file msdp-trace
```

3. (Optional) Configure the maximum number of trace files.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file files 5
```

4. (Optional) Configure the maximum size of each trace file.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file size 1m
```

5. (Optional) Enable unrestricted file access.

```
[edit protocols msdp group groupa traceoptions]
user@host# set file world-readable
```

6. Configure tracing flags. Suppose you are troubleshooting issues with the source-active cache for **groupa**. The following example shows how to trace messages associated with the group address.

```
[edit protocols msdp group groupa traceoptions]
user@host# set flag source-active | match 230.0.0.3
```

7. View the trace file.

```
user@host> file list /var/log
user@host> file show /var/log/msdp-trace
```



- Related Documentation**
- [Understanding MSDP on page 4767](#)
  - *Tracing and Logging Junos OS Operations* in the *Junos OS Administration Library for Routing Devices*

## Configuring the Interface to Accept Traffic from a Remote Source

You can configure an incoming interface to accept traffic from a remote source. A remote source is a source that is not on the same subnet as the incoming interface. This enables the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. You do not need to disable RPF checking, but you do need to ensure that the best path to reach the remote source is through the incoming interface.

In this sample configuration, the incoming interface (**ge-1/3/0**) is on a provider edge (PE) router on the receiver side of a multicast VPN.

To accept traffic from a remote source:

1. Edit the incoming interface.

```
[edit protocols pim interface ge-1/3/0.0]
user@host# set accept-remote-source
```

2. If the incoming interface is not the only way to reach the remote source, ensure that the best path to reach the remote source is through the incoming interface. One way to do this is to use AS path prepending on the other possible routes.

```
[edit policy-options policy-statement as-path-prepend term prepend]
user@host# set from route-filter 192.168.0.0/16 orlonger
user@host# set from route-filter 172.16.0.0/16 orlonger
user@host# set then as-path-prepend "1 1 1 1"
```

Another way to do this might be to configure a static route on the receiver side PE router to the source.

3. After the configuration is committed, use the **show pim statistics** and **show msdp source** commands to verify that the interface is accepting traffic from the remote source.

- Related Documentation**
- *Example: Allowing MBGP MVPN Remote Sources*
  - *Understanding Prepending AS Numbers to BGP AS Paths* in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
  - [show msdp source on page 5107](#) in the [CLI Explorer](#)
  - [show pim statistics on page 5189](#) in the [CLI Explorer](#)

## Example: Configuring MSDP

Configure a router to act as a PIM sparse-mode rendezvous point and an MSDP peer:

```
[edit]
routing-options {
  interface-routes {
```

```
    rib-group ifrg;
  }
  rib-groups {
    ifrg {
      import-rib [inet.0 inet.2];
    }
    mcrg {
      export-rib inet.2;
      import-rib inet.2;
    }
  }
}
protocols {
  bgp {
    group lab {
      type internal;
      family any;
      neighbor 192.168.6.18 {
        local-address 192.168.6.17;
      }
    }
  }
  pim {
    dense-groups {
      224.0.1.39/32;
      224.0.1.40/32;
    }
    rib-group mcrg;
    rp {
      local {
        address 192.168.1.1;
      }
    }
    interface all {
      mode sparse-dense;
      version 1;
    }
  }
  msdp {
    rib-group mcrg;
    group lab {
      peer 192.168.6.18 {
        local-address 192.168.6.17;
      }
    }
  }
}
```

### Example: Configuring MSDP with Active Source Limits and Mesh Groups

This example shows how to configure MSDP to filter source-active messages and limit the flooding of source-active messages.

- [Requirements on page 4881](#)
- [Overview on page 4881](#)

- [Configuration on page 4885](#)
- [Verification on page 4886](#)

---

## Requirements

Before you begin:

- Configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.
- Configure an interior gateway protocol or static routing. See the *Junos OS Routing Protocols Library for Routing Devices*.
- Enable PIM sparse mode. See [“PIM Overview” on page 4737](#).
- Configure the router as a PIM sparse-mode RP. See [“Configuring Local PIM RPs” on page 4795](#).

---

## Overview

A router interested in MSDP messages, such as an RP, might have to process a large number of MSDP messages, especially source-active messages, arriving from other routers. Because of the potential need for a router to examine, process, and create state tables for many MSDP packets, there is a possibility of an MSDP-based denial-of-service (DoS) attack on a router running MSDP. To minimize this possibility, you can configure the router to limit the number of source active messages the router accepts. Also, you can configure a threshold for applying random early discard (RED) to drop some but not all MSDP active source messages. Beginning with Junos OS 12.2, you can optionally configure a warning threshold so the device can log warning messages in the system log when a certain number of source-active messages have been received. It is helpful to review the system log messages for troubleshooting purposes and to detect if an excessive amount of source-active messages have been received. These log messages convey when the configured message limit has been exceeded, when the configured warning threshold has been exceeded, and when the number of messages drop below the configured warning threshold.

By default, the router accepts 25,000 source active messages before ignoring the rest. The limit can be from 1 through 1,000,000. The limit is applied to both the number of messages and the number of MSDP peers.

By default, the router accepts 24,000 source-active messages before applying the RED profile to prevent a possible DoS attack. This number can also range from 1 through 1,000,000. The next 1000 messages are screened by the RED profile and the accepted messages processed. If you configure no drop profiles (as this example does not), RED is still in effect and functions as the primary mechanism for managing congestion. In the default RED drop profile, when the packet queue fill-level is 0 percent, the drop probability is 0 percent. When the fill-level is 100 percent, the drop probability is 100 percent.



**NOTE:** The router ignores source-active messages with encapsulated TCP packets. Multicast does not use TCP; segments inside source-active messages are most likely the result of worm activity.

---

The number configured for the threshold must be less than the number configured for the maximum number of active MSDP sources.

The warning threshold is a percentage of maximum number of MSDP source-active messages received, so you must configure the source-active message limit to configure a warning threshold. The range for the warning threshold is 1 through 100 percent. You can further specify the amount of time (in seconds) between the log messages. The range is 6 through 32,767 seconds.

You can configure an active source limit globally, for a group, or for a peer. If active source limits are configured at multiple levels of the hierarchy (as shown in this example), all are applied.

You can configure an active source limit for an address range as well as for a specific peer. A per-source active source limit uses an IP prefix and prefix length instead of a specific address. You can configure more than one per-source active source limit. The longest match determines the limit.

Per-source active source limits can be combined with active source limits at the peer, group, and global (instance) hierarchy level. Per-source limits are applied before any other type of active source limit. Limits are tested in the following order:

- Per-source
- Per-peer or group
- Per-instance

An active source message must “pass” all limits established before being accepted. For example, if a source is configured with an active source limit of 10,000 active multicast groups and the instance is configured with a limit of 5000 (and there are no other sources or limits configured), only 5000 active source messages are accepted from this source.

MSDP mesh groups are groups of peers configured in a full-mesh topology that limits the flooding of source-active messages to neighboring peers. Every mesh group member must have a peer connection with every other mesh group member. When a source-active message is received from a mesh group member, the source-active message is always accepted but is not flooded to other members of the same mesh group. However, the source-active message is flooded to non-mesh group peers or members of other mesh groups. By default, standard flooding rules apply if **mesh-group** is not specified.



**CAUTION:** When configuring MSDP mesh groups, you must configure all members the same way. If you do not configure a full mesh, excessive flooding of source-active messages can occur.

---

A common application for MSDP mesh groups is peer-reverse-path-forwarding (peer-RPF) check bypass. For example, if there are two MSDP peers inside an autonomous system (AS), and only one of them has an external MSDP session to another AS, the internal MSDP peer often rejects incoming source-active messages relayed by the peer with the external link. Rejection occurs because the external MSDP peer must be reachable by

the internal MSDP peer through the next hop toward the source in another AS, and this next-hop condition is not certain. To prevent rejections, configure an MSDP mesh group on the internal MSDP peer so it always accepts source-active messages.



**NOTE:** An alternative way to bypass the peer-RPF check is to configure a default peer. In networks with only one MSDP peer, especially stub networks, the source-active message always needs to be accepted. An MSDP default peer is an MSDP peer from which all source-active messages are accepted without performing the peer-RPF check. You can establish a default peer at the peer or group level by including the **default-peer** statement.

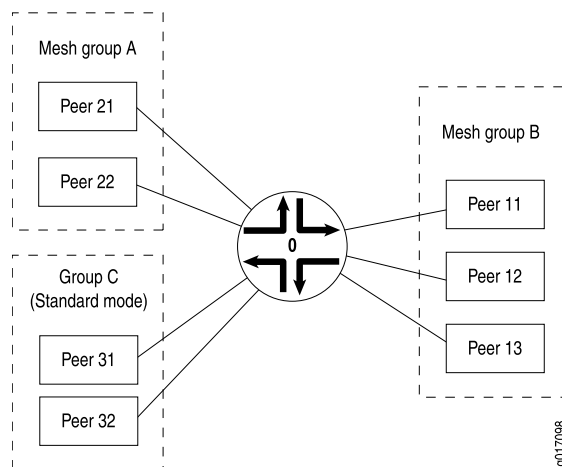
Table 391 on page 4883 explains how flooding is handled by peers in this example.

Figure 173 on page 4883 illustrates source-active message flooding between different mesh groups and peers within the same mesh group.

**Table 391: Source-Active Message Flooding Explanation**

Source-Active Message Received From	Source-Active Message Flooded To	Source-Active Message Not Flooded To
Peer 21	Peer 11, Peer 12, Peer 13, Peer 31, Peer 32	Peer 22
Peer 11	Peer 21, Peer 22, Peer 31, Peer 32	Peer 12, Peer 13
Peer 31	Peer 21, Peer 22, Peer 11, Peer 12, Peer 13, Peer 32	—

**Figure 173: Source-Active Message Flooding**



This example includes the following settings:

- **active-source-limit maximum 10000**—Applies a limit of 10,000 active sources to all other peers.
- **active-source-limit log-warning 80**—(Optional) Applies a warning threshold of 80 percent. In this example, the active source maximum is 10,000, so the device will start logging warning messages once it receives 8,000 active source messages.
- **active-source-limit log-interval 20**—(Optional) Applies a 20 second waiting period between system log messages.
- **data-encapsulation disable**—On an RP router using MSDP, disables the default encapsulation of multicast data received in MSDP register messages inside MSDP source-active messages.

MSDP data encapsulation mainly concerns bursty sources of multicast traffic. Sources that send only one packet every few minutes have trouble with the timeout of state relationships between sources and their multicast groups (S,G). Routers lose data while they attempt to reestablish (S,G) state tables. As a result, multicast register messages contain data, and this data encapsulation in MSDP source-active messages can be turned on or off through configuration.

By default, MSDP data encapsulation is enabled. An RP running MSDP takes the data packets arriving in the source's register message and encapsulates the data inside an MSDP source-active message.

However, data encapsulation creates both a multicast forwarding cache entry in the **inet.1** table (this is also the forwarding table) and a routing table entry in the **inet.4** table. Without data encapsulation, MSDP creates only a routing table entry in the **inet.4** table. In some circumstances, such as the presence of Internet worms or other forms of DoS attack, the router's forwarding table might fill up with these entries. To prevent the forwarding table from filling up with MSDP entries, you can configure the router not to use MSDP data encapsulation. However, if you disable data encapsulation, the router ignores and discards the encapsulated data. Without data encapsulation, multicast applications with bursty sources having transmit intervals greater than about 3 minutes might not work well.

- **group MSDP-group local-address 10.1.2.3**—Specifies the address of the local router (this router).
- **group MSDP-group mode mesh-group**—Specifies that all peers belonging to the **MSDP-group** are mesh group members.
- **group MSDP-group peer 10.10.10.10**—Prevents the sending of source-active messages to neighboring peer 10.10.10.10.
- **group MSDP-group peer 10.10.10.10 active-source-limit maximum 7500**—Applies a limit of 7500 active sources to MSDP peer 10.10.10.10 in group **MSDP-group**.
- **peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000**—Applies a threshold of 4000 active sources and a limit of 5000 active sources to MSDP peer 10.0.0.1.
- **source 10.1.0.0/16 active-source-limit maximum 500**—Applies a limit of 500 active sources to any source on the 10.1.0.0/16 network.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols msdp data-encapsulation disable
set protocols msdp active-source-limit maximum 10000
set protocols msdp active-source-limit log-warning 80
set protocols msdp active-source-limit log-interval 20
set protocols msdp peer 10.0.0.1 active-source-limit maximum 5000
set protocols msdp peer 10.0.0.1 active-source-limit threshold 4000
set protocols msdp source 10.1.0.0/16 active-source-limit maximum 500
set protocols msdp group MSDP-group mode mesh-group
set protocols msdp group MSDP-group local-address 10.1.2.3
set protocols msdp group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure MSDP source active routes and mesh groups:

1. (Optional) Disable data encapsulation.

```
[edit protocols msdp]
user@host# set data-encapsulation disable
```

2. Configure the active source limits.

```
[edit protocols msdp]
user@host# set peer 10.0.0.1 active-source-limit maximum 5000 threshold 4000
user@host# set group MSDP-group peer 10.10.10.10 active-source-limit maximum
7500
user@host# set active-source-limit maximum 10000
user@host# set source 10.1.0.0/16 active-source-limit maximum 500
```

3. (Optional) Configure the threshold at which warning messages are logged and the amount of time between log messages.

```
[edit protocols msdp]
user@host# set active-source-limit log-warning 80
user@host# set active-source-limit log-interval 20
```

4. Configure the mesh group.

```
[edit protocols msdp]
user@host# set group MSDP-group mode mesh-group
user@host# set group MSDP-group peer 10.10.10.10
user@host# set group MSDP-group local-address 10.1.2.3
```

5. If you are done configuring the device, commit the configuration.

```
[edit routing-instances]
user@host# commit
```

## Results

Confirm your configuration by entering the **show protocols** command.

```
user@host# show protocols
msdp {
  data-encapsulation disable;
  active-source-limit {
    maximum 10000;
    log-warning 80;
    log-interval 20;
  }
  peer 10.0.0.1 {
    active-source-limit {
      maximum 5000;
      threshold 4000;
    }
  }
  source 10.1.0.0/16 {
    active-source-limit {
      maximum 500;
    }
  }
  group MSDP-group {
    mode mesh-group;
    local-address 10.1.2.3;
    peer 10.10.10.10 {
      active-source-limit {
        maximum 7500;
      }
    }
  }
}
```

---

## Verification

To verify the configuration, run the following commands:

- [show msdp source-active](#)
- [show msdp statistics](#)

### Related Documentation

- [Example: Configuring MSDP in a Routing Instance](#)
- [Filtering MSDP SA Messages on page 4769](#)
- [Configuring RED Drop Profiles in the Class of Service Feature Guide for Routing Devices](#)
- [Configuring Local PIM RPs on page 4795](#)

## Example: Configuring PIM Anycast With or Without MSDP

When you configure anycast RP, you bypass the restriction of having one active rendezvous point (RP) per multicast group, and instead deploy multiple RPs for the same group range. The RP routers share one unicast IP address. Sources from one RP are known to



other RPs that use the Multicast Source Discovery Protocol (MSDP). Sources and receivers use the closest RP, as determined by the interior gateway protocol (IGP).

You can use anycast RP within a domain to provide redundancy and RP load sharing. When an RP stops operating, sources and receivers are taken to a new RP by means of unicast routing.

You can configure anycast RP to use PIM and MSDP for IPv4, or PIM alone for both IPv4 and IPv6 scenarios. Both are discussed in this section.

We recommend a static RP mapping with anycast RP over a bootstrap router and auto-RP configuration because it provides all the benefits of a bootstrap router and auto-RP without the complexity of the BSR and auto-RP mechanisms.

All systems on a subnet must run the same version of PIM.

The default PIM version can be version 1 or version 2, depending on the mode you are configuring. PIMv1 is the default RP mode (at the **[edit protocols pim rp static address address]** hierarchy level). However, PIMv2 is the default for interface mode (at the **[edit protocols pim interface interface-name]** hierarchy level). Explicitly configured versions override the defaults. This example explicitly configures PIMv2 on the interfaces.

The following example shows an anycast RP configuration for the RP routers, first with MSDP and then using PIM alone, and for non-RP routers.

1. For a network using an RP with MSDP, configure the RP using the **lo0** loopback interface, which is always up. Include the **address** statement and specify the unique and routable router ID and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement for the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```
interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32;
        primary;
        address 198.58.3.253/32;
      }
    }
  }
}
```

2. Specify the RP address. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```

protocols {
  pim {
    rp {
      local {
        family inet;
        address 198.58.3.253;
      }
      interface all {
        mode sparse;
        version 2;
      }
      interface fxp0.0 {
        disable;
      }
    }
  }
}

```

3. Configure MSDP peering. Include the **peer** statement to configure the address of the MSDP peer at the **[edit protocols msdp]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP peering, include the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```

protocols {
  msdp {
    peer 198.58.3.250 {
      local-address 198.58.3.254;
    }
  }
}

```



**NOTE:** If you need to configure a PIM RP for both IPv4 and IPv6 scenarios, perform Step 4 and Step 5. Otherwise, go to Step 6.

4. Configure an RP using the **lo0** loopback interface, which is always up. Include the **address** statement to specify the unique and routable router address and the RP address at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level. In this example, the router ID is **198.58.3.254** and the shared RP address is **198.58.3.253**. Include the **primary** statement on the first address. Including the **primary** statement selects the router's primary address from all the preferred addresses on all interfaces.

```

interfaces {
  lo0 {
    description "PIM RP";
    unit 0 {
      family inet {
        address 198.58.3.254/32 {
          primary;
        }
        address 198.58.3.253/32;
      }
    }
  }
}

```

```
}
```

5. Include the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, include the **mode** statement to set the mode to **sparse**, and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

Include the **anycast-pim** statement to configure anycast RP without MSDP (for example, if IPv6 is used for multicasting). The other RP routers that share the same IP address are configured using the **rp-set** statement. There is one entry for each RP, and the maximum that can be configured is 15. For each RP, specify the routable IP address of the router and whether MSDP source active (SA) messages are forwarded to the RP.

MSDP configuration is not necessary for this type of IPv4 anycast RP configuration.

```
protocols {
  pim {
    rp {
      local {
        family inet {
          address 198.58.3.253;
          anycast-pim {
            rp-set {
              address 198.58.3.240;
              address 198.58.3.241 forward-msdp-sa;
            }
            local-address 198.58.3.254; #If not configured, use lo0 primary
          }
        }
      }
    }
  }
  interface all {
    mode sparse;
    version 2;
  }
  interface fxp0.0 {
    disable;
  }
}
```

6. Configure the non-RP routers. The anycast RP configuration for a non-RP router is the same whether MSDP is used or not. Specify a static RP by adding the address at the **[edit protocols pim rp static]** hierarchy level. Include the **version** statement at the **[edit protocols pim rp static address]** hierarchy level to specify PIM version 2.

```
protocols {
  pim {
    rp {
      static {
        address 198.58.3.253 {
          version 2;
        }
      }
    }
  }
}
```

```
    }  
  }  
}
```

7. Include the **mode** statement at the **[edit protocols pim interface all]** hierarchy level to specify sparse mode on all interfaces. Then include the **version** statement at the **[edit protocols pim rp interface all mode]** to configure all interfaces for PIM version 2. When configuring all interfaces, exclude the **fxp0.0** management interface by including the **disable** statement for that interface.

```
protocols {  
  pim {  
    interface all {  
      mode sparse;  
      version 2;  
    }  
    interface fxp0.0 {  
      disable;  
    }  
  }  
}
```

## Configuring a PIM Anycast RP Router with MSDP

Add the **address** statement at the **[edit protocols pim rp local]** hierarchy level to specify the RP address (the same address as the secondary **lo0** interface).

For all interfaces, use the **mode** statement to set the mode to **sparse** and the **version** statement to specify PIM version 2 at the **[edit protocols pim rp local interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface.

```
protocols {  
  pim {  
    rp {  
      local {  
        family inet;  
        address 198.58.3.253;  
      }  
      interface all {  
        mode sparse;  
        version 2;  
      }  
      interface fxp0.0 {  
        disable;  
      }  
    }  
  }  
}
```

To configure MSDP peering, add the **peer** statement to configure the address of the MSDP peer at the **[edit protocols mspf]** hierarchy level. For MSDP peering, use the unique, primary addresses instead of the anycast address. To specify the local address for MSDP

peering, add the **local-address** statement at the **[edit protocols msdp peer]** hierarchy level.

```
protocols {
  msdp {
    peer 198.58.3.250 {
      local-address 198.58.3.254;
    }
  }
}
```


## Source-Specific Multicast

- [Example: Configuring PIM SSM on a Network on page 4891](#)
- [Example: Configuring an SSM-Only Domain on page 4892](#)
- [Example: Configuring SSM Mapping on page 4893](#)
- [Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895](#)
- [Example: Configuring SSM Maps for Different Groups to Different Sources on page 4899](#)

### Example: Configuring PIM SSM on a Network

The following example shows how PIM SSM is configured between a receiver and a source in the network illustrated in [Figure 174 on page 4891](#).

Figure 174: Network on Which to Configure PIM SSM



This example shows how to configure the IGMP version to IGMPv3 on all receiving host interfaces.

1. Enable IGMPv3 on all host-facing interfaces, and disable IGMP on the **fxp0.0** interface on Router 1.

```
user@router1# set protocols igmp interface all version 3
user@router1# set protocols igmp interface fxp0.0 disable
```



**NOTE:** When you configure IGMPv3 on a router, hosts on interfaces configured with IGMPv2 cannot join the source tree.

2. After the configuration is committed, use the **show configuration protocol igmp** command to verify the IGMP protocol configuration.

```
user@router1> show configuration protocol igmp
```

```
[edit protocols igmp]
interface all {
  version 3;
}
interface fxp0.0 {
  disable;
}
```

3. Use the **show igmp interface** command to verify that IGMP interfaces are configured.

```
user@router1> show igmp interface
Interface      State    Querier      Timeout Version Groups
fe-0/0/0.0     Up      198.58.3.245  213      3      0
fe-0/0/1.0     Up      198.58.3.241  220      3      0
fe-0/0/2.0     Up      198.58.3.237  218      3      0
Configured Parameters:
IGMP Query Interval (1/10 secs): 1250
IGMP Query Response Interval (1/10 secs): 100
IGMP Last Member Query Interval (1/10 secs): 10
IGMP Robustness Count: 2
Derived Parameters:
IGMP Membership Timeout (1/10 secs): 2600
IGMP Other Querier Present Timeout (1/10 secs): 2550
```

4. Use the **show pim join extensive** command to verify the PIM join state on Router 2 and Router 3 (the upstream routers).

```
user@router2> show pim join extensive
232.1.1.1      10.4.1.2      sparse
  Upstream interface: fe-1/1/3.0
  Upstream State: Local Source
  Keepalive timeout: 209
  Downstream Neighbors:
    Interface: so-1/0/2.0
      10.10.71.1      State: Join  Flags: S    Timeout: 209
```

5. Use the **show pim join extensive** command to verify the PIM join state on Router 1 (the router connected to the receiver).

```
user@router1> show pim join extensive
232.1.1.1      10.4.1.2      sparse
  Upstream interface: so-1/0/2.0
  Upstream State: Join to Source
  Keepalive timeout: 209
  Downstream Neighbors:
    Interface: fe-0/2/3.0
      10.3.1.1      State: Join  Flags: S    Timeout: Infinity
```

## Example: Configuring an SSM-Only Domain

Deploying an SSM-only domain is much simpler than deploying an ASM domain because it only requires a few configuration steps. Enable PIM sparse mode on all interfaces by adding the **mode** statement at the **[edit protocols pim interface all]** hierarchy level. When configuring all interfaces, exclude the **fxp0.0** management interface by adding the **disable** statement for that interface. Then configure IGMPv3 on all host-facing interfaces by adding the **version** statement at the **[edit protocols igmp interface interface-name]** hierarchy level.

In the following example, the host-facing interface is **fe-0/1/2**:

```
[edit]
protocols {
  pim {
    interface all {
      mode sparse;
      version 2;
    }
    interface fxp0.0 {
      disable;
    }
  }
  igmp {
    interface fe-0/1/2 {
      version 3;
    }
  }
}
```

### Example: Configuring SSM Mapping

SSM mapping does not require that all hosts support IGMPv3. SSM mapping translates IGMPv1 or IGMPv2 membership reports to an IGMPv3 report. This enables hosts running IGMPv1 or IGMPv2 to participate in SSM until the hosts transition to IGMPv3.

SSM mapping applies to all group addresses that match the policy, not just those that conform to SSM addressing conventions (232/8 for IPv4, ff30::/32 through ff3F::/32 for IPv6).

We recommend separate SSM maps for IPv4 and IPv6 if both address families require SSM support. If you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv4 context (using IGMP), only the IPv4 addresses in the list are used. If there are no such addresses, no action is taken. Similarly, if you apply an SSM map containing both IPv4 and IPv6 addresses to an interface in an IPv6 context (using MLD), only the IPv6 addresses in the list are used. If there are no such addresses, no action is taken.

In this example, you create a policy to match the group addresses that you want to translate to IGMPv3. Then you define the SSM map that associates the policy with the source addresses where these group addresses are found. Finally, you apply the SSM map to one or more IGMP (for IPv4) or MLD (for IPv6) interfaces.

1. Create an SSM policy named **ssm-policy-example**. The policy terms match the IPv4 SSM group address 232.1.1.1/32 and the IPv6 SSM group address ff35::1/128. All other addresses are rejected.

```
user@router1# set policy-options policy-statement ssm-policy-example term A from
route-filter 232.1.1.1/32 exact
user@router1# set policy-options policy-statement ssm-policy-example term A then
accept
user@router1# set policy-options policy-statement ssm-policy-example term B from
route-filter ff35::1/128 exact
user@router1# set policy-options policy-statement ssm-policy-example term B then
accept
```

2. After the configuration is committed, use the **show configuration policy-options** command to verify the policy configuration.

```
user@host> show configuration policy-options

[edit policy-options]
policy-statement ssm-policy-example {
  term A {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then accept;
  }
  term B {
    from {
      route-filter ff35::1/128 exact;
    }
    then accept;
  }
  then reject;
}
```

The group addresses must match the configured policy for SSM mapping to occur.

3. Define two SSM maps, one called **ssm-map-ipv6-example** and one called **ssm-map-ipv4-example**, by applying the policy and configuring the source addresses as a multicast routing option.

```
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv6-example source
fec0::1 fec0::12
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example policy
ssm-policy-example
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
10.10.10.4
user@host# set routing-options multicast ssm-map ssm-map-ipv4-example source
192.168.43.66
```

4. After the configuration is committed, use the **show configuration routing-options** command to verify the policy configuration.

```
user@host> show configuration routing-options

[edit routing-options]
multicast {
  ssm-map ssm-map-ipv6-example {
    policy ssm-policy-example;
    source [ fec0::1 fec0::12 ];
  }
  ssm-map ssm-map-ipv4-example {
    policy ssm-policy-example;
    source [ 10.10.10.4 192.168.43.66 ];
  }
}
```

We recommend separate SSM maps for IPv4 and IPv6.



5. Apply SSM maps for IPv4-to-IGMP interfaces and SSM maps for IPv6-to-MLD interfaces:

```
user@host# set protocols igmp interface fe-0/1/0.0 ssm-map ssm-map-ipv4-example
user@host# set protocols mld interface fe-0/1/1.0 ssm-map ssm-map-ipv6-example
```

6. After the configuration is committed, use the **show configuration protocol** command to verify the IGMP and MLD protocol configuration.

```
user@router1> show configuration protocol
```

```
[edit protocols]
igmp {
  interface fe-0/1/0.0 {
    ssm-map ssm-map-ipv4-example;
  }
}
mld {
  interface fe-0/1/1.0 {
    ssm-map ssm-map-ipv6-example;
  }
}
```

7. Use the **show igmp interface** and the **show mld interface** commands to verify that the SSM maps are applied to the interfaces.

```
user@host> show igmp interface fe-0/1/0.0
Interface: fe-0/1/0.0
Querier: 192.168.224.28
State:      Up Timeout:    None Version:  2 Groups:  2
SSM Map: ssm-map-ipv4-example

user@host> show mld interface fe-0/1/1.0
Interface: fe-0/1/1.0
Querier: fec0:0:0:0:1::12
State:      Up Timeout:    None Version:  2 Groups:  2
SSM Map: ssm-map-ipv6-example
```

## Example: Configuring Source-Specific Multicast Groups with Any-Source Override

This example shows how to extend source-specific multicast (SSM) group operations beyond the default IP address range of 232.0.0.0 through 232.255.255.255. This example also shows how to accept any-source multicast (ASM) join messages (\*G) for group addresses that are within the default or configured range of SSM groups. This allows you to support a mix of any-source and source-specific multicast groups simultaneously.

- [Requirements on page 4895](#)
- [Overview on page 4896](#)
- [Configuration on page 4897](#)
- [Verification on page 4899](#)

### Requirements


Before you begin, configure the router interfaces. See the *Junos OS Network Interfaces Library for Routing Devices*.

## Overview

To deploy SSM, configure PIM sparse mode on all routing device interfaces and issue the necessary SSM commands, including specifying IGMPv3 or MLDv2 on the receiver's LAN. If PIM sparse mode is not explicitly configured on both the source and group members interfaces, multicast packets are not forwarded. Source lists, supported in IGMPv3 and MLDv2, are used in PIM SSM. Only sources that are specified send traffic to the SSM group.


In a PIM SSM-configured network, a host subscribes to an SSM channel (by means of IGMPv3 or MLDv2) to join group G and source S (see [Figure 175 on page 4896](#)). The directly connected PIM sparse-mode router, the receiver's designated router (DR), sends an (S,G) join message to its reverse-path forwarding (RPF) neighbor for the source. Notice in [Figure 175 on page 4896](#) that the RP is not contacted in this process by the receiver, as would be the case in normal PIM sparse-mode operations.

**Figure 175: Receiver Sends Messages to Join Group G and Source S**




The (S,G) join message initiates the source tree and then builds it out hop by hop until it reaches the source. In [Figure 176 on page 4896](#), the source tree is built across the network to Router 3, the last-hop router connected to the source.

**Figure 176: Router 3 (Last-Hop Router) Joins the Source Tree**



Using the source tree, multicast traffic is delivered to the subscribing host (see [Figure 177 on page 4897](#)).


Figure 177: (S,G) State Is Built Between the Source and the Receiver



SSM can operate in include mode or in exclude mode. In exclude mode the receiver specifies a list of sources that it does not want to receive the multicast group traffic from. The routing device forwards traffic to the receiver from any source except the sources specified in the exclusion list. The receiver accepts traffic from any sources except the sources specified in the exclusion list.

This example works with the simple RPF topology shown in [Figure 178 on page 4897](#).

Figure 178: Simple RPF Topology



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 10.255.72.46
set protocols pim rp local group-ranges 239.0.0.0/24
set protocols pim interface fe-1/0/0.0 mode sparse
set protocols pim interface lo0.0 mode sparse
set routing-options multicast ssm-groups 232.0.0.0/8
set routing-options multicast ssm-groups 239.0.0.0/8
set routing-options multicast asm-override-ssm
```

#### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure an RPF policy:

1. Configure OSPF.

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface fxp0.0 disable
user@host# set area 0.0.0.0 interface all
```

2. Configure PIM sparse mode.

```
[edit protocols pim]
user@host# set rp local address 10.255.72.46
user@host# set rp local group-ranges 239.0.0.0/24
user@host# set interface fe-1/0/0.0 mode sparse
user@host# set interface lo0.0 mode sparse
```

3. Configure additional SSM groups.

```
[edit routing-options]
user@host# set ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ]
```

4. Configure the RP to accept ASM join messages for groups within the SSM address range.

```
[edit routing-options]
user@host# set multicast asm-override-ssm
```

5. If you are done configuring the device, commit the configuration.

```
user@host# commit
```

### Results

Confirm your configuration by entering the **show protocols** and **show routing-options** commands.

```
user@host# show protocols
ospf {
  area 0.0.0.0 {
    interface fxp0.0 {
      disable;
    }
    interface all;
  }
}
pim {
  rp {
    local {
      address 10.255.72.46;
      group-ranges {
        239.0.0.0/24;
      }
    }
  }
}
interface fe-1/0/0.0 {
  mode sparse;
}
interface lo0.0 {
  mode sparse;
}
}

user@host# show routing-options
multicast {
  ssm-groups [ 232.0.0.0/8 239.0.0.0/8 ];
  asm-override-ssm;
```

```
}

```

Verification

To verify the configuration, run the following commands:

- `show igmp group`
- `show igmp statistics`
- `show pim join`

Related Documentation

- [Source-Specific Multicast Groups Overview on page 4769](#)

Example: Configuring SSM Maps for Different Groups to Different Sources

- [Multiple SSM Maps and Groups for Interfaces on page 4899](#)
- [Example: Configuring Multiple SSM Maps Per Interface on page 4899](#)

Multiple SSM Maps and Groups for Interfaces

You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces.

Example: Configuring Multiple SSM Maps Per Interface

This example shows how to assign more than one SSM map to an IGMP interface.

- [Requirements on page 4899](#)
- [Overview on page 4899](#)
- [Configuration on page 4900](#)
- [Verification on page 4901](#)

Requirements

This example requires Junos OS Release 11.4 or later.

Overview

In this example, you configure a routing policy, `POLICY-ipv4-example1`, that maps multicast group join messages over an IGMP logical interface to IPv4 multicast source addresses based on destination IP address as follows:

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 1	232.1.1.1	10.10.10.4, 192.168.43.66

Routing Policy Name	Multicast Group Join Messages for a Route Filter at This Destination Address	Multicast Source Addresses
POLICY-ipv4-example1 term 2	232.1.1.2	10.10.10.5, 192.168.43.67

You apply routing policy POLICY-ipv4-example1 to IGMP logical interface fe-0/1/0.0.

### Configuration

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure this example, perform the following task:

#### CLI Quick Configuration

To quickly configure this example, copy the following configuration commands into a text file, remove any line breaks, and then paste the commands into the CLI at the **[edit]** hierarchy level.

```
set policy-options policy-statement POLICY-ipv4-example1 term 1 from route-filter
  232.1.1.1/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  10.10.10.4
set policy-options policy-statement POLICY-ipv4-example1 term 1 then ssm-source
  192.168.43.66
set policy-options policy-statement POLICY-ipv4-example1 term 1 then accept
set policy-options policy-statement POLICY-ipv4-example1 term 2 from route-filter
  232.1.1.2/32 exact
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  10.10.10.5
set policy-options policy-statement POLICY-ipv4-example1 term 2 then ssm-source
  192.168.43.67
set policy-options policy-statement POLICY-ipv4-example1 term 2 then accept
set protocols igmp interface fe-0/1/0.0 ssm-map-policy POLICY-ipv4-example1
```

#### Step-by-Step Procedure

To configure multiple SSM maps per interface:

1. Configure protocol-independent routing options for route filter 232.1.1.1, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```
[edit policy-options policy-statement POLICY-ipv4-example1 term 1]
user@host# set from route-filter 232.1.1.1/32 exact
user@host# set then ssm-source 10.10.10.4
user@host# set then ssm-source 192.168.43.66
user@host# set then accept
```

2. Configure protocol-independent routing options for route filter 232.1.1.2, and specify the multicast source addresses to which matching multicast groups are to be mapped.

```
[edit policy-options policy-statement POLICY-ipv4-example1 term 2]
user@host# set from route-filter 232.1.1.2/32 exact
user@host# set then ssm-source 10.10.10.5
```

```
user@host# set then ssm-source 192.168.43.67
user@host# set then accept
```

3. Apply the policy map POLICY-ipv4-example1 to IGMP logical interface fe-0/1/1/0.

```
[edit protocols igmp interface fe-0/1/0.0]
user@host# set ssm-map-policy POLICY-ipv4-example1
```

**Results** After the configuration is committed, confirm the configuration by entering the **show policy-options** and **show protocols** configuration mode commands. If the command output does not display the intended configuration, repeat the instructions in this procedure to correct the configuration.

```
user@host#> show policy-options
policy-statement POLICY-ipv4-example1 {
  term 1 {
    from {
      route-filter 232.1.1.1/32 exact;
    }
    then {
      ssm-source [ 10.10.10.4 192.168.43.66 ];
      accept;
    }
  }
  term 2 {
    from {
      route-filter 232.1.1.2/32 exact;
    }
    then {
      ssm-source [ 10.10.10.5 192.168.43.67 ];
      accept;
    }
  }
}

user@host# show protocols
igmp {
  interface fe-0/1/0.0 {
    ssm-map-policy POLICY-ipv4-example1;
  }
}
```

### Verification

Confirm that the configuration is working properly.

- [Displaying Information About IGMP-Enabled Interfaces on page 4901](#)
- [Displaying the PIM Groups on page 4902](#)
- [Displaying the Entries in the IP Multicast Forwarding Table on page 4902](#)

### Displaying Information About IGMP-Enabled Interfaces

**Purpose** Verify that the SSM map policy POLICY-ipv4-example1 is applied to logical interface fe-0/1/0.0.

**Action** Use the [show igmp interface](#) operational mode command for the IGMP logical interface to which you applied the SSM map policy.

```
user@host> show igmp interface
Interface: fe-0/1/0.0
  Querier: 10.111.30.1
  State:      Up Timeout:      None Version:  2 Groups:      2
  SSM Map Policy: POLICY-ipv4-example1;
```

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

```
Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

The command output displays the name of IGMP logical interface (fe-0/1/0.0), the address of the routing device that has been elected to send membership queries and group information.

#### *Displaying the PIM Groups*

**Purpose** Verify the Protocol Independent Multicast (PIM) source and group pair (S,G) entries.

**Action** Use the [show pim join extensive 232.1.1.1](#) operational mode command to display the PIM source and group pair (S,G) entries for the 232.1.1.1 group.

#### *Displaying the Entries in the IP Multicast Forwarding Table*

**Purpose** Verify that the IP multicast forwarding table displays the mroute state.

**Action** Use the [show multicast route extensive](#) operational mode command to display the entries in the IP multicast forwarding table to verify that the **Route state** is active and that the **Forwarding state** is forwarding.

**Related Documentation**

- *Example: Configuring Source-Specific Multicast*
- *Example: Configuring Source-Specific Draft-Rosen 7 Multicast VPNs*

---

## PIM Configuration Statements

- [address \(Anycast RPs\) on page 4905](#)
- [address \(Local RPs\) on page 4905](#)
- [address \(Static RPs\) on page 4906](#)
- [algorithm on page 4907](#)
- [anycast-pim on page 4908](#)
- [assert-timeout on page 4909](#)



- [authentication \(Protocols PIM\) on page 4910](#)
- [bfd-liveness-detection \(Protocols PIM\) on page 4911](#)
- [bootstrap on page 4912](#)
- [bootstrap-export on page 4913](#)
- [bootstrap-import on page 4914](#)
- [bootstrap-priority on page 4915](#)
- [detection-time \(BFD for PIM\) on page 4916](#)
- [disable \(PIM\) on page 4917](#)
- [dr-election-on-p2p on page 4918](#)
- [dr-register-policy on page 4918](#)
- [embedded-rp on page 4919](#)
- [export \(Protocols PIM Bootstrap\) on page 4920](#)
- [export \(Protocols PIM\) on page 4920](#)
- [family \(Bootstrap\) on page 4921](#)
- [family \(Protocols PIM\) on page 4922](#)
- [family \(Local RP\) on page 4923](#)
- [group \(RPF Selection\) on page 4924](#)
- [group-ranges on page 4925](#)
- [hello-interval \(Protocols PIM\) on page 4926](#)
- [hold-time \(Protocols PIM\) on page 4927](#)
- [import \(Protocols PIM Bootstrap\) on page 4928](#)
- [import \(Protocols PIM\) on page 4929](#)
- [infinity on page 4930](#)
- [interface on page 4931](#)
- [join-load-balance on page 4932](#)
- [join-prune-timeout on page 4933](#)
- [key-chain \(Protocols PIM\) on page 4934](#)
- [local on page 4935](#)
- [local-address \(Protocols PIM\) on page 4936](#)
- [loose-check on page 4937](#)
- [maximum-rps on page 4938](#)
- [minimum-interval \(PIM BFD Liveness Detection\) on page 4939](#)
- [minimum-interval \(PIM BFD Transmit Interval\) on page 4940](#)
- [minimum-receive-interval on page 4941](#)
- [mode \(Protocols PIM\) on page 4942](#)
- [multiplier on page 4942](#)
- [neighbor-policy on page 4943](#)

- [next-hop \(PIM RPF Selection\) on page 4943](#)
- [no-adaptation \(PIM BFD Liveness Detection\) on page 4944](#)
- [override-interval on page 4945](#)
- [pim on page 4946](#)
- [prefix-list \(PIM RPF Selection\) on page 4949](#)
- [priority \(Bootstrap\) on page 4950](#)
- [priority \(PIM Interfaces\) on page 4951](#)
- [priority \(PIM RPs\) on page 4952](#)
- [propagation-delay on page 4953](#)
- [reset-tracking-bit on page 4954](#)
- [rib-group \(Protocols PIM\) on page 4955](#)
- [rp on page 4956](#)
- [rp-register-policy on page 4958](#)
- [rp-set on page 4959](#)
- [rpf-selection on page 4960](#)
- [source \(PIM RPF Selection\) on page 4961](#)
- [spt-threshold on page 4962](#)
- [static \(Protocols PIM\) on page 4963](#)
- [threshold \(PIM BFD Detection Time\) on page 4964](#)
- [threshold \(PIM BFD Transmit Interval\) on page 4965](#)
- [transmit-interval \(PIM BFD Liveness Detection\) on page 4966](#)
- [traceoptions \(Protocols PIM\) on page 4967](#)
- [version \(BFD\) on page 4970](#)
- [version \(PIM\) on page 4971](#)
- [wildcard-source \(PIM RPF Selection\) on page 4972](#)

## address (Anycast RPs)

<b>Syntax</b>	<code>address <i>address</i> &lt;forward-msdp-sa&gt;;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b> ], [edit protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b> (inet   inet6) <b>anycast-pim rp-set</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.
<b>Options</b>	<b><i>address</i></b> —RP address in an RP set.  <b><i>forward-msdp-sa</i></b> —(Optional) Forward MSDP SAs to this address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## address (Local RPs)

<b>Syntax</b>	<code>address <i>address</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit protocols pim <b>rp local family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the local rendezvous point (RP) address.
<b>Options</b>	<b><i>address</i></b> —Local RP address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 4795</a></li> </ul>

## address (Static RPs)

---

<b>Syntax</b>	<pre>address address {     group-ranges {         destination-ip-prefix &lt;/prefix-length&gt;;     }     override;     version version; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems logical-system-name protocols pim rp static], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols     pim rp static], [edit protocols pim static], [edit routing-instances routing-instance-name protocols pim rp static]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure static rendezvous point (RP) addresses. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p>
<b>Options</b>	<p><b>address</b>—Static RP address.</p> <p><b>Default:</b> 224.0.0.0/4</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4797</a></li></ul>

## algorithm

---

<b>Syntax</b>	<code>algorithm <i>algorithm-name</i>;</code>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the algorithm to use for BFD authentication.
<b>Options</b>	<p><b><i>algorithm-name</i></b>—Name of algorithm to use for BFD authentication:</p> <ul style="list-style-type: none"> <li>• <b>simple-password</b>—Plain-text password. One to 16 bytes of plain text. One or more passwords can be configured.</li> <li>• <b>keyed-md5</b>—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive rates greater than 100 ms.</li> <li>• <b>meticulous-keyed-md5</b>—Meticulous keyed Message Digest 5 hash algorithm.</li> <li>• <b>keyed-sha-1</b>—Keyed Secure Hash Algorithm I for sessions with transmit and receive rates greater than 100 ms.</li> <li>• <b>meticulous-keyed-sha-1</b>—Meticulous keyed Secure Hash Algorithm I.</li> </ul>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding Bidirectional Forwarding Detection Authentication for PIM</i></li> <li>• <a href="#">Configuring BFD Authentication for PIM on page 4823</a></li> <li>• <a href="#">authentication on page 4910</a></li> </ul>

## anycast-pim

---

<b>Syntax</b>	<pre>anycast-pim {   rp-set {     address address &lt;forward-msdp-sa&gt;;   } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit protocols pim <b>rp local family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure properties for anycast RP using PIM.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 4799</a></li></ul>

## assert-timeout

---

<b>Syntax</b>	<code>assert-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds.
<b>Options</b>	<b><i>seconds</i></b> —Time for routing device to wait before another assert message cycle. <b>Range:</b> 5 through 210 seconds <b>Default:</b> 180 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the PIM Assert Timeout on page 4815</a></li> </ul>

## authentication (Protocols PIM)

---

<b>Syntax</b>	<pre>authentication {   algorithm <i>algorithm-name</i>;   key-chain <i>key-chain-name</i>;   loose-check; }</pre>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> family (inet   inet6) bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface family (inet   inet6) <i>interface-name</i> bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the algorithm, security keychain, and level of authentication for BFD sessions running on PIM interfaces.  The remaining statements are explained separately.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 4823</a></li><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM</a></li><li>• <a href="#">bfd-liveness-detection on page 4911</a></li><li>• <a href="#">key-chain (Protocols PIM) on page 4934</a></li><li>• <a href="#">loose-check on page 4937</a></li></ul>



## bfd-liveness-detection (Protocols PIM)

<b>Syntax</b>	<pre> bfd-liveness-detection {   authentication {     algorithm <i>algorithm-name</i>;     key-chain <i>key-chain-name</i>;     loose-check;   }   detection-time {     threshold <i>milliseconds</i>;   }   minimum-interval <i>milliseconds</i>;   minimum-receive-interval <i>milliseconds</i>;   multiplier <i>number</i>;   no-adaptation;   transmit-interval {     minimum-interval <i>milliseconds</i>;     threshold <i>milliseconds</i>;   }   version (0   1   automatic); } </pre>
<b>Hierarchy Level</b>	<p>[edit protocols pim interface <i>interface-name</i> <i>family</i> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <i>family</i> (inet   inet6)]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.1.</p> <p><b>authentication</b> option introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure bidirectional forwarding detection (BFD) timers and authentication for PIM.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4821</a></li> <li>• <a href="#">Configuring BFD Authentication for PIM on page 4823</a></li> </ul>

## bootstrap

---

<b>Syntax</b>	<pre>bootstrap {     family (inet   inet6) {         export [ <i>policy-names</i> ];         import [ <i>policy-names</i> ];         priority <i>number</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b> ], [edit protocols pim <b>rp</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure parameters to control bootstrap routers and messages.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li></ul>

## bootstrap-export

---

<b>Syntax</b>	<code>bootstrap-export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>],</p> <p>[edit protocols pim <a href="#">rp</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Apply one or more export policies to control outgoing PIM bootstrap messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> <li>• <a href="#">bootstrap-import on page 4914</a></li> </ul>

## bootstrap-import

---

<b>Syntax</b>	<code>bootstrap-import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>pim <a href="#">rp</a>],</code> <code>[edit protocols pim <a href="#">rp</a>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp</a>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more import policies to control incoming PIM bootstrap messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">bootstrap-export on page 4913</a></li></ul>

## bootstrap-priority

<b>Syntax</b>	<code>bootstrap-priority <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>],</p> <p>[edit protocols pim <i>rp</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.
<b>Options</b>	<p><i>number</i>—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 0</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> </ul>

## detection-time (BFD for PIM)

---

Syntax	<pre>detection-time {     threshold milliseconds; }</pre>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Enable BFD failure detection. The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the <b>clear bfd adaptation</b> command to return BFD interval timers to their configured values. The <b>clear bfd adaptation</b> command is hitless, meaning that the command does not affect traffic flow on the routing device.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li><li>• <a href="#">bfd-liveness-detection on page 4911</a></li><li>• <a href="#">threshold on page 4964</a></li></ul>

## disable (PIM)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	<pre> [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> protocols pim <b>family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   pim <b>rp local family</b> (inet   inet6)], [edit protocols pim], [edit protocols pim <b>family</b> (inet   inet6)], [edit protocols pim interface <i>interface-name</i>], [edit protocols pim <b>rp local family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>family</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>disable</b> statement extended to the <b>[family]</b> hierarchy level in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Explicitly disable PIM at the protocol, interface or family hierarchy levels.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Disabling PIM on page 4782</li> <li><i>disable (PIM Graceful Restart)</i></li> </ul>

## dr-election-on-p2p

---

<b>Syntax</b>	dr-election-on-p2p;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable PIM designated router (DR) election on point-to-point (P2P) links.
<b>Default</b>	No PIM DR election is performed on point-to-point links.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Designated Router Election on Point-to-Point Links on page 4785</a></li></ul>

## dr-register-policy

---

<b>Syntax</b>	dr-register-policy [ <i>policy-names</i> ];
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <i>rp</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i> ], [edit protocols pim <i>rp</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim <i>rp</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more policies to control outgoing PIM register messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 4813</a></li><li>• <a href="#">rp-register-policy on page 4958</a></li></ul>



## embedded-rp

<b>Syntax</b>	<pre> embedded-rp {   group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;;   }   maximum-rps limit; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],</p> <p>[edit protocols pim <b>rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure properties for embedded IP version 6 (IPv6) RPs.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring PIM Embedded RP for IPv6</i></li> </ul>

## export (Protocols PIM Bootstrap)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)], [edit protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp bootstrap family</a> (inet   inet6)]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more export policies to control outgoing PIM bootstrap messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4</a></li><li>• <a href="#">Configuring PIM Bootstrap Properties for IPv4 or IPv6</a></li><li>• <a href="#">import (Protocols PIM Bootstrap) on page 4928</a></li></ul>

## export (Protocols PIM)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more export policies to control outgoing PIM join and prune messages. PIM join and prune filters can be applied to PIM-SM and PIM-SSM messages. PIM join and prune filters cannot be applied to PIM-DM messages.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Filtering Outgoing PIM Join Messages on page 4811</a></li></ul>

## family (Bootstrap)

<b>Syntax</b>	<pre>family (inet   inet6) {     export [ <i>policy-names</i> ];     import [ <i>policy-names</i> ];     priority <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b>],</p> <p>[edit protocols pim <b>rp bootstrap</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Configure which IP protocol type bootstrap properties to apply.
<b>Options</b>	<p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li> <li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li> </ul>

## family (Protocols PIM)

---

<b>Syntax</b>	family (inet   inet6) { disable; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS 11.3 for the QFX Series.
<b>Description</b>	Enable the PIM protocol for the specified family.
<b>Options</b>	<b>inet</b> —Enable the PIM protocol for the IP version 4 (IPv4) address family.  <b>inet6</b> —Enable the PIM protocol for the IP version 6 (IPv6) address family.  The remaining statement is explained separately.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Disabling PIM on page 4782</a></li><li>• <i>disable (PIM Graceful Restart)</i></li><li>• <a href="#">disable (PIM) on page 4917</a></li></ul>

## family (Local RP)

<b>Syntax</b>	<pre>family (inet   inet6) {     disable;     address address;     anycast-pim {         local-address address;         rp-set {             address address &lt;forward-msdp-sa&gt;;         }     }     group-ranges {         destination-ip-prefix &lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b>],</p> <p>[edit protocols pim <b>rp local</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Configure which IP protocol type local RP properties to apply.
<b>Options</b>	<p><b>inet</b>—Apply IP version 4 (IPv4) local RP properties.</p> <p><b>inet6</b>—Apply IPv6 local RP properties.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 4795</a></li> </ul>

## group (RPF Selection)

---

<b>Syntax</b>	<pre>group group-address{   source source-address {     next-hop next-hop-address;   }   wildcard-source {     next-hop next-hop-address;   } }</pre>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> edit protocols pim rpf-selection]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the PIM group address for which you configure RPF selection <a href="#">group (RPF Selection)</a> .
<b>Default</b>	By default, PIM RPF selection is not configured.
<b>Options</b>	<b>group-address</b> —PIM group address for which you configure RPF selection.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring PIM RPF Selection</i></li></ul>

## group-ranges

<b>Syntax</b>	<pre>group-ranges {     destination-ip-prefix&lt;/prefix-length&gt;; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp embedded-rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp embedded-rp</b>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <b>rp embedded-rp</b>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp embedded-rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p>
<b>Description</b>	Configure the address ranges of the multicast groups for which this routing device can be a rendezvous point (RP).
<b>Default</b>	The routing device is eligible to be the RP for all IPv4 or IPv6 groups (224.0.0.0/4 or FF70::/12 to FFF0::/12).
<b>Options</b>	<b><i>destination-ip-prefix&lt;/prefix-length&gt;</i></b> —Addresses or address ranges for which this routing device can be an RP.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 4795</a> in the <i>Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Configuring PIM Embedded RP for IPv6</i> in the <i>Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>

## hello-interval (Protocols PIM)

---

<b>Syntax</b>	hello-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ], [edit protocols pim interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Specify how often the routing device sends PIM hello packets out of an interface.
<b>Options</b>	<b><i>seconds</i></b> —Length of time between PIM hello packets. <b>Range:</b> 0 through 255 <b>Default:</b> 30 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">hold-time on page 4927</a></li><li>• <a href="#">Modifying the PIM Hello Interval on page 4778</a></li></ul>



## hold-time (Protocols PIM)

<b>Syntax</b>	<code>hold-time seconds;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit protocols pim <b>rp local family</b> (inet   inet6)],</p> <p>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional RP addresses introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 13.3 for the PTX5000 router.</p>
<b>Description</b>	Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).
<b>Options</b>	<p><b>seconds</b>—Hold time.</p> <p><b>Range:</b> 0 through 255</p> <p><b>Default:</b> 150 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Local PIM RPs on page 4795</a> in the <i>Multicast Protocols Feature Guide for Routing Devices</i></li> <li>• <i>Example: Configuring Bidirectional PIM</i></li> </ul>

## import (Protocols PIM Bootstrap)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)], [edit protocols pim <b>rp bootstrap</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more import policies to control incoming PIM bootstrap messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">export (Protocols PIM Bootstrap) on page 4920</a></li></ul>

## import (Protocols PIM)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],</p> <p>[edit protocols pim],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Apply one or more policies to routes being imported into the routing table from PIM. Use the <b>import</b> statement to filter PIM join messages and prevent them from entering the network.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Filtering Incoming PIM Join Messages on page 4812</a></li> </ul>

## infinity

---

<b>Syntax</b>	<code>infinity [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">spt-threshold</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>pim <a href="#">spt-threshold</a>],</code> <code>[edit protocols pim <a href="#">spt-threshold</a>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">spt-threshold</a>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the <b>infinity</b> statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM SPT Threshold Policy on page 4818</a></li></ul>

## interface

<b>Syntax</b>	<pre> <b>interface</b> (all   <i>interface-name</i>) {     <b>disable</b>;     family (inet   inet6) {         <b>disable</b>;     }     <b>hello-interval</b> <i>seconds</i>;     mode (dense   sparse   sparse-dense);     <b>neighbor-policy</b> [ <i>policy-names</i> ];     <b>override-interval</b> <i>milliseconds</i>;     <b>priority</b> <i>number</i>;     <b>propagation-delay</b> <i>milliseconds</i>;     <b>reset-tracking-bit</b>;     <b>version</b> <i>version</i>; } </pre>
<b>Hierarchy Level</b>	[edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable PIM on an interface and configure interface-specific properties.
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">PIM on Aggregated Interfaces on page 4740</a></li> </ul>

## join-load-balance

---

<b>Syntax</b>	<pre>join-load-balance {     automatic; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable load balancing of PIM join messages across interfaces and routing devices.
<b>Options</b>	<b>automatic</b> —Enables automatic load balancing of PIM join messages. When a new interface or neighbor is introduced into the network, ECMP joins are redistributed with minimal disruption to traffic.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring PIM Make-Before-Break Join Load Balancing</i></li><li>• <a href="#">Configuring PIM Join Load Balancing on page 4787</a></li><li>• <i>clear pim join-distribution</i> in the <a href="#">CLI Explorer</a></li></ul>

## join-prune-timeout

---

<b>Syntax</b>	join-prune-timeout <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the timeout for the join state. If the periodic join refresh message is not received before the timeout expires, the join state is removed.
<b>Options</b>	<b>seconds</b> —Number of seconds to wait for the periodic join message to arrive. <b>Range:</b> 210 through 240 seconds <b>Default:</b> 210 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Modifying the Join State Timeout on page 4790</a></li> </ul>

## key-chain (Protocols PIM)

---

<b>Syntax</b>	<code>key-chain <i>key-chain-name</i>;</code>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> family {inet   inet6} bfd-liveness-detection authentication], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> family {inet   inet6} bfd-liveness-detection authentication]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement modified in Junos OS Release 12.2 to include <b>family</b> in the hierarchy level.
<b>Description</b>	Specify the security keychain to use for BFD authentication.
<b>Options</b>	<b><i>key-chain-name</i></b> —Name of the security keychain to use for BFD authentication. The name is a unique integer between <b>0</b> and <b>63</b> . This must match one of the keychains in the <b>authentication-key-chains</b> statement at the [edit security] hierarchy level.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD Authentication for PIM on page 4823</a></li><li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM authentication on page 4910</a></li></ul>



## local

<b>Syntax</b>	<pre> local {   disable;   address address;   family (inet   inet6) {     disable;     address address;     anycast-pim {       local-address address;       rp-set {         address address &lt;forward-msdp-sa&gt;;       }     }     group-ranges {       destination-ip-prefix &lt;/prefix-length&gt;;     }     hold-time seconds;     override;     priority number;   }   group-ranges {     destination-ip-prefix &lt;/prefix-length&gt;;   }   hold-time seconds;   override;   priority number; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],</p> <p>[edit protocols pim <b>rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>The remaining statements are explained separately.</p>
<b>Description</b>	Configure the routing device's RP properties.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Local PIM RPs on page 4795</a></li> </ul>

## local-address (Protocols PIM)

---

<b>Syntax</b>	<code>local-address <i>address</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code> <code>[edit protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6) <b>anycast-pim</b>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the routing device local address for the anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.
<b>Options</b>	<b><i>address</i></b> —Anycast RP IPv4 or IPv6 address, depending on <b>family</b> configuration.
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 4799</a></li></ul>

## loose-check

---

<b>Syntax</b>	loose-check;
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection authentication]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Specify loose authentication checking on the BFD session. Use loose authentication for transitional periods only when authentication might not be configured at both ends of the BFD session.</p> <p>By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure <i>loose checking</i>. When loose checking is configured, packets are accepted without authentication being checked at each end of the session.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD Authentication for PIM on page 4823</a></li> <li>• <a href="#">Understanding Bidirectional Forwarding Detection Authentication for PIM authentication on page 4910</a></li> </ul>

## maximum-rps

---

<b>Syntax</b>	<code>maximum-rps <i>limit</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <a href="#">rp embedded-rp</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a> ], [edit protocols pim <a href="#">rp embedded-rp</a> ], [edit routing-instances <i>routing-instance-name</i> protocols pim <a href="#">rp embedded-rp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Limit the number of RPs that the routing device acknowledges.
<b>Options</b>	<i>limit</i> —Number of RPs. <b>Range:</b> 1 through 500 <b>Default:</b> 100
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring PIM Embedded RP for IPv6</i></li></ul>

## minimum-interval (PIM BFD Liveness Detection)

<b>Syntax</b>	<code>minimum-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the minimum interval after which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the <b>transmit-interval</b> <b>minimum-interval</b> and <b>minimum-receive-interval</b> statements.
<b>Options</b>	<b><i>milliseconds</i></b> —Minimum transmit and receive interval. <b>Range:</b> 1 through 255,000 milliseconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4821</a></li> </ul>

## minimum-interval (PIM BFD Transmit Interval)

---

<b>Syntax</b>	<code>minimum-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the minimum interval after which the local routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the <b>minimum-interval</b> statement at the [edit protocols pim interface <i>interface-name</i> bfd-liveness-detection] hierarchy level.
<b>Options</b>	<i>milliseconds</i> —Minimum transmit interval value. <b>Range:</b> 1 through 255,000



**NOTE:** The threshold value specified in the **threshold** statement must be greater than the value specified in the **minimum-interval** statement for the **transmit-interval** statement.

---

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li><li>• <a href="#">bfd-liveness-detection on page 4911</a></li><li>• <a href="#">minimum-interval on page 4939</a></li><li>• <a href="#">threshold on page 4965</a></li></ul>

## minimum-receive-interval

<b>Syntax</b>	minimum-receive-interval <i>milliseconds</i> ;
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the minimum interval after which the local routing device must receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the <b>minimum-interval</b> statement at the [edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ] hierarchy level.
<b>Options</b>	<i>milliseconds</i> —Minimum receive interval. <b>Range:</b> 1 through 255,000 milliseconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4821</a></li> </ul>

## mode (Protocols PIM)

---

<b>Syntax</b>	mode (dense   sparse   sparse-dense);
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure PIM to operate in sparse, dense, or sparse-dense mode.



**NOTE:** The QFX Series does not support dense or sparse-dense mode.

---

<b>Options</b>	<b>dense</b> —Operate in dense mode.  <b>sparse</b> —Operate in sparse mode.  <b>sparse-dense</b> —Operate in sparse-dense mode. <b>Default:</b> sparse
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## multiplier

---

<b>Syntax</b>	multiplier <i>number</i> ;
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.
<b>Options</b>	<b>number</b> —Number of hello packets. <b>Range:</b> 1 through 255 <b>Default:</b> 3
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li></ul>



## neighbor-policy

<b>Syntax</b>	<code>neighbor-policy [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ], [edit protocols pim interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply a PIM interface-level policy to filter neighbor IP addresses.
<b>Options</b>	<i>policy-name</i> —Name of the policy that filters neighbor IP addresses.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Interface-Level PIM Neighbor Policies on page 4810</a></li> </ul>

## next-hop (PIM RPF Selection)

<b>Syntax</b>	<code>next-hop <i>next-hop-address</i>;</code>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the specific next-hop address for the PIM group source.
<b>Options</b>	<i>next-hop-address</i> —Specific next-hop address for the PIM group source.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM RPF Selection</a></li> </ul>

## no-adaptation (PIM BFD Liveness Detection)

---

<b>Syntax</b>	no-adaptation;
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure BFD sessions not to adapt to changing network conditions. We recommend that you <i>do not</i> disable BFD adaptation unless it is preferable to have BFD adaptation disabled in your network.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li><li>• <a href="#">bfd-liveness-detection on page 4911</a></li></ul>

## override-interval

<b>Syntax</b>	<code>override-interval <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim],  [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],  [edit protocols pim],  [edit protocols pim interface <i>interface-name</i>],  [edit routing-instances <i>routing-instance-name</i> protocols pim]  [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Set the maximum time in milliseconds to delay sending override join messages for a multicast network that has join suppression enabled. When a router or switch sees a prune message for a join it is currently suppressing, it waits for the interval specified by the override timer before it sends an override join message.
<b>Options</b>	<p>This is a random timer with a value in milliseconds.</p> <p><b>Range:</b> 0 through maximum override value</p> <p><b>Default:</b> 2000 milliseconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Enabling Join Suppression on page 4791</a></li> <li>• <a href="#">propagation-delay on page 4953</a></li> <li>• <a href="#">reset-tracking-bit on page 4954</a></li> </ul>

## pim

---

```
Syntax  pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        accept-remote-source;
        disable;
        family (inet | inet6) {
            disable;
        }
        hello-interval seconds;
        mode (dense | sparse | sparse-dense);
        neighbor-policy [ policy-names ];
        override-interval milliseconds;
        priority number;
        propagation-delay milliseconds;
        reset-tracking-bit;
        version version;
    }
    join-load-balance;
    join-prune-timeout;
    nonstop-routing;
    override-interval milliseconds;
    propagation-delay milliseconds;
    reset-tracking-bit;
    rib-group group-name;
    rp {
        auto-rp {
            (announce | discovery | mapping);
            (mapping-agent-election | no-mapping-agent-election);
        }
        bootstrap {
            family (inet | inet6) {
                export [ policy-names ];
                import [ policy-names ];
                priority number;
            }
        }
    }
    bootstrap-import [ policy-names ];
    bootstrap-export [ policy-names ];
```

```

bootstrap-priority number;
dr-register-policy [ policy-names ];
embedded-rp {
    group-ranges {
        destination-ip-prefix </prefix-length>;
    }
    maximum-rps limit;
}
local {
    family (inet | inet6) {
        address address;
        anycast-pim {
            disable;
            rp-set {
                address address <forward-msdp-sa>;
            }
            local-address address;
        }
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        hold-time seconds;
        priority number;
    }
}
rp-register-policy [ policy-names ];
spt-threshold {
    infinity [ policy-names ];
}
static {
    address address {
        group-ranges {
            version version;
            destination-ip-prefix </prefix-length>;
        }
    }
}
rpf-selection {
    group group-address {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
    prefix-list prefix-list-addresses {
        source source-address {
            next-hop next-hop-address;
        }
        wildcard-source {
            next-hop next-hop-address;
        }
    }
}
traceoptions {

```

```
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
  }  
  tunnel-devices [ mt-fpc/pic/port ];  
}
```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>family</b> statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Enable PIM on the routing device.  The statements are explained separately.
<b>Default</b>	PIM is disabled on the routing device.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## prefix-list (PIM RPF Selection)

<b>Syntax</b>	<pre> prefix-list <i>prefix-list-addresses</i> {   source <i>source-address</i> {     next-hop <i>next-hop-address</i>;   }   wildcard-source {     next-hop <i>next-hop-address</i>;   } } </pre>
<b>Hierarchy Level</b>	<p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> wildcard-source],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> source <i>source-address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> wildcard-source]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	(Optional) Configure a list of prefixes (addresses) for multiple PIM groups.
<b>Options</b>	<p><b><i>prefix-list-addresses</i></b>—List of prefixes (addresses) for multiple PIM groups.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>view-level—To view this statement in the configuration.</p> <p>control-level—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>

## priority (Bootstrap)

---

<b>Syntax</b>	<code>priority <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)], [edit protocols pim <b>rp bootstrap</b> (inet   inet6)], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp bootstrap</b> (inet   inet6)]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the routing device's likelihood to be elected as the bootstrap router.
<b>Options</b>	<b>number</b> —Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority. <b>Range:</b> 0 through a 32-bit number <b>Default:</b> 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring PIM Bootstrap Properties for IPv4</i></li><li>• <i>Configuring PIM Bootstrap Properties for IPv4 or IPv6</i></li><li>• <a href="#">bootstrap-priority on page 4915</a></li></ul>



## priority (PIM Interfaces)

<b>Syntax</b>	<code>priority <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ], [edit protocols pim interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the routing device's likelihood to be elected as the designated router.
<b>Options</b>	<b><i>number</i></b> —Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority. <b>Range:</b> 0 through 4294967295 <b>Default:</b> 1 (Each routing device has an equal probability of becoming the DR.)
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Interface Priority for PIM Designated Router Selection on page 4784</a></li> </ul>

## priority (PIM RPs)

---

<b>Syntax</b>	<code>priority <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols pim rp bidirectional address <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols pim</code> <code>rp bidirectional address <i>address</i>],</code> <code>[edit protocols pim rp bidirectional address <i>address</i>],</code> <code>[edit protocols pim <b>rp local family</b> (inet   inet6)],</code> <code>[edit routing-instances <i>instance-name</i> protocols pim rp bidirectional address <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp local family</b> (inet   inet6)]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Support for bidirectional RP addresses introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 13.3 for the PTX5000 router.
<b>Description</b>	For PIM-SM, configure this routing device's priority for becoming an RP.  For bidirectional PIM, configure this RP address' priority for becoming an RP.  The bootstrap router uses this field when selecting the list of candidate rendezvous points to send in the bootstrap message. A smaller number increases the likelihood that the routing device or RP address becomes the RP. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.
<b>Options</b>	<b><i>number</i></b> —Priority for becoming an RP. A lower value corresponds to a higher priority. <b>Range:</b> 0 through 255 <b>Default:</b> 1
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Local PIM RPs on page 4795</a> in the <i>Multicast Protocols Feature Guide for Routing Devices</i></li><li>• <i>Example: Configuring Bidirectional PIM</i></li></ul>

## propagation-delay

<b>Syntax</b>	<code>propagation-delay <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit protocols pim],</code> <code>[edit protocols pim interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols pim],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>pim interface <i>interface-name</i>]</code>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.1.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Set a delay for implementing a PIM prune message on the upstream routing device on a multicast network for which join suppression has been enabled. The routing device waits for the prune pending period to detect whether a join message is currently being suppressed by another routing device.
<b>Options</b>	<p><b><i>milliseconds</i></b>—Interval for the prune pending timer, which is the sum of the <b>propagation-delay</b> value and the <b>override-interval</b> value.</p> <p><b>Range:</b> 250 through 2000 milliseconds</p> <p><b>Default:</b> 500 milliseconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Enabling Join Suppression on page 4791</a></li> <li>• <a href="#">override-interval on page 4945</a></li> <li>• <a href="#">reset-tracking-bit on page 4954</a></li> </ul>

## reset-tracking-bit

---

<b>Syntax</b>	reset-tracking-bit;
<b>Hierarchy Level</b>	[edit protocols pim], [edit protocols pim interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Change the value of a tracking bit (T-bit) field in the LAN prune delay hello option from the default of 1 to 0, which enables join suppression for a multicast interface. When the network starts receiving multiple identical join messages, join suppression triggers a random timer with a value of 66 through 84 milliseconds ( $1.1 \times$ periodic through $1.4 \times$ periodic, where periodic is 60 seconds). This creates an interval during which no identical join messages are sent. Eventually, only one of the identical messages is sent. Join suppression is triggered each time identical messages are sent for the same join.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Enabling Join Suppression on page 4791</a></li><li>• <a href="#">override-interval on page 4945</a></li><li>• <a href="#">propagation-delay on page 4953</a></li></ul>

## rib-group (Protocols PIM)

---

<b>Syntax</b>	<pre> rib-group {     inet <i>group-name</i>;     inet6 <i>group-name</i>; } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Associate a routing table group with PIM.
<b>Options</b>	<i>table-name</i> —Name of the routing table. The name must be one that you defined with the <b>rib-groups</b> statement at the <b>[edit routing-options]</b> hierarchy level.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring a Dedicated PIM RPF Routing Table</i></li> </ul>

## rp

---

```
Syntax  rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bidirectional {
        address address {
            group-ranges {
                destination-ip-prefix </prefix-length>;
            }
            hold-time seconds;
            priority number;
        }
    }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix </prefix-length>;
        }
        maximum-rps limit;
    }
    group-rp-mapping {
        family (inet | inet6) {
            log-interval seconds;
            maximum limit;
            threshold value;
        }
    }
    log-interval seconds;
    maximum limit;
    threshold value;
}
local {
    family (inet | inet6) {
        disable;
        address address;
        anycast-pim {
            local-address address;
            address address <forward-msdp-sa>;
            rp-set {
            }
        }
    }
}
```

```

    }
    group-ranges {
        destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    override;
    priority number;
}
}
register-limit {
    family (inet | inet6) {
        log-interval seconds;
        maximum limit;
        threshold value;
    }
}
log-interval seconds;
maximum limit;
threshold value;
}
}
register-probe-time register-probe-time;
}
rp-register-policy [ policy-names ];
static {
    address address {
        override;
        version version;
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
    }
}
}
}

```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	If you do not include the <b>rp</b> statement, the routing device can never become the RP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding PIM Sparse Mode on page 4741](#)

---

## rp-register-policy

---

<b>Syntax</b>	rp-register-policy [ <i>policy-names</i> ];
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b> ], [edit protocols pim <b>rp</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Apply one or more policies to control incoming PIM register messages.
<b>Options</b>	<i>policy-names</i> —Name of one or more import policies.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Register Message Filters on a PIM RP and DR on page 4813</a></li><li>• <a href="#">dr-register-policy on page 4918</a></li></ul>



## rp-set

<b>Syntax</b>	<pre>rp-set {   address address &lt;forward-msdp-sa&gt;; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>local family</b> (inet   inet6) <b>anycast-pim</b>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring PIM Anycast With or Without MSDP on page 4799</a></li> </ul>

## rpf-selection

---

**Syntax**    `rpf-selection {  
              group group-address {  
                  source source-address {  
                      next-hop next-hop-address;  
                  }  
                  wildcard-source {  
                      next-hop next-hop-address;  
                  }  
              }  
              prefix-list prefix-list-addresses {  
                  source source-address {  
                      next-hop next-hop-address;  
                  }  
                  wildcard-source {  
                      next-hop next-hop-address;  
                  }  
              }  
          }`

**Hierarchy Level**    [edit routing-instances *routing-instance-name* protocols pim]

**Release Information**    Statement introduced in JUNOS Release 10.4.  
Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description**    Configure the PIM RPF next-hop neighbor for a specific group and source for a VRF routing instance.

The remaining statements are explained separately.

**Default**    If you omit the **rpf-selection** statement, PIM RPF checks typically choose the best path determined by the unicast protocol for all multicast flows.

**Options**    **source-address**—Specific source address for the PIM group.

**Required Privilege Level**    view-level—To view this statement in the configuration.  
control-level—To add this statement to the configuration.

**Related Documentation**    • *Example: Configuring PIM RPF Selection*

## source (PIM RPF Selection)

---

<b>Syntax</b>	<code>source source-address {     next-hop next-hop-address; }</code>
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the source address for the PIM group.
<b>Options</b>	<b>source-address</b> —Specific source address for the PIM group.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring PIM RPF Selection</i></li> </ul>

## spt-threshold

---

<b>Syntax</b>	spt-threshold { infinfinity [ <i>policy-names</i> ]; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim], [edit protocols pim], [edit routing-instances <i>routing-instance-name</i> protocols pim]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the PIM SPT Threshold Policy on page 4818</a></li></ul>

## static (Protocols PIM)

<b>Syntax</b>	<pre>static {   address address {     group-ranges {       destination-ip-prefix&lt;/prefix-length&gt;;     }     override;     version version;   } }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>],</p> <p>[edit protocols pim <b>rp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure static RP addresses. The default static RP address is 224.0.0.0/4. To configure other addresses, include one or more <b>address</b> statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.</p> <p>For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Static PIM RP Address on the Non-RP Routing Device on page 4797</a></li> </ul>

## threshold (PIM BFD Detection Time)

---

<b>Syntax</b>	<code>threshold <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection detection-time]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.



**NOTE:** The threshold value must be equal to or greater than the transmit interval.

The threshold time must be equal to or greater than the value specified in the [minimum-interval](#) or the [minimum-receive-interval](#) statement.

---

<b>Options</b>	<i>milliseconds</i> —Value for the detection time adaptation threshold. <b>Range:</b> 1 through 255,000
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li><li>• <a href="#">bfd-liveness-detection on page 4911</a></li><li>• <a href="#">detection-time on page 4916</a></li><li>• <a href="#">minimum-interval on page 4939</a></li><li>• <a href="#">minimum-receive-interval on page 4941</a></li></ul>

## threshold (PIM BFD Transmit Interval)

<b>Syntax</b>	<code>threshold <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection transmit-interval]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent.
<b>Options</b>	<i>milliseconds</i> —Value for the transmit interval adaptation threshold. <b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ )



**NOTE:** The threshold value specified in the `threshold` statement must be greater than the value specified in the `minimum-interval` statement for the `transmit-interval` statement.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for PIM on page 4821</a></li> <li>• <a href="#">bfd-liveness-detection on page 4911</a></li> </ul>

## transmit-interval (PIM BFD Liveness Detection)

---

Syntax	<pre>transmit-interval {     minimum-interval milliseconds;     threshold milliseconds; }</pre>
Hierarchy Level	[edit protocols pim interface <i>interface-name</i> bfd-liveness-detection], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches. Support for BFD authentication introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Specify the transmit interval for the <b>bfd-liveness-detection</b> statement. The negotiated transmit interval for a peer is the interval between the sending of BFD packets to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li><li>• <a href="#">bfd-liveness-detection on page 4911</a></li><li>• <a href="#">threshold on page 4965</a></li><li>• <a href="#">minimum-interval on page 4940</a></li><li>• <a href="#">minimum-receive-interval on page 4941</a></li></ul>



## traceoptions (Protocols PIM)

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim],          [edit protocols pim],          [edit routing-instances <i>routing-instance-name</i> protocols pim]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 9.0 for EX Series switches.          Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	<p>Configure PIM tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>
<b>Default</b>	The default PIM trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the <b>[edit routing-options]</b> hierarchy level.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>pim-log</b> file.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files  <b>Default:</b> 2 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>PIM Tracing Flags</b></p> <ul style="list-style-type: none"> <li>• <b>assert</b>—Assert messages</li> <li>• <b>bidirectional-df-election</b>—Bidirectional PIM designated-forwarder (DF) election events</li> </ul>

- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache
- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 0 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring PIM Trace Options on page 4780</a></li> <li>• <a href="#">Tracing DVMRP Protocol Traffic</a></li> <li>• <a href="#">Tracing MSDP Protocol Traffic on page 4877</a></li> <li>• <a href="#">Configuring PIM Trace Options on page 4780</a></li> </ul>
------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## version (BFD)

---

<b>Syntax</b>	version (0   1   automatic);
<b>Hierarchy Level</b>	[edit protocols piminterface <i>interface-name</i> <b>bfd-liveness-detection</b> ], [edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i> <b>bfd-liveness-detection</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the bidirectional forwarding detection (BFD) protocol version that you want to detect.
<b>Options</b>	Configure the BFD version to detect: <b>1</b> (BFD version 1) or <b>automatic</b> (autodetect the BFD version) <b>Default:</b> automatic
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring BFD for PIM on page 4821</a></li></ul>

## version (PIM)

<b>Syntax</b>	<code>version <i>version</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit protocols pim interface <i>interface-name</i>],</p> <p>[edit protocols pim <b>rp static address</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols pim <b>rp static address</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Specify the version of PIM.
<b>Options</b>	<p><b>version</b>—PIM version number.</p> <p><b>Range:</b> 1 or 2</p> <p><b>Default:</b> PIMv1 for rendezvous point (RP) mode (at the [edit protocols pim rp static address <i>address</i>] hierarchy level). PIMv2 for interface mode (at the [edit protocols pim interface <i>interface-name</i>] hierarchy level).</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling PIM Sparse Mode on page 4786</a></li> <li>• <a href="#">Configuring PIM Dense Mode Properties</a></li> <li>• <a href="#">Configuring PIM Sparse-Dense Mode Properties</a></li> </ul>

## wildcard-source (PIM RPF Selection)

---

<b>Syntax</b>	wildcard-source { <a href="#">next-hop</a> <i>next-hop-address</i> ; }
<b>Hierarchy Level</b>	[edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection group <i>group-address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols pim rpf-selection prefix-list <i>prefix-list-addresses</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Use a wildcard for the multicast source instead of (or in addition to) a specific multicast source.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring PIM RPF Selection</a></li></ul>

## IGMP Configuration Statements

---

- [accounting \(Protocols IGMP\)](#) on page 4973
- [accounting \(Protocols IGMP Interface\)](#) on page 4973
- [asm-override-ssm](#) on page 4974
- [disable \(Protocols IGMP\)](#) on page 4974
- [exclude \(Protocols IGMP\)](#) on page 4975
- [group \(Protocols IGMP\)](#) on page 4976
- [group-count \(Protocols IGMP\)](#) on page 4977
- [group-increment \(Protocols IGMP\)](#) on page 4977
- [group-limit \(IGMP\)](#) on page 4978
- [group-policy \(Protocols IGMP\)](#) on page 4979
- [igmp](#) on page 4980
- [immediate-leave \(Protocols IGMP\)](#) on page 4982
- [interface \(Protocols IGMP\)](#) on page 4983
- [maximum-transmit-rate \(Protocols IGMP\)](#) on page 4984
- [oif-map \(IGMP Interface\)](#) on page 4984
- [passive \(IGMP\)](#) on page 4985
- [promiscuous-mode \(Protocols IGMP\)](#) on page 4986

- [query-interval \(Protocols IGMP\) on page 4986](#)
- [query-last-member-interval \(Protocols IGMP\) on page 4987](#)
- [query-response-interval \(Protocols IGMP\) on page 4988](#)
- [robust-count \(Protocols IGMP\) on page 4989](#)
- [source \(Protocols IGMP\) on page 4990](#)
- [source-count \(Protocols IGMP\) on page 4991](#)
- [source-increment \(Protocols IGMP\) on page 4991](#)
- [static \(Protocols IGMP\) on page 4992](#)
- [traceoptions \(Protocols IGMP\) on page 4993](#)
- [version \(Protocols IGMP\) on page 4995](#)

## accounting (Protocols IGMP)

<b>Syntax</b>	accounting;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ], [edit protocols <a href="#">igmp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Enable the collection of IGMP join and leave event statistics on the system.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Recording IGMP Join and Leave Events on page 4843</a></li> </ul>

## accounting (Protocols IGMP Interface)

<b>Syntax</b>	(accounting   no-accounting);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface</a> <i>interface-name</i> ], [edit protocols <a href="#">igmp interface</a> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Enable or disable the collection of IGMP join and leave event statistics for an interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Recording IGMP Join and Leave Events on page 4843</a></li> </ul>

## asm-override-ssm

---

<b>Syntax</b>	asm-override-ssm;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895</a></li></ul>

## disable (Protocols IGMP)

---

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ], [edit protocols <b>igmp interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Disable IGMP on the system.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Disabling IGMP on page 4847</a></li></ul>



---


## exclude (Protocols IGMP)

---

<b>Syntax</b>	exclude;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3.
<b>Description</b>	Configure the static group to operate in exclude mode. In exclude mode all sources except the address configured are accepted for the group. If this statement is not included, the group operates in include mode.
<b>Required Privilege Level</b>	view-level—To view this statement in the configuration. control-level—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li></ul>

## group (Protocols IGMP)

---

Syntax	<pre>group <i>multicast-group-address</i> {     exclude;     group-count <i>number</i>;     group-increment <i>increment</i>;     source <i>ip-address</i> {         source-count <i>number</i>;         source-increment <i>increment</i>;     } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp interface interface-name static</a> ], [edit protocols <a href="#">igmp interface interface-name static</a> ]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.
<hr/>	
<div> <b>NOTE:</b> You must specify a unique address for each group.</div> <hr/>	
The remaining statements are explained separately.	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li></ul>

## group-count (Protocols IGMP)

<b>Syntax</b>	<code>group-count <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the number of static groups to be created.
<b>Options</b>	<i>number</i> —Number of static groups. <b>Range:</b> 1 through 512
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li> </ul>

## group-increment (Protocols IGMP)

<b>Syntax</b>	<code>group-increment <i>increment</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group</b> <i>mcast-group-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the number of times the address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.
<b>Options</b>	<i>increment</i> —Number of times the address should be incremented. <b>Default:</b> 0.0.0.1 <b>Range:</b> 0.0.0.1 through 255.255.255.255
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li> </ul>

## group-limit (IGMP)

---

<b>Syntax</b>	<code>group-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ], [edit protocols <b>igmp interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.</p> <p>To confirm the configured group limit on the interface, use the <b>show igmp interface</b> command.</p>
<b>Default</b>	By default, there is no limit to the number of multicast groups that can join the interface.
<b>Options</b>	<b>limit</b> —group limit value for the interface. <b>Range:</b> 1 through 32767
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Limiting the Number of IGMP Multicast Group Joins on Logical Interfaces on page 4844</a></li><li>• <i>group-threshold</i></li><li>• <i>log-interval</i></li></ul>

---

## group-policy (Protocols IGMP)

---

<b>Syntax</b>	<code>group-policy [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> ], [edit protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	When this statement is enabled on a router running IGMP version 2 (IGMPv2) or version 3 (IGMPv3), after the router receives an IGMP report, the router compares the group against the specified group policy and performs the action configured in that policy (for example, rejects the report).
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Filtering Unwanted IGMP Reports at the IGMP Interface Level on page 4832</a></li></ul>

## igmp

---

**Syntax**    `igmp {`  
              `accounting;`  
              `interface interface-name {`  
                  `disable;`  
                  `(accounting | no-accounting);`  
                  `group-limit limit;`  
                  `group-policy [ policy-names ];`  
                  `group-threshold`  
                  `immediate-leave;`  
                  `log-interval`  
                  `oif-map map-name;`  
                  `passive;`  
                  `promiscuous-mode;`  
                  `ssm-map ssm-map-name;`  
                  `ssm-map-policy ssm-map-policy-name;`  
                  `static {`  
                      `group multicast-group-address {`  
                          `exclude;`  
                          `group-count number;`  
                          `group-increment increment;`  
                          `source ip-address {`  
                              `source-count number;`  
                              `source-increment increment;`  
                          `}`  
                      `}`  
                  `}`  
                  `version version;`  
              `}`  
              `query-interval seconds;`  
              `query-last-member-interval seconds;`  
              `query-response-interval seconds;`  
              `robust-count number;`  
              `traceoptions {`  
                  `file filename <files number> <size size> <world-readable | no-world-readable>;`  
                  `flag flag <flag-modifier> <disable>;`  
              `}`  
              `}`

**Hierarchy Level**    `[edit logical-systems logical-system-name protocols],`  
                          `[edit protocols]`

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.  
                              Statement introduced in Junos OS Release 12.3R2 for EX Series switches.


**Description**    Enable IGMP on the router or switch. IGMP must be enabled for the router or switch to receive multicast packets.

The remaining statements are explained separately.

<b>Default</b>	IGMP is disabled on the router or switch. IGMP is automatically enabled on all broadcast interfaces when you configure Protocol Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP).
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP on page 4828</a></li></ul>

## immediate-leave (Protocols IGMP)

---

<b>Syntax</b>	<code>immediate-leave;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>Starting in Junos OS Release 9.3, both IGMP version 2 and IGMP version 3 do host tracking when the <b>immediate-leave</b> statement is configured. This means that the multicast group leaves only when the last host leaves. The routing device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p>
	<div> <b>NOTE:</b> Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.</div>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.



- Related Documentation**
- [Specifying Immediate-Leave Host Removal for IGMP on page 4831](#)

## interface (Protocols IGMP)

<b>Syntax</b>	<pre> interface <i>interface-name</i> {     disable;     (accounting   no-accounting);     group-limit <i>limit</i>;     group-policy [ <i>policy-names</i> ];     immediate-leave;     oif-map <i>map-name</i>;     passive;     promiscuous-mode;     ssm-map <i>ssm-map-name</i>;     ssm-map-policy <i>ssm-map-policy-name</i>;     static {         group <i>mcast-group-address</i> {             exclude;             group-count <i>number</i>;             group-increment <i>increment</i>;             source <i>ip-address</i> {                 source-count <i>number</i>;                 source-increment <i>increment</i>;             }         }     }     version <i>version</i>; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b>],</p> <p>[edit protocols <b>igmp</b>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Enable IGMP on an interface and configure interface-specific properties.</p>
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP on page 4828</a></li> </ul>

## maximum-transmit-rate (Protocols IGMP)

---


<b>Syntax</b>	maximum-transmit-rate <i>packets-per-second</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols igmp], [edit protocols igmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Limit the transmission rate of IGMP packets
<b>Options</b>	<b>packets-per-second</b> —Maximum number of IGMP packets transmitted in one second by the routing device. <b>Range:</b> 1 through 10000 <b>Default:</b> 500 packets
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Limiting the Maximum IGMP Message Rate on page 4836</a></li></ul>

## oif-map (IGMP Interface)

---

<b>Syntax</b>	oif-map <i>map-name</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ], [edit protocols <b>igmp interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Associates an outgoing interface (OIF) map to the IGMP interface. The OIF map is a routing policy statement that can contain multiple terms.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast with Subscriber VLANs</a></li></ul>

## passive (IGMP)

<b>Syntax</b>	<code>passive &lt;allow-receive&gt; &lt;send-general-query&gt; &lt;send-group-query&gt;;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. <b>allow-receive</b> , <b>send-general-query</b> , and <b>send-group-query</b> options were added in Junos OS Release 10.0. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify that IGMP run on the interface and either not send and receive control traffic or selectively send and receive control traffic such as IGMP reports, queries, and leaves.
<div>  <p><b>NOTE:</b> You can selectively activate up to two out of the three available options for the <b>passive</b> statement while keeping the other functions passive (inactive). Activating all three options would be equivalent to not using the <b>passive</b> statement.</p> </div>	
<b>Options</b>	<p><b>allow-receive</b>—Enables IGMP to receive control traffic on the interface.</p> <p><b>send-general-query</b>—Enables IGMP to send general queries on the interface.</p> <p><b>send-group-query</b>—Enables IGMP to send group-specific and group-source-specific queries on the interface.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Multicast with Subscriber VLANs</i></li> <li>• <a href="#">Enabling IGMP on page 4828</a></li> </ul>

## promiscuous-mode (Protocols IGMP)

---

<b>Syntax</b>	<code>promiscuous-mode;</code>
<b>Hierarchy Level</b>	[edit dynamic-profiles <i>profile-name</i> protocols igmp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.2 for dynamic profiles. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify that the interface accepts IGMP reports from hosts on any subnetwork. Note that when enabling promiscuous-mode, all routing devices on the ethernet segment must be configured with the promiscuous mode statement. Otherwise, only the interface configured with lowest IPv4 address acts as the querier for IGMP for this Ethernet segment.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Dynamic DHCP Client Access to a Multicast Network</a></li><li>• <a href="#">Accepting IGMP Messages from Remote Subnetworks on page 4833</a></li></ul>

## query-interval (Protocols IGMP)

---

<b>Syntax</b>	<code>query-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ], [edit protocols <b>igmp</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify how often the querier routing device sends general host-query messages.
<b>Options</b>	<i>seconds</i> —Time interval. <b>Range:</b> 1 through 1024 <b>Default:</b> 125 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Host-Query Message Interval on page 4830</a></li><li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 4987</a></li><li>• <a href="#">query-response-interval (Protocols IGMP) on page 4988</a></li></ul>

## query-last-member-interval (Protocols IGMP)

---

<b>Syntax</b>	query-last-member-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ], [edit protocols <a href="#">igmp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify how often the querier routing device sends group-specific query messages.
<b>Options</b>	<b>seconds</b> —Time interval, in fractions of a second or seconds. <b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 999999 <b>Default:</b> 1 second
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Modifying the IGMP Last-Member Query Interval on page 4831</a></li> <li>• <a href="#">query-interval (Protocols IGMP) on page 4986</a></li> <li>• <a href="#">query-response-interval (Protocols IGMP) on page 4988</a></li> </ul>

## query-response-interval (Protocols IGMP)

---

<b>Syntax</b>	<code>query-response-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ], [edit protocols <a href="#">igmp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify how long the querier routing device waits to receive a response to a host-query message from a host.
<b>Options</b>	<b><i>seconds</i></b> —The query response interval must be less than the query interval. <b>Range:</b> 1 through 1024 <b>Default:</b> 10 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Modifying the IGMP Query Response Interval on page 4834</a></li><li>• <a href="#">query-interval (Protocols IGMP) on page 4986</a></li><li>• <a href="#">query-last-member-interval (Protocols IGMP) on page 4987</a></li></ul>

## robust-count (Protocols IGMP)

---

<b>Syntax</b>	<code>robust-count <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">igmp</a> ], [edit protocols <a href="#">igmp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.
<b>Options</b>	<i>number</i> —Robustness variable. <b>Range:</b> 2 through 10 <b>Default:</b> 2
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Modifying the IGMP Robustness Variable on page 4835</a></li> </ul>

## source (Protocols IGMP)

---

<b>Syntax</b>	<code>source <i>ip-address</i> {     <i>source-count</i> <i>number</i>;     <i>source-increment</i> <i>increment</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ], [edit protocols <b>igmp</b> <b>interface</b> <i>interface-name</i> <b>static group</b> <i>multicast-group-address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.
<b>Options</b>	<i>ip-address</i> —IPv4 unicast address.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li></ul>



## source-count (Protocols IGMP)

<b>Syntax</b>	<code>source-count <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the number of multicast source addresses that should be accepted for each static group created.
<b>Options</b>	<i>number</i> —Number of source addresses. <b>Default:</b> 1 <b>Range:</b> 1 through 1024
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li> </ul>

## source-increment (Protocols IGMP)

<b>Syntax</b>	<code>source-increment <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> <b>static group multicast-group-address</b> <i>source</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the number of times the multicast source address should be incremented for each static group created. The increment is specified in dotted decimal notation similar to an IPv4 address.
<b>Options</b>	<i>increment</i> —Number of times the source address should be incremented. <b>Default:</b> 0.0.0.1 <b>Range:</b> 0.0.0.1 through 255.255.255.255
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling IGMP Static Group Membership on page 4836</a></li> </ul>

## static (Protocols IGMP)

---

**Syntax**    static {  
              group *multicast-group-address* {  
                  exclude;  
                  group-count *number*;  
                  group-increment *increment*;  
                  source *ip-address* {  
                      source-count *number*;  
                      source-increment *increment*;  
                  }  
              }  
          }  
      }

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols **igmp** interface *interface-name*],  
                          [edit protocols **igmp** interface *interface-name*]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Test multicast forwarding on an interface without a receiver host.

The **static** statement simulates IGMP joins on a routing device statically on an interface without any IGMP hosts. It is supported for both IGMPv2 and IGMPv3 joins. This statement is especially useful for testing multicast forwarding on an interface without a receiver host.



**NOTE:** To prevent joining too many groups accidentally, the **static** statement is not supported with the **interface all** statement.

---

The remaining statements are explained separately.

**Required Privilege Level**    routing and trace—To view this statement in the configuration.  
                                  routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**    • [Enabling IGMP Static Group Membership on page 4836](#)

## traceoptions (Protocols IGMP)

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> ], [edit protocols <b>igmp</b> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure IGMP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>To trace the paths of multicast packets, use the <b>mtrace</b> command.</p>
<b>Default</b>	The default IGMP trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the file <b>igmp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 2 files</p> <p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>IGMP Tracing Flags</b></p> <ul style="list-style-type: none"> <li><b>leave</b>—Leave group messages (for IGMP version 2 only).</li> <li><b>mtrace</b>—Mtrace packets. Use the <b>mtrace</b> command to troubleshoot the software.</li> <li><b>packets</b>—All IGMP packets.</li> </ul>

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow users to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Tracing IGMP Protocol Traffic on page 4845](#)

## version (Protocols IGMP)

<b>Syntax</b>	<code>version <i>version</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp interface</b> <i>interface-name</i> ], [edit protocols <b>igmp interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the version of IGMP.
<b>Options</b>	<b>version</b> —IGMP version number. <b>Range:</b> 1, 2, or 3 <b>Default:</b> IGMP version 2
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Changing the IGMP Version on page 4829</a></li> </ul>

## IGMP Snooping Configuration Statements

- [data-forwarding on page 4997](#)
- [disable \(IGMP Snooping\) on page 4997](#)
- [group \(IGMP Snooping\) on page 4998](#)
- [group-limit \(IGMP and MLD Snooping\) on page 4999](#)
- [groups \(Multicast VLAN Registration\) on page 5000](#)
- [host-only-interface on page 5001](#)
- [igmp-querier on page 5002](#)
- [igmp-snooping on page 5003](#)

- [immediate-leave \(Bridge Domains\) on page 5004](#)
- [install \(Multicast VLAN Registration\) on page 5005](#)
- [interface \(Bridge Domains\) on page 5006](#)
- [interface \(IGMP Snooping\) on page 5007](#)
- [l2-querier on page 5007](#)
- [multicast-router-interface \(IGMP Snooping\) on page 5008](#)
- [proxy \(Multicast VLAN Registration\) on page 5008](#)
- [query-interval \(Bridge Domains\) on page 5009](#)
- [query-last-member-interval \(Bridge Domains\) on page 5010](#)
- [query-response-interval \(Bridge Domains\) on page 5011](#)
- [receiver on page 5012](#)
- [robust-count \(IGMP Snooping\) on page 5012](#)
- [source \(Multicast VLAN Registration\) on page 5013](#)
- [source-address on page 5013](#)
- [source-address \(IGMP Querier\) on page 5014](#)
- [source-vlans on page 5014](#)
- [static \(IGMP Snooping\) on page 5015](#)
- [traceoptions \(IGMP Snooping\) on page 5016](#)
- [version \(IGMP Snooping\) on page 5018](#)
- [vlan \(IGMP Snooping\) on page 5019](#)

## data-forwarding

<b>Syntax</b>	<pre>data-forwarding {   receiver {     source-vlans <i>vlan-list</i>;     install;   }   source {     groups <i>group-prefix</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	<p>Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMP version 2 (IGMPv2) mode.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Example: Configuring Multicast VLAN Registration</i></li> <li><i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li> </ul>

## disable (IGMP Snooping)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Disable IGMP snooping on all interfaces in a VLAN.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Example: Configuring IGMP Snooping on page 4849</a></li> <li><a href="#">Configuring IGMP Snooping on page 4848</a></li> </ul>

## group (IGMP Snooping)

---

<b>Syntax</b>	<code>group <i>ip-address</i>;</code>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> <b>interface</b> <i>interface-name</i> <b>static</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a static multicast group using a valid IP multicast address.
<b>Default</b>	None.
<b>Options</b>	<i>ip-address</i> —IP address of the multicast group receiving data on an interface.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping vlans on page 5103</a></li><li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li></ul>



## group-limit (IGMP and MLD Snooping)

<b>Syntax</b>	<code>group-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <a href="#">interface</a> <i>interface-name</i>]</p> <p>[edit protocols igmp-snooping vlan <a href="#">interface</a>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>
<b>Description</b>	Configure a limit for the number of multicast groups (or [S,G] channels in IGMPv3) allowed on an interface. After this limit is reached, new reports are ignored and all related flows are not flooded on the interface.
<b>Default</b>	By default, there is no limit to the number of multicast groups joining an interface.
<b>Options</b>	<i>limit</i> —a 32-bit number for the limit on the interface.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>

## groups (Multicast VLAN Registration)

---

<b>Syntax</b>	<code>groups group-prefix;</code>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding source]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Specify the IP address range of the multicast VLAN (MVLAN) source interfaces.
<b>Default</b>	Disabled
<b>Options</b>	<i>group-prefix</i> —IP address range of the source group. Each MVLAN must have exactly one <b>groups</b> statement. If there are multiple MVLANs on the switch, their group ranges must be unique.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul>

## host-only-interface

<b>Syntax</b>	host-only-interface;
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <b>interface</b> <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <b>interface</b> <i>interface-name</i>]</p> <p>[edit protocols igmp-snooping vlan <b>interface</b>]</p> <p>[edit protocols mld-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>]</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols mld-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Support at the [edit protocols mld-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>] and the [edit routing-instances <i>routing-instance-name</i> protocols mld-snooping vlan <i>vlan-id</i> interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 13.3 for EX Series switches.</p>
<b>Description</b>	Configure an interface as a host-facing interface. IGMP and MLD queries received on these interfaces are dropped.
<b>Default</b>	The interface can either be a host-side or multicast-routing device interface.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> <li>• <i>multicast-router-interface</i></li> </ul>

## igmp-querier

---

<b>Syntax</b>	igmp-querier <a href="#">source-address</a> <i>source address</i> ;
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure a QFabric Node device to be an IGMP querier. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that Node is always the IGMP querier on the network.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li><li>• <a href="#">show igmp-snooping vlans on page 5103</a></li><li>• <a href="#">show configuration protocols igmp on page 5083</a></li></ul>

## igmp-snooping

<b>Syntax</b>	<pre> igmp-snooping {   vlan <i>vlan-id</i> {     immediate-leave;     interface <i>interface-name</i> {       group-limit <i>limit</i>;       host-only-interface;       immediate-leave;       multicast-router-interface;       static {         group <i>ip-address</i> {           source <i>ip-address</i>;         }       }     }   }   l2-querier {     source-address <i>ip-address</i>;   }   proxy {     source-address <i>ip-address</i>;   }   query-interval <i>seconds</i>;   query-last-member-interval <i>seconds</i>;   query-response-interval <i>seconds</i>;   robust-count <i>number</i>;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt;;   } } </pre>
<b>Hierarchy Level</b>	[edit protocols]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Enable IGMP snooping on the router or switch.
<b>Default</b>	IGMP snooping is disabled on the router or switch.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding IGMP Snooping</i></li> <li>• <i>IGMP Snooping in MC-LAG Active-Active on MX Series Routers Overview</i></li> </ul>

## immediate-leave (Bridge Domains)

<b>Syntax</b>	<code>immediate-leave;</code>
<b>Hierarchy Level</b>	<pre>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <i>interface</i> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <i>interface</i> <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping <i>interface</i> <i>interface-name</i>] [edit protocols igmp-snooping vlan <i>interface</i>]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>
<b>Description</b>	<p>The immediate leave setting is useful for minimizing the leave latency of IGMP memberships. When this setting is enabled, the routing device leaves the multicast group immediately after the last host leaves the multicast group.</p> <p>The immediate-leave setting enables host tracking, meaning that the device keeps track of the hosts that send join messages. This allows IGMP to determine when the last host sends a leave message for the multicast group.</p> <p>When the immediate leave setting is enabled, the device removes an interface from the forwarding-table entry without first sending IGMP group-specific queries to the interface. The interface is pruned from the multicast tree for the multicast group specified in the IGMP leave message. The immediate leave setting ensures optimal bandwidth management for hosts on a switched network, even when multiple multicast groups are being used simultaneously.</p> <p>When immediate leave is disabled and one host sends a leave group message, the routing device first sends a group query to determine if another receiver responds. If no receiver responds, the routing device removes all hosts on the interface from the multicast group. Immediate leave is disabled by default for both IGMP version 2 and IGMP version 3.</p>



**NOTE:** Although host tracking is enabled for IGMPv2 and MLDv1 when you enable immediate leave, use immediate leave with these versions only when there is one host on the interface. The reason is that IGMPv2 and MLDv1 use a report suppression mechanism whereby only one host on an interface sends a group join report in response to a membership query. The other interested hosts suppress their reports. The purpose of this mechanism is to avoid a flood of reports for the same group. But it also interferes with host tracking, because the routing device only knows about the one interested host and does not know about the others.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring IGMP Snooping*

---

## install (Multicast VLAN Registration)

---

**Syntax** install;

**Hierarchy Level** [edit protocols igmp-snooping vlan (all | *vlan-name*) data-forwarding receiver]

**Release Information** Statement introduced in Junos OS Release 9.6 for EX Series switches.  
Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Install forwarding entries in the multicast receiver VLAN. By default, the multicast VLAN (MVLAN) installs forwarding entries for MVLAN groups only.

**Default** Disabled

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Multicast VLAN Registration*
- *Configuring Multicast VLAN Registration (CLI Procedure)*

## interface (Bridge Domains)

---

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     <i>group-limit limit</i>;     <i>host-only-interface</i>;     multicast-router-interface;     static {         group <i>ip-address</i> {             source <i>ip-address</i>;         }     } }</pre>
<b>Hierarchy Level</b>	[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> ], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping], [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols vlan <i>vlan-id</i> igmp-snooping] [edit protocols igmp-snooping vlan],
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
<b>Description</b>	Enable IGMP snooping on an interface and configure interface-specific properties.
<b>Options</b>	<p><i>interface-name</i>—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify <b>all</b>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring IGMP Snooping</i></li></ul>



## interface (IGMP Snooping)

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     multicast-router-interface;     static {         group <i>ip-address</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Enable IGMP snooping on an interface and configure interface-specific properties.</p> <p>The remaining statements are explained separately.</p>
<b>Options</b>	<i>interface-name</i> —Name of the interface.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5103</a></li> </ul>

## l2-querier

<b>Syntax</b>	<pre>l2-querier {     source-address <i>ip-address</i>; }</pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan],
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Configure the switch to be an IGMP querier. Use the <b>source-address</b> statement to configure the source address to use for IGMP snooping queries.
<b>Options</b>	<p><b>seconds</b>—Time interval.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 125 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	

## multicast-router-interface (IGMP Snooping)

---

<b>Syntax</b>	multicast-router-interface;
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> <b>interface</b> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure an interface to forward IGMP messages to multicast routers.
<b>Default</b>	Disabled. If this statement is disabled, the interface drops IGMP messages it receives.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping vlans on page 5103</a></li><li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li></ul>

## proxy (Multicast VLAN Registration)

---

<b>Syntax</b>	proxy source-address <i>ip-address</i> ;
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Specify that the VLAN operate in proxy mode. The proxy option is supported only for a VLAN acting as a data-forwarding source.
<b>Default</b>	Disabled
<b>Options</b>	<b>source-address <i>ip-address</i></b> —IP address of the source VLAN to act as proxy.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure)</a></li></ul>

## query-interval (Bridge Domains)

<b>Syntax</b>	<code>query-interval seconds;</code>
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface interface-name</a>]</p> <p>[edit protocols igmp-snooping <a href="#">vlan</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>
<b>Description</b>	Configure the interval for host-query message timeouts.
<b>Options</b>	<p><b>seconds</b>—Time interval.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 125 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping</a></li> <li>• <a href="#">query-last-member-interval (Bridge Domains) on page 5010</a></li> <li>• <a href="#">query-response-interval (Bridge Domains) on page 5011</a></li> </ul>

## query-last-member-interval (Bridge Domains)

---

<b>Syntax</b>	<code>query-last-member-interval seconds;</code>
<b>Hierarchy Level</b>	<code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code> <code>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface</a> <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface</a> <i>interface-name</i>]</code> <code>[edit protocols igmp-snooping <a href="#">vlan</a>],</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
<b>Description</b>	Configure the interval for group-specific query timeouts.
<b>Options</b>	<b>seconds</b> —Time interval, in fractions of a second or seconds. <b>Range:</b> 0.1 through 0.9, then in 1-second intervals 1 through 1024 <b>Default:</b> 1 second
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping</a></li><li>• <a href="#">query-interval on page 5009</a></li><li>• <a href="#">query-response-interval on page 5011</a></li></ul>

## query-response-interval (Bridge Domains)

<b>Syntax</b>	query-response-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	<p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</p> <p>[edit bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping <a href="#">interface interface-name</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> protocols igmp-snooping vlan <i>vlan-id</i> <a href="#">interface interface-name</a>]</p> <p>[edit protocols igmp-snooping <a href="#">vlan</a>],</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p>
<b>Description</b>	Specify how long to wait to receive a response to a specific query message from a host.
<b>Options</b>	<p><b><i>seconds</i></b>—Time interval. This interval should be less than the host-query interval.</p> <p><b>Range:</b> 1 through 1024</p> <p><b>Default:</b> 10 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping</a></li> <li>• <a href="#">query-interval (Bridge Domains) on page 5009</a></li> <li>• <a href="#">query-last-member-interval (Bridge Domains) on page 5010</a></li> </ul>

## receiver

---

<b>Syntax</b>	<pre>receiver {     source-vlans <i>vlan-list</i>;     install; }</pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).  The remaining statements are explained separately.
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Multicast VLAN Registration</i></li><li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li></ul>

## robust-count (IGMP Snooping)

---

<b>Syntax</b>	<pre>robust-count <i>number</i>;</pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. Configure the length of each interval using the <b>query-interval</b> statement.
<b>Default</b>	2 intervals
<b>Options</b>	<b><i>number</i></b> —Number of intervals the switch waits before timing out a multicast group. <b>Range:</b> 2 through 10
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li><li>• <a href="#">show igmp-snooping vlans on page 5103</a></li></ul>

## source (Multicast VLAN Registration)

<b>Syntax</b>	source { groups group-prefix; }
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   vlan-name) data-forwarding]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure a VLAN to be a multicast source VLAN (MVLAN).  The remaining statement is explained separately.
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Multicast VLAN Registration</i></li> <li>• <i>Configuring Multicast VLAN Registration (CLI Procedure)</i></li> </ul>

## source-address

<b>Syntax</b>	source-address ip-address;
<b>Hierarchy Level</b>	[edit bridge-domains bridge-domain-name protocols igmp-snooping proxy], [edit bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id proxy], [edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping proxy], [edit routing-instances routing-instance-name bridge-domains bridge-domain-name protocols igmp-snooping vlan vlan-id proxy]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
<b>Description</b>	Specify the IP address to use as the source for IGMP snooping reports in proxy mode. Reports are sent with address 0.0.0.0 as the source address unless there is a source address configured. You can also use this statement to configure the source address to use for IGMP snooping queries.
<b>Options</b>	ip-address—IP address to use as the source for proxy-mode IGMP snooping reports.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring IGMP Snooping</i></li> </ul>

## source-address (IGMP Querier)

---

<b>Syntax</b>	<code>src-address source address;</code>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> <b>igmp-querier</b> ] [edit protocols igmp-snooping vlan <i>vlan-name</i> <b>l2-querier</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure the address that the switch uses as the source address in the IGMP queries that it sends. If there are any multicast routers on the same local network, make sure the source address for the IGMP querier is greater (a higher number) than the IP addresses for those routers on the network. This ensures that switch is always the IGMP querier on the network.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li><li>• <a href="#">show igmp-snooping vlans on page 5103</a></li><li>• <a href="#">show configuration protocols igmp on page 5083</a></li></ul>

## source-vlans

---

<b>Syntax</b>	<code>source-vlans vlan-list;</code>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   <i>vlan-name</i> ) data-forwarding receiver]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.6 for EX Series switches. Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Specify a list of multicast VLANs (MVLANS) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANS must be in proxy mode or none of them can be in proxy mode.
<b>Default</b>	Disabled
<b>Options</b>	<i>vlan-list</i> —Names of the MVLANS.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multicast VLAN Registration</a></li><li>• <a href="#">Configuring Multicast VLAN Registration (CLI Procedure)</a></li></ul>



---

## static (IGMP Snooping)

---

<b>Syntax</b>	<pre>static {     group ip-address; }</pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan <i>vlan-name</i> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Statically define multicast groups on an interface.</p> <p>The remaining statement is explained separately.</p>
<b>Default</b>	No multicast groups are statically defined.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li><li>• <a href="#">show igmp-snooping vlans on page 5103</a></li></ul>

## traceoptions (IGMP Snooping)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;size <i>size</i>&gt; &lt;replace&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> (detail   disable   receive   send); }</pre>
<b>Hierarchy Level</b>	For platforms without ELS:  [edit protocols igmp-snooping]  For platforms with ELS:  [edit protocols igmp-snooping vlan]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define tracing operations for IGMP snooping.
<b>Default</b>	The <b>traceoptions</b> feature is disabled by default.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached (<b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li><li>• <b>general</b>—Trace general IGMP snooping protocol events.</li><li>• <b>krt</b>—Trace communication over routing sockets.</li><li>• <b>nexthop</b>— Trace next-hop related events.</li><li>• <b>normal</b>—Trace normal IGMP snooping protocol events.</li><li>• <b>packets</b>—Trace all IGMP packets.</li><li>• <b>policy</b>—Trace policy processing.</li><li>• <b>query</b>—Trace IGMP membership query messages.</li><li>• <b>report</b>—Trace membership report messages.</li></ul>

- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.
- **vlan**—Trace VLAN related events.

**no-stamp**—(Optional) Do not time stamp trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**size size** —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option. Use **xk** to specify KB, **xm** to specify MB, or **xg** to specify gigabytes.

**Range:** 10 KB through 1 gigabytes

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring IGMP Snooping on page 4849</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> </ul>

## version (IGMP Snooping)

---

<b>Syntax</b>	<code>version number;</code>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping vlan (all   <i>vlan-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify the IGMP version for the IGMP general query that the switch sends to hosts when an interface comes up. The configured IGMP version affects only the version of the general queries sent by a switch. It does not affect the version of IGMP messages that the switch can snoop. For example, If the switch is configured for IGMP version 1 (IGMPv1), it can snoop IGMPv2 and IGMPv3 messages.
<b>Default</b>	If you do not configure the <b>version</b> statement, the default is IGMPv2.
<b>Options</b>	<b>version</b> —IGMP version number. <b>Range:</b> 1 and 2.



**NOTE:** IGMP v3 snooping is not supported.

---

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring IGMP Snooping (CLI Procedure)</a></li><li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li></ul>

## vlan (IGMP Snooping)

<b>Syntax</b>	<pre> vlan <i>vlan-id</i> {     immediate-leave;     interface <i>interface-name</i> {         group-limit <i>limit</i>;         host-only-interface;         multicast-router-interface;         static {             group <i>mcast-group-address</i> {                 source <i>ip-address</i>;             }         }     }     proxy {         source-address <i>ip-address</i>;     }     query-interval <i>seconds</i>;     query-last-member-interval <i>seconds</i>;     query-response-interval <i>seconds</i>;     robust-count <i>number</i>; } </pre>
<b>Hierarchy Level</b>	[edit protocols igmp-snooping ],
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 13.2 for the QFX series.
<b>Description</b>	Configure IGMP snooping parameters for a particular VLAN.
<b>Default</b>	By default, IGMP snooping options apply to all VLANs.
<b>Options</b>	<p><i>vlan-id</i>—Apply the parameters to this VLAN.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring VLAN-Specific IGMP Snooping Parameters on page 4849</a></li> <li>• <a href="#">igmp-snooping on page 5003</a></li> </ul>

## MSDP Configuration Statements

- [active-source-limit on page 5021](#)
- [authentication-key on page 5022](#)
- [data-encapsulation on page 5023](#)
- [default-peer on page 5024](#)

- [disable \(Protocols MSDP\) on page 5025](#)
- [export \(Protocols MSDP\) on page 5026](#)
- [group \(Protocols MSDP\) on page 5027](#)
- [import \(Protocols MSDP\) on page 5028](#)
- [local-address \(Protocols MSDP\) on page 5029](#)
- [maximum \(MSDP Active Source Messages\) on page 5030](#)
- [mode \(Protocols MSDP\) on page 5031](#)
- [msdp on page 5032](#)
- [peer \(Protocols MSDP\) on page 5034](#)
- [rib-group \(Protocols MSDP\) on page 5035](#)
- [source \(Protocols MSDP\) on page 5036](#)
- [threshold \(MSDP Active Source Messages\) on page 5037](#)
- [traceoptions \(Protocols MSDP\) on page 5038](#)

## active-source-limit

<b>Syntax</b>	<pre>active-source-limit {     log-interval <i>seconds</i>;     log-warning <i>value</i>;     <b>maximum</b> <i>number</i>;     <b>threshold</b> <i>number</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp source</b> <i>ip-address/prefix-length</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp source</b> <i>ip-address/prefix-length</i>], [edit protocols <b>msdp</b>], [edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit protocols <b>msdp peer</b> <i>address</i>], [edit protocols <b>msdp source</b> <i>ip-address/prefix-length</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp source</b> <i>ip-address/prefix-length</i>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Limit the number of active source messages the routing device accepts.
<b>Default</b>	If you do not include this statement, the router accepts any number of MSDP active source messages.
<b>Options</b>	The options are explained separately.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li> </ul>

## authentication-key

---

<b>Syntax</b>	<code>authentication-key peer-key;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems logical-system-name protocols msdp group group-name peer address],</code> <code>[edit logical-systems logical-system-name protocols msdp peer address],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols</code> <code>msdp group group-name peer address],</code> <code>[edit logical-systems logical-system-name routing-instances routing-instance-name protocols</code> <code>msdp peer address],</code> <code>[edit protocols msdp group group-name peer address],</code> <code>[edit protocols msdp peer address],</code> <code>[edit routing-instances routing-instance-name protocols msdp group group-name peer</code> <code>address],</code> <code>[edit routing-instances routing-instance-name protocols msdp peer address]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Associate a Message Digest 5 (MD5) signature option authentication key with an MSDP peering session.
<b>Default</b>	If you do not include this statement, the routing device accepts any valid MSDP messages from the peer address.
<b>Options</b>	<b>peer-key</b> —MD5 authentication key. The peer key can be a text string up to 16 letters and digits long. Strings can include any ASCII characters with the exception of (, ), &, and [. If you include spaces in an MSDP authentication key, enclose all characters in quotation marks (" ").
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP in a Routing Instance</i></li></ul>



## data-encapsulation

---

<b>Syntax</b>	data-encapsulation (disable   enable);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ], [edit protocols <a href="#">msdp</a> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure a rendezvous point (RP) using MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP source-active messages.
<b>Default</b>	If you do not include this statement, the RP encapsulates multicast data.
<b>Options</b>	<b>disable</b> —(Optional) Do not use MSDP data encapsulation. <b>enable</b> —Use MSDP data encapsulation. <b>Default:</b> <b>enable</b>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li> </ul>

## default-peer

---

<b>Syntax</b>	default-peer;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>], [edit protocols <b>msdp</b>], [edit protocols <b>msdp group</b> <i>group-name</i>], [edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit protocols <b>msdp peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Establish this peer as the default MSDP peer and accept source-active messages from the peer without the usual peer-reverse-path-forwarding (peer-RPF) check.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li></ul>

## disable (Protocols MSDP)

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Explicitly disable MSDP.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Disabling MSDP</li> </ul>

## export (Protocols MSDP)

---

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>],</code> <code>[edit protocols <a href="#">msdp</a>],</code> <code>[edit protocols <a href="#">msdp group</a> <i>group-name</i>],</code> <code>[edit protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code> <code>[edit protocols <a href="#">msdp peer</a> <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp group</a> <i>group-name</i> <a href="#">peer</a> <i>address</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp peer</a> <i>address</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Apply one or more policies to routes being exported from the routing table into MSDP.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring MSDP in a Routing Instance</i></li><li>• <a href="#">import on page 5028</a></li></ul>

## group (Protocols MSDP)

```
Syntax  group group-name {
        disable;
        export [ policy-names ];
        import [ policy-names ];
        local-address address;
        mode (mesh-group | standard);
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
        peer address; {
            disable;
            active-source-limit {
                maximum number;
                threshold number;
            }
            authentication-key peer-key;
            default-peer;
            export [ policy-names ];
            import [ policy-names ];
            local-address address;
            traceoptions {
                file filename <files number> <size size> <world-readable | no-world-readable>;
                flag flag <flag-modifier> <disable>;
            }
        }
    }
```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [msdp](#)],  
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
[msdp](#)],  
 [edit protocols [msdp](#)],  
 [edit routing-instances *routing-instance-name* protocols [msdp](#)]

**Release Information** Statement introduced before Junos OS Release 7.4.  
 Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Define an MSDP peer group. MSDP peers within groups share common tracing options, if present and not overridden for an individual peer with the [peer](#) statement. To configure multiple MSDP groups, include multiple **group** statements.

By default, the group's options are identical to the global MSDP options. To override the global options, include group-specific options within the **group** statement.

The group must contain at least one peer.

**Options** *group-name*—Name of the MSDP group.

The remaining statements are explained separately.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*

---

## import (Protocols MSDP)

---

**Syntax** `import [ policy-names ];`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols [msdp](#)],  
[edit logical-systems *logical-system-name* protocols [msdp group](#) *group-name*],  
[edit logical-systems *logical-system-name* protocols [msdp group](#) *group-name* [peer](#) *address*],  
[edit logical-systems *logical-system-name* protocols [msdp peer](#) *address*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [msdp](#)],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [msdp group](#) *group-name*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [msdp group](#) *group-name* [peer](#) *address*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols [msdp peer](#) *address*],  
[edit protocols [msdp](#)],  
[edit protocols [msdp group](#) *group-name*],  
[edit protocols [msdp group](#) *group-name* [peer](#) *address*],  
[edit protocols [msdp peer](#) *address*],  
[edit routing-instances *routing-instance-name* protocols [msdp](#)],  
[edit routing-instances *routing-instance-name* protocols [msdp group](#) *group-name*],  
[edit routing-instances *routing-instance-name* protocols [msdp group](#) *group-name* [peer](#) *address*],  
[edit routing-instances *routing-instance-name* protocols [msdp peer](#) *address*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description** Apply one or more policies to routes being imported into the routing table from MSDP.

**Options** *policy-names*—Name of one or more policies.

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring MSDP in a Routing Instance*
- [export on page 5026](#)

## local-address (Protocols MSDP)

<b>Syntax</b>	<code>local-address address;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Configure the local end of an MSDP session. You must configure at least one peer for MSDP to function. When configuring a peer, you must include this statement. This address is used to accept incoming connections to the peer and to establish connections to the remote peer.
<b>Options</b>	<b>address</b> —IP address of the local end of the connection.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>

## maximum (MSDP Active Source Messages)

---

<b>Syntax</b>	<code>maximum <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ], [edit protocols <a href="#">msdp active-source-limit</a> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure the maximum number of MSDP active source messages the router accepts.
<b>Options</b>	<i>number</i> —Maximum number of active source messages. <b>Range:</b> 1 through 1,000,000 <b>Default:</b> 25,000
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li><li>• <a href="#">threshold (MSDP Active Source Messages) on page 5037</a></li></ul>



## mode (Protocols MSDP)

<b>Syntax</b>	mode (mesh-group   standard);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ], [edit protocols <b>msdp group</b> <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure groups of peers in a full mesh topology to limit excessive flooding of source-active messages to neighboring peers. The default flooding mode is <b>standard</b> .
<b>Default</b>	If you do not include this statement, default flooding is applied.
<b>Options</b>	<b>mesh-group</b> —Group of peers that are mesh group members.  <b>standard</b> —Use standard MSDP source-active flooding rules. <b>Default:</b> standard
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li> </ul>

## msdp

```

Syntax  msdp {
        disable;
        active-source-limit {
            log-interval seconds;
            log-warning value;
            maximum number;
            threshold number;
        }
        data-encapsulation (disable | enable);
        export [ policy-names ];
        group group-name {
            ... group-configuration ...
        }
        hold-time seconds;
        import [ policy-names ];
        local-address address;
        keep-alive seconds;
        peer address {
            ... peer-configuration ...
        }
        rib-group group-name;
        source ip-prefix </prefix-length> {
            active-source-limit {
                maximum number;
                threshold number;
            }
        }
        sa-hold-time seconds;
        traceoptions {
            file filename <files number> <size size> <world-readable | no-world-readable>;
            flag flag <flag-modifier> <disable>;
        }
        group group-name {
            disable;
            export [ policy-names ];
            import [ policy-names ];
            local-address address;
            mode (mesh-group | standard);
            peer address {
                ... same statements as at the [edit protocols msdp peer address] hierarchy level shown
                just following ...
            }
            traceoptions {
                file filename <files number> <size size> <world-readable | no-world-readable>;
                flag flag <flag-modifier> <disable>;
            }
        }
        peer address {
            disable;
            active-source-limit {
                maximum number;
                threshold number;
            }
        }
    }

```

```

    }
    authentication-key peer-key;
    default-peer;
    export [ policy-names ];
    import [ policy-names ];
    local-address address;
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
}

```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.4 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Enable MSDP on the router or switch. You must also configure at least one peer for MSDP to function.
<b>Default</b>	MSDP is disabled on the router or switch.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring MSDP in a Routing Instance</i></li> </ul>

## peer (Protocols MSDP)

---

**Syntax**    `peer address {  
              disable;  
              active-source-limit {  
                  maximum number;  
                  threshold number;  
              }  
              authentication-key peer-key;  
              default-peer;  
              export [ policy-names ];  
              import [ policy-names ];  
              local-address address;  
              traceoptions {  
                  file filename <files number> <size size> <world-readable | no-world-readable>;  
                  flag flag <flag-modifier> <disable>;  
              }  
          }`

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols **msdp**],  
                          [edit logical-systems *logical-system-name* protocols **msdp group** *group-name*],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                              **msdp**],  
                          [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols  
                              **msdp group** *group-name*],  
                          [edit protocols **msdp**],  
                          [edit protocols **msdp group** *group-name*],  
                          [edit routing-instances *routing-instance-name* protocols **msdp**],  
                          [edit routing-instances *routing-instance-name* protocols **msdp group** *group-name*]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 12.1 for the QFX Series.

**Description**    Define an MSDP peering relationship. An MSDP routing device must know which routing devices are its peers. You define the peer relationships explicitly by configuring the neighboring routing devices that are the MSDP peers of the local routing device. After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. To configure multiple MSDP peers, include multiple **peer** statements.

By default, the peer's options are identical to the global or group-level MSDP options. To override the global or group-level options, include peer-specific options within the **peer (Protocols MSDP)** statement.

At least one peer must be configured for MSDP to function. You must configure **address** and **local-address**.

**Options**    **address**—Name of the MSDP peer.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • *Example: Configuring MSDP in a Routing Instance*

## rib-group (Protocols MSDP)

<b>Syntax</b>	<code>rib-group group-name;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp</a> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ], [edit protocols <a href="#">msdp</a> ], [edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Associate a routing table group with MSDP.
<b>Options</b>	<b>group-name</b> —Name of the routing table group. The name must be one that you defined with the <b>rib-groups</b> statement at the [edit routing-options] hierarchy level.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <i>Example: Configuring MSDP in a Routing Instance</i>

## source (Protocols MSDP)

---

<b>Syntax</b>	<pre>source ip-address &lt;/prefix-length&gt; {     active-source-limit {         maximum number;         threshold number;     } }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems logical-system-name protocols msdp], [edit logical-systems logical-system-name routing-instances routing-instance-name protocols msdp], [edit protocols msdp], [edit routing-instances routing-instance-name protocols msdp]</pre>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Limit the number of active source messages the routing device accepts from sources in this address range.
<b>Default</b>	If you do not include this statement, the routing device accepts any number of MSDP active source messages.
<b>Options</b>	The other statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li></ul>

## threshold (MSDP Active Source Messages)

<b>Syntax</b>	<code>threshold <i>number</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <a href="#">msdp active-source-limit</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a>],</p> <p>[edit protocols <a href="#">msdp active-source-limit</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <a href="#">msdp active-source-limit</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure the random early detection (RED) threshold for MSDP active source messages. This number must be less than the configured or default maximum.</p>
<b>Options</b>	<p><i>number</i>—RED threshold for active source messages.</p> <p><b>Range:</b> 1 through 1,000,000</p> <p><b>Default:</b> 24,000</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring MSDP with Active Source Limits and Mesh Groups on page 4880</a></li> <li>• <a href="#">maximum (MSDP Active Source Messages) on page 5030</a></li> </ul>

## traceoptions (Protocols MSDP)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp</b>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit protocols <b>msdp peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp</b>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp group</b> <i>group-name</i> <b>peer</b> <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols <b>msdp peer</b> <i>address</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure MSDP tracing options.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>
<b>Default</b>	<p>The default MSDP trace options are those inherited from the routing protocol's <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.</p>
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file</b> <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place tracing output in the <b>msdp-log</b> file.</p> <p><b>files</b> <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p>



If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

**Range:** 2 through 1000 files

**Default:** 2 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

#### MSDP Tracing Flags

- **keepalive**—Keepalive messages
- **packets**—All MSDP packets
- **route**—MSDP changes to the routing table
- **source-active**—Source-active packets
- **source-active-request**—Source-active request packets
- **source-active-response**—Source-active response packets

#### Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

**flag-*modifier***—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-stamp**—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

**no-world-readable**—(Optional) Do not allow any user to read the log file.

**replace**—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing MSDP Protocol Traffic on page 4877</a></li></ul>
------------------------------	--------------------------------------------------------------------------------------------------------------

---

## Source-Specific Multicast Configuration Statements

---

- [asm-override-ssm on page 5041](#)
- [policy \(SSM Maps\) on page 5042](#)
- [ssm-groups on page 5043](#)
- [ssm-map \(Protocols IGMP\) on page 5044](#)
- [ssm-map \(Routing Options Multicast\) on page 5044](#)
- [ssm-map-policy \(IGMP\) on page 5045](#)

## asm-override-ssm

---

<b>Syntax</b>	asm-override-ssm;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Enable the routing device to accept any-source multicast join messages (*G) for group addresses that are within the default or configured range of source-specific multicast groups.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895</a></li> </ul>

## policy (SSM Maps)

---

<b>Syntax</b>	<code>policy [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i> ], [edit routing-instances <i>routing-instance-name</i> routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i> ], [edit routing-options multicast <b>ssm-map</b> <i>ssm-map-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Apply one or more policies to an SSM map.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies for SSM mapping.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To view this statement in the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 4893</a></li></ul>

## ssm-groups

<b>Syntax</b>	<code>ssm-groups [ <i>ip-addresses</i> ];</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
<b>Description</b>	<p>Configure source-specific multicast (SSM) groups.</p> <p>By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the <b>ssm-groups</b> statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the <b>ssm-groups</b> statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.</p> <p>IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.</p>
<b>Options</b>	<i>ip-addresses</i> —List of one or more additional SSM group addresses separated by a space.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Source-Specific Multicast Groups with Any-Source Override on page 4895</a></li> </ul>

## ssm-map (Protocols IGMP)

---

<b>Syntax</b>	<code>ssm-map ssm-map-name;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Apply an SSM map to an IGMP interface.
<b>Options</b>	<i>ssm-map-name</i> —Name of SSM map.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 4893</a></li></ul>

## ssm-map (Routing Options Multicast)

---

<b>Syntax</b>	<code>ssm-map ssm-map-name {     <b>policy</b> [ <i>policy-names</i> ];     <b>source</b> [ <i>addresses</i> ]; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Configure SSM mapping.
<b>Options</b>	<i>ssm-map-name</i> —Name of the SSM map.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Mapping on page 4893</a></li></ul>

---

## ssm-map-policy (IGMP)

---

<b>Syntax</b>	<code>ssm-map-policy <i>ssm-map-policy-name</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols <b>igmp</b> interface <i>interface-name</i> ], [edit protocols <b>igmp</b> interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Apply an SSM map policy to an IGMP interface.
<b>Options</b>	<i>ssm-map-policy-name</i> —Name of SSM map policy.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SSM Maps for Different Groups to Different Sources on page 4899</a></li></ul>





# Administration

- [Routine Monitoring on page 5047](#)
- [Monitoring Commands for Multicast Protocols on page 5048](#)

## Routine Monitoring

- [Monitoring IGMP Snooping on page 5047](#)
- [Verifying the IGMP Snooping Group Timeout Value on page 5048](#)

### Monitoring IGMP Snooping

- Purpose** Use the monitoring feature to view status and information about the IGMP snooping configuration.
- Action** To display IGMP snooping details in the CLI, enter the following commands:
  - `show igmp-snooping vlans`
  - `show igmp-snooping statistics`
  - `show igmp-snooping route`
  - `show igmp-snooping membership`
- Meaning** [Table 392 on page 5047](#) summarizes the IGMP snooping details displayed.

Table 392: Summary of IGMP Snooping Output Fields

Field	Values
IGMP Snooping Monitor	
VLAN	VLAN for which IGMP snooping is enabled.
Interfaces	Interface connected to a multicast router.
Groups	Number of the multicast groups learned by the VLAN.
MRouters	Multicast router.
Receivers	Multicast receiver.

Table 392: Summary of IGMP Snooping Output Fields (*continued*)

Field	Values
IGMP Route Information	
VLAN	VLAN for which IGMP snooping is enabled.
Next-Hop	Next hop assigned by the switch after performing the route lookup.
Group	Multicast groups learned by the VLAN.

- Related Documentation**
- [IGMP Snooping Overview on page 4761](#)
  - [Example: Configuring IGMP Snooping on page 4849](#)
  - [Configuring IGMP Snooping on page 4848](#)
  - [Changing the IGMP Snooping Group Timeout Value](#)

## Verifying the IGMP Snooping Group Timeout Value

**Purpose** Verify that the IGMP snooping group timeout value has been changed correctly from its default value.

**Action** Display the IGMP snooping membership information, which contains the group timeout value that was derived from the IGMP configuration:

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.1
Receiver count: 1, Flags: <v2-hosts>
ge-0/0/15.0 Uptime: 00:00:05 timeout: 510
```

**Meaning** The IGMP snooping group timeout value determines how long a switch waits to receive an IGMP query from a multicast router before removing a multicast group from its multicast cache table. When you enable IGMP snooping, the default IGMP snooping group timeout value of 260 seconds is applied to all VLANs, which means that the switch waits 260 seconds to receive an IGMP query before removing a multicast group from its multicast cache table. You can change the timeout value by using the **robust-count** option.

- Related Documentation**
- [Changing the IGMP Snooping Group Timeout Value](#)

## Monitoring Commands for Multicast Protocols

- [clear igmp membership](#)
- [clear igmp-snooping membership](#)
- [clear igmp statistics](#)
- [clear igmp-snooping statistics](#)

- `clear msdp cache`
- `clear msdp statistics`
- `clear multicast bandwidth-admission`
- `clear multicast scope`
- `clear multicast sessions`
- `clear multicast statistics`
- `clear pim join`
- `clear pim register`
- `clear pim statistics`
- `mtrace`
- `mtrace from-source`
- `mtrace monitor`
- `mtrace to-gateway`
- `show configuration protocols igmp`
- `show igmp group`
- `show igmp interface`
- `show igmp statistics`
- `show igmp-snooping membership`
- `show igmp-snooping route`
- `show igmp-snooping statistics`
- `show igmp-snooping vlans`
- `show msdp`
- `show msdp source`
- `show msdp source-active`
- `show msdp statistics`
- `show multicast flow-map`
- `show multicast interface`
- `show multicast mrinfo`
- `show multicast next-hops`
- `show multicast pim-to-igmp-proxy`
- `show multicast pim-to-mld-proxy`
- `show multicast route`
- `show multicast rpf`
- `show multicast scope`
- `show multicast sessions`
- `show multicast usage`
- `show pim bootstrap`

- [show pim interfaces](#)
- [show pim join](#)
- [show pim neighbors](#)
- [show pim rps](#)
- [show pim source](#)
- [show pim statistics](#)
- [show system statistics igmp](#)
- [test msdp](#)

## clear igmp membership

<b>List of Syntax</b>	<a href="#">Syntax on page 5051</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5051</a>
<b>Syntax</b>	<pre>clear igmp membership &lt;group address-range&gt; &lt;interface interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>clear igmp membership &lt;group address-range&gt; &lt;interface interface-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Clear Internet Group Management Protocol (IGMP) group members.
<b>Options</b>	<p><b>none</b>—Clear all IGMP members on all interfaces and for all address ranges.</p> <p><b>group address-range</b>—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is <b>224.2/16</b>. If you omit the destination prefix length, the default is <b>/32</b>.</p> <p><b>interface interface-name</b>—(Optional) Clear all IGMP group members on an interface.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show igmp group on page 5085</a></li> <li>• <a href="#">show igmp interface on page 5089</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear igmp membership on page 5051</a> <a href="#">clear igmp membership interface on page 5052</a> <a href="#">clear igmp membership group on page 5053</a>
<b>Output Fields</b>	See <a href="#">show igmp group</a> for an explanation of output fields.

## Sample Output

### clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.2.127.253	10.1.128.1	186
so-0/0/0	224.2.127.254	10.1.128.1	186
so-0/0/0	239.255.255.255	10.1.128.1	187
so-0/0/0	224.1.127.255	10.1.128.1	188
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

```
user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

#### clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.2.127.253	10.1.128.1	210
so-0/0/0	239.255.255.255	10.1.128.1	210
so-0/0/0	224.1.127.255	10.1.128.1	215
so-0/0/0	224.2.127.254	10.1.128.1	216
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

```
user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

## clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.2.127.253	10.1.128.1	210
so-0/0/0	239.255.255.255	10.1.128.1	210
so-0/0/0	224.1.127.255	10.1.128.1	215
so-0/0/0	224.2.127.254	10.1.128.1	216
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

```
user@host> clear igmp membership group 239.225/16
```

```
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.1.127.255	10.1.128.1	231
so-0/0/0	224.2.127.254	10.1.128.1	233
so-0/0/0	224.2.127.253	10.1.128.1	236
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

## clear igmp-snooping membership

---

<b>Syntax</b>	<code>clear igmp-snooping membership</code> <code>&lt;vlan <i>vlan-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear IGMP snooping membership information.
<b>Options</b>	<code>vlan <i>vlan-name</i></code> —(Optional) Name of the VLAN.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping membership on page 5096</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear igmp-snooping membership on page 5054</a>

### Sample Output

clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```



## clear igmp statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 5055</a> <a href="#">Syntax (EX Series Switches) on page 5055</a>
<b>Syntax</b>	clear igmp statistics <interface <i>interface-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	clear igmp statistics <interface <i>interface-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear Internet Group Management Protocol (IGMP) statistics.
<b>Options</b>	<b>none</b> —Clear IGMP statistics on all interfaces.  <b>interface <i>interface-name</i></b> —(Optional) Clear IGMP statistics for the specified interface only.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show igmp statistics on page 5093</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear igmp statistics on page 5055</a>
<b>Output Fields</b>	See <a href="#">show igmp statistics</a> for an explanation of output fields.

## Sample Output

### clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report    0            0        0
DVMRP                   19784        35476     0
PIM V1                  18310         0        0
Cisco Trace             0            0        0
V2 Membership Report    0            0        0
Group Leave             0            0        0
Mtrace Response         0            0        0

```

Mtrace Request	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	0	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			0
IGMP v3 mode not applicable for SSM			0

IGMP Global Statistics

Bad Length	0
Bad Checksum	0
Bad Receive If	0
Rx non-local	1227

user@host> clear igmp statistics

user@host> show igmp statistics

IGMP packet statistics for all interfaces

IGMP Message type	Received	Sent	Rx errors
Membership Query	0	0	0
V1 Membership Report	0	0	0
DVMRP	0	0	0
PIM V1	0	0	0
Cisco Trace	0	0	0
V2 Membership Report	0	0	0
Group Leave	0	0	0
Mtrace Response	0	0	0
Mtrace Request	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	0	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			0
IGMP v3 mode not applicable for SSM			0
IGMP Global Statistics			
Bad Length	0		
Bad Checksum	0		
Bad Receive If	0		
Rx non-local	0		

## clear igmp-snooping statistics

---

<b>Syntax</b>	<code>clear igmp-snooping statistics</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear IGMP snooping statistics.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show igmp-snooping statistics on page 5101</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear igmp-snooping statistics on page 5057</a>

### Sample Output

#### clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

## clear msdp cache

---

<b>Syntax</b>	<code>clear msdp cache</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;peer <i>peer-address</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.
<b>Options</b>	<b>none</b> —Clear entries in the MSDP source-active cache for all instances, logical systems, and peers.  <b>instance <i>instance-name</i></b> —(Optional) Clear entries for a specific MSDP instance.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>peer <i>peer-address</i></b> —(Optional) Clear the MSDP source-active cache entries learned from a specific peer.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show msdp source-active on page 5109</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear msdp cache on page 5058</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear msdp cache

```
user@host> clear msdp cache
```

## clear msdp statistics

---

<b>Syntax</b>	clear msdp statistics <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )> <peer <i>peer-address</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Clear Multicast Source Discovery Protocol (MSDP) peer statistics.
<b>Options</b>	<p><b>none</b>—Clear MSDP statistics for all peers.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Clear the statistics for the specified peer.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show msdp statistics on page 5112</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear msdp statistics on page 5059</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear msdp statistics

```
user@host> clear msdp statistics
```

## clear multicast bandwidth-admission

---

<b>Syntax</b>	<pre>clear multicast bandwidth-admission &lt;group <i>group-address</i>&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;source <i>source-address</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Reapply IP multicast bandwidth admissions.
<b>Options</b>	<p><b>none</b>—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p><b>group <i>group-address</i></b>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p><b>inet</b>—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p><b>inet6</b>—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none"><li>• If the interface is congested, and it was admitted previously, it is removed.</li><li>• If the interface was rejected previously, the <b>clear multicast bandwidth-admission</b> command enables the interface to be admitted as long as enough bandwidth exists on the interface.</li><li>• If you do not specify an interface, issuing the <b>clear multicast bandwidth-admission</b> command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.</li></ul> <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p><b>source <i>source-address</i></b>—(Optional) Use with the <b>group</b> option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p>
<b>Required Privilege Level</b>	clear

**Related Documentation** • [show multicast interface on page 5118](#)

**List of Sample Output** [clear multicast bandwidth-admission on page 5061](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear multicast bandwidth-admission](#)

```
user@host> clear multicast bandwidth-admission
```

## clear multicast scope

---

<b>List of Syntax</b>	<a href="#">Syntax on page 5062</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5062</a>
<b>Syntax</b>	<pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>clear multicast scope &lt;inet   inet6&gt; &lt;interface <i>interface-name</i>&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear IP multicast scope statistics.
<b>Options</b>	<p><b>none</b>—(Same as <b>logical-system all</b>) Clear multicast scope statistics.</p> <p><b>inet</b>—(Optional) Clear multicast scope statistics for IPv4 family addresses.</p> <p><b>inet6</b>—(Optional) Clear multicast scope statistics for IPv6 family addresses.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear multicast scope statistics on a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show multicast scope on page 5140</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear multicast scope on page 5062</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear multicast scope

```
user@host> clear multicast scope
```



## clear multicast sessions

<b>List of Syntax</b>	<a href="#">Syntax on page 5063</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5063</a>
<b>Syntax</b>	clear multicast sessions <logical-system (all   <i>logical-system-name</i> )> < <i>regular-expression</i> >
<b>Syntax (EX Series Switch and the QFX Series)</b>	clear multicast sessions < <i>regular-expression</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear IP multicast sessions.
<b>Options</b>	<p><b>none</b>—(Same as <b>logical-system all</b>) Clear multicast sessions.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>regular-expression</i></b>—(Optional) Clear only multicast sessions that contain the specified regular expression.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show multicast sessions on page 5142</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear multicast sessions on page 5063</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear multicast sessions

```
user@host> clear multicast sessions
```

## clear multicast statistics

---

<b>List of Syntax</b>	<a href="#">Syntax on page 5064</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5064</a>
<b>Syntax</b>	<pre>clear multicast statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>clear multicast statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear IP multicast statistics.
<b>Options</b>	<p><b>none</b>—Clear multicast statistics for all supported address families on all interfaces.</p> <p><b>inet</b>—(Optional) Clear multicast statistics for IPv4 family addresses.</p> <p><b>inet6</b>—(Optional) Clear multicast statistics for IPv6 family addresses.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear multicast statistics for the specified instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear multicast statistics on a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><a href="#">show multicast statistics</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear multicast statistics on page 5064</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear multicast statistics

```
user@host> clear multicast statistics
```

## clear pim join

**List of Syntax**    [Syntax on page 5065](#)  
                           [Syntax \(EX Series Switch and the QFX Series\) on page 5065](#)

**Syntax**    clear pim join  
                   <group-address>  
                   <bidirectional | dense | sparse>  
                   <exact>  
                   <inet | inet6>  
                   <instance *instance-name*>  
                   <logical-system (all | *logical-system-name*)>  
                   <rp *ip-address/prefix* | source *ip-address/prefix*>  
                   <sg | star-g>

**Syntax (EX Series Switch and the QFX Series)**    clear pim join  
                                                                   <group-address>  
                                                                   <dense | sparse>  
                                                                   <exact>  
                                                                   <inet | inet6>  
                                                                   <instance *instance-name*>  
                                                                   <rp *ip-address/prefix* | source *ip-address/prefix*>  
                                                                   <sg | star-g>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   **inet6** and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.3 for the QFX Series.  
                                   Multiple new filter options introduced in Junos OS Release 13.2.

**Description**    Clear the Protocol Independent Multicast (PIM) join and prune states.

**Options**    **none**—Clear the PIM join and prune states for all groups, family addresses, and instances.

**group-address**—(Optional) Clear the PIM join and prune states for a group address.

**bidirectional | dense | sparse**—(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.

**exact**—(Optional) Clear only the group that exactly matches the specified group address.

**inet | inet6**—(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Clear the entries for a specific PIM-enabled routing instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**rp *ip-address/prefix* | source *ip-address/prefix***—(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

**sg | star-g**—(Optional) Clear PIM (S,G) or (\*,G) entries.

**Additional Information** The **clear pim join** command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

**Required Privilege Level** clear

**Related Documentation**

- [show pim join on page 5153](#)

**List of Sample Output**

- [clear pim join on page 5066](#)
- [clear pim join inet6 on page 5066](#)
- [clear pim join inet6 star-g on page 5066](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear pim join

```
user@host> clear pim join
Cleared 8 Join/Prune states
```

### clear pim join inet6

```
user@host> clear pim join inet6
Cleared 4 Join/Prune states
```

### clear pim join inet6 star-g

```
user@host> clear pim join inet6 star-g
Cleared 1 Join/Prune states
```

## clear pim register

<b>List of Syntax</b>	<a href="#">Syntax on page 5067</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5067</a> <a href="#">Syntax (PTX Series) on page 5067</a>
<b>Syntax</b>	clear pim register <inet   inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switch and the QFX Series)</b>	clear pim register <inet   inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
<b>Syntax (PTX Series)</b>	clear pim register <inet   inet6> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Clear Protocol Independent Multicast (PIM) register message counters.
<b>Options</b>	<p><b>none</b>—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM register message counters for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Additional Information</b>	The <b>clear pim register</b> command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.
<b>Required Privilege Level</b>	clear

**Related Documentation** • [show pim statistics on page 5189](#)

**List of Sample Output** [clear pim register on page 5068](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear pim register

```
user@host> clear pim register
```

## clear pim statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 5069</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5069</a>
<b>Syntax</b>	<pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>clear pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Clear Protocol Independent Multicast (PIM) statistics.
<b>Options</b>	<p><b>none</b>—Clear PIM statistics for all family addresses, instances, and interfaces.</p> <p><b>inet   inet6</b>—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Clear PIM statistics for a specific interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Additional Information</b>	The <b>clear pim statistics</b> command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show pim statistics on page 5189</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear pim statistics on page 5070</a>
<b>Output Fields</b>	See <a href="#">show pim statistics</a> for an explanation of output fields.

## Sample Output

### clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0      0
Register               0             0      0
Register Stop          0             0      0
Join Prune             0             0      0
Bootstrap              0             0      0
Assert                 0             0      0
Graft                  0             0      0
Graft Ack              0             0      0
Candidate RP           0             0      0
V1 Query               2111          4222      0
V1 Register            0             0      0
V1 Register Stop       0             0      0
V1 Join Prune          14200         13115      0
V1 RP Reachability     0             0      0
V1 Assert              0             0      0
V1 Graft               0             0      0
V1 Graft Ack           0             0      0
PIM statistics summary for all interfaces:
Unknown type           0
V1 Unknown type        0
Unknown Version        0
Neighbor unknown      0
Bad Length             0
Bad Checksum           0
Bad Receive If         0
Rx Intf disabled       2007
Rx V1 Require V2       0
Rx Register not RP     0
RP Filtered Source     0
Unknown Reg Stop       0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...
```

```
user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0      0
Register               0             0      0
Register Stop          0             0      0
Join Prune             0             0      0
Bootstrap              0             0      0
Assert                 0             0      0
Graft                  0             0      0
Graft Ack              0             0      0
Candidate RP           0             0      0
V1 Query               1             0      0
V1 Register            0             0      0
...
```





## mtrace

<b>Syntax</b>	<code>mtrace source</code> <logical-system <i>logical-system-name</i> > <routing-instance <i>routing-instance-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 9.5 for SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices. Command introduced in Junos OS Release 11.3 for the QFX Series. Command introduced in Junos OS Release 12.3 for the PTX Series.
<b>Description</b>	Display trace information about an IP multicast path.
<b>Options</b>	<b>source</b> —Source hostname or address.  <b>logical-system (<i>logical-system-name</i>)</b> —(Optional) Perform this operation on a logical system.  <b>routing-instance <i>routing-instance-name</i></b> —(Optional) Trace a particular routing instance.
<b>Additional Information</b>	The <b>mtrace</b> command for multicast traffic is similar to the <b>traceroute</b> command used for unicast traffic. Unlike <b>traceroute</b> , <b>mtrace</b> traces traffic backwards, from the receiver to the source.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">mtrace source on page 5074</a>
<b>Output Fields</b>	<a href="#">Table 393 on page 5072</a> describes the output fields for the <b>mtrace</b> command. Output fields are listed in the approximate order in which they appear.

**Table 393: mtrace Output Fields**

Field Name	Field Description
<b>Mtrace from</b>	IP address of the receiver.
<b>to</b>	IP address of the source.
<b>via group</b>	IP address of the multicast group (if any).
<b>Querying full reverse path</b>	Indicates the full reverse path query has begun.
<b><i>number-of-hops</i></b>	Number of hops from the source to the named router or switch.
<b><i>router-name</i></b>	Name of the router or switch for this hop.
<b><i>address</i></b>	Address of the router or switch for this hop.

Table 393: mtrace Output Fields (*continued*)

Field Name	Field Description
<i>protocol</i>	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

## Sample Output

### mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

## mtrace from-source

**Syntax** `mtrace from-source source source`  
`<brief | detail>`  
`<extra-hops extra-hops>`  
`<group group>`  
`<interval interval>`  
`<loop>`  
`<max-hops max-hops>`  
`<max-queries max-queries>`  
`<multicast-response | unicast-response>`  
`<no-resolve>`  
`<no-router-alert>`  
`<response response>`  
`<routing-instance routing-instance-name>`  
`<ttl tll>`  
`<wait-time wait-time>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, Junos OS returns additional information, such as packet rates and losses.

**Options** **brief | detail**—(Optional) Display the specified level of output.

**extra-hops *extra-hops***—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

**group *group***—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

**interval *interval***—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

**loop**—(Optional) Loop indefinitely, displaying rate and loss statistics.

**max-hops *max-hops***—(Optional) Maximum hops to trace toward the source. The range of values is **0** through **255**. The default value is **32** hops.

**max-queries *max-queries***—(Optional) Maximum number of query attempts for any hop. The range of values is 1 through **32**. The default is **3**.

**multicast-response**—(Optional) Always request the response using multicast.

**no-resolve**—(Optional) Do not attempt to display addresses symbolically.

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance** *routing-instance-name*—(Optional) Trace a particular routing instance.

**source** *source*—Source hostname or address.

**ttl** *tll*—(Optional) IP time-to-live (TTL) value. You can specify a number between 0 and 255. Local queries to the multicast group use a value of 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time** *wait-time*—(Optional) Number of seconds to wait for a response. The default value is 3.

**Required Privilege Level** view

**List of Sample Output** [mtrace from-source on page 5077](#)

**Output Fields** [Table 394 on page 5076](#) describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

**Table 394: mtrace from-source Output Fields**

Field Name	Field Description
<b>Mtrace from</b>	IP address of the receiver.
<b>to</b>	IP address of the source.
<b>via group</b>	IP address of the multicast group (if any).
<b>Querying full reverse path</b>	Indicates the full reverse path query has begun.
<b>number-of-hops</b>	Number of hops from the source to the named router or switch.
<b>router-name</b>	Name of the router or switch for this hop.
<b>address</b>	Address of the router or switch for this hop.
<b>protocol</b>	Protocol used (for example, PIM).
<b>Round trip time</b>	Average round-trip time, in milliseconds (ms).
<b>total ttl of</b>	Time-to-live (TTL) threshold.
<b>source</b>	Source address.
<b>Response Dest</b>	Response destination address.
<b>Overall</b>	Average packet rate for all traffic at each hop.

Table 394: mtrace from-source Output Fields (*continued*)

Field Name	Field Description
<b>Packet Statistics for Traffic From</b>	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.
<b>Receiver</b>	IP address receiving the multicast.
<b>Query source</b>	IP address sending the mtrace query.

## Sample Output

### mtrace from-source

```

user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-2 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
-3 hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source      Response Dest    Overall    Packet Statistics For Traffic From
192.1.4.2   192.1.1.2  Packet    192.1.4.2 To 225.1.1.1
      v      ___/ rtt    2 ms      Rate      Lost/Sent = Pct  Rate
192.1.2.1
192.1.3.2   routerC.lab.mycompany.net
      v      ^      ttl    2              0/0    = --    0 pps
192.1.4.1
192.1.2.2   routerB.lab.mycompany.net
      v      \__  ttl    3              ?/0              0 pps
192.1.1.2   192.1.1.2
Receiver      Query Source

```

## mtrace monitor

<b>Syntax</b>	mtrace monitor
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Listen passively for IP multicast responses. To exit the <b>mtrace monitor</b> command, type Ctrl+c.
<b>Options</b>	<b>none</b> —Trace the master instance.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">mtrace monitor on page 5079</a>
<b>Output Fields</b>	<a href="#">Table 395 on page 5078</a> describes the output fields for the <b>mtrace monitor</b> command. Output fields are listed in the approximate order in which they appear.

**Table 395: mtrace monitor Output Fields**

Field Name	Field Description
<b>Mtrace query at</b>	Date and time of the query.
<b>by</b>	Address of the host issuing the query.
<b>resp to</b>	Response destination.
<b>qid</b>	Query ID number.
<b>packet from...to</b>	IP address of the query source and default group destination.
<b>from...to</b>	IP address of the multicast source and the response address.
<b>via group</b>	IP address of the group to trace.
<b>mxhop</b>	Maximum hop setting.



## Sample Output

### mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

## mtrace to-gateway

---

**Syntax** `mtrace to-gateway gateway gateway`  
`<brief | detail>`  
`<extra-hops extra-hops>`  
`<group group>`  
`<interface interface-name>`  
`<interval interval>`  
`<loop>`  
`<max-hops max-hops>`  
`<max-queries max-queries>`  
`<multicast-response | unicast-response>`  
`<no-resolve>`  
`<no-router-alert>`  
`<response response>`  
`<routing-instance routing-instance-name>`  
`<tll ttl>`  
`<unicast-response>`  
`<wait-time wait-time>`

**Release Information** Command introduced before Junos OS Release 7.4.  
Command introduced in Junos OS Release 9.0 for EX Series switches.  
Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Display trace information about a multicast path from this router or switch to a gateway router or switch.

**Options** `gateway gateway`—Send the trace query to a gateway multicast address.

`brief | detail`—(Optional) Display the specified level of output.

`extra-hops extra-hops`—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between **0** and **255**.

`group group`—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

`interface interface-name`—(Optional) Source address for sending the trace query.

`interval interval`—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10**.

`loop`—(Optional) Loop indefinitely, displaying rate and loss statistics.

`max-hops max-hops`—(Optional) Maximum hops to trace toward the source. You can specify a number between **0** and **255**. The default value is **32**.

`max-queries max-queries`—(Optional) Maximum number of query attempts for any hop. You can specify a number between **0** and **255**. The default value is **3**.

`multicast-response`—(Optional) Always request the response using multicast.

`no-resolve`—(Optional) Do not attempt to display addresses symbolically.

**no-router-alert**—(Optional) Do not use the router-alert IP option.

**response *response***—(Optional) Send trace response to a host or multicast address.

**routing-instance *routing-instance-name***—(Optional) Trace a particular routing instance.

**ttl *tll***—(Optional) IP time-to-live value. You can specify a number between 0 and 225.

Local queries to the multicast group use TTL 1. Otherwise, the default value is 127.

**unicast-response**—(Optional) Always request the response using unicast.

**wait-time *wait-time***—(Optional) Number of seconds to wait for a response. The default value is 3.

**Required Privilege Level** view

**List of Sample Output** [mtrace to-gateway on page 5081](#)

**Output Fields** [Table 396 on page 5081](#) describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

**Table 396: mtrace to-gateway Output Fields**

Field Name	Field Description
<b>Mtrace from</b>	IP address of the receiver.
<b>to</b>	IP address of the source.
<b>via group</b>	IP address of the multicast group (if any).
<b>Querying full reverse path</b>	Indicates the full reverse path query has begun.
<b><i>number-of-hops</i></b>	Number of hops from the source to the named router or switch.
<b><i>router-name</i></b>	Name of the router or switch for this hop.
<b><i>address</i></b>	Address of the router or switch for this hop.
<b><i>protocol</i></b>	Protocol used (for example, PIM).
<b>Round trip time</b>	Average round-trip time, in milliseconds (ms).
<b>total ttl of</b>	Time-to-live (TTL) threshold.

## Sample Output

### mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief
```

```
Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
```

```
Querying full reverse path... * *
 0 routerA.lab.mycompany.net (192.1.1.2)
-1 routerA.lab.mycompany.net (192.1.1.2) PIM thresh^ 1
-2 routerB.lab.mycompany.net (192.1.2.2) PIM thresh^ 1
-3 routerC.lab.mycompany.net (192.1.3.2) PIM thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

## show configuration protocols igmp

<b>Syntax</b>	show configuration protocols igmp
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display Internet Group Management Protocol (IGMP) information.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">IGMP Snooping Overview on page 4761</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show configuration protocols igmp on page 5083</a>
<b>Output Fields</b>	<a href="#">Table 397 on page 5083</a> describes the output fields for the <b>show configuration protocols igmp</b> command that relate to IGMP querying.

**Table 397: show igmp group Output Fields**

Field Name	Field Description	Level of Output
accounting	Enables notification for join and leave events.	All levels
igmp-querier	Configured source address for the IGMP querier.	All levels
interface	Name of the interface that receives IGMP membership reports.	All levels
query-interval	Interval at which the IGMP querier sends general host-query messages to solicit membership information.	All levels
query-response-interval	How long the IGMP querier waits to receive a response from a query message before sending another query.	All levels
src-address	Source address of IGMP queries.	
version	IGMP version.	All levels

## Sample Output

### show configuration protocols igmp

```

user@switch> show configuration protocols igmp
query-interval 150;
query-response-interval 50;
accounting;
interface vlan.43 {
  version 2;
}
igmp-querier {

```

```
src-address 10.0.0.2;  
}
```

## show igmp group

<b>List of Syntax</b>	<a href="#">Syntax on page 5085</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5085</a>
<b>Syntax</b>	<pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show igmp group &lt;brief   detail&gt; &lt;group-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display Internet Group Management Protocol (IGMP) group membership information.
<b>Options</b>	<p><b>none</b>—Display standard information about membership for all IGMP groups.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group-name</b>—(Optional) Display group membership for the specified IP address only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear igmp membership on page 5051</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show igmp group (Include Mode) on page 5086</a> <a href="#">show igmp group (Exclude Mode) on page 5087</a> <a href="#">show igmp group brief on page 5087</a> <a href="#">show igmp group detail on page 5087</a>
<b>Output Fields</b>	<p><a href="#">Table 397 on page 5083</a> describes the output fields for the <b>show igmp group</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 398: show igmp group Output Fields**

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the interface that received the IGMP membership report. A name of <b>local</b> indicates that the local routing device joined the group itself.	All levels
<b>Group</b>	Group address.	All levels
<b>Group Mode</b>	Mode the SSM group is operating in: <b>Include</b> or <b>Exclude</b> .	All levels

Table 398: show igmp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Source</b>	Source address.	All levels
<b>Source timeout</b>	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	<b>detail</b>
<b>Last reported by</b>	Address of the host that last reported membership in this group.	All levels
<b>Timeout</b>	Time remaining until the group membership is removed.	<b>brief none</b>
<b>Group timeout</b>	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	<b>detail</b>
<b>Type</b>	Type of group membership: <ul style="list-style-type: none"> <li>• <b>Dynamic</b>—Host reported the membership.</li> <li>• <b>Static</b>—Membership is configured.</li> </ul>	All levels

## Sample Output

### show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```



```

Group: 224.0.0.22
Source: 0.0.0.0
Last reported by: Local
Timeout: 0 Type: Dynamic

```

### show igmp group (Exclude Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
  Source: 0.0.0.0
  Last reported by: Local
  Timeout: 0 Type: Dynamic
  Group: 224.0.0.22
  Source: 0.0.0.0
  Last reported by: Local
  Timeout: 0 Type: Dynamic

```

### show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

### show igmp group detail

```

user@host> show igmp group detail
Interface: t1-0/1/0.0
  Group: 232.1.1.1
  Group mode: Include
  Source: 10.0.0.2
  Source timeout: 12
  Last reported by: 10.9.5.2
  Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
  Group mode: Include
  Source: 10.0.0.3
  Source timeout: 12
  Last reported by: 10.9.5.2
  Group timeout: 0 Type: Dynamic
  Group: 232.1.1.1
  Group mode: Include
  Source: 10.0.0.4
  Source timeout: 12
  Last reported by: 10.9.5.2
  Group timeout: 0 Type: Dynamic
  Group: 232.1.1.2
  Group mode: Include
  Source: 10.0.0.4
  Source timeout: 12
  Last reported by: 10.9.5.2
  Group timeout: 0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
  Group mode: Exclude

```

```
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout:      0 Type: Dynamic
Group: 224.0.0.22
Group mode: Exclude
Source: 0.0.0.0
Source timeout: 0
Last reported by: Local
Group timeout:      0 Type: Dynamic
```

## show igmp interface

<b>List of Syntax</b>	<a href="#">Syntax on page 5089</a> <a href="#">Syntax (EX Series Switches and the QFX Series) on page 5089</a>
<b>Syntax</b>	<pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Syntax (EX Series Switches and the QFX Series)</b>	<pre>show igmp interface &lt;brief   detail&gt; &lt;interface-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.
<b>Options</b>	<p><b>none</b>—Display standard information about all IGMP-enabled interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface-name</b>—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear igmp membership on page 5051</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show igmp interface on page 5091</a> <a href="#">show igmp interface brief on page 5092</a> <a href="#">show igmp interface detail on page 5092</a> <a href="#">show igmp interface &lt;interface-name&gt; on page 5092</a>
<b>Output Fields</b>	<p><a href="#">Table 399 on page 5089</a> describes the output fields for the <b>show igmp interface</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 399: show igmp interface Output Fields**

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels

Table 399: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the interface: <b>Up</b> or <b>Down</b> .	All levels
<b>SSM Map Policy</b>	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
<b>Timeout</b>	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
<b>Version</b>	IGMP version being used on the interface: <b>1</b> , <b>2</b> , or <b>3</b> .	All levels
<b>Groups</b>	Number of groups on the interface.	All levels
<b>Group limit</b>	Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.	All levels
<b>Group threshold</b>	Configured threshold at which a warning message is generated.  This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
<b>Group log-interval</b>	Time (in seconds) between consecutive log messages.	All levels
<b>Immediate Leave</b>	State of the immediate leave option: <ul style="list-style-type: none"> <li><b>On</b>—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface.</li> <li><b>Off</b>—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.</li> </ul>	All levels
<b>Promiscuous Mode</b>	State of the promiscuous mode option: <ul style="list-style-type: none"> <li><b>On</b>—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces.</li> <li><b>Off</b>—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.</li> </ul>	All levels
<b>Passive</b>	State of the passive mode option: <ul style="list-style-type: none"> <li><b>On</b>—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves.</li> <li><b>Off</b>—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves.</li> </ul> <p>The <b>passive</b> statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the <b>on</b> state declaration:</p> <ul style="list-style-type: none"> <li><b>send-general-query</b>—The interface sends general queries.</li> <li><b>send-group-query</b>—The interface sends group-specific and group-source-specific queries.</li> <li><b>allow-receive</b>—The interface receives control traffic.</li> </ul>	All levels

Table 399: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>OIF map</b>	Name of the OIF map (if configured) associated with the interface.	All levels
<b>SSM map</b>	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels
<b>Configured Parameters</b>	Information configured by the user: <ul style="list-style-type: none"> <li>• <b>IGMP Query Interval</b>—Interval (in seconds) at which this router sends membership queries when it is the querier.</li> <li>• <b>IGMP Query Response Interval</b>—Time (in seconds) that the router waits for a report in response to a general query.</li> <li>• <b>IGMP Last Member Query Interval</b>—Time (in seconds) that the router waits for a report in response to a group-specific query.</li> <li>• <b>IGMP Robustness Count</b>—Number of times the router retries a query.</li> </ul>	All levels
<b>Derived Parameters</b>	Derived information: <ul style="list-style-type: none"> <li>• <b>IGMP Membership Timeout</b>—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.</li> <li>• <b>IGMP Other Querier Present Timeout</b>—Time (in seconds) that the router waits for the IGMP querier to send a query.</li> </ul>	All levels

## Sample Output

### show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
  Querier: 10.111.30.1
  State:      Up Timeout:  None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
  Querier: 10.111.10.1
  State:      Up Timeout:  None Version:  2 Groups:    2
  SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
  Querier: 10.111.20.1
  State:      Up Timeout:  None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

### show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 5091](#).

### show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 5091](#).

### show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout:      None Version: 3 Groups:      1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```

## show igmp statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 5093</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5093</a>
<b>Syntax</b>	<pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show igmp statistics &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display Internet Group Management Protocol (IGMP) statistics.
<b>Options</b>	<p><b>none</b>—Display IGMP statistics for all interfaces.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP statistics about the specified interface only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear igmp statistics on page 5055</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show igmp statistics on page 5094</a> <a href="#">show igmp statistics interface on page 5095</a>
<b>Output Fields</b>	<p><a href="#">Table 400 on page 5093</a> describes the output fields for the <b>show igmp statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 400: show igmp statistics Output Fields**

Field Name	Field Description
IGMP packet statistics	Heading for IGMP packet statistics for all interfaces or for the specified interface name.

Table 400: show igmp statistics Output Fields (*continued*)

Field Name	Field Description
<b>IGMP Message type</b>	<p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> <li>• <b>Membership Query</b>—Number of membership queries sent and received.</li> <li>• <b>V1 Membership Report</b>—Number of version 1 membership reports sent and received.</li> <li>• <b>DVMRP</b>—Number of DVMRP messages sent or received.</li> <li>• <b>PIM V1</b>—Number of PIM version 1 messages sent or received.</li> <li>• <b>Cisco Trace</b>—Number of Cisco trace messages sent or received.</li> <li>• <b>V2 Membership Report</b>—Number of version 2 membership reports sent or received.</li> <li>• <b>Group Leave</b>—Number of group leave messages sent or received.</li> <li>• <b>Mtrace Response</b>—Number of Mtrace response messages sent or received.</li> <li>• <b>Mtrace Request</b>—Number of Mtrace request messages sent or received.</li> <li>• <b>Domain Wide Report</b>—Number of domain-wide reports sent or received.</li> <li>• <b>V3 Membership Report</b>—Number of version 3 membership reports sent or received.</li> <li>• <b>Other Unknown types</b>—Number of unknown message types received.</li> <li>• <b>IGMP v3 unsupported type</b>—Number of messages received with unknown and unsupported IGMP version 3 message types.</li> <li>• <b>IGMP v3 source required for SSM</b>—Number of IGMP version 3 messages received that contained no source.</li> <li>• <b>IGMP v3 mode not applicable for SSM</b>—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).</li> </ul>
<b>Received</b>	Number of messages received.
<b>Sent</b>	Number of messages sent.
<b>Rx errors</b>	Number of received packets that contained errors.
<b>IGMP Global Statistics</b>	<p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> <li>• <b>Bad Length</b>—Number of messages received with length errors so severe that further classification could not occur.</li> <li>• <b>Bad Checksum</b>—Number of messages received with a bad IP checksum. No further classification was performed.</li> <li>• <b>Bad Receive If</b>—Number of messages received on an interface not enabled for IGMP.</li> <li>• <b>Rx non-local</b>—Number of messages received from senders that are not local.</li> <li>• <b>Timed out</b>—Number of groups that timed out as a result of not receiving an explicit leave message.</li> <li>• <b>Rejected Report</b>—Number of reports dropped because of the IGMP group policy.</li> <li>• <b>Total Interfaces</b>—Number of interfaces configured to support IGMP.</li> </ul>

## Sample Output

### show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report     0            0        0

```



DVMRP	0	0	0
PIM V1	0	0	0
Cisco Trace	0	0	0
V2 Membership Report	0	0	0
Group Leave	0	0	0
Mtrace Response	0	0	0
Mtrace Request	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	0	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			0
IGMP v3 mode not applicable for SSM			0
IGMP Global Statistics			
Bad Length	0		
Bad Checksum	0		
Bad Receive If	0		
Rx non-local	1227		
Timed out	0		
Rejected Report	0		
Total Interfaces	2		

#### show igmp statistics interface

```

user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received      Sent  Rx errors
Membership Query        0           230      0
V1 Membership Report    0           0        0

```

## show igmp-snooping membership

<b>Syntax</b>	<pre>show igmp-snooping membership &lt;brief   detail&gt; &lt;interface <i>interface-name</i>&gt; &lt;vlan <i>vlan-id</i>   <i>vlan-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>IGMPv3 output introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Display IGMP snooping membership information.
<b>Options</b>	<p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display IGMP snooping information for the specified interface.</p> <p><b>vlan <i>vlan-id</i>   <i>vlan-name</i></b>—(Optional) Display IGMP snooping information for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5047</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> <li>• <a href="#">show igmp-snooping route on page 5099</a></li> <li>• <a href="#">show igmp-snooping statistics on page 5101</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5103</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show igmp-snooping membership on page 5097</a></p> <p><a href="#">show igmp-snooping membership detail on page 5098</a></p>
<b>Output Fields</b>	<p><a href="#">Table 401 on page 5096</a> lists the output fields for the <b>show igmp-snooping membership</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 401: show igmp-snooping membership Output Fields**

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Interfaces	Interfaces assigned to the VLAN.	All
Tag	Numerical identifier of the VLAN.	<b>detail</b>

Table 401: show igmp-snooping membership Output Fields (*continued*)

Field Name	Field Description	Level of Output
Router interfaces	Names of multicast router interfaces.	<b>detail</b>
• static or dynamic	Whether the multicast router interface is <b>static</b> or <b>dynamic</b> .	<b>detail</b>
• Uptime	For static interfaces, length of time since the interface was configured as a multicast router interface; for dynamic interfaces, length of time since the first query was received on the interface.	<b>detail</b>
• timeout	Query timeout in seconds.	<b>detail</b>
Group	IP multicast address of the multicast group.	<b>detail</b>
Receiver count	Number of interfaces that have membership in a multicast group.	<b>detail</b>
Flags	IGMP version of the host sending a join message.	<b>detail</b>
Uptime	Length of time a multicast group has been active on the interface.	<b>detail</b>
timeout	Time (in seconds) left until the entry for the multicast group is removed.	All
Last reporter	Last host to report membership for the multicast group.	<b>detail</b>
Include source	Source addresses from which multicast streams are allowed based on IGMPv3 reports.	<b>detail</b>

## Sample Output

### show igmp-snooping membership

```

user@switch> show igmp-snooping membership
VLAN: v1
  224.1.1.1      *           258 secs
    Interfaces: ge-0/0/0.0
  224.1.1.3      *           258 secs
    Interfaces: ge-0/0/0.0
  224.1.1.5      *           258 secs
    Interfaces: ge-0/0/0.0
  224.1.1.7      *           258 secs

```

```
Interfaces: ge-0/0/0.0
224.1.1.9      *           258 secs
Interfaces: ge-0/0/0.0
224.1.1.11     *           258 secs
Interfaces: ge-0/0/0.0
```

### show igmp-snooping membership detail

```
user@switch> show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
Group: 225.0.0.2
Receiver count: 1, Flags: <V3-hosts>
  ge-0/0/15.0 Uptime: 00:00:11 timeout: 248 Last reporter: 10.2.10.16
  Include source: 1.2.1.1, 1.3.1.1
VLAN: v44 Tag: 44 (Index: 5)
Group: 225.0.0.1
Receiver count: 1, Flags: <V2-hosts>
  ge-0/0/21.0 Uptime: 00:00:02 timeout: 257
VLAN: v110 Tag: 110 (Index: 4)
Router interfaces:
  ge-0/0/3.0 static Uptime: 00:08:45
  ge-0/0/2.0 static Uptime: 00:08:45
  ge-0/0/4.0 dynamic Uptime: 00:16:41 timeout: 254
Group: 225.0.0.3
Receiver count: 1, Flags: <V3-hosts>
  ge-0/0/5.0 Uptime: 00:00:19 timeout: 259
Group: 225.1.1.1
Receiver count: 1, Flags: <V2-hosts>
  ge-0/0/5.0 Uptime: 00:22:43 timeout: 96
Group: 225.2.2.2
Receiver count: 1, Flags: <V2-hosts Static>
  ge-0/0/5.0 Uptime: 00:23:13
```

## show igmp-snooping route

<b>Syntax</b>	<pre>show igmp-snooping route &lt;brief   detail&gt; &lt;ethernet-switching &lt;brief   detail   vlan (vlan-id   vlan-name )&gt;&gt; &lt;inet &lt;brief   detail   vlan vlan-name&gt;&gt; &lt;vlan vlan-name&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display IGMP snooping route information.
<b>Options</b>	<p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>ethernet-switching</b>—(Optional) Display Ethernet switching information.</p> <p><b>inet</b>—(Optional) Display <b>inet</b> information.</p> <p><b>vlan vlan-name</b>—(Optional) Display route information for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5047</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> <li>• <a href="#">show igmp-snooping statistics on page 5101</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5103</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show igmp-snooping route on page 5100</a></p> <p><a href="#">show igmp-snooping route vlan v1 on page 5100</a></p>
<b>Output Fields</b>	<p><a href="#">Table 402 on page 5099</a> lists the output fields for the <b>show igmp-snooping route</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 402: show igmp-snooping route Output Fields**

Field Name	Field Description
Table	(For internal use only. Value is always 0.)
VLAN	Name of the VLAN.
Group	Multicast group address.
Interfaces	Interfaces on which IGMP packets were snooped.
Next-hop	ID associated with the next-hop device.

## Sample Output

### show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN          Group          Next-hop
V11           224.1.1.1, *      533
               Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN          Group          Next-hop
v12           224.1.1.3, *      534
               Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

### show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN          Group          Next-hop
v1           224.1.1.1, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.3, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.5, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.7, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.9, *      1266
               Interfaces: ge-0/0/0.0
v1           224.1.1.11, *     1266
               Interfaces: ge-0/0/0.0
```

## show igmp-snooping statistics

<b>Syntax</b>	<b>show igmp-snooping statistics</b>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display IGMP snooping statistics.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5047</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> <li>• <a href="#">show igmp-snooping route on page 5099</a></li> <li>• <a href="#">show igmp-snooping vlans on page 5103</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show igmp-snooping statistics on page 5102</a>
<b>Output Fields</b>	Table 403 on page 5101 lists the output fields for the <b>show igmp-snooping statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 403: show igmp-snooping statistics Output Fields**

Field Name	Field Description
Bad length	IGMP packet has illegal or bad length.
Bad checksum	IGMP or IP checksum is incorrect.
Invalid interface	Packet was received through an invalid interface.
Not local	Number of packets received from senders that are not local.
Receive unknown	Unknown IGMP type.
Timed out	Number of timeouts for all multicast groups.
IGMP Type	Type of IGMP message ( <b>Queries</b> , <b>Reports</b> , <b>Leaves</b> , or <b>Other</b> ).
Received	Number of IGMP packets received.
Transmitted	Number of IGMP packets transmitted.
Recv Errors	Number of general receive errors.

## Sample Output

### show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 58
```

IGMP Type	Received	Transmitted	Recv Errors
Queries:	74295	0	0
Reports:	18148423	0	16333523
Leaves:	0	0	0
Other:	0	0	0



## show igmp-snooping vlans

<b>Syntax</b>	<code>show igmp-snooping vlans</code> <code>&lt;brief   detail&gt;</code> <code>&lt;vlan <i>vlan-id</i>   <i>vlan-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display IGMP snooping VLAN information.
<b>Options</b>	<p><b>none</b>—Display general parameters.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>vlan <i>vlan-id</i>   vlan <i>vlan-number</i></b>—(Optional) Display VLAN information for the specified VLAN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring IGMP Snooping on page 5047</a></li> <li>• <a href="#">Configuring IGMP Snooping on page 4848</a></li> <li>• <a href="#">show igmp-snooping route on page 5099</a></li> <li>• <a href="#">show igmp-snooping statistics on page 5101</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show igmp-snooping vlans on page 5104</a></p> <p><a href="#">show igmp-snooping vlans vlan on page 5104</a></p> <p><a href="#">show igmp-snooping vlans vlan detail on page 5104</a></p>
<b>Output Fields</b>	Table 404 on page 5103 lists the output fields for the <b>show igmp-snooping vlans</b> command. Output fields are listed in the approximate order in which they appear.

**Table 404: show igmp-snooping vlans Output Fields**

Field Name	Field Description	Level of Output
<b>VLAN</b>	Name of the VLAN.	All levels
<b>IGMP-L2-Querier</b>	Source address for IGMP snooping queries (if switch is an IGMP querier)	All levels
<b>Interfaces</b>	Number of interfaces in the VLAN.	All levels
<b>Groups</b>	Number of groups in the VLAN.	All levels
<b>MRouters</b>	Number of multicast routers associated with the VLAN.	All levels
<b>Receivers</b>	Number of host receivers in the VLAN.	All levels

Table 404: show igmp-snooping vlans Output Fields (*continued*)

Field Name	Field Description	Level of Output
Tag	Numerical identifier of the VLAN.	detail
tagged   untagged	Interface participates in a tagged (802.1Q) or untagged (native) VLAN.	detail
vlan-interface	Internal VLAN interface identifier.	detail
Membership timeout	Membership timeout value.	detail
Querier timeout	Timeout value for interfaces dynamically marked as router or switch interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface.	detail
Interface	Name of the interface.	detail
Reporters	Number of dynamic groups on an interface.	detail

## Sample Output

### show igmp-snooping vlans

```

user@switch> show igmp-snooping vlans
VLAN      Interfaces Groups MRouters Receivers
default   0           0      0         0
v1         11          50      0         0
v10        1           0      0         0
v11        1           0      0         0
v180       3           0      1         0
v181       3           0      0         0
v182       3           0      0         0

```

### show igmp-snooping vlans vlan

```

user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN      Interfaces Groups MRouters Receivers
v10       1           0      0         0

```

### show igmp-snooping vlans vlan detail

```

user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
      Interface: ge-0/0/10.0, tagged, Groups: 0
IGMP-L2-Querier: Stopped, SourceAddress: 10.10.1.2

```

## show msdp

<b>Syntax</b>	show msdp <brief   detail> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )> <peer <i>peer-address</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Display Multicast Source Discovery Protocol (MSDP) information.
<b>Options</b>	<p><b>none</b>—Display standard MSDP information for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display information about the specified peer only.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show msdp source on page 5107</a></li> <li>• <a href="#">show msdp source-active on page 5109</a></li> <li>• <a href="#">show msdp statistics on page 5112</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show msdp on page 5106</a> <a href="#">show msdp brief on page 5106</a> <a href="#">show msdp detail on page 5106</a>
<b>Output Fields</b>	Table 405 on page 5105 describes the output fields for the <b>show msdp</b> command. Output fields are listed in the approximate order in which they appear.

**Table 405: show msdp Output Fields**

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels
State	Status of the MSDP connection: <b>Listen</b> , <b>Established</b> , or <b>Inactive</b> .	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels

Table 405: show msdp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> .	All levels
Peer Connect Retries	Number of peer connection retries.	detail
State timer expires	Number of seconds before another message is sent to a peer.	detail
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	detail
SA accepted	Number of entries in the source-active cache accepted from the peer.	detail
SA received	Number of entries in the source-active cache received by the peer.	detail

## Sample Output

### show msdp

```

user@host> show msdp
Peer address   Local address   State      Last up/down Peer-Group SA Count
198.32.8.193   198.32.8.195   Established 5d 19:25:44 North23 120/150
198.32.8.194   198.32.8.195   Established 3d 19:27:27 North23 300/345
198.32.8.196   198.32.8.195   Established 5d 19:39:36 North23 10/13
198.32.8.197   198.32.8.195   Established 5d 19:32:27 North23 5/6
198.32.8.198   198.32.8.195   Established 3d 19:33:04 North23 2305/3000

```

### show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 5106](#).

### show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

## show msdp source

---

<b>Syntax</b>	<pre>show msdp source &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;source-address&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).
<b>Options</b>	<p><b>none</b>—Display standard MSDP source information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>source-address</b>—(Optional) IP address and optional prefix length. Display information for the specified source address only.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show msdp on page 5105</a></li> <li>• <a href="#">show msdp source-active on page 5109</a></li> <li>• <a href="#">show msdp statistics on page 5112</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show msdp source on page 5108</a>

**Output Fields** Table 406 on page 5108 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

**Table 406: show msdp source Output Fields**

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: <ul style="list-style-type: none"> <li>• <b>Configured</b>—Source-active limit explicitly configured for this source.</li> <li>• <b>Dynamic</b>—Source-active limit established when this source was discovered.</li> </ul>
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

## Sample Output

**show msdp source**

```

user@host> show msdp source
Source address /Len  Type      Maximum  Threshold  Exceeded
0.0.0.0        /0    Configured    5         none       0
10.1.0.0       /16   Configured    500       none       0
10.1.1.1       /32   Configured    10000     none       0
10.1.1.2       /32   Dynamic       6936     none       0
10.1.5.5       /32   Dynamic       500      none      123
10.2.1.1       /32   Dynamic        2         none       0

```

## show msdp source-active

<b>Syntax</b>	<pre>show msdp source-active &lt;brief   detail&gt; &lt;group <i>group</i>&gt; &lt;instance <i>instance-name</i>&gt; &lt;local&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;originator <i>originator</i>&gt; &lt;peer <i>peer-address</i>&gt; &lt;source <i>source-address</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
<b>Description</b>	Display the Multicast Source Discovery Protocol (MSDP) source-active cache.
<b>Options</b>	<p><b>none</b>—Display standard MSDP source-active cache information for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>group <i>group</i></b>—(Optional) Display source-active cache information for the specified group.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified instance.</p> <p><b>local</b>—(Optional) Display all source-active caches originated by this router.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>originator <i>originator</i></b>—(Optional) Display information about the peer that originated the source-active cache entries.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display the source-active cache of the specified peer.</p> <p><b>source <i>source-address</i></b>—(Optional) Display the source-active cache of the specified source.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show msdp on page 5105</a></li> <li>• <a href="#">show msdp source on page 5107</a></li> <li>• <a href="#">show msdp statistics on page 5112</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show msdp source-active on page 5110</a></p> <p><a href="#">show msdp source-active brief on page 5110</a></p> <p><a href="#">show msdp source-active detail on page 5111</a></p> <p><a href="#">show msdp source-active source on page 5111</a></p>
<b>Output Fields</b>	Table 407 on page 5110 describes the output fields for the <b>show msdp source-active</b> command. Output fields are listed in the approximate order in which they appear.

Table 407: show msdp source-active Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured.
Flags	Flags: Accept, Reject, or Filtered.

## Sample Output

### show msdp source-active

```

user@host> show msdp source-active
Group address  Source address Peer address  Originator  Flags
230.0.0.0      192.168.195.46 local        10.255.14.30 Accept
230.0.0.1      192.168.195.46 local        10.255.14.30 Accept
230.0.0.2      192.168.195.46 local        10.255.14.30 Accept
230.0.0.3      192.168.195.46 local        10.255.14.30 Accept
230.0.0.4      192.168.195.46 local        10.255.14.30 Accept

```

### show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 5110](#).



### show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 5110](#).

### show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
```

```
Global active source limit exceeded: 0
```

```
Global active source limit maximum: 25000
```

```
Global active source limit threshold: 24000
```

```
Global active source limit log-warning: 100
```

```
Global active source limit log interval: 0
```

Group address	Source address	Peer address	Originator	Flags
226.2.2.1	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.3	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.4	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.5	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.7	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.10	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.11	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.13	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.14	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.15	192.168.215.246	10.255.182.140	10.255.182.140	Accept

## show msdp statistics

<b>Syntax</b>	show msdp statistics <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )> <peer <i>peer-address</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Display statistics about Multicast Source Discovery Protocol (MSDP) peers.
<b>Options</b>	<p><b>none</b>—Display statistics about all MSDP peers for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics about a specific MSDP instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>peer <i>peer-address</i></b>—(Optional) Display statistics about a particular MSDP peer.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear msdp statistics on page 5059</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show msdp statistics on page 5114</a> <a href="#">show msdp statistics peer on page 5114</a>
<b>Output Fields</b>	Table 408 on page 5112 describes the output fields for the <b>show msdp statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 408: show msdp statistics Output Fields**

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Peer	Address of peer.

Table 408: show msdp statistics Output Fields (*continued*)

Field Name	Field Description
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.
Keepalive messages sent	Number of keepalive messages sent.
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.

Table 408: show msdp statistics Output Fields (*continued*)

Field Name	Field Description
Error messages received	Number of error messages received.

## Sample Output

### show msdp statistics

```

user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

### show msdp statistics peer

```

user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
  Last State Change: 8:19:23 (00:01:08)
  Last message received from peer: 8:20:05 (00:00:26)
  RPF Failures: 0
  Remote Closes: 0
  Peer Timeouts: 0
  SA messages sent: 17
  SA messages received: 16
  SA request messages sent: 0
  SA request messages received: 0
  SA response messages sent: 0
  SA response messages received: 0
  Active source exceeded: 20
  Active source Maximum: 10
  Active source threshold: 8
  Active source log-warning: 60
  Active source log-interval: 120
  Keepalive messages sent: 0

```

Keepalive messages received: 0  
Unknown messages received: 0  
Error messages received: 0

## show multicast flow-map

<b>List of Syntax</b>	<a href="#">Syntax on page 5116</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5116</a>
<b>Syntax</b>	<pre>show multicast flow-map &lt;brief   detail&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show multicast flow-map &lt;brief   detail&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.2.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display configuration information about IP multicast flow maps.
<b>Options</b>	<p><b>none</b>—Display configuration information about IP multicast flow maps on all systems.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast flow-map on page 5117</a> <a href="#">show multicast flow-map detail on page 5117</a>
<b>Output Fields</b>	<p><a href="#">Table 409 on page 5116</a> describes the output fields for the <b>show multicast flow-map</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 409: show multicast flow-map Output Fields**

Field Name	Field Description	Levels of Output
<b>Name</b>	Name of the flow map.	All levels
<b>Policy</b>	Name of the policy associated with the flow map.	All levels
<b>Cache-timeout</b>	Cache timeout value assigned to the flow map.	All levels
<b>Bandwidth</b>	Bandwidth setting associated with the flow map.	All levels
<b>Adaptive</b>	Whether or not adaptive mode is enabled for the flow map.	none
<b>Flow-map</b>	Name of the flow map.	<b>detail</b>

Table 409: show multicast flow-map Output Fields (*continued*)

Field Name	Field Description	Levels of Output
<b>Adaptive Bandwidth</b>	Whether or not adaptive mode is enabled for the flow map.	<b>detail</b>
<b>Redundant Sources</b>	Redundant sources defined for the same destination group.	<b>detail</b>

## Sample Output

### show multicast flow-map

```

user@host> show multicast flow-map
Instance: master
Name          Policy          Cache timeout    Bandwidth Adaptive
map2          policy2         never            2000000 no
map1          policy1         60 seconds      2000000 no

```

## Sample Output

### show multicast flow-map detail

```

user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
  Policy:          policy1
  Cache Timeout:   600 seconds
  Bandwidth:       2000000
  Adaptive Bandwidth: yes
  Redundant Sources: 11.11.11.11
  Redundant Sources: 11.11.11.12
  Redundant Sources: 11.11.11.13

```

## show multicast interface

<b>List of Syntax</b>	<a href="#">Syntax on page 5118</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5118</a>
<b>Syntax</b>	show multicast interface <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switch and the QFX Series)</b>	show multicast interface
<b>Release Information</b>	Command introduced in Junos OS Release 8.3. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display bandwidth information about IP multicast interfaces.
<b>Options</b>	none—Display all interfaces that have multicast configured.  logical-system (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast interface on page 5119</a>
<b>Output Fields</b>	<a href="#">Table 410 on page 5118</a> describes the output fields for the <b>show multicast interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 410: show multicast interface Output Fields**

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface.  <b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.  This field does not appear in the output when the no QoS adjustment feature is disabled.



Table 410: show multicast interface Output Fields (*continued*)

Field Name	Field Description
<b>Local bandwidth deduction (bps)</b>	<p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p><b>NOTE:</b> Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
<b>Reverse OIF mapping</b>	<p>State of the reverse OIF mapping feature (<b>on</b> or <b>off</b>).</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
<b>Reverse OIF mapping no QoS adjustment</b>	<p>State of the no QoS adjustment feature (<b>on</b> or <b>off</b>) for interfaces that are using reverse OIF mapping.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
<b>Leave timer</b>	<p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
<b>No QoS adjustment</b>	<p>State (<b>on</b>) of the no QoS adjustment feature when this feature is enabled.</p> <p><b>NOTE:</b> This field does not appear in the output when the no QoS adjustment feature is disabled.</p>

## Sample Output

### show multicast interface

```

user@host> show multicast interface
Interface           Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3            10000000                0
fe-0/0/3.210        10000000                -2000000
fe-0/0/3.220        100000000              100000000
fe-0/0/3.230        20000000               18000000
fe-0/0/2.200        100000000              100000000

```

## show multicast mrinfo

<b>Syntax</b>	<code>show multicast mrinfo</code> <code>&lt;host&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display configuration information about IP multicast networks, including neighboring multicast router addresses.
<b>Options</b>	<b>none</b> —Display configuration information about all multicast networks.  <b>host</b> —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast mrinfo on page 5121</a>
<b>Output Fields</b>	<a href="#">Table 411 on page 5120</a> describes the output fields for the <b>show multicast mrinfo</b> command. Output fields are listed in the approximate order in which they appear.

Table 411: show multicast mrinfo Output Fields

Field Name	Field Description
<i>source-address</i>	Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.
<i>ip-address-1—&gt;ip-address-2</i>	Queried router interface address and directly attached neighbor interface address, respectively.
<i>(name or ip-address)</i>	Name or IP address of neighbor.
<i>[metric/threshold/type/flags]</i>	Neighbor's multicast profile: <ul style="list-style-type: none"> <li><b>metric</b>—Always has a value of 1, because <b>mrinfo</b> queries the directly connected interfaces of a device.</li> <li><b>threshold</b>—Multicast threshold time-to-live (TTL). The range of values is 0 through 255.</li> <li><b>type</b>—Multicast connection type: <b>pim</b> or <b>tunnel</b>.</li> <li><b>flags</b>—Flags for this route: <ul style="list-style-type: none"> <li><b>querier</b>—Queried router is the designated router for the neighboring session.</li> <li><b>leaf</b>—Link is a leaf in the multicast network.</li> <li><b>down</b>—Link status indicator.</li> </ul> </li> </ul>

## Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

## show multicast next-hops

---

<b>List of Syntax</b>	<a href="#">Syntax on page 5122</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5122</a>
<b>Syntax</b>	<pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show multicast next-hops &lt;brief   detail&gt; &lt;identifier-number&gt; &lt;inet   inet6&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p><b>detail</b> option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>
<b>Description</b>	Display the entries in the IP multicast next-hop table.
<b>Options</b>	<p><b>none</b>—Display standard information about all entries in the multicast next-hop table for all supported address families.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p>When you include the <b>detail</b> option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form <b>fe-0/1/2.0-(1048574)</b> where <b>1048574</b> is the next-hop ID number.</p> <p><b>identifier-number</b>—(Optional) Show a particular next hop by ID number. The range of values is 1 through <b>65,535</b>.</p> <p><b>inet   inet6</b>—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast next-hops on page 5123</a> <a href="#">show multicast next-hops (Bidirectional PIM on page 5123</a> <a href="#">show multicast next-hops brief on page 5124</a> <a href="#">show multicast next-hops detail on page 5124</a>

**Output Fields** Table 412 on page 5123 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

**Table 412: show multicast next-hops Output Fields**

Field Name	Field Description
<b>Family</b>	Protocol family (such as <b>INET</b> ).
<b>ID</b>	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.
<b>Refcount</b>	Number of cache entries that are using this next hop.
<b>KRefcount</b>	Kernel reference count for the next hop.
<b>Downstream interface</b>	Interface names associated with each multicast next-hop ID.
<b>Incoming interface list</b>	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.

## Sample Output

### show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
262142      4          2 so-1/0/0.0
262143      2          1 mt-1/1/0.49152
262148      2          1 mt-1/1/0.32769
```

### show multicast next-hops (Bidirectional PIM)

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
2097151      8          4 ge-0/0/1.0

Family: INET6
ID      Refcount  KRefcount Downstream interface
2097157      2          1 ge-0/0/1.0

Family: Incoming interface list
ID      Refcount  KRefcount Downstream interface
513      5          2 lo0.0
           ge-0/0/1.0
514      5          2 lo0.0
           ge-0/0/1.0
           xe-4/1/0.0
515      3          1 lo0.0
           ge-0/0/1.0
           xe-4/1/0.0
544      1          0 lo0.0
           xe-4/1/0.0
```

### show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 5123](#).

### show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefCount Downstream interface
1048577      2          1 fe-0/1/2.0-(1048574)
              ge-0/2/3.0-(1048576)
```

## show multicast pim-to-igmp-proxy

<b>List of Syntax</b>	<a href="#">Syntax on page 5125</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5125</a>
<b>Syntax</b>	<pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show multicast pim-to-igmp-proxy &lt;instance <i>instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 9.6 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3.</p> <p><b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.
<b>Options</b>	<p><b>none</b>—Display configuration information about PIM-to-IGMP message translation for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring PIM-to-IGMP and PIM-to-MLD Message Translation</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show multicast pim-to-igmp-proxy on page 5126</a> <a href="#">show multicast pim-to-igmp-proxy instance on page 5126</a>
<b>Output Fields</b>	<p><a href="#">Table 413 on page 5125</a> describes the output fields for the <b>show multicast pim-to-igmp-proxy</b> command. Output fields are listed in the order in which they appear.</p>

**Table 413: show multicast pim-to-igmp-proxy Output Fields**

Field Name	Field Description
<b>Instance</b>	Routing instance. Default instance is <b>master</b> (inet.0 routing table).
<b>Proxy state</b>	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> .

Table 413: show multicast pim-to-igmp-proxy Output Fields (*continued*)

Field Name	Field Description
<i>interface-name</i>	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

## Sample Output

### show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

### show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```



## show multicast pim-to-mld-proxy

<b>List of Syntax</b>	<a href="#">Syntax on page 5127</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5127</a>
<b>Syntax</b>	show multicast pim-to-mld-proxy <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switch and the QFX Series)</b>	show multicast pim-to-mld-proxy <instance <i>instance-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 9.6 for EX Series switches. <b>instance</b> option introduced in Junos OS Release 10.3. <b>instance</b> option introduced in Junos OS Release 10.3 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.
<b>Options</b>	<b>none</b> —Display configuration information about PIM-to-MLD message translation for all routing instances.  <b>instance</b> <i>instance-name</i> —(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance.  <b>logical-system</b> (all   <i>logical-system-name</i> )—(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast pim-to-mld-proxy on page 5128</a> <a href="#">show multicast pim-to-mld-proxy instance on page 5128</a>
<b>Output Fields</b>	Table 414 on page 5127 describes the output fields for the <b>show multicast pim-to-mld-proxy</b> command. Output fields are listed in the order in which they appear.

**Table 414: show multicast pim-to-mld-proxy Output Fields**

Field Name	Field Description
<b>Proxy state</b>	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: <b>enabled</b> or <b>disabled</b> .
<i>interface-name</i>	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

## Sample Output

### show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

### show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

## show multicast route

**List of Syntax**    [Syntax on page 5129](#)  
                           [Syntax \(EX Series Switch and the QFX Series\) on page 5129](#)

**Syntax**    show multicast route  
                   <brief | detail | extensive | summary>  
                   <active | all | inactive>  
                   <group *group*>  
                   <inet | inet6>  
                   <instance *instance name*>  
                   <logical-system (all | *logical-system-name*)>  
                   <*regular-expression*>  
                   <source-prefix *source-prefix*>

**Syntax (EX Series Switch and the QFX Series)**    show multicast route  
                                                                           <brief | detail | extensive | summary>  
                                                                           <active | all | inactive>  
                                                                           <group *group*>  
                                                                           <inet | inet6>  
                                                                           <instance *instance name*>  
                                                                           <*regular-expression*>  
                                                                           <source-prefix *source-prefix*>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                                   Command introduced in Junos OS Release 9.0 for EX Series switches.  
                                   inet6 and **instance** options introduced in Junos OS Release 10.0 for EX Series switches.  
                                   Command introduced in Junos OS Release 11.3 for the QFX Series.  
                                   Support for bidirectional PIM added in Junos OS Release 12.1.

**Description**    Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

**Options**    **none**—Display standard information about all entries in the multicast forwarding table for all routing instances.

**brief | detail | extensive | summary**—(Optional) Display the specified level of output.

**active | all | inactive**—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

**group *group***—(Optional) Display the cache entries for a particular group.

**inet | inet6**—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**regular-expression**—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.

**source-prefix source-prefix**—(Optional) Display the cache entries for a particular source prefix.

**Required Privilege Level** view

**Related Documentation**

- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*

**List of Sample Output**

- [show multicast route on page 5131](#)
- [show multicast route \(Bidirectional PIM\) on page 5132](#)
- [show multicast route brief on page 5132](#)
- [show multicast route detail on page 5133](#)
- [show multicast route extensive \(Bidirectional PIM\) on page 5133](#)
- [show multicast route extensive \(Multicast-Only Fast Reroute\) on page 5134](#)
- [show multicast route instance <instance-name> on page 5134](#)
- [show multicast route summary on page 5135](#)

**Output Fields** [Table 415 on page 5130](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

**Table 415: show multicast route Output Fields**

Field Name	Field Description	Level of Output
family	IPv4 address family ( <b>INET</b> ) or IPv6 address family ( <b>INET6</b> ).	All levels
Group	Group address.  For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Upstream rpf interface list	When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded.	All levels
Number of outgoing interfaces	Total number of outgoing interfaces for each (S,G) entry.	<b>extensive</b>

Table 415: show multicast route Output Fields (*continued*)

Field Name	Field Description	Level of Output
Session description	Name of the multicast session.	detail extensive
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays <b>Forwarding statistics are not available</b> .  <b>NOTE:</b> On QFX Series switches, this field does not report valid statistics.	detail extensive
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the <b>show multicast nexthops</b> command.	detail extensive
Incoming interface list ID	For bidirectional PIM, incoming interface list identifier.  Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	detail extensive
Upstream protocol	The protocol that maintains the active multicast forwarding route for this group or source.  When the <b>show multicast route extensive</b> command is used with the <b>display-origin-protocol</b> option, the field name is only <b>Protocol</b> and not <b>Upstream Protocol</b> . However, this field also displays the protocol that installed the active route.	detail extensive
Route type	Type of multicast route. Values can be (S,G) or (*,G).	summary
Route state	Whether the group is <b>Active</b> or <b>Inactive</b> .	summary extensive
Route count	Number of multicast routes.	summary
Forwarding state	Whether the prefix is pruned or forwarding.	extensive
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of <b>never</b> indicates a permanent forwarding entry. A value of <b>forever</b> indicates routes that do not have keepalive times.	extensive
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	extensive
Uptime	Time since the creation of a multicast route.	extensive

## Sample Output

### show multicast route

```
user@host> show multicast route
Family: INET
```

```
Group: 228.0.0.0
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
    so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
    so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.70.15/32
Upstream interface: so-1/0/0.0
Downstream interface list:
    mt-1/1/0.1081344

Family: INET6
```

#### show multicast route (Bidirectional PIM)

```
user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0
Downstream interface list:
    ge-0/0/1.0

Group: 224.1.3.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
    ge-0/0/1.0

Group: 225.1.1.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0
Downstream interface list:
    ge-0/0/1.0

Group: 225.1.3.0/24
Source: *
Incoming interface list:
    lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
    ge-0/0/1.0
Family: INET6
```

#### show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see [show multicast route on page 5131](#) or [show multicast route \(Bidirectional PIM\) on page 5132](#).

**show multicast route detail**

```

user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Unknown
  Statistics: 8 kbps, 100 pps, 45272 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 13404 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.1081344
  Session description: Administratively Scoped
  Statistics: 46 kbps, 1000 pps, 921077 packets

  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6

```

**show multicast route extensive (Bidirectional PIM)**

```

user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 2097153
  Incoming interface list ID: 585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

Group: 224.1.3.0/24

```

```
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
```

Family: INET6

### show multicast route extensive (Multicast-Only Fast Reroute)

```
user@host> show multicast route extensive
```

Instance: master Family: INET

```
Group: 225.1.1.1
Source: 10.0.0.1/32
Upstream rpf interface list:
  fe-1/2/13.0 (P) fe-1/2/14.0 (B)
Downstream interface list:
  fe-1/2/15.0
Session description: Unknown
Forwarding statistics are not available
RPF Next-hop ID: 836
Next-hop ID: 1048585
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: 171 seconds
Wrong incoming interface notifications: 0
Uptime: 00:03:09
```

### show multicast route instance <instance-name>

```
user@host> show multicast route instance v1 extensive
```

Instance: v1 Family: INET

```
Group: 224.1.1.1
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Group: 224.1.1.2
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Group: 224.1.1.3
```



```
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3
```

```
Instance: v1 Family: INET6
```

#### show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET
```

Route type	Route state	Route count
(S,G)	Active	2
(S,G)	Inactive	3

```
Instance: master Family: INET6
```

## show multicast rpf

---

<b>List of Syntax</b>	<a href="#">Syntax on page 5136</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5136</a>
<b>Syntax</b>	<pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;prefix&gt; &lt;summary&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show multicast rpf &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;prefix&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display information about multicast reverse-path-forwarding (RPF) calculations.
<b>Options</b>	<p><b>none</b>—Display RPF calculation information for all supported address families.</p> <p><b>inet   inet6</b>—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>prefix</b>—(Optional) Display the RPF calculation information for the specified prefix.</p> <p><b>summary</b>—(Optional) Display a summary of all multicast RPF information.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast rpf on page 5137</a> <a href="#">show multicast rpf inet6 on page 5138</a> <a href="#">show multicast rpf prefix on page 5139</a> <a href="#">show multicast rpf summary on page 5139</a>

**Output Fields** Table 416 on page 5137 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

**Table 416: show multicast rpf Output Fields**

Field Name	Field Description
<b>Instance</b>	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
<b>Source prefix</b>	Prefix and length of the source as it exists in the multicast forwarding table.
<b>Protocol</b>	How the route was learned.
<b>Interface</b>	Upstream RPF interface.  <b>NOTE:</b> The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured.
<b>Neighbor</b>	Upstream RPF neighbor.  <b>NOTE:</b> The displayed neighbor information does not apply to bidirectional PIM. This is because the <b>show multicast rpf</b> command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the <b>show pim join extensive</b> command when bidirectional PIM is configured.

## Sample Output

### show multicast rpf

```

user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
  Protocol: Static

10.255.14.132/32
  Protocol: Direct
  Interface: lo0.0

10.255.245.91/32
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0

```

Neighbor: 192.168.14.254

192.168.0.0/16  
Protocol: Static  
Interface: fxp0.0  
Neighbor: 192.168.14.254

192.168.14.0/24  
Protocol: Direct  
Interface: fxp0.0

192.168.14.132/32  
Protocol: Local

192.168.195.20/30  
Protocol: Direct  
Interface: so-1/1/1.0

192.168.195.22/32  
Protocol: Local

192.168.195.36/30  
Protocol: IS-IS  
Interface: so-1/1/1.0  
Neighbor: 192.168.195.21

### show multicast rpf inet6

```
user@host> show multicast rpf inet6
```

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128  
Protocol: Direct  
Interface: lo0.0

::10.255.245.91/128  
Protocol: IS-IS  
Interface: so-1/1/1.0  
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126  
Protocol: Direct  
Interface: so-1/1/1.0

::192.168.195.22/128  
Protocol: Local

::192.168.195.36/126  
Protocol: IS-IS  
Interface: so-1/1/1.0  
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126  
Protocol: Direct  
Interface: fe-2/2/0.0

::192.168.195.77/128  
Protocol: Local

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

#### show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

#### show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

## show multicast scope

<b>List of Syntax</b>	<a href="#">Syntax on page 5140</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5140</a>
<b>Syntax</b>	<pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show multicast scope &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display administratively scoped IP multicast information.
<b>Options</b>	<p><b>none</b>—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p><b>inet   inet6</b>—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast scope on page 5141</a> <a href="#">show multicast scope inet on page 5141</a> <a href="#">show multicast scope inet6 on page 5141</a>
<b>Output Fields</b>	<p><a href="#">Table 417 on page 5140</a> describes the output fields for the <b>show multicast scope</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 417: show multicast scope Output Fields**

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.

Table 417: show multicast scope Output Fields (*continued*)

Field Name	Field Description
Resolve Rejects	Number of kernel resolve rejects.

## Sample Output

### show multicast scope

```
user@host> show multicast scope
```

Scope name	Group Prefix	Interface	Resolve Rejects
232-net	232.232.0.0/16	fe-0/0/0.1	0
local	239.255.0.0/16	fe-0/0/0.1	0
local	ff05::/16	fe-0/0/0.1	0
larry	ff05::1234/128	fe-0/0/0.1	0

### show multicast scope inet

```
user@host> show multicast scope inet
```

Scope name	Group Prefix	Interface	Resolve Rejects
232-net	232.232.0.0/16	fe-0/0/0.1	0
local	239.255.0.0/16	fe-0/0/0.1	0

### show multicast scope inet6

```
user@host> show multicast scope inet6
```

Scope name	Group Prefix	Interface	Resolve Rejects
local	ff05::/16	fe-0/0/0.1	0
larry	ff05::1234/128	fe-0/0/0.1	0

## show multicast sessions

<b>List of Syntax</b>	<a href="#">Syntax on page 5142</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5142</a>
<b>Syntax</b>	show multicast sessions <brief   detail   extensive> <logical-system (all   <i>logical-system-name</i> )> < <i>regular-expression</i> >
<b>Syntax (EX Series Switch and the QFX Series)</b>	show multicast sessions <brief   detail   extensive> < <i>regular-expression</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display information about announced IP multicast sessions.
<b>Options</b>	<b>none</b> —Display standard information about all multicast sessions for all routing instances.  <b>brief   detail   extensive</b> —(Optional) Display the specified level of output.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b><i>regular-expression</i></b> —(Optional) Display information about announced sessions that match a UNIX-style regular expression.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast sessions on page 5143</a> <a href="#">show multicast sessions regular-expression detail on page 5143</a>
<b>Output Fields</b>	Table 418 on page 5142 describes the output fields for the <b>show multicast sessions</b> command. Output fields are listed in the approximate order in which they appear.

**Table 418: show multicast sessions Output Fields**

Field Name	Field Description
<i>session-name</i>	Name of the known announced multicast sessions.



## Sample Output

### show multicast sessions

```

user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.

```

### show multicast sessions regular-expression detail

```

user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2

```

Attribute: rtpmap:104 L16/22050

1 matching sessions.

## show multicast usage

<b>List of Syntax</b>	<a href="#">Syntax on page 5145</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5145</a>
<b>Syntax</b>	<pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show multicast usage &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.
<b>Options</b>	<p><b>none</b>—Display multicast usage information for all supported address families for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show multicast usage on page 5146</a> <a href="#">show multicast usage brief on page 5146</a> <a href="#">show multicast usage instance on page 5146</a> <a href="#">show multicast usage detail on page 5147</a>
<b>Output Fields</b>	<p><a href="#">Table 419 on page 5146</a> describes the output fields for the <b>show multicast usage</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 419: show multicast usage Output Fields

Field Name	Field Description
<b>Instance</b>	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
<b>Group</b>	Group address.
<b>Sources</b>	Number of sources.
<b>Packets</b>	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays <b>unavailable</b> .
<b>Bytes</b>	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays <b>unavailable</b> .
<b>Prefix</b>	IP address.
<b>/len</b>	Prefix length.
<b>Groups</b>	Number of multicast groups.

## Sample Output

### show multicast usage

```

user@host> show multicast usage
Group          Sources  Packets      Bytes
228.0.0.0      1        52847      4439148
239.1.1.1      2        13450      1125530

Prefix         /len  Groups  Packets      Bytes
10.255.14.144  /32   2        66254      5561304
10.255.70.15   /32   1         43        3374...
```

### show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 5146](#).

### show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group          Sources  Packets      Bytes
224.2.127.254  1        5538      509496
224.0.1.39     1         13         624
224.0.1.40     1         13         624

Prefix         /len  Groups  Packets      Bytes
192.168.195.34 /32   1        5538      509496
10.255.14.30   /32   1         13         624
```

```
10.255.245.91 /32 1 13 624
...
```

### show multicast usage detail

```
user@host> show multicast usage detail
Group          Sources Packets          Bytes
228.0.0.0      1          53159          4465356
  Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356
239.1.1.1      2          13450          1125530
  Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156
  Source: 10.255.70.15  /32 Packets: 43 Bytes: 3374
```

```
Prefix          /len Groups Packets          Bytes
10.255.14.144   /32 2          66566          5587512
  Group: 228.0.0.0      Packets: 53159 Bytes: 4465356
  Group: 239.1.1.1      Packets: 13407 Bytes: 1122156
10.255.70.15    /32 1          43             3374
  Group: 239.1.1.1      Packets: 43 Bytes: 3374
```

## show pim bootstrap

<b>List of Syntax</b>	<a href="#">Syntax on page 5148</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5148</a>
<b>Syntax</b>	<pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show pim bootstrap &lt;instance <i>instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
<b>Description</b>	For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.
<b>Options</b>	<p><b>none</b>—Display PIM bootstrap router information for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show pim bootstrap on page 5149</a> <a href="#">show pim bootstrap instance on page 5149</a>
<b>Output Fields</b>	<p><a href="#">Table 420 on page 5148</a> describes the output fields for the <b>show pim bootstrap</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 420: show pim bootstrap Output Fields**

Field Name	Field Description
<b>Instance</b>	Name of the routing instance.
<b>BSR</b>	Bootstrap router.
<b>Pri</b>	Priority of the routing device as elected to be the bootstrap router.
<b>Local address</b>	Local routing device address.
<b>Pri</b>	Local routing device address priority to be elected as the bootstrap router.

Table 420: show pim bootstrap Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	Local routing device election state: <b>Candidate</b> , <b>Elected</b> , or <b>Ineligible</b> .
<b>Timeout</b>	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

## Sample Output

### show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

BSR	Pri	Local address	Pri	State	Timeout
None	0	10.255.71.46	0	InEligible	0
feco:1:1:1:1:0:aff:785c	34	feco:1:1:1:1:0:aff:7c12	0	InEligible	0

### show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

BSR	Pri	Local address	Pri	State	Timeout
None	0	192.168.196.105	0	InEligible	0

## show pim interfaces

<b>List of Syntax</b>	<a href="#">Syntax on page 5150</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5150</a>
<b>Syntax</b>	<pre>show pim interfaces &lt;inet   inet6&gt; &lt;instance (instance-name   all)&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show pim interfaces &lt;inet   inet6&gt; &lt;instance (instance-name   all)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p>
<b>Description</b>	Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.
<b>Options</b>	<p><b>none</b>—Display interface information for all family addresses for the main instance.</p> <p><b>inet   inet6</b>—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (instance-name   all)</b>—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show pim interfaces on page 5151</a>
<b>Output Fields</b>	<p><a href="#">Table 421 on page 5150</a> describes the output fields for the <b>show pim interfaces</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 421: show pim interfaces Output Fields

Field Name	Field Description
<b>Instance</b>	Name of the routing instance.
<b>Name</b>	Interface name.
<b>State</b>	State of the interface. The state also is displayed in the <b>show interfaces</b> command.



Table 421: show pim interfaces Output Fields (*continued*)

Field Name	Field Description
<b>Mode</b>	<p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> <li>• <b>B</b>—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers.</li> <li>• <b>S</b>—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic.</li> <li>• <b>Dense</b>—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.)</li> <li>• <b>Sparse-Dense</b>—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as <b>dense</b> is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as <b>sparse</b> is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.)</li> </ul> <p>When sparse-dense mode is configured, the output includes both <b>S</b> and <b>D</b>. When bidirectional-sparse mode is configured, the output includes <b>S</b> and <b>B</b>. When bidirectional-sparse-dense mode is configured, the output includes <b>B</b>, <b>S</b>, and <b>D</b>.</p>
<b>IP</b>	Version number of the address family on the interface: <b>4</b> (IPv4) or <b>6</b> (IPv6).
<b>V</b>	PIM version running on the interface: 1 or 2.
<b>State</b>	<p>State of PIM on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Active</b>—Bidirectional mode is enabled on the interface and on all PIM neighbors.</li> <li>• <b>DR</b>—Designated router.</li> <li>• <b>NotCap</b>—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol.</li> <li>• <b>NotDR</b>—Not the designated router.</li> <li>• <b>P2P</b>—Point to point.</li> </ul>
<b>NbrCnt</b>	Number of neighbors that have been seen on the interface.
<b>JoinCnt(sg)</b>	Number of (s,g) join messages that have been seen on the interface.
<b>JointCnt(*g)</b>	Number of (*g) join messages that have been seen on the interface.
<b>DR address</b>	Address of the designated router.

## Sample Output

### show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```

Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name	Stat	Mode	IP	V	State	NbrCnt	JoinCnt(sg/*g)	DR address
ge-0/3/0.0	Up	S	4	2	NotDR,NotCap	1	0/0	40.0.0.3
ge-0/3/3.50	Up	S	4	2	DR,NotCap	1	9901/100	50.0.0.2
ge-0/3/3.51	Up	S	4	2	DR,NotCap	1	0/0	51.0.0.2
pe-1/2/0.32769	Up	S	4	2	P2P,NotCap	0	0/0	

show pim join

List of Syntax	<a href="#">Syntax on page 5153</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5153</a>
Syntax	<pre>show pim join &lt;brief   detail   extensive   summary&gt; &lt;bidirectional   dense   sparse&gt; &lt;exact&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;range&gt; &lt;rp <i>ip-address/prefix</i>   source <i>ip-address/prefix</i>&gt; &lt;sg   star-g&gt;</pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show pim join &lt;brief   detail   extensive   summary&gt; &lt;dense   sparse&gt; &lt;exact&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;range&gt; &lt;rp <i>ip-address/prefix</i>   source <i>ip-address/prefix</i>&gt; &lt;sg   star-g&gt;</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>summary</b> option introduced in Junos OS Release 9.6.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Multiple new filter options introduced in Junos OS Release 13.2.</p>
Description	<p>Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.</p> <p>For bidirectional PIM, display information about PIM group ranges (*G-range) for each active bidirectional RP group range, in addition to each of the joined (*G) routes.</p>
Options	<p><b>none</b>—Display the standard information about PIM groups for all supported family addresses for all routing instances.</p> <p><b>brief   detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>bidirectional   dense   sparse</b>—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.</p> <p><b>exact</b>—(Optional) Display information about only the group that exactly matches the specified group address.</p>

**inet | inet6**—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

**instance *instance-name***—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

**range**—(Optional) Address range of the group, specified as *prefix/prefix-length*.

**rp *ip-address/prefix* | source *ip-address/prefix***—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

**sg | star-g**—(Optional) Display information about PIM (S,G) or (\*,G) entries.

**Required Privilege Level**

view

**Related Documentation**

- [clear pim join on page 5065](#)
- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*
- *Example: Configuring Bidirectional PIM*
- *Example: Configuring PIM State Limits*

**List of Sample Output**

[show pim join summary on page 5158](#)  
[show pim join \(PIM Sparse Mode\) on page 5158](#)  
[show pim join \(Bidirectional PIM\) on page 5159](#)  
[show pim join inet6 on page 5159](#)  
[show pim join inet6 star-g on page 5160](#)  
[show pim join instance <instance-name> on page 5160](#)  
[show pim join detail on page 5160](#)  
[show pim join extensive \(PIM Sparse Mode\) on page 5161](#)  
[show pim join extensive \(Bidirectional PIM\) on page 5162](#)  
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 5163](#)  
[show pim join instance <instance-name> extensive on page 5163](#)  
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 5164](#)  
[show pim join extensive \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 5165](#)  
[show pim join summary on page 5167](#)  
[show pim join \(PIM Sparse Mode\) on page 5167](#)  
[show pim join \(Bidirectional PIM\) on page 5167](#)  
[show pim join inet6 on page 5168](#)  
[show pim join inet6 star-g on page 5168](#)  
[show pim join instance <instance-name> on page 5168](#)  
[show pim join detail on page 5169](#)

[show pim join extensive \(PIM Sparse Mode\) on page 5169](#)  
[show pim join extensive \(Bidirectional PIM\) on page 5170](#)  
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 5171](#)  
[show pim join instance <instance-name> extensive on page 5172](#)  
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 5172](#)  
[show pim join extensive \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 5173](#)

**Output Fields** Table 422 on page 5155 describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

**Table 422: show pim join Output Fields**

Field Name	Field Description	Level of Output
<b>Instance</b>	Name of the routing instance.	<b>brief detail extensive summary none</b>
<b>Family</b>	Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).	<b>brief detail extensive summary none</b>
<b>Route type</b>	Type of multicast route: (S,G) or (*G).	<b>summary</b>
<b>Route count</b>	Number of (S,G) routes and number of (*G) routes.	<b>summary</b>
<b>R</b>	Rendezvous Point Tree.	<b>brief detail extensive none</b>
<b>S</b>	Sparse.	<b>brief detail extensive none</b>
<b>W</b>	Wildcard.	<b>brief detail extensive none</b>
<b>Group</b>	Group address.	<b>brief detail extensive none</b>
<b>Bidirectional group prefix length</b>	For bidirectional PIM, length of the IP prefix for RP group ranges.	All levels
<b>Source</b>	Multicast source: <ul style="list-style-type: none"> <li>• * (wildcard value)</li> <li>• <i>ipv4-address</i></li> <li>• <i>ipv6-address</i></li> </ul>	<b>brief detail extensive none</b>
<b>RP</b>	Rendezvous point for the PIM group.	<b>brief detail extensive none</b>

Table 422: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Flags</b>	<p>PIM flags:</p> <ul style="list-style-type: none"> <li>• <b>bidirectional</b>—Bidirectional mode entry.</li> <li>• <b>dense</b>—Dense mode entry.</li> <li>• <b>rptree</b>—Entry is on the rendezvous point tree.</li> <li>• <b>sparse</b>—Sparse mode entry.</li> <li>• <b>spt</b>—Entry is on the shortest-path tree for the source.</li> <li>• <b>wildcard</b>—Entry is on the shared tree.</li> </ul>	<b>brief detail extensive none</b>
<b>Upstream interface</b>	<p>RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*,G).</p> <p>For bidirectional PIM, <b>RP Link</b> means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>	<b>brief detail extensive none</b>
<b>Upstream neighbor</b>	<p>Information about the upstream neighbor: <b>Direct</b>, <b>Local</b>, <b>Unknown</b>, or a specific IP address.</p> <p>For bidirectional PIM, <b>Direct</b> means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>	<b>extensive</b>
<b>Active upstream interface</b>	<p>When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.</p>	<b>extensive</b>
<b>Active upstream neighbor</b>	<p>On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.</p>	<b>extensive</b>

Table 422: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
MoFRR Backup upstream interface	<p>The MoFRR upstream interface that is used when the primary path fails.</p> <p>When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.</p>	<b>extensive</b>
<b>Upstream state</b>	<p>Information about the upstream interface:</p> <ul style="list-style-type: none"> <li>• <b>Join to RP</b>—Sending a join to the rendezvous point.</li> <li>• <b>Join to Source</b>—Sending a join to the source.</li> <li>• <b>Local RP</b>—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point.</li> <li>• <b>Local Source</b>—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device.</li> <li>• <b>Prune to RP</b>—Sending a prune to the rendezvous point.</li> <li>• <b>Prune to Source</b>—Sending a prune to the source.</li> </ul> <p><b>NOTE:</b> RP group range entries have <b>None</b> in the <b>Upstream state</b> field because RP group ranges do not trigger actual PIM join messages between routing devices.</p>	<b>extensive</b>
<b>Downstream neighbors</b>	<p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> <li>• <b>Interface</b>—Interface name for the downstream neighbor.</li> </ul> <p>A pseudo PIM-SM interface appears for all IGMP-only interfaces.</p> <p>A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p> <ul style="list-style-type: none"> <li>• <b>Interface address</b>—Address of the downstream neighbor.</li> <li>• <b>State</b>—Information about the downstream neighbor: <b>join</b> or <b>prune</b>.</li> <li>• <b>Flags</b>—PIM join flags: <b>R (RPtree)</b>, <b>S (Sparse)</b>, <b>W (Wildcard)</b>, or <b>zero</b>.</li> <li>• <b>Uptime</b>—Time since the downstream interface joined the group.</li> <li>• <b>Time since last Join</b>—Time since the last join message was received from the downstream interface.</li> <li>• <b>Time since last Prune</b>—Time since the last prune message was received from the downstream interface.</li> </ul>	<b>extensive</b>

Table 422: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Number of downstream interfaces</b>	Total number of outgoing interfaces for each (S,G) entry.	<b>extensive</b>
<b>Assert Timeout</b>	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	<b>extensive</b>
<b>Keepalive timeout</b>	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, <b>Keepalive timeout</b> is <b>Infinity</b> .	<b>extensive</b>
<b>Uptime</b>	Time since the creation of (S,G) or (*G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*G) state.	<b>extensive</b>
<b>Bidirectional accepting interfaces</b>	<p>Interfaces on the routing device that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (<b>DF Winner</b>), or the interface is the reverse path forwarding (RPF) interface toward the RP (<b>RPF</b>).</p>	<b>extensive</b>

## Sample Output

### show pim join summary

```

user@host> show pim join summary
Instance: PIM.master Family: INET

Route type      Route count
(s,g)           2
(*,g)           1

Instance: PIM.master Family: INET6

```

### show pim join (PIM Sparse Mode)

```

user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1

```



```

Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join (Bidirectional PIM)

```

user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join inet6

```

user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
  Source: *
  RP: ::46.0.0.13
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

```

```
Group: ff04::e000:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
  Source: ::1.1.1.2
  Flags: sparse
  Upstream interface: unknown (no neighbor)
```

#### show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
  Source: *
  RP: ::46.0.0.13
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
```

#### show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 235.1.1.2
  Source: 192.168.195.74
  Flags: sparse,spt
  Upstream interface: at-0/3/1.0

Group: 235.1.1.2
  Source: 192.168.195.169
  Flags: sparse
  Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
```

```

Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join extensive (PIM Sparse Mode)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: so-1/0/0.0
    10.111.10.2 State: Join Flags: SRW Timeout: 174
    Uptime: 00:03:49 Time since last Join: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: SRW Timeout: Infinity
    Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local Source, Local RP
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
  Interface: so-1/0/0.0
    10.111.10.2 State: Join Flags: S Timeout: 174
    Uptime: 00:03:49 Time since last Prune: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: S Timeout: Infinity
    Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
Upstream neighbor: 10.111.10.2
Upstream state: Local RP, Join to Source

```

```
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
  Interface: Pseudo-GMP
    fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
  Interface: so-1/0/0.0 (pruned)
    10.111.10.2 State: Prune Flags: SR Timeout: 174
    Uptime: 00:03:49 Time since last Prune: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: S Timeout: Infinity
    Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

#### show pim join extensive (Bidirectional PIM)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0 (RPF)
    Interface: lo0.0 (DF Winner)
  Number of downstream interfaces: 0

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0 (RPF)
    Interface: lo0.0 (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join RW Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join RW Timeout: 184
  Number of downstream interfaces: 2

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
```

```

Upstream neighbor: Direct
Upstream state: Local RP
Uptime: 00:03:49
Bidirectional accepting interfaces:
  Interface: ge-0/0/1.0    (RPF)
  Interface: lo0.0         (DF Winner)
  Interface: xe-4/1/0.0    (DF Winner)
Number of downstream interfaces: 0

```

```

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.3
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0    (RPF)
    Interface: lo0.0         (DF Winner)
    Interface: xe-4/1/0.0    (DF Winner)
  Number of downstream interfaces: 0

```

### show pim join instance <instance-name> extensive

```

user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: mt-1/1/0.32768
      10.10.47.101 State: Join Flags: SRW Timeout: 156
      Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 1

Group: 235.1.1.2
  Source: 192.168.195.74
  Flags: sparse,spt
  Upstream interface: at-0/3/1.0
  Upstream neighbor: 10.111.30.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 156
  Uptime: 00:14:52

```

```
Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52
```

#### show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:55
Downstream neighbors:
  Interface: Pseudo-MLDP
  Interface: lt-1/2/0.25
    1.2.5.2 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
  Interface: Pseudo-MLDP

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
  Interface: Pseudo-MLDP

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:25
Downstream neighbors:
```

Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2  
Source: abcd::1:2:7:7  
Flags: sparse,spt  
Upstream interface: lt-1/2/0.27  
Upstream neighbor: Direct  
Upstream state: Local Source  
Keepalive timeout:  
Uptime: 11:27:26  
Downstream neighbors:  
Interface: Pseudo-MLDP

### show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
Source: *
RP: 1.1.1.1
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 11:31:33
Downstream neighbors:
  Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: SRW Timeout: Infinity
    Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:32
Downstream neighbors:
  Interface: so-0/1/3.0
    192.168.92.9 State: Join Flags: S Timeout: Infinity
    Uptime: 11:31:30 Time since last Join: 11:31:30
  Downstream neighbors:
    Interface: fe-1/3/0.0
      192.168.209.9 State: Join Flags: S Timeout: Infinity
      Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
```

```
Uptime: 11:31:32
Downstream neighbors:
  Interface: so-0/1/3.0
    192.168.92.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:30 Time since last Join: 11:31:30
Downstream neighbors:
  Interface: lt-1/2/0.14
    1.1.4.4 State: Join Flags: S Timeout: 177
    Uptime: 11:30:33 Time since last Join: 00:00:33
Downstream neighbors:
  Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:32
Downstream neighbors:
  Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:30
Downstream neighbors:
  Interface: so-0/1/3.0
    192.168.92.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 11:31:32
Downstream neighbors:
  Interface: fe-1/3/0.0
    fe80::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
    Uptime: 11:31:32 Time since last Join: 11:31:32
```



## Sample Output

### show pim join summary

```
user@host> show pim join summary
Instance: PIM.master Family: INET

Route type          Route count
(s,g)               2
(*,g)              1

Instance: PIM.master Family: INET6
```

### show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

### show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
Bidirectional group prefix length: 24
Source: *
```

```
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

### show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: ff04::e000:101
Source: ::1.1.1.1
Flags: sparse
Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
Source: ::1.1.1.1
Flags: sparse
Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
Source: ::1.1.1.2
Flags: sparse
Upstream interface: unknown (no neighbor)
```

### show pim join inet6 star-g

```
user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local
```

### show pim join instance <instance-name>

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
```

```

Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join detail

```

user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### show pim join extensive (PIM Sparse Mode)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: so-1/0/0.0
    10.111.10.2 State: Join Flags: SRW Timeout: 174
    Uptime: 00:03:49 Time since last Join: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: SRW Timeout: Infinity

```

Uptime: 00:03:49 Time since last Join: 00:01:49  
Number of downstream interfaces: 2

Group: 239.1.1.1  
Source: 10.255.14.144  
Flags: sparse,spt  
Upstream interface: Local  
Upstream neighbor: Local  
Upstream state: Local Source, Local RP  
Keepalive timeout: 344  
Uptime: 00:03:49  
Downstream neighbors:  
  Interface: so-1/0/0.0  
    10.111.10.2 State: Join Flags: S Timeout: 174  
    Uptime: 00:03:49 Time since last Prune: 00:01:49  
  Interface: mt-1/1/0.32768  
    10.10.47.100 State: Join Flags: S Timeout: Infinity  
    Uptime: 00:03:49 Time since last Prune: 00:01:49  
Number of downstream interfaces: 2

Group: 239.1.1.1  
Source: 10.255.70.15  
Flags: sparse,spt  
Upstream interface: so-1/0/0.0  
Upstream neighbor: 10.111.10.2  
Upstream state: Local RP, Join to Source  
Keepalive timeout: 344  
Uptime: 00:03:49  
Downstream neighbors:  
  Interface: Pseudo-GMP  
    fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0  
  Interface: so-1/0/0.0 (pruned)  
    10.111.10.2 State: Prune Flags: SR Timeout: 174  
    Uptime: 00:03:49 Time since last Prune: 00:01:49  
  Interface: mt-1/1/0.32768  
    10.10.47.100 State: Join Flags: S Timeout: Infinity  
    Uptime: 00:03:49 Time since last Prune: 00:01:49  
Number of downstream interfaces: 3

Instance: PIM.master Family: INET6  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

### show pim join extensive (Bidirectional PIM)

user@host> show pim join extensive  
Instance: PIM.master Family: INET  
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0  
Bidirectional group prefix length: 24  
Source: \*  
RP: 10.10.13.2  
Flags: bidirectional,rptree,wildcard  
Upstream interface: ge-0/0/1.0  
Upstream neighbor: 10.10.1.2  
Upstream state: None  
Uptime: 00:03:49  
Bidirectional accepting interfaces:  
  Interface: ge-0/0/1.0 (RPF)  
  Interface: lo0.0 (DF Winner)  
Number of downstream interfaces: 0

```

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join   RW   Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

#### show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

```

**show pim join instance <instance-name> extensive**

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 00:03:49
Downstream neighbors:
  Interface: mt-1/1/0.32768
    10.10.47.101 State: Join Flags: SRW Timeout: 156
    Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 1

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0
Upstream neighbor: 10.111.30.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52
```

**show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)**

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:55
Downstream neighbors:
  Interface: Pseudo-MLDP
  Interface: lt-1/2/0.25
    1.2.5.2 State: Join Flags: S Timeout: Infinity
    Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
```

```

Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:25
Downstream neighbors:
    Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
    Interface: Pseudo-MLDP

```

#### show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
Source: 10.0.0.1
Flags: sparse,spt
Active upstream interface: fe-1/2/13.0
Active upstream neighbor: 10.0.0.9
MoFRR Backup upstream interface: fe-1/2/14.0
MoFRR Backup upstream neighbor: 10.0.0.21
Upstream state: Join to Source, No Prune to RP
Keepalive timeout: 354
Uptime: 00:00:06
Downstream neighbors:

```

```
Interface: fe-1/2/15.0
  10.0.0.13 State: Join Flags: S   Timeout: Infinity
    Uptime: 00:00:06 Time since last Join: 00:00:06
Number of downstream interfaces: 1
```



## show pim neighbors

<b>List of Syntax</b>	<a href="#">Syntax on page 5175</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5175</a>
<b>Syntax</b>	<pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (<i>instance-name</i>   all)&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show pim neighbors &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance (<i>instance-name</i>   all)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the <b>instance all</b> option added in Junos OS Release 12.1.</p>
<b>Description</b>	Display information about Protocol Independent Multicast (PIM) neighbors.
<b>Options</b>	<p><b>none</b>—(Same as <b>brief</b>) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance (<i>instance-name</i>   all)</b>—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show pim neighbors on page 5177</a> <a href="#">show pim neighbors brief on page 5177</a> <a href="#">show pim neighbors instance on page 5177</a> <a href="#">show pim neighbors detail on page 5177</a> <a href="#">show pim neighbors detail (With BFD) on page 5178</a>
<b>Output Fields</b>	<p><a href="#">Table 423 on page 5176</a> describes the output fields for the <b>show pim neighbors</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 423: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
<b>Instance</b>	Name of the routing instance.	All levels
<b>Interface</b>	Interface through which the neighbor is reachable.	All levels
<b>Neighbor addr</b>	Address of the neighboring PIM routing device.	All levels
<b>IP</b>	IP version: 4 or 6.	All levels
<b>V</b>	PIM version running on the neighbor: 1 or 2.	All levels
<b>Mode</b>	PIM mode of the neighbor: <b>Sparse</b> , <b>Dense</b> , <b>SparseDense</b> , or <b>Unknown</b> . When the neighbor is running PIM version 2, this mode is always <b>Unknown</b> .	All levels
<b>Option</b>	Can be one or more of the following: <ul style="list-style-type: none"> <li>• <b>B</b>—Bidirectional Capable.</li> <li>• <b>H</b>—Hello Option Holdtime.</li> <li>• <b>G</b>—Generation Identifier.</li> <li>• <b>P</b>—Hello Option DR Priority.</li> <li>• <b>L</b>—Hello Option LAN Prune Delay.</li> </ul>	<b>brief</b> none
<b>Uptime</b>	Time the neighbor has been operational since the PIM process was last initialized, in the format <b>dd:hh:mm:ss ago</b> for less than a week and <b>nwnd:hh:mm:ss ago</b> for more than a week.	All levels
<b>Address</b>	Address of the neighboring PIM routing device.	<b>detail</b>
<b>BFD</b>	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: <b>Enabled</b> , <b>Operational state is up</b> , or <b>Disabled</b> .	<b>detail</b>
<b>Hello Option Holdtime</b>	Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.	<b>detail</b>
<b>Hello Default Holdtime</b>	Default holdtime and the time remaining if the <b>holdtime</b> option is not in the received hello message.	<b>detail</b>
<b>Hello Option DR Priority</b>	Designated router election priority. The range of values is 0 through 255.	<b>detail</b>
<b>Hello Option Generation ID</b>	9-digit or 10-digit number used to tag hello messages.	<b>detail</b>
<b>Hello Option Bi-Directional PIM supported</b>	Neighbor can process bidirectional PIM messages.	<b>detail</b>
<b>Hello Option LAN Prune Delay</b>	Time to wait before the neighbor receives prune messages, in the format <b>delay nnn ms override nnnn ms</b> .	<b>detail</b>

Table 423: show pim neighbors Output Fields (*continued*)

Field Name	Field Description	Level of Output
Join Suppression supported	Neighbor is capable of join suppression.	detail
Rx Join	Information about joins received from the neighbor. <ul style="list-style-type: none"> <li><b>Group</b>—Group addresses in the join message.</li> <li><b>Source</b>—Address of the source in the join message.</li> <li><b>Timeout</b>—Time for which the join is valid.</li> </ul>	detail

## Sample Output

### show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface      IP V Mode      Option      Uptime Neighbor addr
so-1/0/0.0      4 2            HPLG        00:07:10 10.111.10.2

```

### show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 5177](#).

### show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface      IP V Mode      Option      Uptime Neighbor addr
at-0/3/1.0      4 2            HPLG        00:07:54 10.111.30.2
mt-1/1/0.32768  4 2            HPLG        00:07:22 10.10.47.101
so-1/0/1.0      4 2            HPLG        00:07:50 10.111.20.2

```

### show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, ts
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```
Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
  BFD: Disabled
  Hello Option Holdtime: 105 seconds 93 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1734018161
  Hello Option Bi-Directional PIM supported
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
```

Interface: lo0.0

```
Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1997462267
  Hello Option Bi-Directional PIM supported
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                                Join Suppression supported
```

#### show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
```

Instance: PIM.master

Interface: fe-1/0/0.0

```
Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
  Hello Option Holdtime: 65535 seconds
  Hello Option DR Priority: 1
  Hello Option Generation ID: 836607909
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.11.2, IPv4, PIM v2
  BFD: Enabled, Operational state is up
  Hello Default Holdtime: 105 seconds 104 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1907549685
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Interface: fe-1/0/1.0

```
Address: 192.168.12.1, IPv4, PIM v2
  BFD: Disabled
  Hello Default Holdtime: 105 seconds 80 remaining
  Hello Option DR Priority: 1
  Hello Option Generation ID: 1971554705
  Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

## show pim rps

<b>List of Syntax</b>	<a href="#">Syntax on page 5179</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5179</a>
<b>Syntax</b>	<pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show pim rps &lt;brief   detail   extensive&gt; &lt;group-address&gt; &lt;inet   inet6&gt; &lt;instance instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>
<b>Description</b>	Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).
<b>Options</b>	<p><b>none</b>—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>group-address</b>—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance instance-name</b>—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Bidirectional PIM</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show pim rps on page 5182</a> <a href="#">show pim rps brief on page 5182</a> <a href="#">show pim rps &lt;group-address&gt; (Bidirectional PIM) on page 5182</a>

[show pim rps <group-address> \(PIM Dense Mode\) on page 5182](#)  
[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 5182](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 5183](#)  
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 5183](#)  
[show pim rps instance on page 5183](#)  
[show pim rps extensive \(PIM Sparse Mode\) on page 5183](#)  
[show pim rps extensive \(Bidirectional PIM\) on page 5184](#)  
[show pim rps extensive \(PIM Anycast RP in Use\) on page 5184](#)

**Output Fields** [Table 424 on page 5180](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

**Table 424: show pim rps Output Fields**

Field Name	Field Description	Level of Output
<b>Instance</b>	Name of the routing instance.	All levels
<b>Family or Address family</b>	Name of the address family: <b>inet</b> (IPv4) or <b>inet6</b> (IPv6).	All levels
<b>RP address</b>	Address of the rendezvous point.	All levels
<b>Type</b>	Type of RP: <ul style="list-style-type: none"> <li><b>auto-rp</b>—Address of the RP known through the Auto-RP protocol.</li> <li><b>bootstrap</b>—Address of the RP known through the bootstrap router protocol (BSR).</li> <li><b>embedded</b>—Address of the RP known through an embedded RP (IPv6).</li> <li><b>static</b>—Address of RP known through static configuration.</li> </ul>	<b>brief none</b>
<b>Holdtime</b>	How long to keep the RP active, with time remaining, in seconds.	All levels
<b>Timeout</b>	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
<b>Groups</b>	Number of groups currently using this RP.	All levels
<b>Group prefixes</b>	Addresses of groups that this RP can span.	<b>brief none</b>
<b>Learned via</b>	Address and method by which the RP was learned.	<b>detail extensive</b>
<b>Mode</b>	The PIM mode of the RP: bidirectional or sparse.  If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.	All levels
<b>Time Active</b>	How long the RP has been active, in the format <b>hh:mm:ss</b> .	<b>detail extensive</b>

Table 424: show pim rps Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Device Index</b>	Index value of the order in which Junos OS finds and initializes the interface.  For bidirectional RPs, the <b>Device Index</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	<b>detail extensive</b>
<b>Subunit</b>	Logical unit number of the interface.  For bidirectional RPs, the <b>Subunit</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	<b>detail extensive</b>
<b>Interface</b>	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively.  For bidirectional RPs, the <b>Interface</b> output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	<b>detail extensive</b>
<b>Group Ranges</b>	Addresses of groups that this RP spans.	<b>detail extensive</b>  <i>group-address</i>
<b>Active groups using RP</b>	Number of groups currently using this RP.	<b>detail extensive</b>
<b>total</b>	Total number of active groups for this RP.	<b>detail extensive</b>
<b>Register State for RP</b>	Current register state for each group: <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively:</li> <li>• <b>First Hop</b>—PIM-designated routing device that sent the Register message (the source address in the IP header).</li> <li>• <b>RP Address</b>—RP to which the Register message was sent (the destination address in the IP header).</li> <li>• <b>State</b>: <ul style="list-style-type: none"> <li>On the designated router: <ul style="list-style-type: none"> <li>• <b>Send</b>—Sending Register messages.</li> <li>• <b>Probe</b>—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages.</li> <li>• <b>Suppress</b>—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to <b>Probe</b> state.</li> </ul> </li> <li>On the RP: <ul style="list-style-type: none"> <li>• <b>Receive</b>—Receiving Register messages.</li> </ul> </li> </ul> </li> </ul>	<b>extensive</b>
<b>Anycast-PIM rpset</b>	If anycast RP is configured, the addresses of the RPs in the set.	<b>extensive</b>
<b>Anycast-PIM local address used</b>	If anycast RP is configured, the local address used by the RP.	<b>extensive</b>

Table 424: show pim rps Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Anycast-PIM Register State</b>	<p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> <li>• <b>Group</b>—Multicast group address.</li> <li>• <b>Source</b>—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively.</li> <li>• <b>Origin</b>—How the information was obtained: <ul style="list-style-type: none"> <li>• <b>DIRECT</b>—From a local attachment</li> <li>• <b>MSDP</b>—From the Multicast Source Discovery Protocol (MSDP)</li> <li>• <b>DR</b>—From the designated router</li> </ul> </li> </ul>	<b>extensive</b>
<b>RP selected</b>	For sparse mode and bidirectional mode, the identity of the RP for the specified group address.	<i>group-address</i>

## Sample Output

### show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type      Mode   Holdtime Timeout Groups  Group prefixes
10.10.1.3       static   bidir   150     None     2  224.1.3.0/24
                225.1.3.0/24
10.10.13.2      static   bidir   150     None     2  224.1.1.0/24
                225.1.1.0/24

```

### show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 5182](#).

### show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
  11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

### show pim rps <group-address> (PIM Dense Mode)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1

```

### show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```

user@host> show pim rps 224.1.1.1

```



Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75

RP selected: 11.4.12.75

#### show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16  
11.4.12.75 (Bidirectional)

RP selected: (null)

#### show pim rps instance

user@host> show pim rps instance VPN-A

Instance: PIM.VPN-A

Address family INET

RP address	Type	Holdtime	Timeout	Groups	Group prefixes
10.10.47.100	static	0	None	1	224.0.0.0/4

Address family INET6

#### show pim rps extensive (PIM Sparse Mode)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET

RP: 10.255.245.91

Learned via: static configuration

Time Active: 00:05:48

Holdtime: 45 with 36 remaining

Device Index: 122

Subunit: 32768

Interface: pd-6/0/0.32768

Group Ranges:

224.0.0.0/4, 36s remaining

Active groups using RP:

225.1.1.1

total 1 groups active

Register State for RP:

Group	Source	FirstHop	RP Address	State	Timeout
225.1.1.1	192.168.195.78	10.255.14.132	10.255.245.91	Receive	0

**show pim rps extensive (Bidirectional PIM)**

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
    224.1.3.0/24
    225.1.3.0/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
    224.1.1.0/24
    225.1.1.0/24
```

**show pim rps extensive (PIM Anycast RP in Use)**

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
    10.100.111.34
    10.100.111.17
    10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:

```

Group	Source	Origin
224.1.1.1	10.10.95.2	DIRECT
224.1.1.2	10.10.95.2	DIRECT
224.10.10.10	10.10.70.1	MSDP
224.10.10.11	10.10.70.1	MSDP
224.20.20.1	10.10.71.1	DR

```
Address family INET6

Anycast-PIM rpset:
```

```
ab::1
ab::2
Anycast-PIM local address used: cd::1
```

Anycast-PIM Register State:

Group	Source	Origin
::224.1.1.1	::10.10.95.2	DIRECT
::224.1.1.2	::10.10.95.2	DIRECT
::224.20.20.1	::10.10.71.1	DR

## show pim source

---

<b>List of Syntax</b>	<a href="#">Syntax on page 5186</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5186</a>
<b>Syntax</b>	<pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;source-prefix&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show pim source &lt;brief   detail&gt; &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;source-prefix&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.
<b>Options</b>	<p><b>none</b>—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>inet   inet6</b>—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>source-prefix</b>—(Optional) Display the state for source RPF states in the given range.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show pim source on page 5187</a> <a href="#">show pim source brief on page 5187</a> <a href="#">show pim source detail on page 5187</a> <a href="#">show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 5188</a>
<b>Output Fields</b>	<a href="#">Table 425 on page 5187</a> describes the output fields for the <b>show pim source</b> command. Output fields are listed in the approximate order in which they appear.

Table 425: show pim source Output Fields

Field Name	Field Description
<b>Instance</b>	Name of the routing instance.
<b>Source</b>	Address of the source or reverse path.
<b>Prefix/length</b>	Prefix and prefix length for the route used to reach the RPF address.
<b>Upstream Protocol</b>	Protocol toward the source address.
<b>Upstream interface</b>	RPF interface toward the source address.  A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.
<b>Upstream Neighbor</b>	Address of the RPF neighbor used to reach the source address.  The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.

## Sample Output

### show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local

Source 10.255.70.15
  Prefix 10.255.70.15/32
  Upstream interface so-1/0/0.0
  Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

### show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 5187](#).

### show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local
  Active groups:228.0.0.0
    239.1.1.1

```

239.1.1.1

Source 10.255.70.15  
Prefix 10.255.70.15/32  
Upstream interface so-1/0/0.0  
Upstream neighbor 10.111.10.2  
Active groups:239.1.1.1

Instance: PIM.master Family: INET6

#### show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

user@host> show pim source

Instance: PIM.master Family: INET

Source 1.1.1.1  
Prefix 1.1.1.1/32  
Upstream interface Local  
Upstream neighbor Local

Source 1.2.7.7  
Prefix 1.2.7.0/24  
Upstream protocol MLDP  
Upstream interface Pseudo MLDP  
Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11  
Prefix 192.168.219.0/28  
Upstream protocol MLDP  
Upstream interface Pseudo MLDP  
Upstream neighbor MLDP LSP root <1.1.1.2>

Instance: PIM.master Family: INET6

Source abcd::1:2:7:7  
Prefix abcd::1:2:7:0/120  
Upstream protocol MLDP  
Upstream interface Pseudo MLDP  
Upstream neighbor MLDP LSP root <1.1.1.2>

## show pim statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 5189</a> <a href="#">Syntax (EX Series Switch and the QFX Series) on page 5189</a>
<b>Syntax</b>	<pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switch and the QFX Series)</b>	<pre>show pim statistics &lt;inet   inet6&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>inet6</b> and <b>instance</b> options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>
<b>Description</b>	Display Protocol Independent Multicast (PIM) statistics.
<b>Options</b>	<p><b>none</b>—Display PIM statistics.</p> <p><b>inet   inet6</b>—(Optional) Display IPv4 or IPv6 PIM statistics, respectively.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).</p> <p><b>interface <i>interface-name</i></b>—(Optional) Display statistics about the specified interface.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear pim statistics on page 5069</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show pim statistics on page 5196</a> <a href="#">show pim statistics inet interface &lt;interface-name&gt; on page 5198</a> <a href="#">show pim statistics inet6 interface &lt;interface-name&gt; on page 5198</a> <a href="#">show pim statistics instance &lt;instance-name&gt; on page 5199</a> <a href="#">show pim statistics interface &lt;interface-name&gt; on page 5200</a>
<b>Output Fields</b>	<p><a href="#">Table 426 on page 5190</a> describes the output fields for the <b>show pim statistics</b> command.</p> <p>Output fields are listed in the approximate order in which they appear.</p>

Table 426: show pim statistics Output Fields

Field Name	Field Description
Instance	<p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet interface <i>interface-name</i></b></li> <li>• <b>inet6 interface <i>interface-name</i></b></li> <li>• <b>interface <i>interface-name</i></b></li> </ul>
Family	<p>Output is for IPv4 or IPv6 PIM statistics. <b>INET</b> indicates IPv4 statistics, and <b>INET6</b> indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> <li>• <b>inet interface <i>interface-name</i></b></li> <li>• <b>inet6 interface <i>interface-name</i></b></li> <li>• <b>interface <i>interface-name</i></b></li> </ul>
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgment packets.
V2 Candidate RP	PIM version 2 candidate RP packets.
V2 State Refresh	<p>PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.</p> <p>State refresh is an extension to PIM-DM. It not supported in Junos OS.</p>



Table 426: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.

Table 426: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the routing device is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to the source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the routing device has a source address filter configured for the RP.
Rx Unknown Reg Stop	Number of register stop messages received with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the routing device has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.

Table 426: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the routing device has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the routing device has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encaps if	Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the routing device has an RP mismatch.
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.

Table 426: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the <b>maximum-rps</b> statement is exceeded. The <b>maximum-rps</b> statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	<p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the routing device:</p> <ul style="list-style-type: none"> <li>• Multicast Listener Discovery (MLD) report for an embedded RP multicast group address</li> <li>• PIM join message with an embedded RP multicast group address</li> <li>• Static embedded RP multicast group address associated with an interface</li> <li>• Packets sent to an embedded RP multicast group address received on the DR</li> </ul> <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p>
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of received register messages dropped because of a filter configured for PIM register messages.
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.
Rx Bidir Join/Prune on non-Bidir if	Error counter for join and prune messages received on non-bidirectional PIM interfaces.
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.
V4 (S,G) Maximum	Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.

Table 426: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V4 (S,G) Accepted	Number of accepted (S,G) IPv4 multicast routes.
V4 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).
V4 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V6 (S,G) Maximum	Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V6 (S,G) Accepted	Number of accepted (S,G) IPv6 multicast routes.
V6 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).
V6 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V4 (grp-prefix, RP) Maximum	Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V4 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv4 multicast mappings.
V4 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).
V4 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V6 (grp-prefix, RP) Maximum	Maximum number of group-to-RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V6 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv6 multicast mappings.
V6 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).
V6 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.

Table 426: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V4 Register Maximum	Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.  You configure the register limits on the RP.
V4 Register Accepted	Number of accepted IPv4 PIM registers.
V4 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).
V4 Register Log Interval	Time (in seconds) between consecutive log messages.
V6 Register Maximum	Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted.  You configure the register limits on the RP.
V6 Register Accepted	Number of accepted IPv6 PIM registers.
V6 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).
V6 Register Log Interval	Time (in seconds) between consecutive log messages.
(*G) Join drop due to SSM range check	PIM join messages that are dropped because the multicast addresses are outside of the SSM address range of 232.0.0.0 through 232.255.255.255. You can extend the accepted SSM address range by configuring the <a href="#">ssm-groups</a> statement.

## Sample Output

### show pim statistics

```

user@host> show pim statistics
PIM Message type    Received    Sent    Rx errors
V2 Hello            15          32        0
V2 Register          0          362       0
V2 Register Stop    483          0        0
V2 Join Prune       18          518       0
V2 Bootstrap        0           0         0
V2 Assert           0           0         0
V2 Graft            0           0         0
V2 Graft Ack        0           0         0
V2 Candidate RP     0           0         0
V2 State Refresh    0           0         0
V2 DF Election      0           0         0
V1 Query            0           0         0
V1 Register         0           0         0
V1 Register Stop    0           0         0
V1 Join Prune       0           0         0

```

V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

## Global Statistics

Hello dropped on neighbor policy	0
Unknown type	0
V1 Unknown type	0
Unknown Version	0
ipv4 BSR pkt drop due to excessive rate	0
ipv6 BSR pkt drop due to excessive rate	0
Neighbor unknown	0
Bad Length	0
Bad Checksum	0
Bad Receive If	0
Rx Bad Data	0
Rx Intf disabled	0
Rx V1 Require V2	0
Rx V2 Require V1	0
Rx Register not RP	0
Rx Register no route	0
Rx Register no decap if	0
Null Register Timeout	0
RP Filtered Source	0
Rx Unknown Reg Stop	0
Rx Join/Prune no state	0
Rx Join/Prune on upstream if	0
Rx Join/Prune for invalid group	5
Rx Join/Prune messages dropped	0
Rx sparse join for dense group	0
Rx Graft/Graft Ack no state	0
Rx Graft on upstream if	0
Rx CRP not BSR	0
Rx BSR when BSR	0
Rx BSR not RPF if	0
Rx unknown hello opt	0
Rx data no state	0
Rx RP no state	0
Rx aggregate	0
Rx malformed packet	0
Rx illegal TTL	0
Rx illegal destination address	0
No RP	0
No register encap if	0
No route upstream	0
Nexthop Unusable	0
RP mismatch	0
RP mode mismatch	0
RPF neighbor unknown	0
Rx Joins/Prunes filtered	0
Tx Joins/Prunes filtered	0
Embedded-RP invalid addr	0
Embedded-RP limit exceed	0
Embedded-RP added	0

```
Embedded-RP removed                0
Rx Register msgs filtering drop      0
Tx Register msgs filtering drop      0
Rx Bidir Join/Prune on non-Bidir if 0
Rx Bidir Join/Prune on non-DF if    0
(*,G) Join drop due to SSM range check 0
```

## Sample Output

**show pim statistics inet interface <interface-name>**

```
user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET
```

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	4	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
V1 Query	0	0	0
V1 Register	0	0	0
V1 Register Stop	0	0	0
V1 Join Prune	0	0	0
V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

## Sample Output

**show pim statistics inet6 interface <interface-name>**

```
user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6
```

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	4	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
Anycast Register	0	0	0
Anycast Register Stop	0	0	0



## show pim statistics instance &lt;instance-name&gt;

```

user@host> show pim statistics instance VPN-A
PIM Message type      Received      Sent  Rx errors
V2 Hello              31           37      0
V2 Register           0            0      0
V2 Register Stop      0            0      0
V2 Join Prune         0           16      0
V2 Bootstrap          0            0      0
V2 Assert             0            0      0
V2 Graft              0            0      0
V2 Graft Ack          0            0      0
V2 Candidate RP       0            0      0
V2 State Refresh      0            0      0
V2 DF Election        0            0      0
V1 Query              0            0      0
V1 Register           0            0      0
V1 Register Stop      0            0      0
V1 Join Prune         0            0      0
V1 RP Reachability    0            0      0
V1 Assert             0            0      0
V1 Graft              0            0      0
V1 Graft Ack          0            0      0
AutoRP Announce       0            0      0
AutoRP Mapping         0            0      0
AutoRP Unknown type   0            0      0
Anycast Register      0            0      0
Anycast Register Stop 0            0      0

```

## Global Statistics

```

Hello dropped on neighbor policy      0
Unknown type                          0
V1 Unknown type                       0
Unknown Version                       0
Neighbor unknown                      0
Bad Length                            0
Bad Checksum                          0
Bad Receive If                        0
Rx Bad Data                           0
Rx Intf disabled                      0
Rx V1 Require V2                      0
Rx V2 Require V1                      0
Rx Register not RP                    0
Rx Register no route                  0
Rx Register no decap if                0
Null Register Timeout                 0
RP Filtered Source                    0
Rx Unknown Reg Stop                   0
Rx Join/Prune no state                0
Rx Join/Prune on upstream if          0
Rx Join/Prune for invalid group        0
Rx Join/Prune messages dropped         0
Rx sparse join for dense group         0
Rx Graft/Graft Ack no state           0
Rx Graft on upstream if               0
Rx CRP not BSR                        0
Rx BSR when BSR                       0
Rx BSR not RPF if                     0
Rx unknown hello opt                  0
Rx data no state                      0

```

Rx RP no state	0
Rx aggregate	0
Rx malformed packet	0
Rx illegal TTL	0
Rx illegal destination address	0
No RP	0
No register encap if	0
No route upstream	28
Nexthop Unusable	0
RP mismatch	0
RP mode mismatch	0
RPF neighbor unknown	0
Rx Joins/Prunes filtered	0
Tx Joins/Prunes filtered	0
Embedded-RP invalid addr	0
Embedded-RP limit exceed	0
Embedded-RP added	0
Embedded-RP removed	0
Rx Register msgs filtering drop	0
Tx Register msgs filtering drop	0
Rx Bidir Join/Prune on non-Bidir if	0
Rx Bidir Join/Prune on non-DF if	0
V4 (S,G) Maximum	10
V4 (S,G) Accepted	9
V4 (S,G) Threshold	80
V4 (S,G) Log Interval	80
V6 (S,G) Maximum	8
V6 (S,G) Accepted	8
V6 (S,G) Threshold	50
V6 (S,G) Log Interval	100
V4 (grp-prefix, RP) Maximum	100
V4 (grp-prefix, RP) Accepted	5
V4 (grp-prefix, RP) Threshold	80
V4 (grp-prefix, RP) Log Interval	10
V6 (grp-prefix, RP) Maximum	20
V6 (grp-prefix, RP) Accepted	0
V6 (grp-prefix, RP) Threshold	90
V6 (grp-prefix, RP) Log Interval	20
V4 Register Maximum	100
V4 Register Accepted	10
V4 Register Threshold	80
V4 Register Log Interval	10
V6 Register Maximum	20
V6 Register Accepted	0
V6 Register Threshold	90
V6 Register Log Interval	20
(*,G) Join drop due to SSM range check	0

## Sample Output

show pim statistics interface <interface-name>

```

user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type      Received      Sent  Rx errors
V2 Hello                0            3        0
V2 Register             0            0        0
V2 Register Stop        0            0        0

```

V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
V1 Query	0	0	0
V1 Register	0	0	0
V1 Register Stop	0	0	0
V1 Join Prune	0	0	0
V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	3	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

## show system statistics igmp

---

<b>List of Syntax</b>	<a href="#">Syntax on page 5202</a> <a href="#">Syntax (EX Series Switches) on page 5202</a> <a href="#">Syntax (TX Matrix Router) on page 5202</a> <a href="#">Syntax (TX Matrix Plus Router) on page 5202</a>
<b>Syntax</b>	show system statistics igmp
<b>Syntax (EX Series Switches)</b>	show system statistics igmp <all-members> <local> <member <i>member-id</i> >
<b>Syntax (TX Matrix Router)</b>	show system statistics igmp <all-chassis   all-lcc   lcc <i>number</i>   scc>
<b>Syntax (TX Matrix Plus Router)</b>	show system statistics igmp <all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. <b>sfc</b> option introduced for the TX Matrix Plus router in Junos OS Release 9.6. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Display system-wide Internet Group Management Protocol (IGMP) statistics.
<b>Options</b>	<b>none</b> —Display system statistics for IGMP.  <b>all-chassis</b> —(TX Matrix routers and TX Matrix Plus routers only ) (Optional) Display system statistics for IGMP for all the routers in the chassis.  <b>all-lcc</b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs.  <b>all-members</b> —(EX4200 switches only) (Optional) Display IGMP statistics for all members of the Virtual Chassis configuration.  <b>lcc <i>number</i></b> —(TX Matrix routers and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router. Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li></ul>

- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**local**—(EX4200 switches only) (Optional) Display IGMP statistics for the local Virtual Chassis member.

**member *member-id***—(EX4200 switches only) (Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

**scc**—(TX Matrix routers only) (Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).

**sfc *number***—(TX Matrix Plus routers only) (Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace *number* with 0.

**Additional Information** By default, when you issue the **show system statistics igmp** command on the master Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the master Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

**Required Privilege Level** view

**Related Documentation** • [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

**List of Sample Output** [show system statistics igmp on page 5203](#)  
[show system statistics igmp \(EX Series Switches\) on page 5203](#)  
[show system statistics igmp \(TX Matrix Plus Router\) on page 5204](#)

## Sample Output

### show system statistics igmp

```
user@host> show system statistics igmp
igmp:
    17178 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid field(s)
    0 membership reports received
    0 membership reports received with invalid field(s)
    0 membership reports received for groups to which we belong
    0 membership reports sent
```

### show system statistics igmp (EX Series Switches)

```
user@host> show system statistics igmp
```

```
igmp:
  0 messages received
  0 messages received with too few bytes
  0 messages received with bad checksum
  0 membership queries received
  0 membership queries received with invalid fields
  0 membership reports received
  0 membership reports received with invalid fields
  0 membership reports received for groups to which we belong
  0 Membership reports sent
```

### show system statistics igmp (TX Matrix Plus Router)

```
user@host> show system statistics igmp
sfc0-re0:
```

```
-----
igmp:
  0 messages received
  0 messages received with too few bytes
  0 messages received with bad checksum
  0 membership queries received
  0 membership queries received with invalid field(s)
  0 membership reports received
  0 membership reports received with invalid field(s)
  0 membership reports received for groups to which we belong
  0 membership reports sent
```

```
lcc0-re0:
```

```
-----
igmp:
  0 messages received
  0 messages received with too few bytes
  0 messages received with bad checksum
  0 membership queries received
  0 membership queries received with invalid field(s)
  0 membership reports received
  0 membership reports received with invalid field(s)
  0 membership reports received for groups to which we belong
  0 membership reports sent
```

```
lcc1-re0:
```

```
-----
igmp:
  0 messages received
  0 messages received with too few bytes
  0 messages received with bad checksum
  0 membership queries received
  0 membership queries received with invalid field(s)
  0 membership reports received
  0 membership reports received with invalid field(s)
  0 membership reports received for groups to which we belong
  0 membership reports sent
```

```
lcc2-re0:
```

```
-----
igmp:
  0 messages received
  0 messages received with too few bytes
  0 messages received with bad checksum
  0 membership queries received
  0 membership queries received with invalid field(s)
```

```
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent
```

```
lcc3-re0:
```

```
-----
igmp:
```

```
0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent
```

## test msdp

---

<b>Syntax</b>	<code>test msdp (dependent-peers <i>prefix</i>   rpf-peer <i>originator</i>)</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Find Multicast Source Discovery Protocol (MSDP) peers.
<b>Options</b>	<b>dependent-peers <i>prefix</i></b> —Find downstream dependent MSDP peers.  <b>rpf-peer <i>originator</i></b> —Find the MSDP reverse-path-forwarding (RPF) peer for the originator.  <b>instance <i>instance-name</i></b> —(Optional) Find MDSP peers for the specified routing instance.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">test msdp dependent-peers on page 5206</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```



## PART 17

# Security

- [Overview on page 5209](#)
- [Configuration on page 5279](#)
- [Administration on page 5379](#)
- [Troubleshooting on page 5411](#)



## CHAPTER 60

# Overview

- [Firewall Filters on page 5209](#)
- [Policers on page 5241](#)
- [Port Security on page 5251](#)
- [Device Security on page 5272](#)

## Firewall Filters

---

- [Overview of Firewall Filters on page 5209](#)
- [Understanding Filter-Based Forwarding on page 5212](#)
- [Understanding How Firewall Filters Are Evaluated on page 5212](#)
- [Understanding How Firewall Filters Control Packet Flows on page 5214](#)
- [Understanding Firewall Filter Match Conditions on page 5215](#)
- [Firewall Filter Match Conditions and Actions on page 5219](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 5233](#)
- [Understanding Firewall Filter Planning on page 5234](#)
- [Planning the Number of Firewall Filters to Create on page 5236](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 5239](#)
- [Applying Firewall Filters to Interfaces on page 5240](#)

## Overview of Firewall Filters

Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. If a packet is accepted, you can configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). You configure firewall filters to determine whether to accept or discard a packet before it enters or exits any of these:

- Port
- VLAN

- Layer 3 (routed) interface
- Routed VLAN interface (RVI)

An *ingress* firewall filter is applied to packets that are entering an interface or VLAN, and an *egress* firewall filter is applied to packets that are exiting an interface or VLAN.



**NOTE:** Firewall filters are sometimes called *access control lists (ACLs)*.

- [Firewall Filter Types on page 5210](#)
- [Firewall Filter Components on page 5211](#)
- [Firewall Filter Processing on page 5211](#)

---

### Firewall Filter Types

---

The following firewall filter types are supported:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 traffic transiting system ports.
- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, or leave a VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter in both ingress and egress directions on IPv4 or IPv6 Layer 3 (routed) interfaces, routed VLAN interfaces (RVI) and a loopback interface, which filters traffic sent to the switch itself or generated by the switch. (You apply a filter to a loopback interface in the input direction to protect the switch from unwanted traffic. You also might want to apply a filter to a loopback interface in the output direction so that you can set the forwarding class and DSCP bit value for packets that originate on the switch itself. This feature gives you very fine control over the classification of CPU generated packets. For example, you might want to assign different DSCP values and forwarding classes to traffic generated by different routing protocols so the traffic for those protocols can be treated in a differentiated manner by other devices. You can apply a filter to a loopback interface in the output direction starting with Junos OS 13.2X51-D15.)



**NOTE:** You can apply a firewall filter to a management interface (for example, `me0`) on a QFX and EX4600 standalone switch. You cannot apply a firewall filter to a management interface on a QFX3000-G or QFX3000-M system.

- MPLS filter—You can apply a firewall filter to an MPLS interface

To apply a firewall filter:

1. Configure the firewall filter.
2. Apply the firewall filter to a port, VLAN, or router interface.



**NOTE:** You can apply only one firewall filter to a port, VLAN, or interface for a given direction. For example, for interface ge-0/0/6.0, you can apply one filter for the ingress direction and one for the egress direction.

---

## Firewall Filter Components

In a firewall filter, you first define the family address type (ethernet-switching, inet (for IPv4), inet6 (for IPv6), or mpls) and then define one or more terms that specify the filtering criteria and the action to take if a match occurs.

Each term consists of the following components:

- Match conditions—Specify values that a packet must contain to be considered a match. You can specify values for most fields in the IP, TCP, UDP, or ICMP headers. You can also match on interface names.
- Action—Specifies what to do if a packet matches the match conditions. A filter can accept, discard, or reject a matching packet and then perform additional actions, such as counting, classifying, and policing. If no action is specified for a term, the default is to accept the matching packet.

---

## Firewall Filter Processing

If there are multiple terms in a filter, the order of the terms is important. If a packet matches the first term, the switch executes the action defined by that term, and no other terms are evaluated. If the switch does not find a match between the packet and the first term, it compares the packet to the next term. If no match occurs between the packet and the second term, the system continues to compare the packet to each successive term in the filter until a match is found. If the packet does not match any terms in the filter, the switch discards the packet by default.

### Related Documentation

- [Understanding Firewall Filter Planning on page 5234](#)
- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 5239](#)
- [Understanding How Firewall Filters Are Evaluated on page 5212](#)
- [Understanding Firewall Filter Match Conditions on page 5215](#)
- [Overview of Policers on page 5241](#)
- [Configuring Firewall Filters on page 5290](#)

## Understanding Filter-Based Forwarding

You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or other security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment. For example, you might want to ensure that the highest-priority traffic is forwarded over a 40-Gigabit Ethernet link. You might also use filter-based forwarding to obtain more control over load balancing than dynamic routing protocols provide.



**NOTE:** You can create as many as 128 filters or terms that direct packets to a given virtual routing instance.

Filters used for filter-based forwarding consume memory in two ternary content addressable memories (TCAMs), and this affects the number of supported filters. See [“Planning the Number of Firewall Filters to Create” on page 5236](#) and [“Understanding FIP Snooping, FBF, and MVR Filter Scalability” on page 5546](#) for more information. The section *FBF Filter VFP TCAM Consumption* in the latter topic specifically addresses the number of supported filters when using filter-based forwarding.

### Related Documentation

- [Understanding Virtual Router Routing Instances on page 2898](#)
- [Overview of Firewall Filters on page 5209](#)
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 5279](#)

## Understanding How Firewall Filters Are Evaluated

A firewall filter consists of one or more terms, and the order of the terms within a filter is important. Before you configure firewall filters, you should understand how switches evaluate the terms within a filter and how packets are evaluated against the terms.

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.
- If the packet matches all the conditions, and no action is specified in the **then** statement, the default action **accept** is taken.
- If the packet does not match all the conditions, the switch discards it.

When a firewall filter consists of more than one term, the filter is evaluated sequentially:

1. The packet is evaluated against the conditions in the **from** statement in the first term.
2. If the packet matches all the conditions in the term, the action in the **then** statement is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
3. If the packet does not match all the conditions in the term, the packet is evaluated against the conditions in the **from** statement in the second term.

This process continues until the packet matches all the conditions in the **from** statement in one of the subsequent terms or there are no more terms in the filter.

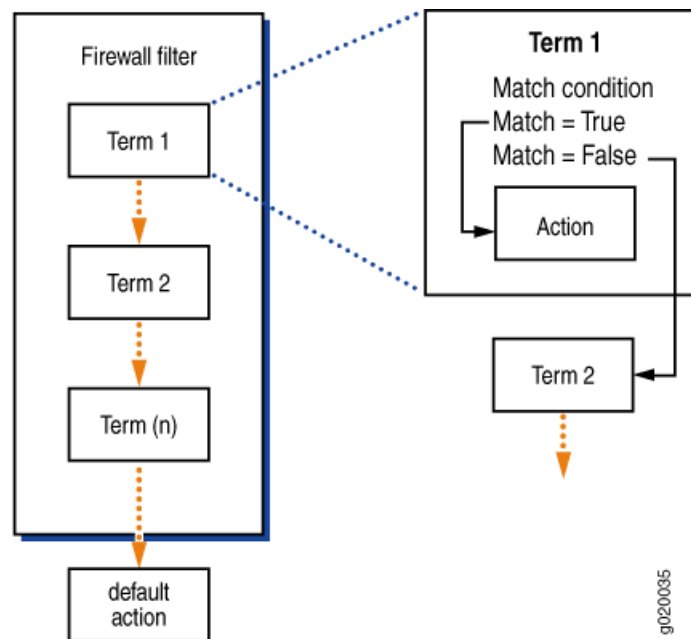
4. If a packet passes through all the terms in the filter without a match, the switch discards it.



**NOTE:** The order of conditions in a **from** statement is not important because a packet must match all the conditions to be considered a match.

Figure 179 on page 5213 shows how switches evaluate the terms within a firewall filter.

Figure 179: Evaluation of Terms Within a Firewall Filter



If you do not include a **from** statement in a term, all packets will match the term and be processed by the **then** statement. If a term does not contain a **then** statement or if an action has not been configured in the **then** statement, the term accepts any matching packets.

Every firewall filter contains an implicit **deny** statement at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {  
    then discard;  
}
```

Consequently, a packet that does not match any of the terms in a firewall filter is discarded. If you configure a filter that has no terms, all packets that pass through the filter are discarded.



**NOTE:** Firewall filtering is supported on packets that are at least 64 bytes long.

**Related  
Documentation**

- [Overview of Firewall Filters on page 5209](#)
- [Understanding Firewall Filter Match Conditions on page 5215](#)
- [Overview of Policers on page 5241](#)
- [Configuring Firewall Filters on page 5290](#)

## Understanding How Firewall Filters Control Packet Flows

A switch supports firewall filters that allow you to control flows of data packets and local packets. *Data packets* transit a switch as they are forwarded from a source to a destination. *Local packets* are destined for or sent by a Routing Engine (they do not transit a switch). Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, or data for administrative protocols such as the Internet Control Message Protocol (ICMP).

Firewall filters affect packet flows entering into or exiting from a switch as follows:


- Ingress firewall filters affect the flow of data packets that are received on switch interfaces. When a switch receives a data packet, the Packet Forwarding Engine in the system that contains the ingress interface determines where to forward the packet by looking in its Layer 2 or Layer 3 forwarding table for the best route to the destination. Data packets are forwarded to an egress interface. Locally destined packets are forwarded to the Routing Engine.
- Egress firewall filters affect data packets that are transiting a switch but do not affect packets sent by the Routing Engine. These filters are applied by the Packet Forwarding Engine in the system that contains the egress interface.

[Figure 180 on page 5215](#) illustrates the application of ingress and egress firewall filters to control the flow of packets through a switch:

1. Ingress firewall filter applied to locally destined packets that are received on switch interfaces and are destined for the Routing Engine.
2. Ingress firewall filter applied to data packets that are received on switch interfaces and will transit the switch.
3. Egress firewall filter applied to data packets that are transiting the switch.



Figure 180: Application of Firewall Filters to Control Packet Flow



#### Related Documentation

- [Understanding Firewall Filter Processing Points for Bridged and Routed Packets on page 5239](#)
- [Understanding How Firewall Filters Are Evaluated on page 5212](#)
- [Configuring Firewall Filters on page 5290](#)

## Understanding Firewall Filter Match Conditions

Before you define terms for firewall filters, you must understand how the conditions in a term are handled and how to specify interface, numeric, address, and bit-field filter match conditions to achieve the desired filter results.

- [Filter Match Conditions on page 5215](#)
- [Numeric Filter Match Conditions on page 5216](#)
- [Interface Filter Match Conditions on page 5216](#)
- [IP Address Filter Match Conditions on page 5217](#)
- [MAC Address Filter Match Conditions on page 5217](#)
- [Bit-Field Filter Match Conditions on page 5218](#)

### Filter Match Conditions

In the **from** statement of a firewall filter term, you specify the conditions that the packet must match for the action in the **then** statement to be taken. All conditions must match for the action to be implemented. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify multiple values for the same condition, a match on any one of those values matches that condition. For example, if you specify multiple IP source addresses using the **source-address** statement, a packet that contains any one of those IP source addresses

matches the condition. In some cases you can specify multiple values for the same condition by enclosing the possible values in square brackets, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set protocol (icmp | udp)
```

In other cases you must enter multiple statements, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.1.1
user@switch# set source-address 10.1.1.2
```

If you specify no match conditions in a term, that term matches all packets.



**NOTE:** Unlike traditional Junos OS firewall filters, you cannot use `except` in a condition statement to negate the condition.

---

### Numeric Filter Match Conditions

---

You can specify numeric filter match conditions that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify the condition and a single value that a field in a packet must contain to be considered a match.

You can specify the numeric value in one of the following ways:

- Single number—A match occurs if the value of the field matches the number. For example, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 23
```

- Text synonym for a single number—A match occurs if the value of the field matches the number that corresponds to the synonym. For example, to match Telnet traffic:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port telnet
```

- To specify multiple values for the same match condition in a filter term, enter each value in its own match statement. For example, a match occurs in the following term if the value of the source port in the packet is 22 or 23.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-port 22
user@switch# set source-port 23
```

### Interface Filter Match Conditions

---

You can specify an interface filter match condition to match an interface on which a packet is received or transmitted. For example, if you apply a filter to a VLAN you might want the filter to match on some interfaces that participate in the VLAN and not match on other interfaces in the VLAN. When you specify the name of the interface, you must include a logical unit.

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/0/6.0
```

In this example, the final character (0) specifies the logical unit. You can include the wildcard (\*) as part of the interface name. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set interface ge-0/*/6.0
user@switch# set interface ge-0/1/*0
user@switch# set interface ge-0/0/6.*
```

Note that you must specify a value or a wildcard for the logical unit.

### IP Address Filter Match Conditions

You can specify an address filter match condition to match an IP source or destination address or prefix in a packet. Specify the address or prefix type and the address or prefix itself. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10.2.1.0/24;
```

If you omit the prefix length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@switch# show
destination-address {
  10.0.0.0/32;
}
```

To specify more than one IP address or prefix in a filter term, enter each address or prefix in its own match statement. For example, a match occurs in the following term if the source address of a packet matches either of the following prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-address 10.1.0.0/16
user@switch# set source-address 10.2.0.0/16
```

### MAC Address Filter Match Conditions

You can specify a MAC address filter match condition to match a source or destination MAC address. You specify the address type and value that a packet must contain to be considered a match.

You can specify the MAC address as six hexadecimal bytes in any of the following formats:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 00:11:22:33:44:55
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 0011.2233.4455
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set destination-mac-address 001122334455
```

Regardless of the formats you use, the system resolves the address to the standard format, in this case 00:11:22:33:44:55.

To specify more than one MAC address in a filter term, enter each MAC address in its own match statement. For example, a match occurs in the following term if the value of the MAC source address matches either of the following addresses:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set source-mac-address 00:11:22:33:44:55
```

```
user@switch# set source-mac-address 00:11:22:33:20:15
```

### Bit-Field Filter Match Conditions

You can specify bit-field filter match conditions to match particular bits within certain fields in Ethernet frames and IP, TCP, UDP, and ICMP headers. You usually specify the field and the bit within the field that must be set in a packet to be considered a match.

In most cases you can use a keyword to specify the bit you want to match on. For example, to match on a TCP SYN packet you can enter **syn**, as in:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags syn
```

You can also enter **0x02** because the SYN bit is the third least-significant bit of the 8-bit tcp-flags field:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags 0x02
```

To match multiple bit-field values, use the logical operators, which are described in [Table 427 on page 5218](#). The operators are listed in order from highest precedence to lowest precedence. Operations are evaluated from left to right.

**Table 427: Actions for Firewall Filters**

Logical Operators	Description
!	Negation
&	Logical AND
	Logical OR

If you use a logical operator, enclose the values in quotation marks and do not include any spaces. For example, the following statement matches the second packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn&ack"
```

To negate a match, precede the value with an exclamation point. For example, the following statement matches only the initial packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-flags "syn&!ack"
```

You can use text synonyms to specify some common bit-field matches. For example, the following statement also matches the initial packet of a TCP handshake:

```
[edit firewall family family-name filter filter-name term term-name from]
user@switch# set tcp-initial
```

#### Related Documentation

- [Overview of Firewall Filters on page 5209](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 5233](#)
- [Firewall Filter Match Conditions and Actions on page 5219](#)
- [Configuring Firewall Filters on page 5290](#)

## Firewall Filter Match Conditions and Actions

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the fields and values that a packet must contain to be considered a match. You can define single or multiple match conditions in *match statements*. You can also include no match statement, in which case the term matches all packets.

When a packet matches a filter, a switch takes the action specified in the term. In addition, you can specify action modifiers to count, mirror, rate-limit, and classify packets. If no match conditions are specified for the term, the switch accepts the packet by default.

This topic describes the various match conditions, actions, and action modifiers that you can define in a firewall filter.

- [Table 428 on page 5219](#) describes the match conditions you can specify when configuring a firewall filter. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. To see a list of all the synonyms for a match condition, type `?` at the appropriate place in a statement.
- [Table 429 on page 5230](#) shows the actions that you can specify in a term.
- [Table 430 on page 5231](#) shows the action modifiers you can use to count, mirror, rate-limit, and classify packets.

**Table 428: Supported Match Conditions for Firewall Filters**

Match Condition	Description	Direction and Interface
<b>arp-type</b>	ARP request packet or ARP reply packet.	Egress and ingress ports.
<b>destination-address</b> <i>ip-address</i>	IP destination address field, which is the address of the final destination node.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.
<b>destination-mac-address</b> <i>mac-address</i>	Destination media access control (MAC) address of the packet.	Ingress ports, VLANs and IPv4 (inet) interfaces.  Egress ports and VLANs.

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>destination-port value</b>	<p>TCP or UDP destination port field. Typically, you specify this match in conjunction with the <b>protocol</b> match statement. For the following well-known ports you can specify text synonyms (the port numbers are also listed):</p> <p><b>afs</b> (1483), <b>bgp</b> (179), <b>biff</b> (512), <b>bootpc</b> (68), <b>bootps</b> (67),</p> <p><b>cmd</b> (514), <b>cvspserver</b> (2401),</p> <p><b>dhcp</b> (67), <b>domain</b> (53),</p> <p><b>eklogin</b> (2105), <b>ekshell</b> (2106), <b>exec</b> (512),</p> <p><b>finger</b> (79), <b>ftp</b> (21), <b>ftp-data</b> (20),</p> <p><b>http</b> (80), <b>https</b> (443),</p> <p><b>ident</b> (113), <b>imap</b> (143),</p> <p><b>kerberos-sec</b> (88), <b>klogin</b> (543), <b>kpasswd</b> (761), <b>krb-prop</b> (754), <b>krbupdate</b> (760), <b>kshell</b> (544),</p> <p><b>ldap</b> (389), <b>login</b> (513),</p> <p><b>mobileip-agent</b> (434), <b>mobileip-mn</b> (435), <b>msdp</b> (639),</p> <p><b>netbios-dgm</b> (138), <b>netbios-ns</b> (137), <b>netbios-ssn</b> (139), <b>nfsd</b> (2049), <b>nntp</b> (119), <b>ntalk</b> (518), <b>ntp</b> (123),</p> <p><b>pop3</b> (110), <b>pptp</b> (1723), <b>printer</b> (515),</p> <p><b>radacct</b> (1813), <b>radius</b> (1812), <b>rip</b> (520), <b>rkinit</b> (2108),</p> <p><b>smtp</b> (25), <b>snmp</b> (161), <b>snmptrap</b> (162), <b>snpp</b> (444), <b>socks</b> (1080), <b>ssh</b> (22), <b>sunrpc</b> (111), <b>syslog</b> (514),</p> <p><b>tacacs-ds</b> (65), <b>talk</b> (517), <b>telnet</b> (23), <b>tftp</b> (69), <b>timed</b> (525),</p> <p><b>who</b> (513),</p> <p><b>xmcp</b> (177),</p> <p><b>zephyr-clt</b> (2103), <b>zephyr-hm</b> (2104)</p>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>destination-port range-optimize</b> <i>range</i>	Match a range of TCP or UDP port ranges while using the available memory more efficiently. Using this condition allows you to configure more firewall filters than if you configure individual destination ports. (Not supported with filter-based forwarding.)	Egress and ingress IPv4 (inet) interfaces.
<b>destination-prefix-list</b> <i>prefix-list</i>	IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the <b>[edit policy-options]</b> hierarchy level.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.
<b>dot1q-tag</b> <i>number</i>	802.1Q VLAN ID field in the Ethernet frame. The tag values can be 1–4094.	Ingress ports and VLANs.  Egress ports and VLANs ( <i>Number</i> must be the VLAN ID of the VLAN you want to match).
<b>dot1q-user-priority</b> <i>number</i>	<p>802.1Q priority field in the Ethernet frame (used for class-of-service priorities). Values can be 0–7.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>best-effort (0)</b>—Best effort</li> <li>• <b>background (1)</b>—Background</li> <li>• <b>standard (2)</b>—Standard or spare</li> <li>• <b>excellent-load (3)</b>—Excellent load</li> <li>• <b>controlled-load (4)</b>—Controlled load</li> <li>• <b>video (5)</b>—Video</li> <li>• <b>voice (6)</b>—Voice</li> <li>• <b>network-control (7)</b>—Network control reserved traffic</li> </ul>	Ingress ports and VLANs.  Egress ports and VLANs.

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>dscp value</b>	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most-significant 6 bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>be</b>—best effort (default)</li> <li>• <b>ef (46)</b>—as defined in <a href="#">RFC 3246</a>, <i>An Expedited Forwarding PHB</i>.</li> <li>• <b>af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38)</b> These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in <a href="#">RFC 2597</a>, <i>Assured Forwarding PHB</i>.</li> <li>• <b>cs0, cs1, cs2, cs3, cs4, cs5, cs6, cs7, cs5</b></li> </ul>	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>



Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>ether-type value</b>	<p>Ethernet type field of a packet. The EtherType value specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>aarp (0x80F3)</b>—EtherType value AARP</li> <li>• <b>appletalk (0x809B)</b>—EtherType value AppleTalk</li> <li>• <b>arp (0x0806)</b>—EtherType value ARP</li> <li>• <b>fcoe (0x8906)</b>—EtherType value FCoE</li> <li>• <b>fip (0x8914)</b>—EtherType value FIP</li> <li>• <b>ipv4 (0x0800)</b>—EtherType value IPv4</li> <li>• <b>ipv6 (0x08DD)</b>—EtherType value IPv6</li> <li>• <b>mpls-multicast (0x8848)</b>—EtherType value MPLS multicast</li> <li>• <b>mpls-unicast (0x8847)</b>—EtherType value MPLS unicast</li> <li>• <b>oam (0x88A8)</b>—EtherType value OAM</li> <li>• <b>ppp (0x880B)</b>—EtherType value PPP</li> <li>• <b>pppoe-discovery (0x8863)</b>—EtherType value PPPoE Discovery Stage</li> <li>• <b>pppoe-session (0x8864)</b>—EtherType value PPPoE Session Stage</li> <li>• <b>sna (0x80D5)</b>—EtherType value SNA</li> </ul>	<p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>
<b>exp</b>	Match on MPLS EXP bits.	<p>Ingress MPLS interfaces.</p> <p>Egress MPLS interfaces.</p>
<b>fragment-flags value</b>	<p>IP fragmentation flags. In place of the numeric value, you can specify one of the following text synonyms (the hexadecimal values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>is-fragment</b></li> <li>• <b>dont-fragment (0x4000)</b></li> <li>• <b>more-fragments (0x2000)</b></li> <li>• <b>reserved (0x8000)</b></li> </ul>	Ingress ports and VLANs.

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>icmp-code value</b>	<p>ICMP code field. Because the meaning of the value depends upon the associated <b>icmp-type</b>, you must specify a value for <b>icmp-type</b> along with a value for <b>icmp-code</b>. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <ul style="list-style-type: none"> <li><i>IPv4</i>: parameter-problem—ip-header-bad (0), required-option-missing (1)</li> <li><i>IPv6</i>: parameter-problem—ip6-header-bad (0), unrecognized-next-header (1), unrecognized-option (2)</li> <li>redirect—redirect-for-network (0), redirect-for-host (1), redirect-for-tos-and-net (2), redirect-for-tos-and-host (3)</li> <li>time-exceeded—ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)</li> <li><i>IPv4</i>: unreachable—network-unreachable (0), host-unreachable (1), protocol-unreachable (2), port-unreachable (3), fragmentation-needed (4), source-route-failed (5), destination-network-unknown (6), destination-host-unknown (7), source-host-isolated (8), destination-network-prohibited (9), destination-host-prohibited (10), network-unreachable-for-TOS (11), host-unreachable-for-TOS (12), communication-prohibited-by-filtering (13), host-precedence-violation (14), precedence-cutoff-in-effect (15)</li> <li><i>IPv6</i>: unreachable—address-unreachable (3), administratively-prohibited (1), no-route-to-destination (0), port-unreachable (4)</li> </ul>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>hop-limitvalue</b>	Match the the specified hop limit or set of hop limits. Specify a single value or a range of values from 0 through 255.	Ingress and egress IPv6 (inet6) interfaces.

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>icmp-type</b> <i>value</i>	<p>ICMP message type field. Typically, you specify this match in conjunction with the <b>protocol</b> match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <p><i>IPv4:</i> echo-reply (0), destination unreachable (3), source-quench (4), redirect (5), echo-request (8), IPv4 (inet)-advertisement (9), IPv4 (inet)-solicit (10), time-exceeded (11), parameter-problem (12), timestamp (13), timestamp-reply (14), info-request (15), info-reply (16), mask-request (17), mask-reply (18)</p> <p><i>IPv6:</i> destination-unreachable (1), packet-too-big (2), time-exceeded (3), parameter-problem (4), echo-request (128), echo-reply (129), membership-query (130), membership-report (131), membership-termination (132), router-solicit (133), router-advertisement (134), neighbor-solicit (135), neighbor-advertisement (136), redirect (137), router-renumbering (138), node-information-request (139), node-information-reply (140)</p> <p>See also <b>icmp-code</b> <i>variable</i>.</p>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>interface</b> <i>interface-name</i>	<p>Interface on which the packet is received, including the logical unit. You can include the wildcard character (*) as part of an interface name or logical unit.</p> <p><b>NOTE:</b> An interface from which a packet is sent cannot be used as a match condition.</p>	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces and IPv6 (inet6) interfaces.</p>
<b>ip-destination-address</b> <i>address</i>	IPv4 address that is the final destination node address for the packet.	Ingress ports and VLANs.
<b>ip6-destination-address</b> <i>address</i>	IPv6 address that is the final destination node address for the packet.	Ingress ports and VLANs. (You cannot simultaneously apply a filter with this match criterion to a Layer 2 port and VLAN that includes that port.)
<b>ip-options</b>	Specify <b>any</b> to create a match if anything is specified in the options field in the IP header.	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>ip-precedence</b> <i>ip-precedence-field</i>	IP precedence field. In place of the numeric field value, you can specify one of the following text synonyms (the field values are also listed): <b>critical-ecp</b> (0xa0), <b>flash</b> (0x60), <b>flash-override</b> (0x80), <b>immediate</b> (0x40), <b>internet-control</b> (0xc0), <b>net-control</b> (0xe0), <b>priority</b> (0x20), or <b>routine</b> (0x00).	Ingress ports, VLANs, and IPv4 (inet) interfaces.  Egress IPv4 (inet) interfaces.
<b>ip-protocol</b> <i>number</i>	IP protocol field.	Ingress ports, VLANs, and IPv4 (inet) interfaces.  Egress IPv4 (inet) interfaces.
<b>ip-source-address</b> <i>address</i>	IPv4 address of the source node sending the packet.	Ingress ports and VLANs.
<b>ip6-source-address</b> <i>address</i>	IPv6 address of the source node sending the packet.	Ingress ports and VLANs. (You cannot simultaneously apply a filter with this match criterion to a Layer 2 port and VLAN that includes that port.)
<b>ip-version</b> <i>address</i>	IP version of the packet. Use this condition to match IPv4 or IPv6 header fields in traffic that arrives on a Layer 2 port or VLAN interface.	Ingress ports and VLANs.
<b>is-fragment</b>	Using this condition causes a match if the More Fragments flag is enabled in the IP header or if the fragment offset is not zero.	Ingress ports, VLANs, and IPv4 (inet) interfaces.  Egress IPv4 (inet) interfaces.
<b>l2-encap-type</b> <i>llc-non-snap</i>	Match on logical link control (LLC) layer packets for non-Subnet Access Protocol (SNAP) Ethernet Encapsulation type.	Ingress ports and VLANs.  Egress ports and VLANs.
<b>label</b>	Match on MPLS label bits.	Ingress MPLS interfaces.  Egress MPLS interfaces.
<b>learn-vlan-id</b> <i>number</i>	VLAN identifier used for MAC learning.	Ingress ports and VLANs.  Egress ports and VLANs.

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>next-header</b>	<p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p><b>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b></p>	<p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>
<b>packet-length</b>	<p>Packet length in bytes. You must enter a value between 0 and 65535.</p>	<p>Ingress ports, VLANs, IPv4 (inet), and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>payload-protocol</b>	<p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p><b>hop-by-hop (0), icmp (1), icmp6 (58), igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b></p>	<p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>
<b>precedence value</b>	<p>IP precedence bits in the type-of-service (ToS) byte in the IP header. (This byte can also be used for the DiffServ DSCP.) In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <ul style="list-style-type: none"> <li>• <b>routine (0)</b></li> <li>• <b>priority (1)</b></li> <li>• <b>immediate (2)</b></li> <li>• <b>flash (3)</b></li> <li>• <b>flash-override (4)</b></li> <li>• <b>critical-ecp (5)</b></li> <li>• <b>internet-control (6)</b></li> <li>• <b>net-control (7)</b></li> </ul>	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>protocol type</b>	<p>IPv4 or IPv6 protocol value. In place of the numeric value, you can specify one of the following text synonyms (the numeric values are also listed):</p> <p><b>hop-by-hop (0), icmp (1), icmp6, igmp (2), ipip (4), tcp (6), egp (8), udp (17), ipv6 (41), routing (43), fragment (44), rsvp (46), gre (47), esp (50), ah (51), icmp6 (58), no-next-header (59), dstopts (60), ospf (89), pim (103), vrrp (112), sctp (132)</b></p>	<p>Ingress ports, VLANs and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>rat-type</b> <b>tech-type-value</b>	<p>Match the radio-access technology (RAT) type specified in the 8-bit Tech-Type field of Proxy Mobile IPv4 (PMIPv4) access technology type extension. The technology type specifies the access technology through which the mobile device is connected to the access network. Specify a single value, a range of values, or a set of values. You can specify a technology type as a numeric value from 0 through 255 or as a system keyword.</p> <ul style="list-style-type: none"> <li>Numeric value 1 matches IEEE 802.3.</li> <li>Numeric value 2 matches IEEE 802.11a/b/g.</li> <li>Numeric value 3 matches IEEE 802.16e</li> <li>Numeric value 4 matches IEEE 802.16m.</li> <li>Text string <b>eutran</b> matches 4G.</li> <li>Text string <b>geran</b> matches 2G.</li> <li>Text string <b>utran</b> matches 3G.</li> <li></li> </ul>	Egress and ingress IPv4 (inet) interfaces.
<b>sample</b>	Sample the packet traffic. Apply this option only if you have enabled traffic sampling.	Egress and ingress IPv4 (inet) interfaces.
<b>source-address</b> <b>ip-address</b>	IP source address field, which is the address of the node that sent the packet.	<p>Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>source-mac-address</b> <i>mac-address</i>	Source media access control (MAC) address of the packet.	<p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>source-port</b> <i>value</i>	TCP or UDP source port. Typically, you specify this match in conjunction with the <b>protocol</b> match statement. In place of the numeric field, you can specify one of the text synonyms listed under <b>destination-port</b> .	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces.
<b>source-port range-optimize</b> <i>range</i>	Match a range of TCP or UDP port ranges while using the available memory more efficiently. Using this condition allows you to configure more firewall filters than if you configure individual source ports. (Not supported with filter-based forwarding.)	Egress and ingress IPv4 (inet) interfaces.
<b>source-prefix-list</b> <i>prefix-list</i>	IP source prefix list. You can define a list of IP address prefixes under a prefix-list alias for frequent use. Define this list at the <b>[edit policy-options]</b> hierarchy level.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces.
<b>tcp-established</b>	Match packets of an established TCP connection. This condition matches packets other than those used to set up a TCP connection—that is, three-way handshake packets are not matched.  When you specify <b>tcp-established</b> , a switch does not implicitly verify that the protocol is TCP. You must also specify the <b>protocol tcp</b> match condition.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces.
<b>tcp-flags</b> <i>value</i>	One or more TCP flags: <ul style="list-style-type: none"> <li>• <b>ack</b> (0x10)</li> <li>• <b>fin</b> (0x01)</li> <li>• <b>push</b> (0x08)</li> <li>• <b>rst</b> (0x04)</li> <li>• <b>syn</b> (0x02)</li> <li>• <b>urgent</b> (0x20)</li> </ul>	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces.
<b>tcp-initial</b>	Match the first TCP packet of a connection. A match occurs when the TCP flag <b>SYN</b> is set and the TCP flag <b>ACK</b> is not set.  When you specify <b>tcp-initial</b> , a switch does not implicitly verify that the protocol is TCP. You must also specify the <b>protocol tcp</b> match condition.	Ingress ports, VLANs, IPv4 (inet) interfaces, and IPv6 (inet6) interfaces.  Egress IPv4 (inet) interfaces.

Table 428: Supported Match Conditions for Firewall Filters (*continued*)

Match Condition	Description	Direction and Interface
<b>traffic-class</b>	<p>8-bit field that specifies the class-of-service (CoS) priority of the packet. The traffic-class field is used to specify a DiffServ code point (DSCP) value. This field was previously used as the type-of-service (ToS) field in IPv4, and, the semantics of this field (for example, DSCP) are identical to those of IPv4.</p> <p>You can specify one of the following text synonyms (the field values are also listed):</p> <p><b>af11 (10), af12 (12), af13 (14), af21 (18), af22 (20), af23 (22), af31 (26), af32 (28), af33 (30), af41 (34), af42 (36), af43 (38), cs0 (0), cs1 (8), cs2 (16), cs3 (24), cs4 (32), cs5 (40), cs6 (48), cs7 (56), ef (46)</b></p>	<p>Ingress ports, VLANs, and IPv6 (inet6) interfaces.</p> <p>Egress IPv6 (inet6) interfaces.</p>
<b>ttl value</b>	IP Time-to-live (TTL) field in decimal. The value can be 1-255.	<p>Ingress IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>user-vlan-1p-priority value</b>	Match on the IEEE 802.1p user priority bits in the customer VLAN tag (the inner tag in a dual-tag frame with 802.1Q VLAN tags). Specify a single value or multiple values from 0 through 7.	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>user-vlan-id number</b>	Match the first VLAN identifier that is part of the payload.	<p>Ingress ports, VLANs, and IPv4 (inet) interfaces.</p> <p>Egress IPv4 (inet) interfaces.</p>
<b>vlan (vlan-name   vlan-id )</b>	VLAN names or ID.	<p>Ingress ports and VLANs.</p> <p>Egress ports and VLANs.</p>

Use **then** statements to define actions that should occur if a packet matches all conditions in a **from** statement. [Table 429 on page 5230](#) shows the actions that you can specify in a term. (If you do not include a **then** statement, the system accepts packets that match the filter.)

Table 429: Actions for Firewall Filters

Action	Description
<b>accept</b>	Accept a packet. This is the default action for packets that match a term.
<b>discard</b>	Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.



Table 429: Actions for Firewall Filters (*continued*)

Action	Description
<b>reject</b> <i>message-type</i>	<p>Discard a packet and send a “destination unreachable” ICMPv4 message (type 3). To log rejected packets, configure the <b>syslog</b> action modifier.</p> <p>You can specify one of the following message types: <b>administratively-prohibited</b> (default), <b>bad-host-tos</b>, <b>bad-network-tos</b>, <b>host-prohibited</b>, <b>host-unknown</b>, <b>host-unreachable</b>, <b>network-prohibited</b>, <b>network-unknown</b>, <b>network-unreachable</b>, <b>port-unreachable</b>, <b>precedence-cutoff</b>, <b>precedence-violation</b>, <b>protocol-unreachable</b>, <b>source-host-isolated</b>, <b>source-route-failed</b>, or <b>tcp-reset</b>.</p> <p>If you specify <b>tcp-reset</b>, the system sends a TCP reset if the packet is a TCP packet; otherwise nothing is sent.</p> <p>If you do not specify a message type, the ICMP notification “destination unreachable” is sent with the default message “communication administratively filtered.”</p> <p><b>NOTE:</b> The <b>reject</b> action is supported on ingress interfaces only.</p>
<b>routing-instance</b> <i>instance-name</i>	Forward matched packets to a virtual routing instance.
<b>vlan</b> <i>VLAN-name</i>	Forward matched packets to a specific VLAN.
	<b>NOTE:</b> The <b>vlan</b> action is supported on ingress interfaces only.

You can also specify the action modifiers listed in [Table 430 on page 5231](#) to count, mirror, rate-limit, and classify packets.

Table 430: Action Modifiers for Firewall Filters

Action Modifier	Description
<b>analyzer</b> <i>analyzer-name</i>	<p>(Non-ELS platforms) Mirror traffic (copy packets) to an analyzer configured at the <b>[edit ethernet-switching-options analyzer]</b> hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p>
<b>count</b> <i>counter-name</i>	Count the number of packets that match the term.
<b>decapsulate</b> [ <b>gre</b>   <i>routing-instance</i> ]	De-encapsulate GRE packets or forward de-encapsulated GRE packets to the specified routing instance

Table 430: Action Modifiers for Firewall Filters (*continued*)

Action Modifier	Description
<b>dscp value</b>	<p>Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most-significant 6 bits of this byte form the DSCP.</p> <p>You can specify DSCP in hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed):</p> <ul style="list-style-type: none"> <li><b>be</b>—best effort (default)</li> <li><b>ef (46)</b>—as defined in <a href="#">RFC 3246</a>, <i>An Expedited Forwarding PHB</i>.</li> <li><b>af11 (10), af12 (12), af13 (14); af21 (18), af22 (20), af23 (22); af31 (26), af32 (28), af33 (30); af41 (34), af42 (36), af43 (38)</b></li> </ul> <p>These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in <a href="#">RFC 2597</a>, <i>Assured Forwarding PHB</i>.</p> <ul style="list-style-type: none"> <li><b>cs0, cs1, cs2, cs3, cs4, cs5, cs6, cs7, cs5</b></li> </ul>
<b>forwarding-class class</b>	<p>Classify the packet in one of the following forwarding classes:</p> <ul style="list-style-type: none"> <li><b>best-effort</b></li> <li><b>fcoe</b></li> <li><b>mcast</b></li> <li><b>network-control</b></li> <li><b>no-loss</b></li> </ul> <p><b>NOTE:</b> To configure a forwarding class, you must also configure loss priority.</p>
<b>log</b>	<p>Log the packet's header information in the Routing Engine. To view this information, enter the <b>show firewall log</b> operational mode command.</p> <p><b>NOTE:</b> The <b>log</b> action modifier is supported on ingress interfaces only.</p>
<b>loss-priority (low   medium-low   medium-high   high)</b>	<p>Set the packet loss priority (PLP).</p> <p><b>NOTE:</b> The <b>loss-priority</b> action modifier is supported on ingress interfaces only.</p> <p><b>NOTE:</b> The <b>loss-priority</b> action modifier is not supported in combination with the <b>policer</b> action.</p>
<b>policer policer-name</b>	<p>Send packets to a policer (for the purpose of applying rate limiting).</p> <p>You can specify a policer for ingress port, VLAN, IPv4 (inet), IPv6 (inet6), and MPLS filters.</p> <p><b>NOTE:</b> The <b>policer</b> action modifier is not supported in combination with the <b>loss-priority</b> action.</p>

Table 430: Action Modifiers for Firewall Filters (*continued*)

Action Modifier	Description
<b>port-mirror</b>	<p>(ELS platforms) Mirror traffic (copy packets) to an output interface configured in a port-mirroring instance at the <b>[edit forwarding-options port-mirroring]</b> hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p>
<b>port-mirror-instance</b> <i>port-mirror-instance-name</i>	<p>(ELS platforms) Mirror traffic to a port-mirroring instance configured at the <b>[edit forwarding-options port-mirroring]</b> hierarchy level.</p> <p>You can specify port mirroring for ingress port, VLAN, and IPv4 (inet) firewall filters only.</p>
<b>syslog</b>	<p>Log an alert for this packet.</p> <p><b>NOTE:</b> The <b>syslog</b> action modifier is supported on ingress interfaces only.</p>
<b>three-color-policer</b> <i>three-color-policer-name</i>	<p>Send packets to a three-color policer (for the purpose of applying rate limiting).</p> <p>You can specify a three-color policer for ingress and egress port, VLAN, IPv4 (inet), IPv6 (inet6), and MPLS filters.</p> <p><b>NOTE:</b> The <b>policer</b> action modifier is not supported in combination with the <b>loss-priority</b> action.</p>

**Related  
Documentation**

- [Understanding Firewall Filter Match Conditions on page 5215](#)
- [Understanding How Firewall Filters Are Evaluated on page 5212](#)
- [Understanding How a Firewall Filter Tests a Protocol on page 5233](#)
- [Overview of Policers on page 5241](#)
- [Understanding Port Mirroring on page 5425](#)
- [Configuring Firewall Filters on page 5290](#)

## Understanding How a Firewall Filter Tests a Protocol

When examining match conditions in a firewall filter, a switch tests only the fields that you specify. It does not implicitly test any fields that you do not explicitly configure. For example, if you specify a match condition of **source-port ssh**, there is no implied test to determine if the protocol is TCP. In this case, the switch considers any packet that has a value of **22** (decimal) in the 2-byte field that follows a *presumed* IP header to be a match. To ensure that the term matches on TCP packets, you also specify a **protocol tcp** match condition.

For the following match conditions, you should explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify **protocol tcp** or **protocol udp**.
- **icmp-code**—Specify **protocol icmp** and **icmp-type**.
- **icmp-type**—Specify **protocol tcp** or **protocol udp**.
- **source-port**—Specify **protocol tcp** or **protocol udp**.
- **tcp-flags**—Specify **protocol tcp**.

**Related  
Documentation**

- [Overview of Firewall Filters on page 5209](#)
- [Understanding Firewall Filter Match Conditions on page 5215](#)
- [Configuring Firewall Filters on page 5290](#)

## Understanding Firewall Filter Planning

Before you create a firewall filter and apply it, determine what you want the filter to accomplish and how to use its match conditions and actions to achieve your goals. It is important that you understand how packets are matched, the default and configured actions of the firewall filter, and where to apply the firewall filter.

You can apply no more than one firewall filter per port, VLAN, or router interface per direction (input and output). For example, for a given port you can apply at most one filter in the input direction and one filter in the output direction. You should try to be conservative in the number of terms (rules) that you include in each firewall filter, because a large number of terms requires longer processing time during a commit operation and can make testing and troubleshooting more difficult.

Before you configure and apply firewall filters, answer the following questions for each of them:

1. What is the purpose of the filter?

For example, the system can drop packets based on header information, rate-limit traffic, classify packets into forwarding classes, log and count packets, or prevent denial-of-service attacks.

2. What are the appropriate match conditions? Determine the packet header fields that the packet must contain for a match. Possible fields include:

- Layer 2 header fields—Source and destination MAC addresses, 802.1Q tag, Ethernet type, or VLAN.
- Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, or TTL type).
- TCP header fields—Source and destination ports and flags.
- ICMP header fields—Packet type and code.

3. What are the appropriate actions to take if a match occurs?

The system can accept, discard, or reject packets.

4. What additional action modifiers might be required?

For example, you can configure the system to mirror (copy) packets to a specified port, count matching packets, apply traffic management, or police packets.

5. On what port, router interface, or VLAN should the firewall filter be applied?

Start with the following basic guidelines:

- If packets entering or leaving a Layer 2 interface (port) need to be filtered, apply the filter at the **[edit family ethernet switching filter]** hierarchy level. This is a port filter.
- If packets entering or leaving any port in a specific VLAN need to be filtered, use a VLAN filter.
- If packets entering or leaving a Layer 3 (routed) interface or routed VLAN interface (RVI) need to be filtered, use a router firewall filter. Apply the filter to the interface at the **[edit family inet]** hierarchy level. You can also apply a router firewall filter on a loopback interface.

Before you choose the interface or VLAN on which to apply a firewall filter, understand how that placement can affect traffic flow to other interfaces. In general, apply a filter close to the source device if the filter matches on source or destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP or UDP port numbers. However, you should apply a filter close to the destination device if the filter matches *only* on a source IP address. When you apply a filter too close to the source device, the filter could prevent that source device from accessing other services that are available on the network.



**NOTE:** Egress firewall filters do not affect the flow of locally generated control packets from the Routing Engine.

6. In which direction should the firewall filter be applied?

You typically configure different actions for traffic entering an interface than you configure for traffic exiting an interface.

7. How many filters should I create?

See “[Planning the Number of Firewall Filters to Create](#)” on [page 5236](#) for information about how many firewall filters you can apply.

**Related Documentation**

- [Overview of Firewall Filters on page 5209](#)
- [Overview of Policers on page 5241](#)
- [Understanding How Firewall Filters Are Evaluated on page 5212](#)
- [Planning the Number of Firewall Filters to Create on page 5236](#)
- [Configuring Firewall Filters on page 5290](#)

## Planning the Number of Firewall Filters to Create

- [Understanding How Many Firewall Filters Are Supported on page 5236](#)
- [Egress Filters on page 5237](#)
- [Avoid Configuring too Many Filters on page 5237](#)
- [Policers can Limit Egress Filters on page 5238](#)
- [Planning for Filter-Specific Policers on page 5239](#)
- [Planning for Filter-Based Forwarding on page 5239](#)

### Understanding How Many Firewall Filters Are Supported

QFX3500, QFX3600, QFX5100, and EX4600 switches, QFabric Node devices, and VCF members support the maximum numbers of firewall filter terms per type of attachment point shown in [Table 431 on page 5236](#).

**Table 431: Supported Firewall Filter Numbers**

Filter Type	QFX3500, QFX3600	QFX5100, EX4600
Ingress	768	1536
Egress	1024	1024

These totals are applied in aggregate. For example, on the QFX3500 and QFX3600 you can apply a total of 768 terms in all your port filters, Layer 3 filters, and VLAN filters that are applied in the input direction and 1024 terms in port filters, Layer 3 filters, and VLAN filters that are applied in the output direction.



**NOTE:** If you want to create more than 512 egress VLAN filters, your first VLAN ID should be 6 and the subsequent VLAN IDs should increase by 1. For example, to create 1024 egress VLAN filters, the first VLAN ID would be 6, the second ID would be 7, and the sequence would continue through VLAN ID 1029. Similarly, if you want to create fewer than 512 egress VLAN filters but want the total number of terms in those filters to exceed 512, you should number your VLAN IDs in the same manner. If you do not use this approach to create your VLAN IDs, the total number of allowed terms or filters will be less than 1024 and might be 512.

The ternary content addressable memory (TCAM) for firewall filters is divided into slices that accommodate 256 terms, and all the terms in a memory slice must be in filters of the same type and applied in the same direction. A memory slice is reserved as soon as you commit a filter. For example, if you create a port filter and apply it in the input direction, a memory slice is reserved that will only store ingress port filters. If you create and apply only one ingress port filter and that filter has only one term, the rest of this slice is unused and is unavailable for other filter types.

Continuing with the above example, assume that you create and apply 256 ingress port filters with one term each so that one memory slice is filled. This leaves two more memory

slices available for ingress filters. (Remember that the maximum number of ingress terms is 768.) If you then create and apply an ingress Layer 3 filter with one term, another memory slice is reserved for ingress Layer 3 filters. As before, the rest of the slice is unused and is unavailable for different filter types. At this point there is one memory slice available for any ingress filter type.

Now assume that you create and apply a VLAN ingress filter. The final memory slice is reserved for VLAN ingress filters. Memory allocation for ingress filters (once again assuming one term per filter) is as follows:

- Slice 1: Filled with 256 ingress port filters. You cannot commit any more ingress port filters.
- Slice 2: Contains one ingress Layer 3 filter with one term. You can commit 255 more terms in ingress Layer 3 filters.
- Slice 3: Contains one ingress VLAN filter with one term. You can commit 255 more terms in ingress VLAN filters.

Here is another example. Assume that you create 257 ingress port filters with one term per filter—that is, you create one more term than a single memory slice can accommodate. When you apply the filters and commit the configuration, the filter memory allocation is:

- Slice 1: Filled with 256 ingress port filters. You cannot apply any more ingress port filters.
- Slice 2: Contains one ingress port filter. You can apply 255 more terms in ingress port filters.
- Slice 3: This slice is unassigned. You can create and apply 256 terms in ingress filters of any type (port, Layer 3, or VLAN), but all the filters must be of the same type.

### Egress Filters

---

All of the preceding principles also apply to egress filters, but four memory slices are used because IPv4 Layer 3 filters and IPv6 Layer 3 filters are stored in separate slices. The memory slices for egress filters are the same size as those for ingress filters, so the maximum number of egress filter terms is therefore 1024.

### Avoid Configuring too Many Filters

---

If you violate any of these restrictions and commit a configuration that is not in compliance, Junos OS rejects the excessive filters. For example, if you configure 300 ingress port filters and 300 ingress Layer 3 filters and try to commit the configuration, Junos OS does the following (again assuming one term per filter):

- Accepts the 300 ingress port filters (storing them in two memory slices).
- Accepts the first 256 ingress Layer 3 filters it processes (storing them in the third memory slice).
- Rejects the remaining 44 ingress Layer 3 filters.



**NOTE:** In this situation, be sure to delete excessive filters (for example, the remaining 44 ingress Layer 3 filters) from the configuration before you reboot the device. If you reboot a device that has a noncompliant configuration, you cannot predict which filters are installed after the reboot. Using the example above, the 44 ingress Layer 3 filters that were originally rejected might be installed, and 44 of the port filters that were originally accepted might be rejected.

---

### Policers can Limit Egress Filters

---

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem from occurring by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.



You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

### Planning for Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented above to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

### Planning for Filter-Based Forwarding

You can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature—called filter-based forwarding—you specify a filter and match criteria and then specify the virtual routing instance to send packets to. Filters used in this way also consume memory in an additional TCAM. See [“Understanding FIP Snooping, FBF, and MVR Filter Scalability” on page 5546](#) for more information. The section *FBF Filter VFP TCAM Consumption* in this topic specifically addresses the number of supported filters when using filter-based forwarding.

#### Related Documentation

- [Overview of Firewall Filters on page 5209](#)
- [Understanding How Firewall Filters Are Evaluated on page 5212](#)
- [Understanding Firewall Filter Planning on page 5234](#)
- [Configuring Firewall Filters on page 5290](#)
- [Understanding Filter-Based Forwarding on page 5212](#)

## Understanding Firewall Filter Processing Points for Bridged and Routed Packets

You apply firewall filters at multiple processing points in the forwarding path. At each processing point, the action to be taken on a packet is determined by the configuration of the filter and the results of the lookup in the forwarding or routing table.

For both bridged (Layer 2) unicast packets and routed (Layer 3) unicast packets, firewall filters are applied in the prescribed order shown below (assuming that each filter is present and a packet is accepted by each one).

Bridged packets:

1. Ingress port filter
2. Ingress VLAN filter
3. Egress VLAN filter
4. Egress port filter

Routed packets:

1. Ingress port firewall filter
2. Ingress VLAN firewall filter (Layer 2 CoS)
3. Ingress router firewall filter (Layer 3 CoS)
4. Egress router firewall filter
5. Egress VLAN firewall filter
6. Egress port filter



**NOTE:** MAC learning occurs before filters are applied, so switches learn the MAC addresses of packets that are dropped by ingress filters.

---

**Related  
Documentation**

- [Overview of Firewall Filters on page 5209](#)
- [Understanding How Firewall Filters Control Packet Flows on page 5214](#)
- [Configuring Firewall Filters on page 5290](#)

## Applying Firewall Filters to Interfaces

For a firewall filter to work, you must apply it to at least one interface. To do this, include the **filter** statement when configuring a logical interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family family-name filter (input |
output) filter-name
```

In the **input** statement, specify a firewall filter to be evaluated when packets are received on the interface. Input filters applied to a loopback interface affect only traffic destined for the Routing Engine.

In the **output** statement, specify a filter to be evaluated when packets exit the interface.



**NOTE:** When you create a loopback interface, it is important to apply an ingress filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface lo0, you include the `apply-groups` statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including lo0 and other loopback interfaces.

#### Related Documentation

- [Configuring Firewall Filters on page 5290](#)

## Policers

- [Overview of Policers on page 5241](#)
- [Understanding Policers with Link Aggregation Groups on page 5246](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 5249](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 5249](#)

## Overview of Policers

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.


- [Policer Overview on page 5241](#)
- [Policer Types on page 5242](#)
- [Policer Actions on page 5243](#)
- [Policer Colors on page 5244](#)
- [Filter-Specific Policers on page 5244](#)
- [Suggested Naming Convention for Policers on page 5244](#)
- [Policer Counters on page 5245](#)
- [Policer Algorithms on page 5245](#)
- [How Many Policers are Supported? on page 5245](#)
- [Policers can Limit Egress Firewall Filters on page 5245](#)

### Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 181 on page 5242](#) illustrates this process.

**Figure 181: Flow of Tricolor Marking Policer Operation**



9017049

After you name and configure a policer, you use it by specifying it as an action in one or more firewall filters.

### Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- **Two-rate three-color marker**—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic

based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 432 on page 5243](#) for information about how metering results are applied for each of these policer types.

### Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 432 on page 5243](#) describes the policer actions.

**Table 432: Policer Actions**

Policer	Marking	Implicit Action	Configurable Action
Single-rate two-color	Green (conforming)	Assign low loss priority	None
	Red (nonconforming)	None	Discard
Single-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the EBS)	Assign high loss priority	Discard
Two-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the PIR and PBS)	Assign high loss priority	Discard



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is discard.

## Policer Colors

---

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

## Filter-Specific Policers

---

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5236](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

## Suggested Naming Convention for Policers

---

We recommend that you use the naming convention ***policertypeTCM#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named ***srTCM1-ca***. The second two-rate, color-blind three-color configured would be named ***trTCM2-cb***. The elements of this naming convention are explained below:

- **sr** (single-rate)
- **tr** (two-rate)
- **TCM** (tricolor marking)
- **1 or 2** (number of marker)

- ca (color-aware)
- cb (color-blind)

### Policer Counters

---

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

### Policer Algorithms

---

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.

### How Many Policers are Supported?

---

You can configure and commit the following numbers of policers on QFX3500 and QFX3600 standalone switches and QFabric Node devices:

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

### Policers can Limit Egress Firewall Filters

---

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms,

1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.

- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

#### Related Documentation

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 5249](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 5249](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)

## Understanding Policers with Link Aggregation Groups

If you apply a policer to a link aggregation group (LAG) on a QFX3500 switch or node, the policer applies to all the interfaces in the LAG in aggregate. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer (by using a firewall filter) to a LAG that has two member interfaces on a single switch or node, the total allowed throughput for both members is 1 Gbps.

If you apply a policer to a LAG that has members on different nodes in a QFabric network Node group or redundant server Node group, the configured rate applies to the interface on each node. For example, if you configure a policer to rate-limit at 1 Gbps and apply the policer to a LAG that has one member on server node A and one member on server node B, the allowed throughput for each member is 1 Gbps, for a total allowed throughput of 2 Gbps.



- Related Documentation**
- [Overview of Policers on page 5241](#)
  - [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)

## Understanding Color-Blind Mode for Single-Rate Tricolor Marking

With the color-blind mode of single-rate tricolor marking, all packets are evaluated against the CBS. If a packet exceeds the CBS, it is evaluated against the EBS. In color-blind mode, the policer supports three loss priorities only: low, medium-high, and high.

Packets that exceed the CBS but are below the EBS are marked yellow (medium-high). Packets that exceed the EBS are marked red (high), as shown in [Table 433 on page 5247](#).

**Table 433: Color-Blind Mode TCM Color-to-PLP Mapping**

Color	PLP	Meaning
Green	<b>low</b>	Conforming.
Yellow	<b>medium-high</b>	Packet exceeds the CIR and CBS but does not exceed the EBS.
Red	<b>high</b>	Packet exceeds the EBS.

- Related Documentation**
- [Overview of Policers on page 5241](#)
  - [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296](#)

## Understanding Color-Aware Mode for Single-Rate Tricolor Marking

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase a preassigned PLP but cannot decrease it.

### Summary of PLP Changes

[Table 434 on page 5247](#) shows how a packet's incoming priority can be modified with single-rate marking.

**Table 434: Color-Aware Mode Single-Rate PLP Mapping**

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
<b>low</b>	CIR, CBS, and EBS	Conforming	<b>low</b>
		Packet exceeds the CIR and CBS but does not exceed the EBS.	<b>medium-high</b>
		Packet exceeds the EBS.	<b>high</b>
<b>medium-low</b>	EBS only	Packet does not exceed the EBS.	<b>medium-low</b>
		Packet exceeds the EBS.	<b>high</b>

Table 434: Color-Aware Mode Single-Rate PLP Mapping (*continued*)

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
medium-high	EBS only	Packet does not exceed the EBS.	medium-high
		Packet exceeds the EBS.	high
high	Not metered by the policer.	All cases.	high

The following sections describe single-rate color-aware PLP mapping in more detail.

#### ***Effect on Green Packets (Low PLP)***

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the PLP unchanged or increase it to medium-high or high, so these packets are therefore metered against both the CBS and the EBS. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, packets remain marked as low PLP.
- If bursts exceed the CBS but not the EBS, some of the packets are marked as medium-high PLP, and some of the packets remain marked as low PLP.
- If bursts exceed the EBS, some of the packets are marked as high PLP, and some of the packets remain marked as low PLP.

#### ***Effect on Yellow Packets (Medium PLP)***

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high, so these packets are therefore metered against the EBS only. For example, if a behavior aggregate or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CBS, the packets remain marked as medium-high PLP.

- If the rate of traffic flow is greater than the CBS but less than the EBS, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the EBS, some of the packets are marked as high PLP and some remain marked as medium-high PLP.

#### ***Effect on Red Packets (High PLP)***

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CBS or the EBS.

#### **Related Documentation**

- [Overview of Policers on page 5241](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296](#)

## **Understanding Color-Blind Mode for Two-Rate Tricolor Marking**

With the color-blind mode of two-rate tricolor marking, all packets are evaluated against the committed information rate (CIR). If a packet exceeds the CIR, it is evaluated against the peak information rate (PIR). Packets that exceed the CIR but are below the PIR are marked yellow (medium-high). Packets that exceed the PIR are marked red (high).

**Table 435: Color-Blind Mode TCM Color-to-PLP Mapping**

Color	PLP	Meaning
Green	low	Packet does not exceed the CIR.
Yellow	medium-high	Packet exceeds the CIR but does not exceed the PIR.
Red	high	Packet exceeds the PIR.

#### **Related Documentation**

- [Overview of Policers on page 5241](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296](#)

## **Understanding Color-Aware Mode for Two-Rate Tricolor Marking**

In color-aware mode, the treatment the packet receives depends on its classification. Marking can increase the preassigned PLP but cannot decrease it

#### **Summary of PLP Changes**

[Table 436 on page 5250](#) shows how a packet's incoming priority can be modified with two-rate marking.

Table 436: Color-Aware Mode Two-Rate PLP Mapping

Incoming PLP	Packet Metered Against	Possible Cases	Outgoing PLP
<b>low</b>	CIR and PIR	Packet does not exceed the CIR.	<b>low</b>
		Packet exceeds the CIR but not the PIR.	<b>medium-high</b>
		Packet exceeds the PIR.	<b>high</b>
<b>medium-low</b>	PIR only	Packet does not exceed the PIR.	<b>medium-low</b>
		Packet exceeds the PIR.	<b>high</b>
<b>medium-high</b>	PIR only	Packet does not exceed the PIR.	<b>medium-high</b>
		Packet exceeds the PIR.	<b>high</b>
<b>high</b>	Not metered by the policer.	All cases.	<b>high</b>

The following sections describe color-aware two-rate PLP mapping in more detail.

#### Effect on Green Packets (Low PLP)

Packets belonging to the green class have already been marked by a classifier with low PLP. The marking policer can leave the packet's PLP unchanged or increase the PLP to medium-high or high. These packets are therefore metered against both the CIR and the PIR. For example, if a behavior aggregate or multifield classifier marks a packet with low PLP and the two-rate TCM policer is in color-aware mode, the output loss priority is as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, some of the packets are marked as medium-high PLP and some of the packets remain marked as low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as low PLP.

#### Effect on Yellow Packets (Medium PLP)

Packets belonging to the yellow class have already been marked by a classifier with medium-low or medium-high PLP. The marking policer can leave the PLP unchanged or increase it to high. These packets are therefore metered against the PIR only. For example, if a behavior aggregate (BA) or multifield classifier marks a packet with medium-low PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-low PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-low PLP.

If a BA or multifield classifier marks a packet with medium-high PLP and the two-rate TCM policer is in color-aware mode, the policer assigns output loss priority as follows:

- If the rate of traffic flow is less than the CIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the CIR but less than the PIR, the packets remain marked as medium-high PLP.
- If the rate of traffic flow is greater than the PIR, some of the packets are marked as high PLP and some of the packets remain marked as medium-high PLP.

#### Effect on Red Packets (High PLP)

Packets belonging to the red class have already been marked by a classifier with high PLP. Because the policer cannot decrease the PLP, it does not change it, and these packets are not metered against the CIR or the PIR.

#### **Related Documentation**

- [Overview of Policers on page 5241](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296](#)

---

## Port Security

- [Overview of Access Port Protection on page 5251](#)
- [Understanding Port Security on page 5254](#)
- [Understanding DHCP Snooping for Port Security on page 5256](#)
- [Understanding DAI for Port Security on page 5263](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 5266](#)
- [Understanding Trusted and Untrusted Ports on page 5268](#)
- [Understanding Trusted DHCP Servers for Port Security on page 5268](#)
- [Understanding DHCP Option 82 for Port Security on page 5269](#)
- [Understanding Static ARP Entries on page 5271](#)

### Overview of Access Port Protection

Port security features can protect a switch against various types of attacks. Protection methods against some common attacks are:

- [Mitigation of Ethernet Switching Table Overflow Attacks on page 5252](#)
- [Mitigation of Rogue DHCP Server Attacks on page 5252](#)

- [Protection Against ARP Spoofing Attacks on page 5253](#)
- [Protection Against DHCP Snooping Database Alteration Attacks on page 5253](#)
- [Protection Against DHCP Starvation Attacks on page 5253](#)

### **Mitigation of Ethernet Switching Table Overflow Attacks**

---

In an overflow attack on an Ethernet switching table, an intruder sends so many requests from new MAC addresses that the table cannot learn all the addresses. The attack forces the switch to send broadcast messages when it needs to send traffic to addresses for which it lacks MAC addresses. In addition to generating unnecessary traffic, the attacker might be able to sniff the broadcast packets.

To mitigate such attacks, you can configure a limit for learned MAC addresses or allow only specific MAC addresses. Use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface or interfaces. By setting the MAC addresses that are explicitly allowed, you ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table.

### **Mitigation of Rogue DHCP Server Attacks**

---

By default, all access ports are untrusted, and all trunk ports are trusted with regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. If someone connects an unauthorized DHCP server to a trusted port, the unauthorized server can start issuing IP addresses and configuration information to the network's DHCP clients. The information provided to the clients by this server can disrupt their network access. The unauthorized server might also assign itself as the default gateway device for the network. An attacker can then sniff the network traffic and perpetrate a man-in-the-middle attack—that is, it misdirects traffic intended for a legitimate network device to a device of its choice.

To mitigate this problem, set the interface to which the unauthorized server is connected as untrusted. That action blocks all ingress DHCP server messages from that interface.



**NOTE:** The switch logs all DHCP server packets that are received on untrusted ports. For example:

```
5 untrusted DHCPOFFER received, interface xe-0/0/2.0[65], vlan v1[10] server  
ip/mac 12.12.12.1/00:00:00:00:01:12 offer ip/client mac  
12.12.12.253/00:AA:BB:CC:DD:01
```

You can use these messages to detect unauthorized DHCP servers on the network.



**NOTE:** If you attach a DHCP server to an access port, you must configure the port as trusted.

---

### Protection Against ARP Spoofing Attacks

---

In ARP spoofing, an attacker sends faked ARP messages on the network. The attacker associates its own MAC address with the IP address of a network device connected to the switch. Any traffic sent to that IP address is instead sent to the attacker. Now the attacker can create various types of problems, including sniffing the packets that were meant for another host and perpetrating man-in-the middle attacks. (In a man-in-the-middle attack, the attacker intercepts messages between two hosts, reads them, and perhaps alters them, all without the original hosts knowing that their communications have been compromised.)

To protect against ARP spoofing on your switch, enable both DHCP snooping and dynamic ARP inspection (DAI). DHCP snooping builds and maintains the DHCP snooping table. That table contains the MAC addresses, IP addresses, lease times, binding types, VLAN information, and interface information for the untrusted interfaces on the switch. DAI uses the information in the DHCP snooping table to validate ARP packets. Invalid ARP packets are blocked, and when they are blocked, a system log message is recorded that includes the type of ARP packet and the sender's IP address and MAC address.

*See Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks.*

### Protection Against DHCP Snooping Database Alteration Attacks

---

In an attack designed to alter the DHCP snooping database, an intruder introduces a DHCP client on one of the switch's untrusted access interfaces that has a MAC address identical to that of a client on another untrusted port. The intruder acquires the DHCP lease, which results in changes to the entries in the DHCP snooping table. Subsequently, what would have been valid ARP requests from the legitimate client are blocked.

To protect against this type of alteration of the DHCP snooping database, configure MAC addresses that are explicitly allowed on the interface. *See Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks.*

### Protection Against DHCP Starvation Attacks

---

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses so that trusted DHCP servers cannot keep up with requests from legitimate DHCP clients. The address space of those servers is completely used up, so they can no longer assign IP addresses and lease times to clients. DHCP requests from those clients are either dropped—that is, the result is a denial of service (DoS)—or directed to a rogue DHCP server set up by the attacker to imitate a legitimate DHCP server.

To protect the switch from DHCP starvation attacks, use the MAC limiting feature. Specify the maximum number of MAC addresses that the switch can learn on the access interfaces to which DHCP clients connect. The DHCP server or servers can then supply only the specified number of IP addresses over each of those interfaces. If a DHCP starvation attack occurs after the maximum number of IP addresses has been assigned, the attack fails.

**Related Documentation**

- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 5266](#)
- [Configuring MAC Limiting](#)
- [Verifying That MAC Limiting Is Working Correctly on page 5385](#)
- [Understanding DHCP Option 82 for Port Security on page 5269](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)
- [Understanding DAI for Port Security on page 5263](#)

## Understanding Port Security

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. Port security features help protect the access ports on your device against the loss of information and productivity that such attacks can cause.

The Juniper Networks Junos operating system (Junos OS) provides features to help secure ports on a device. Ports can be categorized as either trusted or untrusted. You apply policies appropriate to each category to protect ports against various types of attacks.

Basic port security features are enabled in the device's default configuration. You can configure additional features with minimal configuration steps.

Depending on the particular feature, you can configure the feature either on VLANs or bridge domain interfaces.

Port security features supported on switching devices are:

- DHCP snooping—Filters and blocks ingress Dynamic Host Configuration Protocol (DHCP) server messages on untrusted ports; builds and maintains a database of valid bindings between IP addresses and MAC addresses (IP-MAC bindings), which is called the DHCP snooping database.



**NOTE:** DHCP snooping is not enabled in the default configuration of the switching device. DHCP snooping is enabled on a VLAN or bridge domain. The details of enabling DHCP snooping depend on the particular device.

- DHCPv6 snooping—DHCP for IPv6 (DHCPv6) equivalent for DHCP snooping.
- DHCP option 82—Also known as the DHCP Relay Agent information option. This DHCPv4 feature helps protect the switching device against attacks such as spoofing of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- DHCPv6 option 37—Option 37 is the DHCPv6 equivalent of option 82 and is enabled by default when DHCPv6 is enabled on a VLAN.
- Dynamic ARP inspection (DAI)—Prevents Address Resolution Protocol (ARP) spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping



database, and filtering decisions are made on the basis of the results of those comparisons. You enable DAI on a VLAN.

- IPv6 neighbor discovery inspection—Prevents IPv6 address spoofing attacks. Neighbor discovery requests and replies are compared against entries in the DHCPv6 snooping database, and filtering decisions are made on the basis of the results of those comparisons. You enable neighbor discovery inspection on a VLAN.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. With IP source guard enabled, the source IP address in the packet sent from an untrusted access interface is validated against the source MAC address in the DHCP snooping database. The packet is forwarded if the source IP-MAC binding is valid; if the binding is not valid, the packet is discarded. You enable IP source guard on a VLAN or bridge domain.



**NOTE:** IP source guard is not supported on the QFX Series.

- IPv6 source guard—IP source guard for IPv6.



**NOTE:** IPv6 source guard is not supported on the QFX Series.

- MAC limiting—Protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You can enable MAC limiting on an interface.
- MAC move limiting—(Not supported on EX9200) Tracks MAC movement and detects MAC spoofing on access ports. You enable this feature on a VLAN or bridge domain.
- Persistent MAC learning—Also known as sticky MAC. Persistent MAC learning enables interfaces to retain dynamically learned MAC addresses across switch reboots. You enable this feature on an interface.
- RA guard—Examines incoming Router Advertisement (RA) messages and decides whether to forward or block them based on statically configured IPv6/MAC address bindings. If the content of the RA message does not match the bindings, the message is dropped.
- Trusted DHCP server—Configuring the DHCP server on a trusted port protects against rogue DHCP servers sending leases. You enable this feature on an interface (port). By default, access ports are untrusted, and trunk ports are trusted. (Access ports are the switch ports that connect to Ethernet endpoints such as user PCs and laptops, servers, and printers. Trunk ports are the switch ports that connect an Ethernet switch to other switches or to routers.)

#### Related Documentation

- [Security Features for EX Series Switches Overview](#)
- [Understanding DHCP Snooping for Port Security on page 5256](#)
- [Understanding DHCP Snooping for Port Security](#)
- [Understanding IPv6 Neighbor Discovery Inspection](#)

- [Understanding DAI for Port Security on page 5263](#)
- [Understanding IP Source Guard for Port Security on EX Series Switches](#)
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches](#)
- [Understanding DHCP Option 82 for Port Security on Switching Devices](#)

## Understanding DHCP Snooping for Port Security

DHCP snooping enables the switching device, which can be either a switch or a router, to monitor and control DHCP messages received from untrusted devices connected to it. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information and builds and maintains a database of valid bindings between IP addresses and MAC address (IP-MAC bindings) called the DHCP snooping database. Only clients with valid bindings are allowed access to the network.

- [DHCP Snooping Basics on page 5256](#)
- [DHCP Snooping Process on page 5257](#)
- [DHCPv6 Snooping on page 5258](#)
- [Rapid Commit for DHCPv6 on page 5259](#)
- [DHCP Server Access on page 5259](#)
- [Static IP Address Additions to the DHCP Snooping Database on page 5262](#)
- [Snooping DHCP Packets That Have Invalid IP Addresses on page 5262](#)
- [Prioritizing Snooped Packets on page 5263](#)

---

### DHCP Snooping Basics

The Dynamic Host Configuration Protocol (DHCP) allocates IP addresses dynamically, *leasing* addresses to devices so that the addresses can be reused when no longer needed. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port).

By default, all trunk ports on the switch are trusted and all access ports are untrusted for DHCP snooping.

When DHCP snooping is enabled, the lease information from the switching device is used to create the DHCP snooping table, also known as the binding table. The table shows current IP-MAC bindings, as well as lease time, type of binding, names of associated VLANs, and associated interfaces.



**NOTE:** DHCP snooping is disabled in the default configuration of the switching device. You must explicitly enable DHCP snooping by setting `examine-dhcp` at the `[edit ethernet-switching-options secure-access-port]` hierarchy level.

Entries in the DHCP snooping database are updated in these events:

- When a DHCP client releases an IP address (sends a DHCPRELEASE message). In this event, the associated mapping entry is deleted from the database.
- If you move a network device from one VLAN to another. In this event, typically the device needs to acquire a new IP address. Therefore, its entry in the database, including its VLAN ID, is updated.
- When the lease time (timeout value) assigned by the DHCP server expires. In this event, the associated entry is deleted from the database.



**TIP:** By default, the IP-MAC bindings are lost when the switching device is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the `dhcp-snooping-file` statement to store the database file either locally or remotely.

You can configure the switching device to snoop DHCP server responses from particular VLANs only. This prevents spoofing of DHCP server messages.

You configure DHCP snooping per VLAN, not per interface (port). DHCP snooping is disabled by default on switching devices.

### DHCP Snooping Process

The basic process of DHCP snooping consists of the following steps:



**NOTE:** When DHCP snooping is enabled for a VLAN, all DHCP packets sent from the network devices in that VLAN are subjected to DHCP snooping. The final IP-MAC binding occurs when the DHCP server sends DHCPACK to the DHCP client.

1. The network device sends a DHCPDISCOVER packet to request an IP address.
2. The switching device forwards the packet to the DHCP server.
3. The server sends a DHCPOFFER packet to offer an address. If the DHCPOFFER packet is from a trusted interface, the switching device forwards the packet to the DHCP client.
4. The network device sends a DHCPREQUEST packet to accept the IP address. The switching device adds an IP-MAC placeholder binding to the database. The entry is

considered a placeholder until a DHCPACK packet is received from the server. Until then, the IP address could still be assigned to some other host.

5. The server sends a DHCPACK packet to assign the IP address or a DHCPNAK packet to deny the address request.
6. The switching device updates the DHCP snooping database according to the type of packet received:
  - If the switching device receives a DHCPACK packet, it updates lease information for the IP-MAC bindings in its database.
  - If the switching device receives a DHCPNAK packet, it deletes the placeholder.



**NOTE:** The DHCP snooping database is updated only after the DHCPREQUEST packet has been sent.

For general information about the messages that the DHCP client and DHCP server exchange during the assignment of an IP address for the client, see the *Junos OS Administration Library for Routing Devices*.

### DHCPv6 Snooping

DHCPv6 snooping is the equivalent of DHCP snooping for IPv6. The process for DHCPv6 snooping is similar to that for DHCP snooping, but uses different names for the messages exchanged between the client and server to assign IPv6 addresses. [Table 437 on page 5258](#) shows DHCPv6 messages and their DHCP equivalents.

**Table 437: DHCPv6 Messages and Equivalent DHCPv4 Messages**

Sent by	DHCPv6 Messages	Equivalent DHCP Messages
Client	SOLICIT	DHCPDISCOVER
Server	ADVERTISE	DHCPOFFER
Client	REQUEST, RENEW, REBIND	DHCPREQUEST
Server	REPLY	DHCPACK/DHCPNAK
Client	RELEASE	DHCPRELEASE
Client	INFORMATION-REQUEST	DHCPINFORM
Client	DECLINE	DHCPDECLINE
Client	CONFIRM	none
Server	RECONFIGURE	DHCPFORCERENEW
Client	RELAY-FORW, RELAY-REPLY	none

---

## Rapid Commit for DHCPv6

---

DHCPv6 provides for a Rapid Commit option (DHCPv6 option 14), which, when supported by the server and set by the client, shortens the exchange from a four-way relay to a two-message handshake. For more information about enabling the Rapid Commit option, see *Enabling DHCPv6 Rapid Commit Support*.

In the rapid commit process:

1. The DHCPv6 client sends out a SOLICIT message that contains a request that rapid assignment of address, prefix, and other configuration parameters be preferred.
2. If the DHCPv6 server supports rapid assignment, it responds with a REPLY message, which contains the assigned IPv6 address and prefix and other configuration parameters.

---

## DHCP Server Access

---

You can configure a switching device's access to the DHCP server in three ways:

- [Switching Device, DHCP Clients, and DHCP Server Are All on the Same VLAN on page 5259](#)
- [Switching Device Acts as DHCP Server on page 5260](#)
- [Switching Device Acts as Relay Agent on page 5261](#)

### ***Switching Device, DHCP Clients, and DHCP Server Are All on the Same VLAN***

When the switching device, DHCP clients, and DHCP server are *all members of the same VLAN*, the DHCP server can be connected to the switching device in one of two ways:

- The server is directly connected to the same switching device as the one connected to the DHCP clients (the hosts, or network devices, that are requesting IP addresses from the server). The VLAN is enabled for DHCP snooping to protect the untrusted access ports. The trunk port is configured by default as a trusted port. See [Figure 182 on page 5260](#).
- The server is connected to an intermediary switching device (Switching Device 2). The DHCP clients are connected to Switching Device 1, which is connected through a trunk port to Switching Device 2. Switching Device 2 is being used as a transit device. The VLAN is enabled for DHCP snooping to protect the untrusted access ports. The trunk port is configured by default as a trusted port. As shown in [Figure 183 on page 5260](#), ge-0/0/11 is a trusted trunk port.

Figure 182: DHCP Server Connected Directly to a Switching Device





Figure 183: DHCP Server Connected Directly to Switching Device 2, with Switching Device 2 Connected to Switching Device 1 Through a Trusted Trunk Port




### Switching Device Acts as DHCP Server



**NOTE:** The switching device acting as a DHCP server is not supported on the QFX Series.

The switching device itself is configured as a DHCP server; this is known as a *local configuration*. See [Figure 184 on page 5261](#).

Figure 184: Switching Device Is the DHCP Server

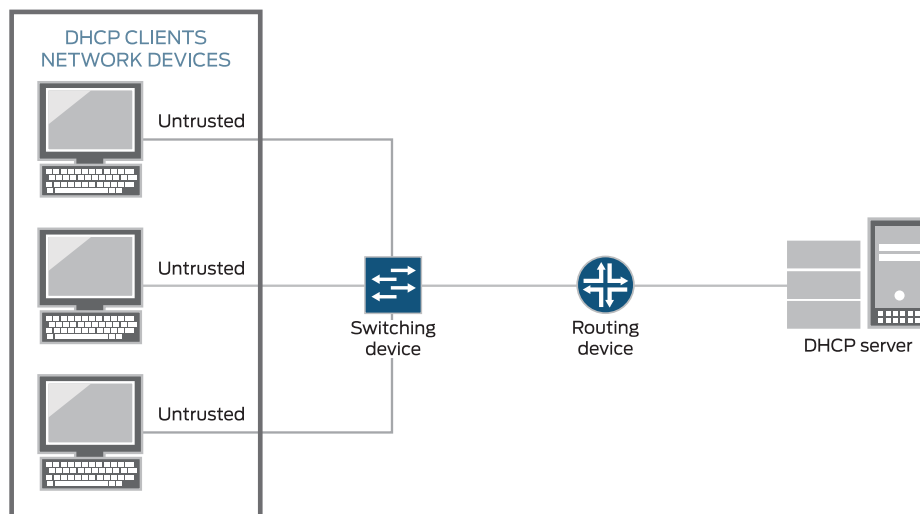
***Switching Device Acts as Relay Agent***

The switching device functions as a relay agent when the DHCP clients or the DHCP server is connected to the device through a Layer 3 interface. The Layer 3 interfaces on the switching device are configured as routed VLAN interfaces (RVIs), which are also known as integrated routing and bridging (IRB) interfaces. The trunk interfaces are trusted by default.

These two scenarios illustrate the switching device acting as a relay agent:

- The DHCP server and clients are in different VLANs.
- The switching device is connected to a router that is in turn connected to the DHCP server. See [Figure 185 on page 5262](#).

**Figure 185: Switching Device Acting as Relay Agent Through Router to DHCP Server**



8042487

### Static IP Address Additions to the DHCP Snooping Database

You can add specific static IP addresses to the database as well as have the addresses dynamically assigned through DHCP snooping. To add static IP addresses, you supply the IP address, the MAC address of the device, the interface on which the device is connected, and the VLAN with which the interface is associated. No lease time is assigned to the entry. The statically configured entry never expires.

### Snooping DHCP Packets That Have Invalid IP Addresses

If you enable DHCP snooping on a VLAN and then devices on that VLAN send DHCP packets that request invalid IP addresses, these invalid IP addresses are stored in the DHCP snooping database until they are deleted when their default timeout is reached. To eliminate this unnecessary consumption of space in the DHCP snooping database, the switching device drops the DHCP packets that request invalid IP addresses, preventing the snooping of these packets. The invalid IP addresses are:

- 0.0.0.0
- 128.0.x.x
- 191.255.x.x
- 192.0.0.x
- 223.255.255.x
- 224.x.x.x
- 240.x.x.x to 255.255.255.255



## Prioritizing Snooped Packets



**NOTE:** Prioritizing snooped packets is not supported on the QFX Series and the EX4600 switch.

You can use class-of-service (CoS) forwarding classes and queues to prioritize DHCP snooped packets for a specified VLAN. This type of configuration places the DHCP snooped packets for that VLAN in a specified egress queue, so that the security procedure does not interfere with the transmission of high-priority traffic. For additional information, see *Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic*.

### Related Documentation

- [Understanding Port Security on page 5254](#)
- [Understanding Trusted DHCP Servers for Port Security on page 5268](#)
- [Enabling a Trusted DHCP Server on an MX Series Router \(CLI Procedure\)](#)
- [Configuring Static IP Addresses for DHCP Bindings on Access Ports for MX Series Routers \(CLI Procedure\)](#)

## Understanding DAI for Port Security

Dynamic ARP inspection (DAI) protects switching devices against ARP spoofing.

DAI inspects Address Resolution Protocol (ARP) packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP spoofing (also known as ARP poisoning or ARP cache poisoning). ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. When an attacker tries to use a forged ARP packet to spoof an address, the switch compares the address with entries in the database. If the media access control (MAC) address or IP address in the ARP packet does not match a valid entry in the DHCP snooping database, the packet is dropped.

ARP packets are sent to the Routing Engine and are rate-limited to protect the switching device from CPU overload.

- [Address Resolution Protocol on page 5263](#)
- [ARP Spoofing on page 5264](#)
- [Dynamic ARP Inspection on page 5264](#)
- [Prioritizing Inspected Packets on page 5265](#)

### Address Resolution Protocol

Sending IP packets on a multi-access network requires mapping an IP address to an Ethernet MAC address.

Ethernet LANs use ARP to map MAC addresses to IP addresses.

The switching device maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

### ARP Spoofing

---

ARP spoofing is one way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switching device sending traffic to the proper network device, it sends the traffic to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that must have gone to another device. The result is that traffic from the switching device is misdirected and cannot reach its proper destination.

One type of ARP spoofing is gratuitous ARP, which is when a network device sends an ARP request to resolve its own IP address. In normal LAN operation, gratuitous ARP messages indicate that two devices have the same MAC address. They are also broadcast when a network interface card (NIC) in a device is changed and the device is rebooted, so that other devices on the LAN update their ARP caches. In malicious situations, an attacker can poison the ARP cache of a network device by sending an ARP response to the device that directs all packets destined for a certain IP address to go to a different MAC address instead.

To prevent MAC spoofing through gratuitous ARP and through other types of spoofing, the switches examine ARP responses through DAI.

### Dynamic ARP Inspection

---

DAI examines ARP requests and responses on the LAN and validates ARP packets. The switch intercepts ARP packets from an access port and validates them against the DHCP snooping database. If no IP-MAC entry in the database corresponds to the information in the ARP packet, DAI drops the ARP packet and the local ARP cache is not updated with the information in that packet. DAI also drops ARP packets when the IP address in the packet is invalid. ARP probe packets are not subjected to dynamic ARP inspection. The switch always forwards such packets.

Junos OS for EX Series switches and the QFX Series uses DAI for ARP packets received on access ports because these ports are untrusted by default. Trunk ports are trusted by default, and therefore ARP packets bypass DAI on them.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

If you set an interface to be a DHCP trusted port, it is also trusted for ARP packets.

**NOTE:**

- If your switching device is an EX Series switch and uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see *Enabling a Trusted DHCP Server (CLI Procedure)* for information about configuring an access interface to be a DHCP trusted port.
- If your switching device is an EX Series switch and is *not* using Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style, see *Enabling a Trusted DHCP Server (CLI Procedure)* for information about configuring an access interface to be a DHCP trusted port.

For packets directed to the switching device to which a network device is connected, ARP queries are broadcast on the VLAN. The ARP responses to those queries are subjected to the DAI check.

For DAI, all ARP packets are trapped to the Packet Forwarding Engine. To prevent CPU overloading, ARP packets destined for the Routing Engine are rate-limited.

If the DHCP server goes down and the lease time for an IP-MAC entry for a previously valid ARP packet runs out, that packet is blocked.

### Prioritizing Inspected Packets



**NOTE:** Prioritizing inspected packets is not supported on the QFX Series and the EX4600 switch.

You can use class-of-service (CoS) forwarding classes and queues to prioritize DAI packets for a specified VLAN. This type of configuration places inspected packets for that VLAN in the egress queue, that you specify, ensuring that the security procedure does not interfere with the transmission of high-priority traffic.

#### Related Documentation

- [Understanding Port Security on page 5254](#)
- [Understanding DHCP Snooping for Port Security on page 5256](#)
- [Example: Configuring Basic Port Security Features](#)
- [Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch](#)
- [Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks](#)
- [Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing](#)
- [Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#)
- [Enabling Dynamic ARP Inspection \(CLI Procedure\)](#)

- [Enabling Dynamic ARP Inspection \(J-Web Procedure\)](#)

## Understanding MAC Limiting and MAC Move Limiting for Port Security

MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on Layer 2 interfaces (ports). MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

- [MAC Limiting on page 5266](#)
- [MAC Move Limiting on page 5266](#)
- [Actions for MAC Limiting on page 5267](#)
- [MAC Addresses That Exceed the MAC Limit or MAC Move Limit on page 5267](#)

### MAC Limiting

---

MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. Junos OS provides two MAC limiting methods:

- Maximum number of MAC addresses—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- Allowed MAC addresses—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.



**NOTE:** If you do not want the system to log messages about invalid MAC addresses received by an interface that has been configured for allowed MAC addresses, disable the logging by configuring the [no-allowed-mac-log](#) statement.

You configure MAC limiting per interface, not per VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface (including tagged-access interfaces) or on all Layer 2 access interfaces.

### MAC Move Limiting

---

MAC move limiting causes the switch to track the number of times a MAC address can move to a new interface (port). It can help to prevent MAC spoofing, and it can also detect and prevent loops.

If a MAC address moves more than the configured number of times within 1 second, the switch performs the configured action. You can configure MAC move limiting to apply to all VLANs or to a specific VLAN.



**CAUTION:** Mac move limiting does not work properly on a QFX5100 switch used as a Node device in a QFabric system. Do not use this feature on a QFX5100 switch in a QFabric system.

### Actions for MAC Limiting

You can choose to have one of the following actions performed when the limit of MAC addresses or the limit of MAC moves is exceeded:

- **drop**—Drop the packet and generate a system log entry. This is the default.
- **log**—Do not drop the packet but generate a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interface and generate an alarm. If you configure the switch with the **port-error-disable** statement, the disabled interface recovers automatically upon expiration of the specified timeout. If this is not configured, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

See descriptions of results of these various action settings in “[Verifying That MAC Limiting Is Working Correctly](#)” on page 5385.

If you set a MAC limit to apply to all interfaces on the switch, you can override that setting for a particular interface by specifying action **none**. See *Configuring the none Action to Override a MAC Limit Applied to All Interfaces (CLI Procedure)*

### MAC Addresses That Exceed the MAC Limit or MAC Move Limit

If you have configured the **port-error-disable** statement, you can view which interfaces are temporarily disabled because the MAC limit or MAC move limit was exceeded. Use the **show ethernet-switching interfaces** command.

The log messages that indicate the MAC limit or MAC move limit has been exceeded include the offending MAC addresses.

#### Related Documentation

- [Understanding Port Security on page 5254](#)
- [Configuring MAC Limiting](#)
- [Configuring MAC Move Limiting \(CLI Procedure\)](#)
- [Verifying That MAC Limiting Is Working Correctly on page 5385](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 5388](#)
- [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)
- [Example: Configuring Basic Port Security Features](#)
- [no-allowed-mac-log on page 5339](#)

## Understanding Trusted and Untrusted Ports

By default, all access ports are untrusted and all trunk ports are trusted in regard to DHCP. Trusted ports allow DHCP servers to provide IP addresses and other information to requesting devices. Untrusted ports drop traffic from DHCP servers to prevent unauthorized servers from providing any configuration information to clients.

If you attach a DHCP server to an access port, you must configure the port as trusted. Before you do so, ensure that the server is physically secure—that is, that access to the server is monitored and controlled.

- Related Documentation**
- [Understanding DHCP Snooping for Port Security on page 5256](#)
  - *Example: Configuring Basic Port Security Features*
  - *Enabling a Trusted Port for DHCP*

## Understanding Trusted DHCP Servers for Port Security

Any interface on the switching device that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

- Related Documentation**
- [Understanding DHCP Snooping for Port Security on page 5256](#)
  - *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
  - *Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing*
  - *Enabling a Trusted DHCP Server (CLI Procedure)*
  - *Enabling a Trusted DHCP Server (CLI Procedure)*

## Understanding DHCP Option 82 for Port Security

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on Ethernet LAN switches send requests for IP addresses in order to access the Internet. The switch forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to perpetrate address spoofing and starvation.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Juniper Networks Junos operating system (Junos OS) implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>.

This topic covers:

- [DHCP Option 82 Processing on page 5269](#)
- [Suboption Components of Option 82 on page 5270](#)
- [Configurations That Support Option 82 on page 5270](#)

### DHCP Option 82 Processing

If DHCP option 82 is enabled on the switch, then when a DHCP client that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See “[Suboption Components of Option 82](#)” on page 5270 for details about option 82 information.

You can enable DHCP option 82 on a single VLAN or on all VLANs on the switch. You can also configure it on Layer 3 interfaces (in routed VLAN interfaces, or RVIs) when the switch is functioning as a relay agent.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards or relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.



**NOTE:** To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If it is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information in setting parameters and it does not echo the information in its response message.

---

### Suboption Components of Option 82

---

When configuring DHCP option 82, you can use the following suboptions:

- **circuit ID**—Identifies the circuit (interface and/or VLAN) on the switch on which the request was received. The circuit ID contains the interface name and/or VLAN name, with the two elements separated by a colon—for example, **xe-0/0/10:vlan1**. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, **xe-0/0/10**.

Use the **prefix** option to add an optional prefix to the circuit ID. If you enable the **prefix** option, the hostname for the switch is used as the prefix; for example, **switch1:xe-0/0/10:vlan1**.

You can also specify that the interface description be used rather than the interface name and that the VLAN ID be used rather than the VLAN name.

- **remote ID**—Identifies the host. By default, the remote ID is the MAC address of the switch. You can specify that the remote ID be the hostname of the switch, the interface description, or a character string of your choice. You can also add an optional prefix to the remote ID.
- **vendor ID**—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value **Juniper** is used. To specify a value, you type a character string.

---

### Configurations That Support Option 82


---

You can use option 82 with the following configurations:

- The DHCP client and the DHCP server are on the same VLAN. In this case the switch forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. For this configuration, you set DHCP option 82 at the **[edit ethernet-switching-options secure-access-port vlan]** hierarchy level.
- The DHCP client or the DHCP server is connected to the switch through a Layer 3 interface and the switch is configured to relay DHCP requests. [Figure 186 on page 5271](#) illustrates a scenario for the switch-as-relay-agent; in this instance, the switch relays requests through a router to the server.



Figure 186: Switch Relays DHCP Requests to Server



For the configuration shown in [Figure 186 on page 5271](#), you set DHCP option 82 at the `[edit forwarding-options helpers bootp]` hierarchy level.

**Related Documentation**

- [Overview of Access Port Protection on page 5251](#)
- [DHCP and BOOTP Relay Overview on page 5431](#)
- `dhcp-option82`
- *Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server*
- *Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*

## Understanding Static ARP Entries

You can create explicit mappings between IP addresses and MAC addresses, which are called static ARP table entries. Unlike dynamically learned ARP entries, static entries do not age out. You might want to create static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically for any reason.

**Related Documentation**

- [Configuring Static ARP Entries on page 1676](#)
- `arp`

## Device Security

---

- [Understanding Storm Control on page 5272](#)
- [Understanding Unicast RPF on page 5273](#)
- [Understanding Unknown Unicast Forwarding on page 5277](#)

### Understanding Storm Control

A traffic storm occurs when broadcast packets prompt receiving devices to broadcast packets in response. This prompts further responses, creating a snowball effect. The switch is flooded with packets, which creates unnecessary traffic that leads to poor performance or even a complete loss of service by some clients. Storm control causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

Storm control is enabled by default on ELS platforms and disabled by default on non-ELS platforms. If storm control is enabled, the default level is 80 percent of the available bandwidth for ingress traffic. You can change the storm control level by configuring it as a specific bandwidth value. (The **level** configuration statement, which allows you to configure the storm control level as a percentage of the combined broadcast and unknown unicast streams, is deprecated and might be removed from future releases. We recommend that you phase out its use and replace it with the **bandwidth** statement.)



NOTE: Storm control is not enabled by default on MX platforms.



NOTE: When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually enforced. For example, if you configure a bandwidth limit of 150 Kbps, storm control enforces a bandwidth limit of 128 Kbps.



NOTE: On an FCoE-FC gateway, storm control must be disabled on all Ethernet interfaces that belong to an FCoE VLAN to prevent FCoE traffic from being dropped. Configuring storm control on an Ethernet interface that is included in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss. After disabling storm control on all interfaces, enable storm control on any interfaces that are not part of an FCoE-FC gateway on which you want to use storm control. However, on an FCoE transit switch, you can enable storm control on interfaces that carry FCoE traffic.



**CAUTION:** The Junos OS allows you to configure a storm control value that exceeds the bandwidth of the interface. If you configure an interface this way, storm control does not drop broadcast or unknown unicast packets even if they consume all the available bandwidth.

To recognize a storm, you must be able to identify when traffic has reached an abnormal level. Suspect a storm when operations begin timing out and network response times slow down. Users might be unable to access expected services. Monitor the percentage of broadcast and unknown unicast traffic in the network when it is operating normally. This data can then be used as a benchmark to determine when traffic levels are too high. You can then configure storm control to set the level at which you want to drop broadcast and unknown unicast traffic.

#### Related Documentation

- *Example: Configuring Storm Control to Prevent Network Outages*
- *Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)*
- *Disabling Storm Control on FCoE Interfaces on an FCoE-FC Gateway*
- [action-shutdown on page 5365](#)
- *interface (Storm Control)*
- [port-error-disable on page 5341](#)
- *storm-control*

## Understanding Unicast RPF

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF.



**NOTE:** On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see [“Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches” on page 5276](#).

This topic covers:

- [Unicast RPF for Switches Overview on page 5274](#)
- [Unicast RPF Implementation on page 5274](#)

- [When to Enable Unicast RPF on page 5275](#)
- [When Not to Enable Unicast RPF on page 5276](#)
- [Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches on page 5276](#)

### **Unicast RPF for Switches Overview**

---

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see [“When to Enable Unicast RPF” on page 5275](#).)

For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at <http://www.ietf.org/rfc/rfc3704.txt>.

### **Unicast RPF Implementation**

---

This section includes:

- [Unicast RPF Packet Filtering on page 5274](#)
- [Bootstrap Protocol \(BOOTP\) and DHCP Requests on page 5274](#)
- [Default Route Handling on page 5275](#)

#### **Unicast RPF Packet Filtering**

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

#### **Bootstrap Protocol (BOOTP) and DHCP Requests**

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The

switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

### Default Route Handling

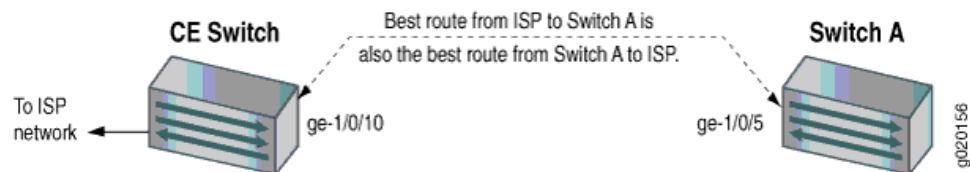
If the best return path to the source is the default route (0.0.0.0) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

### When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in [Figure 187 on page 5275](#). Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 187: Symmetrically Routed Interfaces



Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link



**NOTE:** Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.



**TIP:** Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.


### When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in [Figure 188 on page 5276](#). This means that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

**Figure 188: Asymmetrically Routed Interfaces**



**NOTE:** Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

### Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.



**NOTE:** You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

#### Related Documentation

- *Example: Configuring Unicast RPF on an EX Series Switch*
- [Configuring Unicast RPF \(CLI Procedure\) on page 5303](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 5304](#)

## Understanding Unknown Unicast Forwarding

Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that are traveling in a VLAN to all interfaces that are members of the VLAN. Forwarding this type of traffic can create unnecessary traffic that leads to poor network performance or even a complete loss of network service. This is known as a traffic storm.

To prevent a storm, you can disable the flooding of unknown unicast packets to all VLAN interfaces by configuring one VLAN or all VLANs to forward all unknown unicast traffic to a specific interface. This channels the unknown unicast traffic to a single interface.

#### Related Documentation

- *Configuring Unknown Unicast Forwarding (CLI Procedure)*
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 5305](#)
- *Understanding Storm Control on EX Series Switches*
- *Understanding Storm Control on Switching Devices*
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*
- *Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches*





## CHAPTER 61

# Configuration

- [Firewall and Policer Configuration Examples on page 5279](#)
- [Device Security Configuration Example on page 5288](#)
- [Firewall and Policer Configuration Tasks on page 5290](#)
- [Device Security Configuration Tasks on page 5303](#)
- [Configuration Statements for Firewall Filters on page 5306](#)
- [Configuration Statements for Policers on page 5314](#)
- [Configuration Statements for Port Security on page 5332](#)
- [Configuration Statements for Port Security on page 5344](#)
- [Configuration Statements for Device Security on page 5364](#)

### Firewall and Policer Configuration Examples

---

- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 5279](#)
- [Example: Using Two-Color Policers and Prefix Lists on page 5283](#)
- [Example: Using Policers to Manage Oversubscription on page 5286](#)

#### Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device

You can configure filter-based forwarding by using a firewall filter to forward matched traffic to a specific virtual routing instance.

This example describes how to set up filter-based forwarding:

- [Requirements on page 5279](#)
- [Overview and Topology on page 5280](#)
- [Configuration on page 5280](#)
- [Verification on page 5282](#)

#### Requirements

---

This example requires Junos OS Release 12.2X50-D20 or later.

## Overview and Topology

---

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address of the source application server. Any matching packets are routed to a virtual routing instance that sends the traffic to a security device. In this case, the security device must be able to forward the traffic to the destination application server. For this example, assume that the address of the destination application server is 192.168.0.1.

## Configuration

---

To configure filter-based forwarding:

### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste them into the CLI at the **[edit]** hierarchy level.

```
[edit]
set interfaces xe-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces xe-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter f1 term t1 from source-address 10.1.0.50/32
set firewall family inet filter f1 term t1 from protocol tcp
set interfaces xe-0/0/0 unit 0 family inet filter input f1
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface xe-0/0/3.0
set routing-instances vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
set firewall family inet filter f1 term t1 then routing-instance vrf01
```

### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*. To configure filter-based forwarding:

1. Configure an interface to connect to the application server:  

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet address 10.1.0.1/24
```
2. Configure an interface to connect to the security device:  

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 10.1.3.1/24
```
3. Create a firewall filter that matches packets based on the address of the application server that the traffic will be sent from. Also configure the filter so that it matches only TCP packets:  

```
[edit firewall]
user@switch# set family inet filter f1 term t1 from source-address 10.1.0.50/32
user@switch# set firewall family inet filter f1 term t1 from protocol tcp
```
4. Apply the filter to the interface that connects to the source application server and configure it to match incoming packets:  

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family inet filter input f1
```
5. Create a virtual router:  

```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```
6. Associate the virtual router with the interface that connects to the security device:

- ```
[edit routing-instances]
user@switch# set vrf01 interface xe-0/0/3.0
```
7. Configure the routing information for the virtual routing instance:
- ```
[edit routing-instances]
user@switch# set vrf01 routing-options static route 192.168.0.1/24 next-hop 10.1.3.254
```
8. Set the filter to forward packets to the virtual router:
- ```
[edit firewall]
user@switch# set family inet filter f1 term t1 then routing-instance vrf01
```

### Results

Check the results of the configuration:

```
user@switch> show configuration
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        filter {
          input f1;
        }
        address 10.1.0.1/24;
      }
    }
  }
  xe-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.3.1/24;
      }
    }
  }
}
firewall {
  family inet {
    filter f1 {
      term t1 {
        from {
          source-address {
            10.1.0.50/32;
          }
          protocol tcp;
        }
        then {
          routing-instance vrf01;
        }
      }
    }
  }
}
routing-instances {
  vrf01 {
    instance-type virtual-router;
    interface xe-0/0/1.0;
    routing-options {
```

```

static {
    route 12.34.56.0/24 next-hop 10.1.3.254;
}
}
}
}

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That Filter-Based Forwarding Was Configured on page 5282](#)

### Verifying That Filter-Based Forwarding Was Configured

**Purpose** Verify that filter-based forwarding was properly enabled on the switch.

**Action** 1. Use the `show interfaces filters` command:

```

user@switch> show interfaces filters xe-0/0/0.0
Interface      Admin Link Proto Input Filter      Output Filter
xe-0/0/0.0      up    down inet f1

```

2. Use the `show route forwarding-table` command:

```

user@switch> show route forwarding-table

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          user  1 0:12:f2:21:cf:0 ucst  331  4 me0.0
default          perm  0                rjct   36   3
0.0.0.0/32       perm  0                dscd   34   1
10.1.0.0/24      ifdn  0                rslv   613  1 xe-0/0/0.0
10.1.0.0/32      iddn  0 10.1.0.0        recv   611  1 xe-0/0/0.0
10.1.0.1/32      user  0                rjct   36   3
10.1.0.1/32      intf  0 10.1.0.1        locl   612  2
10.1.0.1/32      iddn  0 10.1.0.1        locl   612  2
10.1.0.255/32    iddn  0 10.1.0.255      bcst   610  1 xe-0/0/0.0
10.1.1.0/26      ifdn  0                rslv   583  1 vlan.0
10.1.1.0/32      iddn  0 10.1.1.0        recv   581  1 vlan.0
10.1.1.1/32      user  0                rjct   36   3
10.1.1.1/32      intf  0 10.1.1.1        locl   582  2
10.1.1.1/32      iddn  0 10.1.1.1        locl   582  2
10.1.1.63/32     iddn  0 10.1.1.63       bcst   580  1 vlan.0
255.255.255.255/32 perm  0                bcst   32   1

```

```

Routing table: vrf01.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct   559  2
0.0.0.0/32       perm  0                dscd   545  1
10.1.3.0/24      ifdn  0                rslv   617  1 xe-0/0/3.0
10.1.3.0/32      iddn  0 10.1.3.0        recv   615  1 xe-0/0/3.0
10.1.3.1/32      user  0                rjct   559  2
192.168.0.1/24   user  0 10.1.3.254      ucst   616  2 xe-0/0/3.0
192.168.0.1/24   user  0 10.1.3.254      ucst   616  2 xe-0/0/3.0
10.1.3.255/32    iddn  0 10.1.3.255      bcst   614  1 xe-0/0/3.0
224.0.0.0/4      perm  0                mdsc   546  1
224.0.0.1/32     perm  0 224.0.0.1       mcst   529  1

```

```

255.255.255.255/32 perm      0                bcst    543      1

Routing table: default.iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm      0                rjct    60       1

Routing table: vrf01.iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm      0                rjct   600       1

```

**Meaning** The output indicates that the filter was created on the interface and that the virtual routing instance is forwarding matching traffic to the correct IP address.

- Related Documentation**
- [Configuring Firewall Filters on page 5290](#)
  - [Understanding Filter-Based Forwarding on page 5212](#)
  - [Understanding Virtual Router Routing Instances on page 2898](#)

## Example: Using Two-Color Policers and Prefix Lists

If you provide specific amounts of bandwidth to internal or external customers, you can use policing to make sure that customers do not consume more bandwidth than they should receive. For example, you might connect many customers to one 10-Gbps interface and want to ensure that none of them congest the interface by using more bandwidth than they have been allotted.

You could accomplish this by creating a two-color policer similar to the following for each customer:

```

firewall {
  policer Limit-Customer-1 {
    if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 150m;
    }
    then discard;
  }
}

```

Creating a policer for each customer is clearly not a scalable solution, however. As an alternative, you can create prefix lists that group classes of customers and then create policers for each prefix list. For example, you could create prefix lists such as **Class-A-Customer-Prefixes**, **Class-B-Customer-Prefixes**, and **Class-C-Customer-Prefixes** (at the **[edit policy-options]** hierarchy level) and create the following corresponding policers:

```

firewall {
  policer Class-A {
    if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 150m;
    }
  }
}

```

```
    }
    then discard;
  }
  policer Class-B {
    if-exceeding {
      bandwidth-limit 75m;
      burst-size-limit 100m;
    }
    then discard;
  }
  policer Class-C {
    if-exceeding {
      bandwidth-limit 50m;
      burst-size-limit 75m;
    }
    then discard;
  }
}
```

You must create filter terms that specify the prefix lists in their **from** statements and the corresponding policers in their **then** statements similar to the following:

```
firewall
family inet {
  filter Class-A-Customers {
    term term-1 {
      from {
        destination-prefix-list {
          Class-A-Customer-Prefixes;
        }
      }
      then policer Class-A;
    }
  }
  filter Class-B-Customers {
    term term-1 {
      from {
        destination-prefix-list {
          Class-B-Customer-Prefixes;
        }
      }
      then policer Class-B;
    }
  }
  filter Class-C-Customers {
    term term-1 {
      from {
        destination-prefix-list {
          Class-C-Customer-Prefixes;
        }
      }
      then policer Class-C;
    }
  }
}
```

Here are the steps to create this firewall configuration:

1. Create the first policer:

```
[edit firewall]
user@switch# set policer Class-A if-exceeding bandwidth-limit 100m burst-size-limit 150m
user@switch# set policer Class-A then discard
```

2. Create the second policer:

```
[edit firewall]
user@switch# set policer Class-B if-exceeding bandwidth-limit 75m burst-size-limit 100m
user@switch# set policer Class-B then discard
```

3. Create the third policer:

```
[edit firewall]
user@switch# set policer Class-C if-exceeding bandwidth-limit 50m burst-size-limit 75m
user@switch# set policer Class-C then discard
```

4. Create a filter for class A customers:

```
[edit firewall]
user@switch# edit family inet filter Class-A-Customers
```

5. Configure the filter to send packets matching the Class-A-Customer-Prefixes prefix list to the Class-A policer:

```
[edit firewall family inet filter Class-A-Customers]
user@switch# set term term-1 from source-prefix-list Class-A-Customers
user@switch# set term term-1 then policer Class-A
```

6. Create a filter for class B customers:

```
[edit firewall]
user@switch# edit family inet filter Class-B-Customers
```

7. Configure the filter to send packets matching the Class-B-Customer-Prefixes prefix list to the Class-B policer:

```
[edit firewall family inet filter Class-B-Customers]
user@switch# set term term-1 from source-prefix-list Class-B-Customers
user@switch# set term term-1 then policer Class-B
```

8. Create a filter for class C customers:

```
[edit firewall]
user@switch# edit family inet filter Class-C-Customers
```

9. Configure the filter to send packets matching the Class-C-Customer-Prefixes prefix list to the Class-C policer:

```
[edit firewall family inet filter Class-C-Customers]
user@switch# set term term-1 from source-prefix-list Class-C-Customers
user@switch# set term term-1 then policer Class-C
```

10. Apply the filters you created to the appropriate interfaces in the output direction.



**NOTE:** Note that the implicit deny statement in this filter will block traffic from any source that does not match one of the prefix lists. If you want the filter to allow this traffic, you must include an explicit term for this purpose.

#### Related Documentation

- [Overview of Policers on page 5241](#)
- [Applying Firewall Filters to Interfaces on page 5240](#)

- *prefix-list*

## Example: Using Policers to Manage Oversubscription

You might want to use a policer when an interface is oversubscribed and you want to control what will happen if congestion occurs. For example, you might have servers connected to a switch as listed in [Table 438 on page 5286](#).

**Table 438: Servers Connected to Switch**

| Server Type                | Connection           | IP Address |
|----------------------------|----------------------|------------|
| Network application server | 1-gigabit interface  | 10.0.0.1   |
| Authentication server      | 1-gigabit interface  | 10.0.0.2   |
| Database server            | 10-gigabit interface | 10.0.0.3   |

In this example, users access services provided by the network application server, which requests information from the database server as appropriate. When it receives a request from a user, the network application server first contacts the authentication server to verify the user's credentials. When a user is authenticated and the network application server provides the requested service, all the packets sent from the database server to the application server must transit the 1-Gigabit Ethernet interface connected to the application server twice—once on ingress to the application server and again on egress to the user.

The sequence of events for a user session is as follows:

1. A user connects to the application server and requests a service.
2. The application server requests the user's credentials and relays them to the authentication server.
3. If the authentication server verifies the credentials, the application server initiates the requested service.
4. The application server requests the files necessary to meet the user's request from the database server.
5. The database server sends the requested files to the application server.
6. The application server includes the requested files in its response to the user.

Traffic from the database server to the application server might congest the 1-gigabit interface to which that the application server is connected. This congestion might prevent the server from responding to requests from users and creating new sessions for them. You can use policing to make sure that this does not occur.



To create this firewall configuration, perform the following steps on the database server:

1. Create a policer to drop traffic from the database server to the application server if it exceeds certain limits:

```
[edit firewall]
user@switch# set policer Database-Egress-Policer if-exceeding bandwidth-limit 400
burst-size-limit 500m
user@switch# set policer Database-Egress-Policer then discard
```

2. Create a filter to examine traffic from the database server to the application server:

```
[edit firewall]
user@switch# edit family inet filter Database-Egress-Filter
```

3. Configure the filter to apply the policer to traffic egressing the database server and destined for the application server:

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-1 from destination-address 10.0.0.1
user@switch# set term term-1 then policer Database-Egress-Policer
```

4. If required, configure a term to allow traffic from the database server to other destinations (otherwise the traffic will be dropped by the implicit deny statement):

```
[edit firewall family inet filter Database-Egress-Filter]
user@switch# set term term-2 then accept
```

Note that omitting a **from** statement causes the term to match all packets, which is the desired behavior.

5. Install the egress filter as an output filter on the database server interface that is connected the application server:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet filter output Database-Egress-Filter
```

Here is how the final configuration would appear:

```
firewall {
  policer Database-Egress-Policer {
    if-exceeding {
      bandwidth-limit 400;
      burst-size-limit 500m;
    }
    then discard;
  }
  family inet {
    filter Database-Egress-Filter {
      term term-1 {
        from {
          destination-address {
            10.0.0.1/24;
          }
        }
        then policer Database-Egress-Policer;
      }
      term term-2 { # If required, include this term so that traffic from the database server
        to other destinations is allowed.
        then accept;
      }
    }
  }
}
```

```
}  
]
```

**Related  
Documentation**

- [Overview of Policers on page 5241](#)

---

## Device Security Configuration Example

---

- [Example: Configuring Storm Control to Prevent Network Outages on page 5288](#)

### Example: Configuring Storm Control to Prevent Network Outages

Storm control enables you to prevent network outages caused by broadcast storms on the LAN. You can configure storm control on to rate-limit broadcast traffic, multicast traffic, and unknown unicast traffic at a specified level and to have packets dropped when the specified traffic level is exceeded, thereby preventing packets from proliferating and degrading the LAN.



**NOTE:** This example uses a Junos OS release that supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Storm Control to Prevent Network Outages*.

- 
- [Requirements on page 5288](#)
  - [Overview and Topology on page 5288](#)
  - [Configuration on page 5289](#)

---

#### Requirements

---

This example uses the following hardware and software components:

- One QFX Series switch running Junos OS with ELS
- Junos OS Release 13.2 or later

---

#### Overview and Topology

---

A storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect and resulting in a broadcast storm that can cause network outages.

You can use storm control to prevent broadcast storms by specifying the amount, also known as the *storm control level*, of broadcast traffic, multicast traffic, and unknown unicast traffic to be allowed on an interface. You specify the storm control level as the traffic rate in kilobits per second (Kbps) of the combined applicable traffic streams or as the percentage of available bandwidth used by the combined applicable traffic streams. On ELS systems, storm control is enabled by default on all interfaces at a level of 80 percent of the available bandwidth.

Storm control monitors the level of applicable incoming traffic and compares it with the level that you specify. If the combined level of the applicable traffic exceeds the specified level, the switch drops packets for the controlled traffic types. As an alternative to having the switch drop packets, you can configure storm control to shut down interfaces or temporarily disable interfaces (see the **action-shutdown** statement or the [recovery-timeout](#) statement) when the storm control level is exceeded.

The topology used in this example consists of one switch connected to various network devices. This example shows how to configure the storm control level on interface xe-0/0/0 by setting the level to a traffic rate of 15,000 Kbps, based on the traffic rate of the combined applicable traffic streams. If the combined traffic exceeds this level, the switch drops packets for the controlled traffic types to prevent a network outage.

### Configuration

#### CLI Quick Configuration

To quickly configure storm control based on the traffic rate in kilobits per second of the combined traffic streams, copy the following command and paste it into the switch terminal window:

```
[edit]
set forwarding-options storm-control-profiles sc-profile all bandwidth-level 15000
set interfaces xe-0/0/0 unit 0 family ethernet-switching storm-control sc-profile
```

#### Step-by-Step Procedure

To configure storm control:

1. Configure a storm control profile, **sc-profile**, and specify the traffic rate in kilobits per second of the combined traffic streams:

```
[edit]
user@switch> set forwarding-options storm-control-profiles sc-profile all bandwidth-level 15000
```

2. Bind the storm control profile, **sc**, to a logical interface:

```
[edit]
user@switch> set interfaces xe-0/0/0 unit 0 family ethernet-switching storm-control sc-profile
```

**Results** Display the results of the configuration:

```
[edit forwarding-options]
user@switch> show storm-control-profiles sc-profile
all {
  bandwidth 15000;
}

[edit]
user@switch> show interfaces xe-0/0/0
unit 0 {
  family ethernet-switching {
    vlan {
      members default;
    }
    storm-control sc-profile;
  }
}
```

- Related Documentation**
- [Understanding Storm Control on page 5272](#)
  - [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\)](#)

---

## Firewall and Policer Configuration Tasks

---

- [Configuring Firewall Filters on page 5290](#)
- [Applying Firewall Filters to Interfaces on page 5293](#)
- [Assigning Forwarding Classes and Loss Priority on page 5294](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)
- [Configuring MPLS Firewall Filters and Policers on page 5299](#)
- [Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch on page 5301](#)

### Configuring Firewall Filters

You can configure firewall filters in a switch to control traffic that enters switch ports or enters and exits VLANs and Layer 3 (routed) interfaces. To use a firewall filter, you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

- [Configuring a Firewall Filter on page 5290](#)
- [Applying a Firewall Filter to a Port on page 5292](#)
- [Applying a Firewall Filter to a VLAN on page 5292](#)
- [Applying a Firewall Filter to a Layer 3 \(Routed\) Interface on page 5293](#)

---

#### Configuring a Firewall Filter

---

To configure a firewall filter:

1. Configure the family address type, filter name, term name, and at least one match condition—for example, match on packets that contain a specific source address:

```
[edit]
user@switch# set firewall family ethernet-switching filter ingress-port-filter term term-one
from source-address 192.0.2.14
```

For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching**. For a firewall filter that is applied to a Layer 3 (routed) interface, specify the family address type **inet**.

The filter and term names can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. Each filter name must be unique. A filter can contain one or more terms, and each term name must be unique within a filter.

2. Configure additional match conditions. For example, match on packets that contain a specific source port:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one from]
user@switch# set source-port 80
```

You can specify one or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term. The **from** statement is optional, but if included in a term, it cannot be empty. If you omit the **from** statement, all packets are considered to match.

3. If you want to apply a firewall filter to multiple interfaces and be able to see counters specific to each interface, configure the **interface-specific** option:

```
[edit firewall family ethernet-switching filter ingress-port-filter]
user@switch# set interface-specific
```

4. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term. You can specify an action and action modifiers:
  - To specify a filter action, for example, to discard packets that match the conditions of the filter term:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one then]
user@switch# set discard
```

You can specify no more than one action (**accept**, **discard**, **reject**, **routing-instance**, or **vlan**) per term.

- To specify action modifiers, for example, to count and classify packets to a forwarding class:

```
[edit firewall family ethernet-switching filter ingress-port-filter term
term-one then]
user@switch# set count counter-one
user@switch# set forwarding-class expedited-forwarding
user@switch# set loss-priority high
```

You can specify any of the following action modifiers in a **then** statement:

- **analyzer analyzer-name**—Mirror port traffic to a specified analyzer, which you must configure at the **[ethernet-switching-options]** level.
- **count counter-name**—Count the number of packets that pass this filter term.



**NOTE:** We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.



**NOTE:** On QFX3500 and QFX3600 switches, filters automatically count packets that have been dropped on ingress because of cyclic redundancy check (CRC) errors.

- **forwarding-class class**—Assign packets to a forwarding class.
- **log**—Log the packet header information in the Routing Engine.
- **loss-priority priority**—Set the priority of dropping a packet.
- **policer policer-name**—Apply rate-limiting to the traffic.
- **syslog**—Log an alert for this packet.

If you omit the **then** statement or do not specify an action, packets that match all the conditions in the **from** statement are accepted. However, you should always explicitly configure an action in the **then** statement. You can include no more than one action statement, but you can use any combination of action modifiers. For an action or action modifier to take effect, all conditions in the **from** statement must match.



**NOTE:** Implicit discard is also applicable to a firewall filter applied to the loopback interface, lo0.

---

### Applying a Firewall Filter to a Port

To apply a firewall filter to an ingress port:

1. Provide a meaningful description of the firewall filter in the configuration of the port to which the filter will be applied:

```
[edit]
user@switch# set interfaces ge-0/0/6 description "filter to limit tcp traffic at trunk port for employee-vlan"
```

2. Apply the filter to the interface, specifying the unit number, family address type, the direction of the filter (for packets entering the port), and the filter name:

```
[edit]
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input ingress-port-filter
```

For firewall filters that are applied to ports, the family address type must be **ethernet-switching**.



**NOTE:** You can apply only one filter to a port for a given direction (ingress or egress).

---

### Applying a Firewall Filter to a VLAN

To apply a firewall filter to a VLAN:

1. Provide a meaningful description of the firewall filter in the configuration of the VLAN to which the filter will be applied:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 description "filter to block rogue devices on employee-vlan"
```

2. Apply firewall filters to filter packets that are entering or exiting the VLAN:

- To apply a filter to match packets that are entering the VLAN:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 filter input ingress-vlan-rogue-block
```

- To apply a firewall filter to match packets that are exiting the VLAN:

```
[edit]
user@switch# set vlans employee-vlan vlan-id 20 filter output egress-vlan-filter
```



**NOTE:** You can apply only one filter to a VLAN for a given direction (ingress or egress).

### Applying a Firewall Filter to a Layer 3 (Routed) Interface

To apply a firewall filter to a Layer 3 routed interface:

1. Provide a meaningful description of the firewall filter in the configuration of the interface to which the filter will be applied:

```
[edit]
user@switch# set interfaces ge-0/1/6 description "filter to count and monitor traffic on layer
3 interface"
```

2. You can apply firewall filters to filter packets that enter or exit a Layer 3 routed interface:

- To apply a firewall filter to filter packets that enter a Layer 3 interface:

```
[edit]
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter input ingress-router-filter
```

- To apply a firewall filter to filter packets that exit a Layer 3 interface:

```
[edit]
user@switch# set interfaces ge-0/1/6 unit 0 family inet filter output egress-router-filter
```

For firewall filters applied to Layer 3 routed interfaces, the family address type must be **inet**.



**NOTE:** You can apply only one filter to an interface for a given direction (ingress or egress).

#### Related Documentation

- [Overview of Firewall Filters on page 5209](#)
- [Firewall Filter Match Conditions and Actions on page 5219](#)
- [Verifying That Firewall Filters Are Operational on page 5382](#)
- [Monitoring Firewall Filter Traffic on page 5379](#)
- [Configuring Port Mirroring](#)

### Applying Firewall Filters to Interfaces

For a firewall filter to work, you must apply it to at least one interface. To do this, include the **filter** statement when configuring a logical interface at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family family-name filter (input |
output) filter-name
```

In the **input** statement, specify a firewall filter to be evaluated when packets are received on the interface. Input filters applied to a loopback interface affect only traffic destined for the Routing Engine.

In the **output** statement, specify a filter to be evaluated when packets exit the interface.



**NOTE:** When you create a loopback interface, it is important to apply an ingress filter to it so the Routing Engine is protected. We recommend that when you apply a filter to the loopback interface `lo0`, you include the `apply-groups` statement. Doing so ensures that the filter is automatically inherited on every loopback interface, including `lo0` and other loopback interfaces.

**Related Documentation**

- [Configuring Firewall Filters on page 5290](#)

## Assigning Forwarding Classes and Loss Priority

You can configure firewall filters to assign packet loss priority (PLP) and forwarding classes so that if congestion occurs, the marked packets can be dropped according to the priority you set. The valid match conditions are one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. In other words, you can set the forwarding class and the PLP for each packet entering or an interface with a specific destination address, source address, IP protocol, source port, destination port, or DSCP.



**NOTE:** Junos OS assigns forwarding classes and PLP on ingress only. Do not use a filter that assigns forwarding classes or PLP as an egress filter.

When tricolor marking is enabled, a switch supports four PLP designations: **low**, **medium-low**, **medium-high**, and **high**. You can also specify any of the forwarding classes listed in [Table 439 on page 5294](#)

**Table 439: Unicast Forwarding Classes**

| Unicast Forwarding Class | For CoS Traffic Type   |
|--------------------------|--|
| <b>be</b>                | Best-effort traffic  |
| <b>no-loss</b>           | Guaranteed delivery for TCP traffic                                |
| <b>fcoe</b>              | Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic |
| <b>nc</b>                | Network-control traffic  |

To assign forwarding classes in firewall filters:

1. Configure the family address type and filter name:



```
[edit]
```

```
user@switch# edit firewall family ethernet-switching filter ingress-filter
```

2. Configure the terms of the filter as appropriate, including the **forwarding-class** and **loss-priority** action modifiers. For example, each of the following terms in the filter examines various packet header fields and assigns the appropriate forwarding class and packet loss priority:

- The term **corp-traffic** matches all IPv4 packets with a 10.1.1.0/24 source address and assigns the packets to forwarding class **no-loss** with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term corp-traffic from source-address 10.1.1.0/24;
user@switch# set term corp-traffic then forwarding-class no-loss
user@switch# set term corp-traffic then loss-priority low
```

- The term **data-traffic** matches all IPv4 packets with a 10.1.2.0/24 source address and assigns the packets to forwarding class **be** (best effort) with a loss priority of **medium-high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from source-address 10.1.2.0/24;
user@switch# set term data-traffic then forwarding-class be
user@switch# set term data-traffic then loss-priority medium-high
```

- Because the loss of network-generated packets can jeopardize proper network operation, the delay of these packets is preferable to discarding these packets. The term **network-traffic** assigns the packets with an IP precedence of **net-control** to forwarding class **nc** (network control) with a loss priority of **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term network-traffic from precedence net-control
user@switch# set term network-traffic then forwarding-class nc
user@switch# set term network-traffic then loss-priority low
```

- The last term **accept-traffic** matches any packets that did not match on any of the preceding terms and assigns the packets to forwarding class **be** with a loss priority of **high**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic then forwarding-class be
user@switch# set term accept-traffic then loss-priority high
```

3. Apply the filter **ingress-filter** to a port, VLAN, or Layer 3 interface. For information about applying the filter, see [“Configuring Firewall Filters” on page 5290](#). (Assigning forwarding classes and PLP is supported only on ingress filters.)

#### Related Documentation

- [Configuring Firewall Filters on page 5290](#)
- [Verifying That Firewall Filters Are Operational on page 5382](#)
- [Monitoring Firewall Filter Traffic on page 5379](#)
- [Overview of Policers on page 5241](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS Forwarding Classes on page 5830](#)

## Configuring Color-Blind Egress Policers for Medium-Low PLP

If you use color-blind mode and want to configure an egress policer that marks packets to have medium-low PLP, you must configure a single-rate two-color policer at the **[edit firewall policer *policer-name*]** hierarchy level, because color-blind mode does not support medium-low priority. For example:

1. Specify the name of the policer, the bandwidth limit in bits per second (bps) to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit]
user@switch# set firewall policer policer-name if-exceeding bandwidth-limit bytes
burst-size-limit bytes
```

2. Specify medium-low loss priority for matching packets:

```
[edit]
user@switch# set firewall policer policer-name then loss-priority medium-low;
```

3. Apply the filter to a port, VLAN, or Layer 3 interface.

### Related Documentation

- [Overview of Policers on page 5241](#)
- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 5249](#)
- [Configuring Firewall Filters on page 5290](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)

## Configuring Two-Color and Three-Color Policers to Control Traffic Rates

You can rate-limit traffic by configuring a policer and specifying it as an action modifier for a term in a firewall filter. By default, if you specify the same policer in multiple terms, Junos OS creates a separate policer instance for each term and applies rate limiting separately for each instance. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, each policer instance enforces a 1-Gbps limit. In this case, the total bandwidth allowed by the filter is 3 Gbps.

You can also configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps.



**NOTE:** You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

---

1. [Configuring Two-Color Policers on page 5297](#)
2. [Configuring Three-Color Policers on page 5297](#)

3. [Specifying Policers in a Firewall Filter Configuration on page 5298](#)
4. [Applying a Firewall Filter That Includes a Policer on page 5298](#)

### Configuring Two-Color Policers

To configure a two-color policer:

1. Specify the name of the policer, the bandwidth limit to control the traffic rate on an interface, and the maximum allowed burst size to control the amount of traffic bursting:

```
[edit firewall]
user@switch# set policer policer-name <filter-specific> if-exceeding bandwidth-limit bps
burst-size-limit bytes
```

The policer name can contain letters, numbers, and hyphens (-) and can have as many as 64 characters.

The range for the bandwidth limit is 32000 (32k) through 102,300,000,000 (102300m) bps.

To determine the value for the burst-size limit, multiply the bandwidth of the interface on which the filter is applied by the amount of time to allow a burst of traffic at that bandwidth to occur and divide the result by 8:

**maximum burst size = (interface bandwidth) X (allowable time for burst) / (8 bits/byte)**

The range for the burst-size limit is 1 through 2,147,450,880 bytes.

2. Specify the policer action to discard or assign a loss priority to packets that exceed the rate limits:

```
[edit firewall policer policer-name]
user@switch# set then (discard | loss-priority low | loss-priority high)
```

### Configuring Three-Color Policers

To configure a three-color policer:

1. Specify the name of the policer and (optionally) whether to automatically discard packets with high loss priority (PLP):

```
[edit firewall]
user@switch# set three-color-policer policer-name
user@switch# set three-color-policer policer-name action loss-priority high then discard
```

2. Specify whether the three-color policer should be single-rate or two-rate and whether it should be color-aware or color-blind:

```
[edit firewall three-color-policer policer-name]
user@switch# set (single-rate | two-rate) (color-aware | color-blind)
```

3. For single-rate three-color policers, configure the CIR, CBS, and EBS:

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
user@switch# set excess-burst-size bytes
```

4. For two-rate three-color policers, configure the CIR, CBS, PIR, and PBS:

```
[edit firewall three-color-policer policer-name single-rate]
user@switch# set committed-information-rate bps
user@switch# set committed-burst-size bytes
```

```
user@switch# set peak-information-rate bps
user@switch# set peak-burst-size bytes
```

### Specifying Policers in a Firewall Filter Configuration

---

To use a two-color policer, configure a filter term that includes the action **policer**:

```
[edit firewall family family-name]
user@switch# set filter filter-name term name then name
```

For example, the following commands apply a two-color policer to all packets sent from 192.0.2.0/24.

```
[edit firewall family family-name]
user@switch# set filter limit—hosts term term1 from source-address 192.0.2.0/24
user@switch# set filter limit—hosts term term1 then policer policer1
```

To use a three-color policer, configure a filter term that includes the action **three-color-policer**:

```
[edit firewall family name]
user@switch# set filter name term name from match-condition
user@switch# set filter name term name then three-color-policer (single-rate | two-rate) name
```

For example, the following commands apply a single-rate three-color policer to all packets received or sent by interface **ge-0/0/6** (depending on whether the filter is an ingress or egress filter).

```
[edit firewall family name]
user@switch# set filter srTCM term term-one from interface ge-0/0/6
user@switch# set filter srTCM term term-one then three-color-policer single-rate srTCM1-ca
```

You must specify whether the three-color policer is single-rate or two-rate, and this must match the policer itself. Otherwise, the configuration listing includes an error message indicating that the three-color policer you referenced in the filter does not exist.

### Applying a Firewall Filter That Includes a Policer

---

A firewall filter that includes one or more policer action modifiers must be applied to a port, VLAN, or Layer 3 interface like any other filter. For information about applying firewall filters, see [“Configuring Firewall Filters” on page 5290](#).



**NOTE:** You can include two-color policer actions on ingress firewall filters only. You can include three-color policer actions on ingress and egress filters.

---

#### Related Documentation

- [Configuring Firewall Filters on page 5290](#)
- [Overview of Policers on page 5241](#)
- [Verifying That Two-Color Policers Are Operational on page 5393](#)
- [Verifying That Three-Color Policers Are Operational on page 5392](#)
- [Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296](#)

## Configuring MPLS Firewall Filters and Policers

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can also configure policers for MPLS LSPs.

The following sections discuss MPLS firewall filters and policers:

- [Configuring MPLS Firewall Filters on page 5299](#)
- [Examples: Configuring MPLS Firewall Filters on page 5300](#)
- [Configuring Policers for LSPs on page 5300](#)

### Configuring MPLS Firewall Filters

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can then apply this filter to a specific interface on input or output. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached. You cannot apply MPLS firewall filters to loopback interfaces.

You can configure the following match conditions for MPLS filters at the **[edit firewall family mpls filter *filter-name* term *term-name* from]** hierarchy level:

- **exp**
- **label**

These **exp** match condition can accept EXP bits in the range 0 through 7. You can configure the following choices:

- A single EXP bit—for example, **exp 3**;
- Several EXP bits—for example, **exp 0, 4**;
- A range of EXP bits—for example, **exp [0-5]**;

The **label** match condition can accept a range of values from 0 to 1048575.

If you do not specify a match criterion (that is, you do not configure the **from** statement and use only the **then** statement with the **count** action keyword), all the MPLS packets passing through the interface on which the filter is applied will be counted.

You also can configure any of the following action keywords at the **[edit firewall family mpls filter *filter-name* term *term-name* then]** hierarchy level:

- **accept**
- **count**
- **discard**
- **policer**
- **three-color-policer**

### Examples: Configuring MPLS Firewall Filters

---

The following examples illustrate how you might configure an MPLS firewall filter and then apply the filter to an interface. This filter is configured to count MPLS packets with EXP bits set to either 0 or 4.

The following shows a configuration for an MPLS firewall filter:

```
[edit firewall]
family mpls {
  filter expf {
    term expt0 {
      from {
        exp 0,4;
      }
      then {
        count counter0;
        accept;
      }
    }
  }
}
```

### Configuring Policers for LSPs

---

MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

#### ***LSP Policer Limitations***

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- LSP policers are not supported on aggregated interfaces.

- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.

**Related  
Documentation**

- [Overview of Policers on page 524](#)

## Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through a network by encapsulating (or tunneling) the packets. GRE tunneling is performed by tunnel endpoints that encapsulate or de-encapsulate traffic.

You can use a firewall filter to de-encapsulate GRE traffic on a QFX5100 switch. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term.



**NOTE:** QFX5100 switches support as many as 512 GRE tunnels, including tunnels created with a firewall filter. That is, you can create a total of 512 GRE tunnels, regardless of which method you use.

This topic describes:

1. [Configuring a Filter to De-encapsulate GRE Traffic on page 5301](#)
2. [Applying the Filter to an Interface on page 5302](#)

### Configuring a Filter to De-encapsulate GRE Traffic

To configure a firewall filter to de-encapsulate GRE traffic on a QFX5100 switch:

1. Create an IPv4 firewall filter and (optionally) specify a source address for the tunnel:

[edit]

```
user@switch# set firewall family inet filter filter-name term term-name from
source-address address
```

You must create an IPv4 filter by using **family inet** because the outer header of a GRE packet must be IPv4. If you specify a source address, it should be an address on a device that will encapsulate traffic into GRE packets.



**NOTE:** To terminate many tunnels from multiple source IP addresses with one firewall term, do not configure a source address. In this case, the filter will de-encapsulate any GRE packets received by the interface that you apply the filter to.

2. Specify a destination address for the tunnel:

[edit]

```
user@switch# set firewall family inet filter filter-name term term-name from
destination-address address
```

This should be an address on an interface of the QFX5100 switch on which you want the tunnel or tunnels to terminate and the GRE packets to be de-encapsulated. You should also configure this address as a tunnel endpoint on all the tunnel source routers that you want to form tunnels with the QFX5100 switch.

3. Specify that the filter should match and accept GRE traffic:

```
[edit ]
user@switch# set firewall family inet filter filter-name term term-name from protocol
gre
```

4. Specify that the filter should de-encapsulate GRE traffic:

```
[edit ]
user@switch# set firewall family inet filter filter-name term term-name then decapsulate
gre
```

Based on the configuration you have performed so far, the switch forwards the de-encapsulated packets by comparing the inner header to the default routing table (*inet0*). If you want the switch to use a virtual routing instance to forward the de-encapsulated packets, perform the following steps:

5. Specify the name of the virtual routing instance:

```
[edit ]
user@switch# set firewall family inet filter filter-name term term-name then decapsulate
routing-instance instance-name
```

6. Specify that the virtual routing instance is a virtual router:

```
[edit ]
user@switch# set routing-instances instance-name instance-type virtual-router
```

7. Specify the interfaces that belong to the virtual router:

```
[edit ]
user@switch# set routing-instances instance-name interface interface-name
```

---

### Applying the Filter to an Interface

After you create the firewall filter, you must also apply it to an interface that will receive GRE traffic. Be sure to apply it in the input direction. For example, enter

```
[edit ]
user@switch# set interfaces interface-name unit logical-unit-number family inet filter
input filter-name
```

Because the outer header of a GRE packet must be IPv4, you must apply the filter to an IPv4 interface and specify **family inet**.

#### Related Documentation

- [Understanding Generic Routing Encapsulation on page 2449](#)
- [Configuring Generic Routing Encapsulation Tunneling on page 2600](#)
- [Configuring Firewall Filters on page 5290](#)



## Device Security Configuration Tasks

- [Configuring Unicast RPF \(CLI Procedure\) on page 5303](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 5304](#)
- [Configuring Unknown Unicast Forwarding \(CLI Procedure\) on page 5305](#)

### Configuring Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Enabling unicast RPF on the switch interfaces filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. When a packet comes into an interface, if that interface is not the best return path to the source, the switch discards the packet. If the incoming interface is the best return path to the source, the switch forwards the packet.



**NOTE:** On EX3200, EX4200, and EX4300 switches, you can enable unicast RPF only globally—that is, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.

Before you begin:

- On an EX8200, EX6200, or QFX Series switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.
- On an EX3200, EX4200, or EX4300 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

[edit interfaces]

```
user@switch# set ge-1/0/10 unit 0 family inet rpf-check
```



**BEST PRACTICE:** On EX3200, EX4200, and EX4300 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On EX3200, EX4200, and EX4300 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback of this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

---

**Related  
Documentation**

- *Example: Configuring Unicast RPF on an EX Series Switch*
- [Verifying Unicast RPF Status on page 5389](#)
- [Disabling Unicast RPF \(CLI Procedure\) on page 5304](#)
- *Troubleshooting Unicast RPF*
- [Understanding Unicast RPF on page 5273](#)

## Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on an EX3200, EX4200, or EX4300 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the **warning: statement not found** message appears. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the EX3200, EX4200, or EX4300 switch.

On EX8200, EX6200, and QFX Series switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

[edit interfaces]

user@switch# **delete** ge-1/0/10 unit 0 family inet **rpf-check**



**NOTE:** On EX3200, EX4200, and EX4300 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

#### Related Documentation

- [Example: Configuring Unicast RPF on an EX Series Switch](#)
- [Verifying Unicast RPF Status on page 5389](#)
- [Configuring Unicast RPF \(CLI Procedure\) on page 5303](#)
- [Understanding Unicast RPF on page 5273](#)

## Configuring Unknown Unicast Forwarding (CLI Procedure)



**NOTE:** This task uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see [Configuring Unknown Unicast Forwarding \(CLI Procedure\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#)

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets to all interfaces associated with a VLAN. Forwarding such traffic to interfaces on the switch can create a security issue.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN out to a specific interface. You can configure each VLAN to divert unknown unicast traffic to different interfaces or use one interface for multiple VLANs.

To configure unknown unicast forwarding options:

- Configure unknown unicast forwarding for a specific VLAN (here, the VLAN name is employee), and specify the interface to which all unknown unicast traffic will be forwarded:

[edit switch-options]

user@switch# **set** unknown-unicast-forwarding vlan *vlan-name* interface *ge-x/y/z.0*

#### Related Documentation

- [Verifying That Unknown Unicast Packets Are Forwarded to a Single Interface](#)
- [Understanding Unknown Unicast Forwarding on page 5277](#)

## Configuration Statements for Firewall Filters

---

- [family](#) on page 5307
- [filter](#) on page 5308
- [filter \(Layer 2 and Layer 3 Interfaces\)](#) on page 5309
- [filter \(VLANs\)](#) on page 5310
- [firewall](#) on page 5311
- [from](#) on page 5312
- [interface-specific](#) on page 5313
- [term](#) on page 5313
- [then \(Filters\)](#) on page 5314

## family

**Syntax**

```
family family-name {
  filter filter-name {
    interface-specific;
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
}
```

**Hierarchy Level** [edit [firewall](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the fields a firewall filter can match on.

**Options** *family-name*—Type of addressing protocol:

- **ethernet-switching**—Filter Layer 2 Ethernet packets and Layer 3 (IP) packets (allows some Layer 3 filtering).
- **inet**—Filter Layer 3 IPv4 packets (provides additional Layer 3 filter options).
- **inet6**—Filter Layer 3 IPv6 packets (provides additional Layer 3 filter options).
- **mpls**—Filter multiprotocol label switched packets.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Firewall Filter Match Conditions and Actions on page 5219](#)
- [Configuring Firewall Filters on page 5290](#)
- [Overview of Firewall Filters on page 5209](#)

## filter

---

**Syntax**    `filter filter-name {  
                  interface-specific;  
                  term term-name {  
                    from {  
                      match-conditions;  
                    }  
                    then {  
                      action;  
                      action-modifiers;  
                    }  
                  }  
                }`

**Hierarchy Level**    [edit `firewall family family-name`]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure firewall filters.

**Options**    *filter-name*—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

The remaining statements are explained separately.

**Required Privilege Level**    firewall—To view this statement in the configuration.  
                                  firewall-control—To add this statement to the configuration.

**Related Documentation**

- [Firewall Filter Match Conditions and Actions on page 5219](#)
- [Configuring Firewall Filters on page 5290](#)
- [Overview of Firewall Filters on page 5209](#)

## filter (Layer 2 and Layer 3 Interfaces)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>filter (input   output) <i>filter-name</i>;</code>  |
| <b>Hierarchy Level</b>          | [ <a href="#">edit interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> <i>family-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Apply a firewall filter to traffic transiting a port or Layer 3 interface.  |
| <b>Default</b>                  | All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.   |
| <b>Options</b>                  | <p><b><i>filter-name</i></b>—Name of a firewall filter defined at the [<a href="#">edit firewall family <i>family-name</i> filter</a>] hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to traffic entering the port or Layer 3 interface.</p> <p><b>output</b>—Apply a firewall filter to traffic exiting the port or Layer 3 interface.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</li> <li><a href="#">Configuring Firewall Filters on page 5290</a></li> <li><a href="#">Overview of Firewall Filters on page 5209</a></li> </ul>   |

## filter (VLANs)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>filter (input   output) <i>filter-name</i>;</code>   |
| <b>Hierarchy Level</b>          | <code>[edit vlans <i>vlan-name</i>]</code><br><code>[edit vlans <i>vlan-name</i> forwarding-options]</code>  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Apply a firewall filter to traffic ingressing or egressing a VLAN.   |
| <b>Default</b>                  | All incoming traffic is accepted unmodified to a VLAN, and all outgoing traffic is sent unmodified from a VLAN.  |
| <b>Options</b>                  | <p><b><i>filter-name</i></b>—Name of a firewall filter defined at the <code>[edit firewall family <i>family-name</i> filter]</code> hierarchy level.</p> <p><b>input</b>—Apply a firewall filter to VLAN ingress traffic.</p> <p><b>output</b>—Apply a firewall filter to VLAN egress traffic.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Overview of Firewall Filters on page 5209</a></li></ul>  |



## firewall

```
Syntax  firewall {
        family family-name {
            filter filter-name {
                interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
            }
        }
        policer policer-name {
            filter-specific;
            if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
        three-color-policer policer-name {
            action {
                loss-priority high then discard;
            }
            single-rate {
                (color-aware | color-blind);
                committed-information-rate bps;
                committed-burst-size bytes;
                excess-burst-size bytes;
            }
            two-rate {
                (color-aware | color-blind);
                committed-information-rate bps;
                committed-burst-size bytes;
                peak-information-rate bps;
                peak-burst-size bytes;
            }
        }
    }
```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure firewall filters and policers.

The remaining statements are explained separately.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions and Actions on page 5219</a></li><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Firewall Filters on page 5209</a></li></ul> |

---

## from

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | from {<br><i>match-conditions</i> ;<br>}  |
| <b>Hierarchy Level</b>          | [edit <b>firewall</b> <b>family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> <b>term</b> <i>term-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Match packet fields to values specified in a match condition. If the <b>from</b> statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the <b>then</b> statement are implemented.   |
| <b>Options</b>                  | <b><i>match-conditions</i></b> —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the <b>then</b> statement to be implemented. |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions and Actions on page 5219</a></li><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Understanding Firewall Filter Match Conditions on page 5215</a></li></ul>   |

## interface-specific

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | interface-specific;   |
| <b>Hierarchy Level</b>          | [edit <b>firewall</b> family <i>family-name</i> <b>filter</b> <i>filter-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure separate counters for each interface to which a filter is applied.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions and Actions on page 5219</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5290</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5209</a></li> </ul> |

## term

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>term <i>term-name</i> {     from {         <i>match-conditions</i>;     }     then {         <i>action</i>;         <i>action-modifiers</i>;     } }</pre>  |
| <b>Hierarchy Level</b>          | [edit <b>firewall</b> family <i>family-name</i> <b>filter</b> <i>filter-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Define a firewall filter term.   |
| <b>Options</b>                  | <p><b><i>term-name</i></b>—Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Firewall Filter Match Conditions and Actions on page 5219</a></li> <li>• <a href="#">Configuring Firewall Filters on page 5290</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5209</a></li> </ul>            |

## then (Filters)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>then {<br/>    action;<br/>    action-modifiers;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit <b>firewall family</b> <i>family-name</i> <b>filter</b> <i>filter-name</i> <b>term</b> <i>term-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure a firewall filter action.   |
| <b>Options</b>                  | <p><b>action</b>—Actions to accept, discard, or forward packets that match all conditions specified in a filter term.</p> <p><b>action-modifiers</b>—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.</p>      |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions and Actions on page 5219</a></li><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Understanding Firewall Filter Match Conditions on page 5215</a></li></ul> |

## Configuration Statements for Policers

---

- [action on page 5315](#)
- [bandwidth-limit on page 5315](#)
- [burst-size-limit on page 5316](#)
- [color-aware on page 5317](#)
- [color-blind on page 5318](#)
- [committed-burst-size on page 5319](#)
- [committed-information-rate on page 5320](#)
- [excess-burst-size on page 5321](#)
- [filter-specific on page 5322](#)
- [firewall on page 5323](#)
- [if-exceeding on page 5324](#)
- [loss-priority high then discard \(Three-Color Policer\) on page 5325](#)
- [peak-burst-size on page 5326](#)
- [peak-information-rate on page 5327](#)
- [policer on page 5328](#)

- [single-rate](#) on page 5329
- [then \(Policers\)](#) on page 5330
- [three-color-policer](#) on page 5331
- [two-rate](#) on page 5332

## action

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>action {<br/>    <a href="#">loss-priority high then discard</a>;<br/>}</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall three-color-policer name</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Discard traffic on a logical interface using tricolor marking policing.   |
| <b>Options</b>                  | The statements are explained separately.  |
| <b>Required Privilege Level</b> | <code>firewall</code> —To view this statement in the configuration.<br><code>firewall-control</code> —To add this statement to the configuration. |

## bandwidth-limit

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>bandwidth-limit <i>bps</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall policer policer-name if-exceeding</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Specify the traffic rate in bits per second.   |
| <b>Options</b>                  | <code><i>bps</i></code> —Traffic rate in bits per second. Specify <code><i>bps</i></code> as a decimal value or as a decimal number followed by one of the abbreviation <code><i>k</i></code> (1000), <code><i>m</i></code> (1,000,000), or <code><i>g</i></code> (1,000,000,000).<br><b>Range:</b> 32000 bps (32 Kbps) through 10,000,000,000 bps (10 Gbps) |
| <b>Required Privilege Level</b> | <code>firewall</code> —To view this statement in the configuration.<br><code>firewall-control</code> —To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates</a> on page 5296</li> <li>• <a href="#">Overview of Policers</a> on page 5241</li> </ul>  |

## burst-size-limit

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>burst-size-limit bytes;</code>   |
| <b>Hierarchy Level</b>          | [edit <code>firewall policer policer-name if-exceeding</code> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Specify the maximum allowed burst size to control the amount of traffic bursting.  |
| <b>Options</b>                  | <b>bytes</b> —Decimal value or a decimal number followed by k (thousand), m (million), or g (giga).<br><b>Range:</b> 1 through 2,147,450,880 bytes (2147 MB)   |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul> |

## color-aware

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | color-aware;   |
| <b>Hierarchy Level</b>          | [edit <b>firewall three-color-policer</b> <i>policer-name</i> single-rate],<br>[edit <b>firewall three-color-policer</b> <i>policer-name</i> two-rate]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure the way preclassified packets are metered. In color-aware mode, the switch can assign a higher packet-loss priority, but cannot assign a lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded its committed information rate (CIR). The switch cannot change the PLP to low even if the packet conforms to the configured CIR of the appropriate interface. On the other hand, if an upstream device assigns low PLP to a packet but the packet exceeds the CIR and committed burst size (CBS) of the switch interface, the switch can increase the PLP to medium-high. |
| <b>Default</b>                  | If you omit the <b>color-aware</b> statement, the default behavior is color-aware mode.  |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Overview of Policers on page 5241</a></li> <li>• <a href="#">Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 5247</a></li> <li>• <a href="#">Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 5249</a></li> <li>• <a href="#">color-blind on page 5318</a></li> </ul>   |


## color-blind

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | color-blind;  |
| <b>Hierarchy Level</b>          | [edit <b>firewall three-color-policer</b> <i>policer-name</i> single-rate],<br>[edit <b>firewall three-color-policer</b> <i>policer-name</i> two-rate]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure the way preclassified packets are metered. In color-blind mode, the switch ignores any preclassification of packets and can assign a higher or lower packet loss priority (PLP). For example, suppose an upstream device assigns medium-high PLP to a packet because the packet exceeded the CIR on the upstream device. The switch can change the PLP to low if the packet conforms to the CIR of the appropriate interface.                             |
| <b>Default</b>                  | If you omit the <b>color-blind</b> statement, the default behavior is color-aware mode.   |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Overview of Policers on page 5241</a></li><li>• <a href="#">Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 5247</a></li><li>• <a href="#">Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 5249</a></li><li>• <a href="#">Configuring Color-Blind Egress Policers for Medium-Low PLP on page 5296</a></li><li>• <a href="#">color-aware on page 5317</a></li></ul> |




## committed-burst-size


|  |   |
|--|---|
| <b>Syntax</b>  | <code>committed-burst-size bytes;</code>  |
| <b>Hierarchy Level</b>   | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate],<br>[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>   | Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with low packet loss priority (green).   |
| <div>  <p><b>NOTE:</b> When you include the <code>committed-burst-size</code> statement in the configuration, you must also include the <code>committed-information-rate</code> statement at the same hierarchy level.</p> </div> |   |
| <b>Options</b>   | <p><b>bytes</b>—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 512 bytes through 268435456 bytes (268 MB)</p> |
| <b>Required Privilege Level</b>  | <p><b>firewall</b>—To view this statement in the configuration.</p> <p><b>firewall-control</b>—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>   |

## committed-information-rate

---

|   |   |
|---|---|
| <b>Syntax</b>   | <code>committed-information-rate <i>bits-per-second</i>;</code>   |
| <b>Hierarchy Level</b>  | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> single-rate],<br>[edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]  |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>  | Configure the guaranteed bandwidth under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green).  |
| <div> <b>NOTE:</b> When you include the <code>committed-information-rate</code> statement in the configuration, you must also include the <code>committed-burst-size</code> statement at the same hierarchy level.</div> |   |
| <b>Options</b>  | <b><i>bits-per-second</i></b> —Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br><b>Range:</b> 32,000 bps through 10,000,000,000 bps (10 gbps) |
| <b>Required Privilege Level</b>   | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul>  |

## excess-burst-size

|  |   |
|--|---|
| <b>Syntax</b>  | <code>excess-burst-size bytes;</code>   |
| <b>Hierarchy Level</b>   | [edit <code>firewall three-color-policer policer-name</code> single-rate]   |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>   | Configure the maximum number of bytes allowed for incoming traffic to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red).             |
| <div>  <p><b>NOTE:</b> When you include the <code>excess-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>committed-information-rate</code> statements at the same hierarchy level.</p> </div> |   |
| <b>Options</b>   | <p><b>bytes</b>—Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 512 bytes through 268435456 bytes (268 MB)</p> |
| <b>Required Privilege Level</b>  | <p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>   |

## filter-specific

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | filter-specific;   |
| <b>Hierarchy Level</b>          | [edit <b>firewall</b> <b>policer</b> <i>policer-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure a policer to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. If you use a filter-specific policer in multiple terms, both of the following are true:</p> <ul style="list-style-type: none"><li>• Traffic is policed at the aggregate rate. For example, if you create a policer that has a bandwidth limit of 100 Mbps and use the policer in two terms, the total allowed bandwidth for both terms is 100 Mbps—not 100 Mbps for each term.</li><li>• The implicit counter counts all the packets are that matched by any of the terms. For example, if you reference the same filter-specific policer in term1 and term2, and term1 matches 1000 packets and term2 matches 500 packets, the implicit counter shows 1500 matches for the policer.</li></ul> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul>   |

## firewall

```

Syntax  firewall {
        family family-name {
            filter filter-name {
                interface-specific;
                term term-name {
                    from {
                        match-conditions;
                    }
                    then {
                        action;
                        action-modifiers;
                    }
                }
            }
        }
        policer policer-name {
            filter-specific;
            if-exceeding {
                bandwidth-limit bps;
                burst-size-limit bytes;
            }
            then {
                policer-action;
            }
        }
        three-color-policer policer-name {
            action {
                loss-priority high then discard;
            }
            single-rate {
                (color-aware | color-blind);
                committed-information-rate bps;
                committed-burst-size bytes;
                excess-burst-size bytes;
            }
            two-rate {
                (color-aware | color-blind);
                committed-information-rate bps;
                committed-burst-size bytes;
                peak-information-rate bps;
                peak-burst-size bytes;
            }
        }
    }

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure firewall filters and policers.

The remaining statements are explained separately.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Firewall Filter Match Conditions and Actions on page 5219</a></li><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Firewall Filters on page 5209</a></li></ul> |

---

## if-exceeding

---


|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>if-exceeding {<br/>    bandwidth-limit <i>bps</i>;<br/>    burst-size-limit <i>bytes</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall policer</a> <i>policer-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Configure policer rate limits.<br><br>The remaining statements are explained separately.   |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul> |

## loss-priority high then discard (Three-Color Policer)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | loss-priority high then discard;   |
| <b>Hierarchy Level</b>          | [edit firewall <b>three-color-policer</b> <i>policer-name</i> <b>action</b> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>For packets with high loss priority, discard the packets. The loss priority setting is not configurable. Include this statement if you do not want the switch to forward packets that have high packet-loss priority.</p> <p>For single-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the committed information rate and the excess burst size.</p> <p>For two-rate three-color policers, Junos OS assigns high loss priority to packets that exceed the peak information rate and the peak burst size.</p> |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>  |


## peak-burst-size

---

|  |  |
|--|--|
| <b>Syntax</b>  | <code>peak-burst-size bytes;</code>  |
| <b>Hierarchy Level</b>   | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> two-rate]   |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>   | Configure the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red).           |
| <div> <b>NOTE:</b> When you include the <code>peak-burst-size</code> statement in the configuration, you must also include the <code>committed-burst-size</code> and <code>peak-information-rate</code> statements at the same hierarchy level.</div> |  |
| <b>Options</b>   | <b>bytes</b> —Number of bytes. You can specify a value in bytes either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).<br><b>Range:</b> 1500 bytes through 100,000,000,000 bytes (100 GB) |
| <b>Required Privilege Level</b>  | <b>firewall</b> —To view this statement in the configuration.<br><b>firewall-control</b> —To add this statement to the configuration.  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul>   |



## peak-information-rate

|  |   |
|--|---|
| <b>Syntax</b>  | <code>peak-information-rate <i>bits-per-second</i>;</code>  |
| <b>Hierarchy Level</b>   | [edit <code>firewall three-color-policer <i>policer-name</i> two-rate</code> ]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>   | Configure the maximum achievable rate. Packets that exceed the committed information rate (CIR) but are below the peak information rate (PIR) are marked with medium-high packet loss priority (yellow). Packets that exceed the PIR are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. |
| <div>  <b>NOTE:</b> When you include the <code>peak-information-rate</code> statement in the configuration, you must also include the <code>committed-information-rate</code> and <code>peak-burst-size</code> statements at the same hierarchy level. </div> |   |
| <b>Options</b>   | <p><b><i>bits-per-second</i></b>—Number of bits per second. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 32,000 bps through 10,000,000,000 bps (10 gbps)</p>               |
| <b>Required Privilege Level</b>  | <p><code>firewall</code>—To view this statement in the configuration.</p> <p><code>firewall-control</code>—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>   |

## policer

---


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>policer <i>policer-name</i> {<br/>    filter-specific;<br/>    if-exceeding {<br/>        bandwidth-limit <i>bps</i>;<br/>        burst-size-limit <i>bytes</i>;<br/>    }<br/>    then {<br/>        <i>policer-action</i>;<br/>    }<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Configure policer rate limits and actions. To activate a policer, you must include the <b>policer</b> action modifier in the <b>then</b> statement in a firewall filter term.</p> <p>Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or across filters—the policer's implicit counter is used to count packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these approaches:</p> <ul style="list-style-type: none"><li>• Configure a unique policer for each term.</li><li>• Configure only one policer, but use a unique, explicit counter in each term.</li></ul> |
| <b>Options</b>                  | <p><b><i>policer-name</i></b>—Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul>  |

## single-rate

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>single-rate {   (color-aware   color-blind);   committed-information-rate <i>bps</i>;   committed-burst-size <i>bytes</i>;   excess-burst-size <i>bytes</i>; }</pre>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>Configure a single-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the EBS are assigned medium-high loss priority (yellow). Packets that exceed the EBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p> |
| <b>Options</b>                  | <b><i>policer-name</i></b> —Name of the three-color policer. Use this name when you apply the policer to an interface.   |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>  |

## then (Policers)

---

|   |   |
|---|---|
| Syntax  | then {<br><i>policer-action</i> ;<br>}  |
| Hierarchy Level   | [edit <b>firewall</b> <b>policer</b> <i>policer-name</i> ]  |
| Release Information   | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| Description   | Configure a policer action.   |
| Options   | <i>policer-action</i> —Allowed policer actions are <b>discard</b> , <b>loss-priority high</b> , and <b>loss-priority low</b> . <b>discard</b> causes the system to drop traffic that exceeds the rate limits defined by the policer. Use <b>loss-priority high</b> to allow the system to forward matching traffic in some cases. |
| <hr/>   |   |
| <div> <b>NOTE:</b> If you specify a policer in an egress firewall filter, the only supported action is <b>discard</b>.</div> <hr/> |   |
| Required Privilege Level  | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.   |
| Related Documentation   | <ul style="list-style-type: none"><li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li><li>• <a href="#">Configuring Firewall Filters on page 5290</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul>  |

## three-color-policer

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> three-color-policer <i>policer-name</i> {   action {     loss-priority high then discard;   }   single-rate {     (color-aware   color-blind);     committed-information-rate <i>bps</i>;     committed-burst-size <i>bytes</i>;     excess-burst-size <i>bytes</i>;   }   two-rate {     (color-aware   color-blind);     committed-information-rate <i>bps</i>;     committed-burst-size <i>bytes</i>;     peak-information-rate <i>bps</i>;     peak-burst-size <i>bytes</i>;   } } </pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall</a> ],<br>[edit logical-systems <i>logical-system-name</i> firewall]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure a three-color policer.  |
| <b>Options</b>                  | <p><b><i>policer-name</i></b>—Name of the three-color policer. Use this name when you apply the policer to an interface.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | firewall—To view this statement in the configuration.<br>firewall-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>   |

## two-rate

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>two-rate {<br/>  (color-aware   color-blind);<br/>  committed-information-rate <i>bps</i>;<br/>  committed-burst-size <i>bytes</i>;<br/>  peak-information-rate <i>bps</i>;<br/>  peak-burst-size <i>bytes</i>;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">firewall three-color-policer</a> <i>policer-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Configure a two-rate three-color policer in which marking is based on the committed information rate (CIR), committed burst size (CBS), peak information rate (PIR), and peak burst size (PBS).</p> <p>Packets that conform to the CIR or the CBS are assigned low loss priority (green). Packets that exceed the CIR and the CBS but are within the PIR or the PBS are assigned medium-high loss priority (yellow). Packets that exceed the PIR and the PBS are assigned high loss priority (red).</p> <p>Green and yellow packets are always forwarded; this action is not configurable. You can configure red packets to be discarded. By default, red packets are forwarded.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>firewall—To view this statement in the configuration.</p> <p>firewall-control—To add this statement to the configuration.</p>  |

## Configuration Statements for Port Security

---

- [circuit-id](#) on page 5333
- [dhcp-snooping-file](#) on page 5334
- [fc-map](#) on page 5335
- [fcoe-trusted](#) on page 5337
- [mac-move-limit](#) on page 5338
- [no-allowed-mac-log](#) on page 5339
- [no-gratuitous-arp-request](#) on page 5340
- [persistent-learning](#) on page 5340
- [port-error-disable](#) on page 5341
- [vendor-id](#) on page 5343
- [write-interval](#) on page 5344

## circuit-id

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> circuit-id {   prefix {     host-name;     logical-system-name;     routing-instance-name;   }   use-interface-description (device   logical);   use-vlan-id; } </pre>   |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS):<br/>[edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security option-82</b> ]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i>) dhcp-option82],<br/>[edit forwarding-options helpers bootp dhcp-option82] ,<br/>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]</li> <li>For MX Series platforms:<br/>[edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security <b>option-82</b>]</li> </ul> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b>] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 14.1 for the MX Series.</p>   |
| <b>Description</b>              | <p>Configure the <b>circuit-id</b> suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (the interface, the VLAN, or both) on which the DHCP request arrived.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Default</b>                  | <p>If DHCP option 82 is enabled on the switch, the circuit ID is supplied by default in the format <i>interface-name:vlan-name</i> or, on a Layer 3 interface, just <i>interface-name</i>.</p>   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> </ul>   |

- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- *Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)*
- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

---

## dhcp-snooping-file

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>dhcp-snooping-file {<br/>    location <i>local_pathname</i>   <i>remote_URL</i>;<br/>    timeout <i>seconds</i>;<br/>    write-interval <i>seconds</i>;<br/>}</pre>            |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options secure-access-port]<br><br>For platforms with ELS:<br><br>[edit system processes] <a href="#">dhcp-service</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Specify a local pathname or remote URL for the DHCP snooping database file to maintain persistence of IP-MAC bindings.<br><br>The remaining statements are explained separately.    |
| <b>Default</b>                  | The IP-MAC bindings in the DHCP snooping database file are not persistent. If the switch is rebooted, the bindings are lost.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li></ul>  |



## fc-map

**Syntax** `fc-map fc-map-value;`

**Hierarchy Level** Original CLI

[edit ethernet-switching options secure-access-port vlan (all | *vlan-name*) examine-fip]

ELS CLI for Platforms that Support FCoE

[edit **v**lans *vlan-name* **f**orwarding-**o**ptions **f**ip-**s**ecurity]



**NOTE:** The `fc-map` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

[edit fc-fabrics *fc-fabric-name* protocols fip]

**Release Information** Statement introduced in Junos OS Release 10.4 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

**Description** Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).

You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (0x0EFC00) than for VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping.

The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN\_Port MAC address.

When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.



**NOTE:** Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

**Options** `fc-map-value`—FC-MAP value, hexadecimal value preceded by “0x”.

**Range:** 0x0EFC00 through 0x0EFCFF


**Default:** 0x0EFC00 for VN2VF\_Port FIP snooping 0x0EFD00 for VN2VN\_Port FIP snooping

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.


**Related Documentation**

- *examine-fip*
- [show fip snooping on page 5746](#)
- *Example: Configuring an FCoE Transit Switch*
- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)

## fcoe-trusted

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | fcoe-trusted;  |
| <b>Hierarchy Level</b>          | Original CLI<br><br>[edit ethernet-switching-options secure-access-port interface <i>interface-name</i> ]<br><br>ELS CLI for Platforms that Support FCoE<br><br>[edit <b>vlangs</b> <i>vlan-name</i> <b>forwarding-options fip-security interface</b> <i>interface-name</i> ]  |
|                                 | <div>  <p><b>NOTE:</b> The <b>fcoe-trusted</b> configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.</p> </div>  |
|                                 | <p>QFX Series that Support FCoE-FC Gateway Configuration</p> <p>[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]</p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p> <p>(QFX Series FCoE-FC gateway) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the <b>fcoe-trusted</b> configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> </ul>  |

## mac-move-limit

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>mac-move-limit <i>limit</i> &lt;fabric-limit <i>limit</i>&gt; action <i>action</i>;</code>  |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options secure-access-port (all   <i>vlan-name</i>)]</pre> <p>For platforms with ELS:</p> <pre>[edit vlans <i>vlan-name</i> switch-options],</pre>  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Specify the number of times a MAC address can move to a new interface (port) in 1 second and the action to be taken by the switch if the MAC address move limit is exceeded.  |
|                                 | <div>  <p><b>CAUTION:</b> Mac move limiting does not work properly on a QFX5100 switch used as a Node device in a QFabric system. Do not use this feature on a QFX5100 switch in a QFabric system.</p> </div>  |
| <b>Default</b>                  | The default move limit is unlimited. The default action is <b>drop</b> .  |
| <b>Options</b>                  | <p><b>fabric-limit</b>—Specify the maximum number of moves in a QFabric system. If you do not specify a fabric limit, the value for <b>mac-move-limit</b> applies to the QFabric system.</p> <p><b>limit</b>—Maximum number of moves to a new interface per second.</p> <p><b>action <i>action</i></b>—(Optional) Action to take when the MAC address move limit is reached:</p> <ul style="list-style-type: none"> <li>• <b>drop</b>—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.</li> <li>• <b>log</b>—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.</li> <li>• <b>none</b>—No action.</li> <li>• <b>shutdown</b>—Disable the interface and generate an alarm. If you have configured the switch with the <b>port-error-disable</b> statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the <b>clear-ethernet-switch-port</b> command.</li> </ul> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>  |

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | <ul style="list-style-type: none"> <li>• <i>mac-limit</i></li> <li>• <i>Example: Configuring Basic Port Security Features</i></li> <li>• <i>Configuring MAC Move Limiting (CLI Procedure)</i></li> <li>• <i>Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)</i></li> </ul> |
|------------------------------|---|

## no-allowed-mac-log

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-allowed-mac-log;  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>• For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port interface (all   <i>interface-name</i>)]</li> <li>• For platforms with ELS:<br/>[edit switch-options interface <i>interface-name</i>]</li> </ul> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Specify that the switch should not log messages when it receives packets from invalid MAC addresses on an interface that has been configured for allowed MAC addresses.  |
| <b>Default</b>                  | The switch logs messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular allowed (specific) MAC addresses.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing—control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding MAC Limiting and MAC Move Limiting for Port Security on page 5266</a></li> <li>• <i>Configuring MAC Limiting</i></li> <li>• <i>allowed-mac</i></li> <li>• <i>mac-limit</i></li> </ul>                     |

## no-gratuitous-arp-request

---


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-gratuitous-arp-request;  |
| <b>Hierarchy Level</b>          | [edit <a href="#">interfaces</a> <i>interface-name</i> ],<br>[edit <a href="#">interfaces</a> <i>interface-range</i> <i>interface-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs). |
| <b>Default</b>                  | Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring IRB Interfaces on page 1675</a></li></ul>   |

## persistent-learning

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | persistent-learning;  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>• For platforms without ELS:<br/>[edit ethernet-switching-options secure-access-port interface (all   <i>interface-name</i>)]</li><li>• For platforms with ELS:<br/>[edit switch-options interface <i>interface-name</i>]</li></ul> |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.4 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.<br>Hierarchy level [edit switch-options interface <i>interface-name</i> ] introduced in Junos OS Release 13.2X50-D10           |
| <b>Description</b>              | Specify that learned MAC addresses persist on the specified interfaces across restarts of the switch and link-down conditions. This feature is also known as sticky MAC.  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Basic Port Security Features</i></li><li>• <i>Configuring Persistent MAC Learning (CLI Procedure)</i></li><li>• <i>Configuring Persistent MAC Learning (CLI Procedure)</i></li></ul>                      |

## port-error-disable

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>port-error-disable {   (disable-timeout <i>seconds</i>   <i>recovery-timeout seconds</i>); }</pre>   |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms without ELS:<br/>[edit ethernet-switching-options]</li> <li>For platforms with ELS:<br/>[edit switch-options ]</li> </ul>  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 on the QFX Series.  |
| <b>Description</b>              | Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and storm control, and allow the interface to recover automatically from the error condition after a specified period of time:   |
|                                 | <div>  <p><b>NOTE:</b> The <b>port-error-disable</b> configuration does not apply to preexisting error conditions. It affects only error conditions that are detected after you enable and commit the <b>port-error-disable</b> statement. To clear a preexisting error condition and restore the interface to service, use the <a href="#">clear ethernet-switching port-error</a> command.</p> </div>  |
|                                 | <ul style="list-style-type: none"> <li>If you enable the <i>mac-limit</i> statement with the <b>shutdown</b> option and also enable the <b>port-error-disable</b> statement, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.</li> <li>If you have enabled the <a href="#">mac-move-limit</a> statement with the <b>shutdown</b> option and you enable the <b>port-error-disable</b> statement, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.</li> <li>If you enable the <i>storm-control</i> statement with the <b>action-shutdown</b> option and you also enable <b>port-error-disable</b>, the switch disables (rather than shuts down) the interface when broadcast traffic and unknown unicast traffic exceed the specified levels.</li> </ul> |
| <b>Default</b>                  | Not enabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing—control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">Understanding MAC Limiting and MAC Move Limiting for Port Security on page 5266</a></li> <li><a href="#">Understanding Storm Control on page 5272</a></li> <li><a href="#">Example: Configuring Storm Control to Prevent Network Outages</a></li> </ul>  |

- *Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)*
- [action-shutdown on page 5365](#)
- *disable-timeout*
- [clear ethernet-switching port-error on page 5397](#)



## vendor-id

|   |  |
|---|--|
| <b>Syntax</b>   | <code>vendor-id &lt;string&gt;;</code>   |
| <b>For Platforms with Enhanced Layer 2 Software (ELS)</b> | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security option-82</b> ]   |
| <b>For Platforms Without ELS</b>                          | [edit ethernet-switching-options secure-access-port vlan (all   <i>vlan-name</i> ) dhcp-option82],<br>[edit forwarding-options helpers bootp dhcp-option82],<br>[edit forwarding-options helpers bootp interface <i>interface-name</i> dhcp-option82]  |
| <b>For MX Series Platforms</b>                            | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security <b>option-82</b> ]   |
| <b>Release Information</b>                                | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 11.3 for the QFX Series.<br>Hierarchy level [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> ] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)<br>Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options <b>dhcp-security</b> ] introduced in Junos OS Release 14.1 for the MX Series.  |
| <b>Description</b>  | Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.   |
| <b>Default</b>  | If <b>vendor-id</b> is not explicitly configured for DHCP option 82, then no vendor ID is set.   |
| <b>Options</b>  | <b>string</b> —(Optional) A single string that designates the vendor ID.<br><br><b>Range:</b> 1–255 characters<br><br><b>Default:</b> If you specify <b>vendor-id</b> with no <b>string</b> value, then the default vendor ID <b>Juniper Networks</b> is configured.   |
| <b>Required Privilege Level</b>                           | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>                              | <ul style="list-style-type: none"> <li><i>Setting Up DHCP Option 82 on an MX Series Router (CLI Procedure)</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch with No Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Example: Setting Up DHCP Option 82 with a Switch as a Relay Agent Between Clients and a DHCP Server</i></li> <li><i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> <li><i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> </ul> |

- *Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)*

## write-interval

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>write-interval <i>seconds</i>;</code>   |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options secure-access-port <a href="#">dhcp-snooping-file</a> ]<br><br>For platforms with ELS:<br><br>[edit system processes] <a href="#">dhcp-service</a> dhcp-snooping-file] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Specify how frequently the switch writes the database entries from memory into the specified DHCP snooping database file.   |
| <b>Default</b>                  | None  |
| <b>Options</b>                  | <i>seconds</i> —Value in seconds.<br><b>Range:</b> 60 through 86400   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li></ul>  |



## Configuration Statements for Port Security

---

- [accept-source-mac on page 5346](#)
- [arp-inspection on page 5348](#)
- [dhcp-security on page 5350](#)
- [dhcp-service on page 5352](#)
- [group \(DHCP Security\) on page 5353](#)
- [interface \(DHCP Security\) on page 5354](#)
- [interface-mac-limit on page 5355](#)
- [no-dhcp-snooping on page 5357](#)
- [no-option-82 on page 5358](#)
- [option-82 on page 5359](#)
- [overrides \(DHCP Security\) on page 5360](#)
- [recovery-timeout on page 5361](#)
- [static-ip on page 5362](#)

- [switch-options on page 5363](#)
- [trusted on page 5364](#)
- [untrusted on page 5364](#)

## accept-source-mac

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> accept-source-mac {     mac-address <i>mac-address</i> {         policer {             input <i>cos-policer-name</i>;             output <i>cos-policer-name</i>;         }     } } </pre>   |
| <b>Hierarchy Level</b>          | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p>  |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>  |
| <b>Description</b>              | <p>For Gigabit Ethernet intelligent queuing (IQ) interfaces only, accept traffic from and to the specified remote media access control (MAC) address.</p> <p>The <b>accept-source-mac</b> statement is equivalent to the <b>source-address-filter</b> statement, which is valid for aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only. To allow the interface to receive packets from specific MAC addresses, include the <b>accept-source-mac</b> statement.</p> <p>On untagged Gigabit Ethernet interfaces, you should not configure the <b>source-address-filter</b> statement and the <b>accept-source-mac</b> statement simultaneously. On tagged Gigabit Ethernet interfaces, you should not configure the <b>source-address-filter</b> statement and the <b>accept-source-mac</b> statement with an identical MAC address specified in both filters.</p> <p>The statements are explained separately.</p> |
|                                 | <p> <b>NOTE:</b> The <b>policer</b> statement is not supported on PTX Series Packet Transport Routers.</p>  |
|                                 | <p> <b>NOTE:</b> On QFX platforms, if you configure source MAC addresses for an interface using the <b>static-mac</b> or <b>persistent-learning</b> statements and later configure a different MAC address for the same interface using the <b>accept-source-mac</b> statement, the MAC addresses that you previously configured for the interface remain in the ethernet-switching table and can still be used to send packets to the interface.</p>   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- *Configuring MAC Address Filtering*
  - *Configuring MAC Address Filtering on PTX Series Packet Transport Routers*
  - *source-filtering*

## arp-inspection

---

**Syntax**    `arp-inspection {  
                  forwarding-class class-name;  
                  }`

- Hierarchy Level**
- For platforms with ELS:  
    `[edit vlans vlan-name forwarding-options dhcp-security],`  
    `[edit forwarding-options dhcp-relay ]`
  - For platforms without ELS:  
    `[edit ethernet-switching-options secure-access-port vlan (all | vlan-name)],`  
    `[edit forwarding-options dhcp-relay ]`

**Release Information**    Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See [“Getting Started with Enhanced Layer 2 Software” on page 43](#) for information about ELS.)  
Statement introduced in Junos OS Release 13.2 for the QFX series.

**Description**    Perform dynamic ARP inspection (DAI) on all VLANs or on the specified VLAN.

When DAI is enabled, the switch logs invalid ARP packets that it receives on each interface, along with the sender’s IP and MAC addresses. ARP probe packets, which have the sender IP address 0.0.0.0, are validated by DAI.



**NOTE:** If you configure DAI at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level:

- DAI can be configured only for a specific VLAN, not for a list or a range of VLAN IDs.
- DHCP snooping is automatically enabled on the specified VLAN.
- The `forwarding-class` statement is not available at the `[edit vlans vlan-name forwarding-options dhcp-security]` hierarchy level.

See *Enabling Dynamic ARP Inspection (CLI Procedure)* for more information about this configuration.

---



**NOTE:** On EX9200 switches, DAI is not supported in an MC-LAG scenario.

---

The remaining statement is explained separately.

**Default**    Disabled.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch</i></li><li>• <i>Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks</i></li><li>• <i>Example: Configuring IP Source Guard and Dynamic ARP Inspection to Protect the Switch from IP Spoofing and ARP Spoofing</i></li><li>• <i>Example: Using CoS Forwarding Classes to Prioritize Snooped Packets in Heavy Network Traffic</i></li><li>• <i>Enabling Dynamic ARP Inspection (CLI Procedure)</i></li><li>• <i>Enabling Dynamic ARP Inspection (J-Web Procedure)</i></li></ul> |

## dhcp-security

---

**Syntax**    dhcp-security {  
              arp-inspection;  
              group *group-name* {  
                  interface *interface-name* {  
                      static-ip *ip-address* {  
                          mac *mac-address*;  
                      }  
                  }  
                  overrides {  
                      no-option82;  
                      trusted;  
                      untrusted;  
                  }  
              }  
              ip-source-guard;  
              neighbor-discovery-inspection;  
              no-dhcp-snooping;  
              option-82 {  
                  circuit-id {  
                      prefix {  
                          host-name;  
                          logical-system-name;  
                          routing-instance-name;  
                      }  
                  use-interface-description (device | logical);  
                  use-vlan-id;  
              }  
              remote-id {  
                  host-name *hostname*;  
                  use-interface-description (device | logical);  
                  mac;  
                  use-string *string*;  
              }  
              vendor-id {  
                  use-string *string*;  
              }  
              }  
              }

**Hierarchy Level**    [edit vlans *vlan-name* forwarding-options]

**Release Information**    Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
Statement introduced in Junos OS Release 13.2 for the QFX series.  
Support for **static-ipv6**, **nd-inspection**, **ipv6-source-guard**, **no-dhcpv6-snooping**, and **no-option-37** introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

**Description**    Configure port security features on the switch. DHCP snooping is enabled automatically if you configure any of the following port security features within this hierarchy:

- Dynamic ARP inspection (DAI)
- IP source guard



- DHCP option 82
- Static IP

For switches that support DHCPv6, both DHCP snooping and DHCPv6 snooping are enabled automatically if you configure any of the features listed above or any of the following IPv6 features:

- IPv6 Neighbor Discovery inspection
- IPv6 source guard
- Static IPv6



**NOTE:** On EX9200 switches, DHCP snooping, DAI, and IP source guard are not supported in an MC-LAG scenario.

The remaining statements are explained separately.

|                                 |   |
|---------------------------------|---|
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Enabling Dynamic ARP Inspection (CLI Procedure)</i></li> <li>• <i>Configuring IP Source Guard (CLI Procedure)</i></li> <li>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> <li>• <i>Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure)</i></li> </ul> |

## dhcp-service

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>dhcp-service {<br/>    dhcp-snooping-file (<i>local_pathname</i>   <i>remote_URL</i>);<br/>    write-interval <i>interval</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit system processes]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series.  |
| <b>Description</b>              | <p>Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can improve performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Performance (CLI Procedure)</i></li></ul>   |

## group (DHCP Security)

```
Syntax  group group-name {
        interface interface-name {
            static-ip ip-address {
                mac mac-address;
            }
            static-ipv6 ip-address {
                mac mac-address;
            }
        }
        overrides {
            no-option37;
            no-option-82;
            trusted;
            untrusted;
        }
    }
```

**Hierarchy Level** [edit vlans *vlan-name* forwarding-options [dhcp-security](#)]

**Release Information** Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.  
Statement introduced in Junos OS Release 13.2 for the QFX series.  
Support for the **static-ipv6** and **no-option37** statements introduced in Junos OS Release 13.2X51-D20 for EX Series switches.

**Description** Specify the name of a group of access interfaces that you want to configure for DHCP security attributes that are different from the attributes set for other interfaces in the VLAN. A group must contain at least one interface.

The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- *Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure)*
- *Enabling a Trusted DHCP Server (CLI Procedure)*
- [Understanding DHCP Snooping for Port Security on page 5256](#)

## interface (DHCP Security)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>interface <i>interface-name</i> {<br/>    <b>static-ip</b> <i>ip-address</i> {<br/>        mac <i>mac-address</i>;<br/>    }<br/>    static-ipv6 <i>ip-address</i> {<br/>        mac <i>mac-address</i>;<br/>    }<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Support for the <b>static-ipv6</b> statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches.   |
| <b>Description</b>              | <p>Configure an interface for a static IPv4 or IPv6 address to MAC address binding (IP-MAC binding) or configure an interface to belong to a group within the VLAN that has DHCP security attributes that are different from the attributes of other interfaces in the VLAN.</p> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure)</i></li><li>• <i>Enabling a Trusted DHCP Server (CLI Procedure)</i></li><li>• <i>Configuring Port Security (CLI Procedure)</i></li></ul>   |

## interface-mac-limit

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | <pre>interface-mac-limit <i>limit</i> {     <b>packet-action</b> drop; }</pre>   |
| <b>Hierarchy Level</b>     | <p>[edit bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit logical-systems <i>logical-system-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options],<br/>         [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>         [edit logical-systems <i>logical-system-name</i> switch-options],<br/>         [edit logical-systems <i>logical-system-name</i> switch-options interface <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options],<br/>         [edit routing-instances <i>routing-instance-name</i> bridge-domains <i>bridge-domain-name</i> bridge-options interface <i>interface-name</i>],<br/>         [edit routing-instances <i>routing-instance-name</i> switch-options],<br/>         [edit routing-instances <i>routing-instance-name</i> switch-options interface <i>interface-name</i>],<br/>         [edit switch-options],<br/>         [edit switch-options interface <i>interface-name</i>],<br/>         [edit switch-options interface <i>interface-name</i>],<br/>         [edit vlans <i>vlan-name</i> switch-options],<br/>         [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>]</p> |
| <b>Release Information</b> | <p>Statement introduced in Junos OS Release 8.4.</p> <p>Support for the <b>switch-options</b> statement added in Junos OS Release 9.2.</p> <p>Support for top-level configuration for the <b>virtual-switch</b> type of routing instance added in Junos OS Release 9.2. In Junos OS Release 9.1 and earlier, the routing instances hierarchy supported this statement only for a VPLS instance or a bridge domain configured within a virtual switch.</p> <p>Support for logical systems added in Junos OS Release 9.6.</p> <p>[edit switch-options], [edit switch-options interface <i>interface-name</i>], [edit vlans <i>vlan-name</i> switch-options], and [edit vlans <i>vlan-name</i> switch-options interface <i>interface-name</i>] hierarchy levels introduced in Junos OS Release 12.3R2 for EX Series switches.</p>   |
| <b>Description</b>         | <p>(MX Series routers or EX Series switches only) Configure a limit to the number of MAC addresses that can be learned from a bridge domain, VLAN, virtual switch, or set of bridge domains or VLANs.</p>  |



**NOTE:** For multichassis link aggregation (MC-LAG) peers in active-active mode, configuring the `interface-mac-limit` statement or changing the `interface-mac-limit` configuration when traffic is flowing can cause the MAC entries to be out of synchronization between the two MC-LAG peers, which might result in flooding. To avoid flooding, you must either halt traffic forwarding and then configure the `interface-mac-limit` statement or use the `commit at configuration` statement to commit the changes at the same time in both the peer nodes.

Alternatively, if flooding does occur, you can clear the bridge MAC table on both the routers by using the `clear bridge mac-table` command. Running this command ensures that the MAC entries are re-learned and in synchronization between both the peers.

---

**Default** For an access port, the default MAC limit is 1024 MAC addresses. For a trunk port, the default MAC limit is 8192 MAC addresses.

**Options** *limit*—Maximum number of MAC addresses learned from an interface.

**Range:** 1 through 131,071 MAC addresses per interface

The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Layer 2 Learning and Forwarding for Bridge Domains Overview](#)
- [Layer 2 Learning and Forwarding for VLANs Overview on page 1526](#)
- [Layer 2 Learning and Forwarding for Bridge Domains Functioning as Switches with Layer 2 Trunk Ports](#)
- [Layer 2 Learning and Forwarding for VLANs Acting as a Switch for a Layer 2 Trunk Port](#)

## no-dhcp-snooping

|  |  |
|--|--|
| <b>Syntax</b>                                  | no-dhcp-snooping;  |
| <b>Hierarchy Level (EX Series, QFX Series)</b> | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security</b> ]   |
| <b>Hierarchy Level (MX Series)</b>             | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]   |
| <b>Release Information</b>                     | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Hierarchy level [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security] introduced in Junos OS Release 14.1 for the MX Series. |
| <b>Description</b>                             | Disable DHCP snooping for the specified VLAN or bridge domain.   |



**NOTE:** Explicitly disabling DHCP snooping also disables any other port security features that you have enabled under [edit vlans *vlan-name* forwarding-options **dhcp-security**], including dynamic ARP inspection (DAI) and IP source guard for the specified VLAN or bridge domain.

There is no configuration statement that explicitly enables DHCP snooping.

**Default** DHCP snooping is not enabled.



**NOTE:** Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style and MX Series routers do not have a configuration statement that explicitly enables DHCP snooping.

DHCP snooping is enabled automatically by Junos OS if any of the following is configured at the [edit vlans *vlan-name* forwarding-options dhcp-security] hierarchy level for EX Series switches or at the [edit bridge-domains *bridge-domain-name* forwarding-options dhcp-security] for MX Series routers:

- DAI
- IP source guard
- Static IP
- DHCP option 82

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

- Related Documentation**
- [Understanding DHCP Snooping for Port Security on page 5256](#)

---

## no-option-82

---

|  |   |
|--|---|
| <b>Syntax</b>                                  | no-option-82;   |
| <b>Hierarchy Level (EX Series, QFX Series)</b> | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security group group-name overrides</a> ]  |
| <b>Hierarchy Level (MX Series)</b>             | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options group group <i>group-name overrides</i> ]   |
| <b>Release Information</b>                     | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series.   |
| <b>Description</b>                             | Configure a specific group of one or more access interfaces within the VLAN or bridge domain <i>not</i> to transmit DHCP option 82 information, even if the VLAN or bridge domain is configured to use option 82.   |
| <b>Required Privilege Level</b>                | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>                   | <ul style="list-style-type: none"><li>• <a href="#">option-82 on page 5359</a></li><li>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li><li>• <i>Understanding DHCP Option 82 for Port Security on Switching Devices</i></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li></ul> |



## option-82

|  |  |
|--|--|
| <b>Syntax</b>                                  | <pre> option-82 {   circuit-id {     prefix (host-name   routing-instance-name);     use-interface-description;     use-vlan-id;   }   remote-id {     host-name;     mac;     use-interface-description;     use-string string;   }   vendor-id {     use-string string;   } } </pre>   |
| <b>Hierarchy Level (EX Series, QFX Series)</b> | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security</a> ]  |
| <b>Hierarchy Level (MX Series)</b>             | [edit bridge-domains <i>bridge-domain-name</i> forwarding-options dhcp-security]   |
| <b>Release Information</b>                     | <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Juos OS Release 14.1 for the MX Series.</p>  |
| <b>Description</b>                             | <p>Have the device insert DHCP option 82 (also known as the DHCP relay agent information option) information in the DHCP request packet header of a DHCP request that it receives from a DHCP client connected to one of its interfaces before it forwards or relays that DHCP request to a DHCP server. The server uses the option 82 information, which provides details about the circuit and host the request came from. However, in formulating the reply, the server does not make any changes to the option 82 information in the packet header. The device receives the reply and then removes the DHCP option 82 information before forwarding the reply to the client.</p> <p>The remaining statements are explained separately.</p> |
| <b>Default</b>                                 | Insertion of DHCP option 82 information is not enabled.  |
| <b>Required Privilege Level</b>                | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>                   | <ul style="list-style-type: none"> <li>• <i>Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)</i></li> <li>• <a href="#">no-option-82 on page 5358</a></li> <li>• <i>Understanding DHCP Option 82 for Port Security on Switching Devices</i></li> <li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li> </ul>  |


- RFC 3046, *DHCP Relay Agent Information Option*, at <http://tools.ietf.org/html/rfc3046>

## overrides (DHCP Security)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | overrides (trusted   untrusted [no-option37   no-option-82];   |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <b>dhcp-security group</b> <i>group-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Support for the <b>no-option37</b> option introduced in Junos OS Release 13.2X51-D20 for EX Series switches.  |
| <b>Description</b>              | Modify selected attributes of a specific interface within a group of interfaces that is configured within a specified VLAN.  |
| <b>Options</b>                  | <b>no-option37</b> —The interface specified in this group does not support DHCPv6 option 37.<br><b>no-option82</b> —The interface specified in this group does not support DHCP option 82.<br><b>trusted</b> —The interface specified in this group is trusted. DHCP snooping does not apply to the trusted interface. Likewise, DAI and IP source guard—even if they are enabled for the VLAN—do not apply to the interface that is configured with the <b>overrides</b> and the <b>trusted</b> options. Access interfaces are untrusted by default.<br><b>untrusted</b> —(Only for EX9200) The interface specified in this group is untrusted. Trunk interface are trusted by default. Access interfaces are untrusted by default. |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Enabling a Trusted DHCP Server (CLI Procedure)</i></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li><li>• <i>Understanding DHCP Option 82 for Port Security on Switching Devices</i></li></ul>  |

## recovery-timeout

|   |   |
|---|---|
| <b>Syntax</b>                                     | <code>recovery-timeout seconds;</code>  |
| <b>Hierarchy Level (EX Series and QFX Series)</b> | [edit interfaces <i>interface-name</i> unit 0 family ethernet-switching]  |
| <b>Hierarchy Level (MX Series)</b>                | [edit interfaces <i>interface-name</i> unit 0 family bridge]  |
| <b>Release Information</b>                        | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series routers.   |
| <b>Description</b>                                | <p>Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, or rate-limiting configuration options for shutting down the interface, and allow the interface to recover automatically from the error condition after the specified period of time:</p> <ul style="list-style-type: none"> <li>• If you have enabled MAC limiting with the <b>shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.</li> <li>• If you have enabled MAC move limiting (not supported on EX9200) with the <b>shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.</li> <li>• If you have enabled storm control with the <b>action-shutdown</b> option and you enable <b>recovery-timeout</b>, the switch disables (rather than shuts down) the interface when applicable traffic exceeds the specified levels. Depending upon the configuration, applicable traffic could include broadcast, unknown unicast, and multicast traffic.</li> </ul> |
|   | <p> <b>NOTE:</b> The <b>recovery-timeout</b> configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after the <b>recovery-timeout</b> statement has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the operational mode command <code>clear ethernet-switching recovery-timeout</code> for EX Series and QFX Series and <code>clear bridge recovery-timeout</code> for MX Series routers.</p>  |
| <b>Default</b>                                    | Not enabled.  |
| <b>Options</b>                                    | <p><b>seconds</b>— Number of seconds that the interface remains in a disabled state due to a port error prior to automatic recovery.</p> <p><b>Range:</b> 10 through 3600</p>   |
| <b>Required Privilege Level</b>                   | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>  |

- Related Documentation**
- *action-shutdown*
  - [Configuring MAC Limiting \(CLI Procedure\) on page 1672](#)
  - *Configuring MAC Move Limiting (CLI Procedure)*
  - *Configuring or Disabling Storm Control (CLI Procedure)*

---

## static-ip

---

- Syntax** `static-ip ip-addresses {  
    vlan vlan-name;  
    mac mac-address;  
}`
- Hierarchy Level**
- For platforms with ELS:  
`[edit vlans vlan-name forwarding-options dhcp-security group group-name interface interface-name]`
  - For platforms without ELS:  
`[edit ethernet-switching-options secure-access-port interface (all | interface-name)]`
- Release Information** Statement introduced in Junos OS Release 9.2 for EX Series switches.  
Hierarchy level `[edit vlans vlan-name forwarding-options dhcp-security]` introduced in Junos OS Release 13.2X50-D10. (See [“Getting Started with Enhanced Layer 2 Software” on page 43](#) for information about ELS.)
- Description** Configure a static IP address to MAC address (IP-MAC) binding to be added to the DHCP snooping database.



**NOTE:** The VLAN is specified at the higher hierarchy level when `static-ip` is configured at `[edit vlans vlan-name forwarding-options dhcp-security group group-name interface interface-name]`.

---

- Options**
- ip-address*—Static IP address assigned to a device connected on the specified interface.
- mac-address*—Static MAC address assigned to a device connected on the specified interface.
- The remaining statements are explained separately.
- Required Privilege Level**
- system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.
- Related Documentation**
- *Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)*
  - *Configuring Static IP Addresses for DHCP and DHCPv6 Bindings on Access Ports (CLI Procedure)*

## switch-options

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> switch-options {   interface <i>interface-name</i> {     interface-mac-limit <i>limit</i> {       packet-action drop;     }     no-mac-learning;     static-mac <i>static-mac-address</i> {       vlan-id <i>number</i>;     }   }   interface-mac-limit <i>limit</i> {     packet-action drop;   }   mac-statistics;   mac-table-size <i>limit</i> {     packet-action drop;   }   no-mac-learning;   service-id <i>number</i>;   vtep-source-interface } </pre> |
| <b>Hierarchy Level</b>          | <pre> [edit <i>number</i>], [edit vlans <i>vlan--name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> vlans   <i>vlan-name</i>], [edit routing-instances <i>routing-instance-name</i> vlans <i>vlan-name</i>] </pre>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches and MX Series routers.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p>   |
| <b>Description</b>              | <p>Configure Layer 2 learning and forwarding properties for a VLAN or a virtual switch.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |

## trusted

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | trusted;  |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security group group-name overrides</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.   |
| <b>Description</b>              | Allow DHCP responses from the specified interface. The interface is not subject to DHCP snooping, even if the VLAN is enabled for DHCP snooping.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling a Trusted DHCP Server (CLI Procedure)</a></li><li>• <a href="#">Understanding Trusted DHCP Servers for Port Security on page 5268</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li></ul> |

## untrusted

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | untrusted;  |
| <b>Hierarchy Level</b>          | [edit vlans <i>vlan-name</i> forwarding-options <a href="#">dhcp-security group group-name overrides</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.  |
| <b>Description</b>              | Override the default behavior of a trunk interface from trusted to untrusted.   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Enabling a Trusted DHCP Server (CLI Procedure)</a></li><li>• <a href="#">Understanding Trusted DHCP Servers for Port Security on page 5268</a></li><li>• <a href="#">Understanding DHCP Snooping for Port Security on page 5256</a></li></ul> |

## Configuration Statements for Device Security

---


- [action-shutdown on page 5365](#)
- [bandwidth-level on page 5366](#)
- [bandwidth-percentage on page 5367](#)
- [interface \(Unknown Unicast Forwarding\) on page 5368](#)
- [no-broadcast on page 5369](#)

- [no-multicast on page 5370](#)
- [no-registered-multicast on page 5371](#)
- [no-unknown-unicast on page 5372](#)
- [no-unregistered-multicast on page 5373](#)
- [rpf-check on page 5374](#)
- [storm-control on page 5375](#)
- [storm-control-profiles on page 5376](#)
- [unknown-unicast-forwarding on page 5377](#)

## action-shutdown


|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | action-shutdown;   |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options storm-control]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <a href="#">storm-control-profiles</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Shut down or disable interfaces when the storm control level is exceeded, as follows: <ul style="list-style-type: none"> <li>• If you set both the <b>action-shutdown</b> and the <b>port-error-disable</b> statements, the affected interfaces are disabled temporarily and recover automatically when the disable timeout expires.</li> <li>• If you set the <b>action-shutdown</b> statement and do not set the <b>port-error-disable</b> statement, the affected interfaces are shut down when the storm control level is exceeded, and they do not recover automatically. You must issue the <b>clear ethernet-switching port-error</b> command to clear the port error and restore the interfaces to service.</li> </ul> |
| <b>Default</b>                  | The <b>action-shutdown</b> feature is disabled. If the storm control level is exceeded, the switch drops broadcast and unknown unicast messages on the specified interfaces.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Storm Control on page 5272</a></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li> <li>• <a href="#">port-error-disable on page 5341</a></li> <li>• <i>disable-timeout</i></li> <li>• <a href="#">clear ethernet-switching port-error on page 5397</a></li> </ul>  |

## bandwidth-level

|   |  |
|---|--|
| <b>Syntax</b>   | <code>bandwidth-level <i>kbps</i>;</code>  |
| <b>Hierarchy Level</b>  | [edit forwarding-options <a href="#">storm-control-profiles</a> <i>profile-name</i> all]   |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for MX Series routers.  |
| <b>Description</b>  | Configure the storm control level as the bandwidth in kilobits per second of the available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams.   |
| <div>  <p><b>NOTE:</b> When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</p> </div> |  |
| <b>Default</b>  | <p>On EX4300 switches—If you do not specify the storm control level using either the <b>bandwidth-level</b> or the <b>bandwidth-percentage</b> statements, the storm control level defaults to 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.</p> <p>On EX9200 switches—Storm control is not enabled by default.</p> <p>On MX Series routers—Storm control is not enabled by default.</p> |
| <b>Options</b>  | <p><b>bandwidth-level <i>kbps</i></b>—Traffic rate in kilobits per second of the combined broadcast, multicast, and unknown unicast traffic streams.</p> <p><b>Range:</b> 100 through 10,000,000</p> <p><b>Default:</b> None</p>   |
| <b>Required Privilege Level</b>   | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">bandwidth-percentage on page 5367</a></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</i></li> <li>• <i>Configuring or Disabling Storm Control (CLI Procedure)</i></li> </ul>  |



## bandwidth-percentage

|   |  |
|---|--|
| <b>Syntax</b>   | <code>bandwidth-percentage <i>percentage</i>;</code>   |
| <b>Hierarchy Level</b>  | [edit forwarding-options <a href="#">storm-control-profiles</a> <i>profile-name</i> all]   |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.1 for MX Series routers.  |
| <b>Description</b>  | Configure the storm control level as the percentage of available bandwidth used by the combined broadcast, multicast, and unknown unicast traffic streams on an interface.<br>The storm control level is configured as part of the storm control profile.  |
| <div>  <p><b>NOTE:</b> When you configure storm control level on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control level of 15,000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15,000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30,000 Kbps of combined broadcast, multicast, and unknown unicast traffic.</p> </div> |  |
| <b>Default</b>  | <p>On EX4300 switches—The storm control level is 80 percent of the available bandwidth used by the combined broadcast, unknown unicast, and multicast traffic streams.</p> <p>On EX9200 switches—Storm control is not enabled by default.</p> <p>On MX Series routers—Storm control is not enabled by default.</p>   |
| <b>Required Privilege Level</b>   | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">bandwidth-level on page 5366</a></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li> <li>• <i>Example: Configuring Storm Control to Prevent Network Outages on MX Series Routers</i></li> <li>• <i>Configuring or Disabling Storm Control (CLI Procedure)</i></li> </ul> |

## interface (Unknown Unicast Forwarding)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>interface <i>interface-name</i>;</code>   |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"><li>For platforms with ELS:<br/>[edit switch-options <b>unknown-unicast-forwarding</b> vlan <i>vlan-name</i>]</li><li>For platforms without ELS:<br/>[edit ethernet-switching-options <b>unknown-unicast-forwarding</b> vlan <i>vlan-name</i>]</li></ul>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Hierarchy level <b>[edit switch-options]</b> introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> |
| <b>Description</b>              | Specify the interface to which unknown unicast packets will be forwarded.   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li><code>show vlans</code></li><li><code>show ethernet-switching table</code></li><li><i>Configuring Unknown Unicast Forwarding (CLI Procedure)</i></li><li><a href="#">Understanding Unknown Unicast Forwarding on page 5277</a></li></ul>  |

---

## no-broadcast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-broadcast;  |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options storm-control interface (all   <i>interface-name</i> )]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <a href="#">storm-control-profiles</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | For interfaces configured for storm control, disable broadcast traffic storm control on the interface.   |
| <b>Default</b>                  | When storm control is enabled on an interface, it is enabled for both unknown unicast traffic and broadcast traffic.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 5272</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li></ul>  |

## no-multicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-multicast;  |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options storm-control interface (all   <i>interface-name</i> )]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <a href="#">storm-control-profiles</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Disable storm control for all multicast traffic (both registered multicast and unregistered multicast) for the specified interface or for all interfaces.  |
| <b>Default</b>                  | Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic.  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 5272</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li></ul>  |

## no-registered-multicast

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-registered-multicast;  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):<br/>[edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)]</li> </ul>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p> |
| <b>Description</b>              | <p>(EX8200 switches only) Disable storm control for registered multicast traffic for the specified interface or for all interfaces.</p> <p>(EX4300 and EX9200 switches only) Exclude storm control for registered multicast traffic from a storm control profile.</p> <p>(MX Series routers only) Exclude storm control for registered multicast traffic from a storm control profile.</p>  |
| <b>Default</b>                  | <p>EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.</p> <p>EX9200 switches—Storm control is not enabled by default.</p> <p>MX Series routers—Storm control is not enabled by default.</p>   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>no-multicast</i></li> <li><a href="#">no-unregistered-multicast on page 5373</a></li> <li><i>Understanding Storm Control on EX Series Switches</i></li> <li><i>Understanding Storm Control on Switching Devices</i></li> </ul>  |

## no-unknown-unicast

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | no-unknown-unicast;  |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options storm-control interface (all   <i>interface-name</i> )]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <a href="#">storm-control-profiles</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | For interfaces configured for storm control, disable unknown unicast traffic storm control on the interface.   |
| <b>Default</b>                  | When storm control is enabled on an interface, it is enabled for both unknown unicast traffic and broadcast traffic.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Storm Control on page 5272</a></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li></ul>  |

## no-unregistered-multicast

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-unregistered-multicast;  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with Enhanced Layer 2 Software (ELS) (EX Series switches and MX Series routers):<br/>[edit forwarding-options <b>storm-control-profiles</b> <i>profile-name</i> all]</li> <li>For platforms without ELS:<br/>[edit ethernet-switching-options storm-control interface (all   <i>interface-name</i>)],</li> </ul>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p> <p>Hierarchy level [edit forwarding-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX series.</p> <p>Statement introduced in Junos OS Release 14.1 for MX Series routers.</p> |
| <b>Description</b>              | <p>(EX8200 switches only) Disable storm control for unregistered multicast traffic for the specified interface or for all interfaces.</p> <p>(EX4300 and EX9200 switches only) Exclude storm control for unregistered multicast traffic from a storm control profile.</p> <p>(MX Series routers) Exclude storm control for unregistered multicast traffic from a storm control profile.</p>   |
| <b>Default</b>                  | <p>EX4300 and EX8200 switches—Storm control is enabled for unknown unicast traffic, multicast traffic, and broadcast traffic. The default storm control level is 80 percent of the available bandwidth used by the combined applicable traffic streams.</p> <p>EX9200 switches—Storm control is not enabled by default.</p> <p>MX Series routers—Storm control is not enabled by default.</p>   |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>no-multicast</i></li> <li><a href="#">no-registered-multicast on page 5371</a></li> <li><i>Understanding Storm Control on EX Series Switches</i></li> <li><i>Understanding Storm Control on Switching Devices</i></li> </ul>  |

## rpf-check

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | rpf-check;  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet],<br>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 9.3 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.   |
| <b>Description</b>              | <p>On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.</p> <p>On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.</p> <p>On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.</p> <p>On QFX Series switches, enable an RPF check on unicast traffic (except ECMP packets) on the selected ingress interfaces.</p> |
| <b>Default</b>                  | Unicast RPF is disabled on all interfaces.  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Unicast RPF on an EX Series Switch</i></li><li>• <a href="#">Configuring Unicast RPF (CLI Procedure) on page 5303</a></li><li>• <a href="#">Disabling Unicast RPF (CLI Procedure) on page 5304</a></li><li>• <a href="#">Understanding Unicast RPF on page 5273</a></li></ul>   |



---

## storm-control

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>storm-control <i>storm-control-profile</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> unit <i>number</i> family ethernet-switching],<br>[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX series.<br>Statement introduced in Junos OS Release 14.1 for the MX Series routers.  |
| <b>Description</b>              | <p>Bind a storm control profile to a logical interface.</p> <p>On switches running ELS software, storm control is enabled by default on all switch interfaces at a level of 80 percent of the combined broadcast and unknown unicast streams. (For the equivalent statement for platforms running non-ELS software, see <i>storm-control</i>.)</p> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li><li>• <i>Understanding Storm Control on Switching Devices</i></li></ul>   |

## storm-control-profiles

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>storm-control-profiles <i>profile-name</i> {<br/>    action-shutdown;<br/>    all {<br/>        bandwidth-level;<br/>        bandwidth-percentage;<br/>        no-broadcast;<br/>        no-multicast;<br/>        no-registered-multicast;<br/>        no-unknown-unicast;<br/>        no-unregistered-multicast;<br/>    }<br/>}</pre> |
| <b>Hierarchy Level</b>          | [edit forwarding-options]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.<br>Statement introduced in Junos OS Release 13.2 for the QFX Series.<br>Statement introduced in Junos OS Release 14.1 for MX Series routers.   |
| <b>Description</b>              | Configure a storm control profile on a switch or router.<br><br>The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Storm Control to Prevent Network Outages on EX Series Switches</i></li><li>• <i>Understanding Storm Control on Switching Devices</i></li></ul>  |

## unknown-unicast-forwarding

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>unknown-unicast-forwarding {   vlan <i>vlan-name</i> {     <i>interface</i> <i>interface-name</i>;   } }</pre>  |
| <b>Hierarchy Level</b>          | <ul style="list-style-type: none"> <li>For platforms with ELS:           <br/>[edit switch-options]         </li> <li>For platforms without ELS:           <br/>[edit ethernet-switching-options]         </li> </ul>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10. (See <a href="#">“Getting Started with Enhanced Layer 2 Software” on page 43</a> for information about ELS.)</p>           |
| <b>Description</b>              | <p>Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Default</b>                  | Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.  |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>show vlans</i></li> <li><i>show ethernet-switching table</i></li> <li><i>Configuring Unknown Unicast Forwarding (CLI Procedure)</i></li> <li><a href="#">Configuring Unknown Unicast Forwarding (CLI Procedure) on page 5305</a></li> <li><a href="#">Understanding Unknown Unicast Forwarding on page 5277</a></li> </ul> |



## CHAPTER 62

# Administration

- [Routine Monitoring on page 5379](#)
- [Monitoring Commands on page 5393](#)

### Routine Monitoring

---

- [Monitoring Firewall Filter Traffic on page 5379](#)
- [Monitoring Port Security on page 5381](#)
- [Verifying That Firewall Filters Are Operational on page 5382](#)
- [Verifying That DAI Is Working Correctly on page 5383](#)
- [Verifying That DHCP Snooping Is Working Correctly on page 5384](#)
- [Verifying That MAC Limiting Is Working Correctly on page 5385](#)
- [Verifying That MAC Move Limiting Is Working Correctly on page 5388](#)
- [Verifying That the Port Error Disable Setting Is Working Correctly on page 5388](#)
- [Verifying Unicast RPF Status on page 5389](#)
- [Verifying That a Trusted DHCP Server Is Working Correctly on page 5392](#)
- [Verifying That Three-Color Policers Are Operational on page 5392](#)
- [Verifying That Two-Color Policers Are Operational on page 5393](#)

### Monitoring Firewall Filter Traffic

You can use operational mode commands to monitor firewall filter traffic.

- [Monitoring Traffic for All Firewall Filters and Policers That Are Configured on page 5379](#)
- [Monitoring Traffic for a Specific Firewall Filter on page 5380](#)
- [Monitoring Traffic for a Specific Policer on page 5380](#)

#### Monitoring Traffic for All Firewall Filters and Policers That Are Configured

**Purpose** Monitor the number of packets and bytes that matched the firewall filters and monitor the number of packets that exceeded policer rate limits:

**Action** Use the **show firewall** operational mode command:

```
user@switch> show firewall
```

```
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes          Packets
counter-employee-web              3348           27
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                      560            10
Policers:
Name                               Packets
icmp-connection-policer          10
tcp-connection-policer           0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall** command displays the names of all firewall filters, counters, and policers that are configured. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits.

---

#### Monitoring Traffic for a Specific Firewall Filter

**Purpose** Monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded policer rate limits.

**Action** Use the **show firewall filter *filter-name*** operational mode command:

```
user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                      560            10
```

**Meaning** The **show firewall filter *filter-name*** command limits the display information to the counters and policers that are defined for the specified filter.

---

#### Monitoring Traffic for a Specific Policer

**Purpose** Monitor the number of packets that exceeded the rate limits of a policer:

**Action** Use the **show firewall policer *policer-name*** operational mode command:

```
user@switch> show firewall policer icmp-connection-policer
Filter: ingress-port-limit-tcp-icmp
Policers:
Name                               Packets
icmp-connection-policer           10
```

**Meaning** The **show firewall policer *policer-name*** command displays the number of packets that exceeded the rate limits for the specified policer.

- Related Documentation**
- [Configuring Firewall Filters on page 5290](#)
  - [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)
  - [Verifying That Firewall Filters Are Operational on page 5382](#)

## Monitoring Port Security

### Purpose



**NOTE:** This topic applies only to the J-Web Application package.

Use the monitoring functionality to view these port security details:

- DHCP snooping database for a VLAN or all VLANs
- ARP inspection details for all interfaces

### Action

To monitor port security in the J-Web interface, select **Monitor > Security > Port Security**.

To monitor and manipulate the DHCP snooping database and ARP inspection statistics in the CLI, enter the following commands:

- **show dhcp snooping binding**
- **clear dhcp snooping binding**—In addition to clearing the whole database, you can clear database entries for specified VLANs or MAC addresses.
- **show arp inspection statistics**
- **clear arp inspection statistics**



**NOTE:** On EX4300 switches, to monitor and manipulate the DHCP snooping database and ARP inspection statistics in the CLI, enter the following commands:

- **show dhcp-security binding**
- **clear dhcp-security binding**—In addition to clearing the whole database, you can clear database entries for specified VLANs or IP Address.
- **show dhcp-security arp inspection statistics**
- **clear arp inspection statistics**

### Meaning

The J-Web Port Security Monitoring page comprises two sections:

- **DHCP Snooping Details**—Displays the DHCP snooping database for all the VLANs for which DHCP snooping is enabled. To view the DHCP snooping database for a specific VLAN, select the specific VLAN from the list.

- **ARP Inspection Details**—Displays the ARP inspection details for all interfaces. The information includes details of the number of packets that passed ARP inspection and the number of packets that failed the inspection. The pie chart graphically represents these statistics when you select an interface. To view ARP inspection statistics for a specific interface, select the interface from the list.

You can use the following options on the page to clear DHCP snooping and ARP inspection details:

- **Clear All**—Clears the DHCP snooping database, either for all VLANs if the option **ALL** has been selected in the Select VLANs list or for the specific VLAN that has been selected in that list.
- **Clear**—Deletes a specific IP address from the DHCP snooping database.

To clear ARP inspection details on the page, click **Clear All** in the ARP inspection details section.



**NOTE:** Clear All button in the ARP inspection details section is not supported on EX4300 switches.

Use the CLI commands to show and clear DHCP snooping database and ARP inspection statistics details.

#### Related Documentation

- *Configuring Port Security (CLI Procedure)*
- *Configuring Port Security (J-Web Procedure)*
- *Example: Configuring Basic Port Security Features*

## Verifying That Firewall Filters Are Operational

**Purpose** Verify that firewall filters are working properly after you apply them to ports, VLANs, or Layer 3 interfaces.

**Action** Use the **show firewall** operational mode command to verify that the firewall filters are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes      Packets
counter-employee-web              0           0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes      Packets
icmp-counter                      560        10
Policers:
Name                               Packets
icmp-connection-policer          10
tcp-connection-policer           0
```



```
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The **show firewall** command displays the names of all firewall filters, counters, and policers that are configured. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. In the above example, the **icmp-counter** in the filter **ingress-port-limit-tcp-icmp** shows that the filter matched 10 packets. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits. The policer **icmp-connection-policer** shows that 10 ICMP packets were policed.

- Related Documentation**
- [Configuring Firewall Filters on page 5290](#)
  - [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)
  - [Monitoring Firewall Filter Traffic on page 5379](#)

## Verifying That DAI Is Working Correctly

**Purpose** Verify that dynamic ARP inspection (DAI) is working on the switch.

**Action** Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface          Packets received  ARP inspection pass  ARP inspection failed
-----
ge-0/0/1.0          7                 5                   2
ge-0/0/2.0         10                10                  0
ge-0/0/3.0         12                12                  0
```

**Meaning** The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

- Related Documentation**
- *Enabling Dynamic ARP Inspection (CLI Procedure)*
  - *Enabling Dynamic ARP Inspection (J-Web Procedure)*
  - *Example: Configuring Basic Port Security Features*
  - *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
  - *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
  - [Monitoring Port Security on page 5381](#)

## Verifying That DHCP Snooping Is Working Correctly

**Purpose** Verify that DHCP snooping is working on the switch and that the DHCP snooping database is correctly populated with both dynamic and static bindings.

**Action** Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
DHCP Snooping Information:
MAC address      IP address  Lease (seconds) Type      VLAN    Interface
00:05:85:3A:82:77 192.0.2.17  600          dynamic employee ge-0/0/1.0
00:05:85:3A:82:79 192.0.2.18  653          dynamic employee ge-0/0/1.0
00:05:85:3A:82:80 192.0.2.19  720          dynamic employee ge-0/0/2.0
00:05:85:3A:82:81 192.0.2.20  932          dynamic employee ge-0/0/2.0
00:05:85:3A:82:83 192.0.2.21  1230         dynamic employee ge-0/0/2.0
00:05:85:27:32:88 192.0.2.22  -            static   data     ge-0/0/4.0
```

**Meaning** When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires. Static IP addresses have no assigned lease time. The statically configured entry never expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

- Related Documentation**
- *Enabling DHCP Snooping (CLI Procedure)*
  - *Enabling DHCP Snooping (J-Web Procedure)*
  - *Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)*
  - *Example: Configuring Basic Port Security Features*
  - *Example: Configuring DHCP Snooping, DAI, and MAC Limiting on a Switch with Access to a DHCP Server Through a Second Switch*
  - *Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks*
  - [Monitoring Port Security on page 5381](#)
  - *Troubleshooting Port Security*

## Verifying That MAC Limiting Is Working Correctly

MAC limiting protects against flooding of the Ethernet switching table by setting a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

Junos OS provides two MAC limiting methods:

- **Maximum number of MAC addresses**—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses can be ignored, dropped, or logged. You can also specify that the interface be shut down or temporarily disabled.
- **Allowed MAC addresses**—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned, and the switch logs an appropriate message. The allowed MAC method binds MAC addresses to a VLAN so that the address is not registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

This topic includes the following tasks:

1. [Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly on page 5385](#)
2. [Verifying That Allowed MAC Addresses Are Working Correctly on page 5386](#)
3. [Verifying That Interfaces Are Shut Down on page 5386](#)
4. [Customizing the Ethernet Switching Table Display to View Information for a Specific Interface on page 5387](#)

### Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly

**Purpose** Verify that MAC limiting for dynamic MAC addresses is working.

**Action** Display the MAC addresses that have been learned. The following sample output shows the results of sending two packets from hosts connected to **xe-1:0/0/1** and five packets from hosts connected to **xe-1:0/0/2**, with both interfaces configured with a MAC limit of **4** and the action **drop**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 7 entries, 6 learned
```

| VLAN          | MAC address       | Type  | Age | Interfaces   |
|---------------|-------------------|-------|-----|--------------|
| employee-vlan | *                 | Flood | -   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:77 | Learn | 0   | xe-1:0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:79 | Learn | 0   | xe-1:0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:80 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:81 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:83 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:85 | Learn | 0   | xe-1:0/0/2.0 |

**Meaning** The output shows that the fifth packet received on the **xe-1:0/0/2** interface was dropped because it exceeded the MAC limit for that interface. The address was not learned, and thus an asterisk (\*) rather than an address appears in the MAC address column in the first line of the sample output.

---

### Verifying That Allowed MAC Addresses Are Working Correctly

**Purpose** Verify that allowed MAC addresses are working.

**Action** Display the MAC cache information after allowed MAC addresses have been configured on an interface. The following sample shows the MAC cache after four allowed MAC addresses had been configured on interface **xe-1:0/0/2** and a fifth MAC address appeared on the interface.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 5 entries, 4 learned
```

| VLAN          | MAC address       | Type  | Age | Interfaces   |
|---------------|-------------------|-------|-----|--------------|
| employee-vlan | 00:05:85:3A:82:80 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:81 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:83 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:85 | Learn | 0   | xe-1:0/0/2.0 |
| employee-vlan | *                 | Flood | -   | xe-1:0/0/2.0 |

**Meaning** Because the fifth address was not allowed it was not learned, and an asterisk (\*) rather than an address appears in the MAC address column in the last line of the sample output.

---

### Verifying That Interfaces Are Shut Down

**Purpose** Verify that an interface is shut down when the MAC limit is exceeded.

**Action** For more information about interfaces that have been shut down because the MAC limit was exceeded, use the **show ethernet-switching interfaces** command.

```
user@switch> show ethernet-switching interfaces
Interface      State  VLAN members      Tag  Tagging  Blocking

bme0.32770     down  mgmt              untagged unblocked
xe-0/0/0.0     down  v1                untagged MAC limit exceeded
xe- 0/0/1.0    up    v1                untagged unblocked
xe-0/0/2.0     up    v1                untagged unblocked
me0.0          up    mgmt              untagged unblocked
```



**NOTE:** You can configure interfaces to recover automatically when the MAC limit has been exceeded by specifying the **port-error-disable** statement with a **disable timeout** value. The switch automatically restores the disabled interface to service when the disable timeout expires. The **port-error-disable** configuration does not apply to preexisting error conditions—it affects only error conditions that are detected after the **port-error-disable** statement has been enabled and the configuration has been committed. To clear a preexisting error condition and restore the interface to service, use the **clear ethernet-switching port-error** command.

### Customizing the Ethernet Switching Table Display to View Information for a Specific Interface

**Purpose** You can use the **show ethernet-switching table** command to view information for a specific interface.

**Action** For example, to display the MAC addresses that have been learned on the **xe-0/0/2** interface, enter:

```
user@switch> show ethernet-switching table interface xe-0/0/2.0
Ethernet-switching table: 1 unicast entries
```

| VLAN | MAC address       | Type  | Age | Interfaces  |
|------|-------------------|-------|-----|-------------|
| v1   | *                 | Flood | -   | All-members |
| v1   | 00:00:06:00:00:00 | Learn | 0   | xe-0/0/2.0  |

**Meaning** The MAC limit value for the **xe-0/0/2** interface had been set to **1**, and the output shows that only one MAC address was learned and added to the MAC cache.

- Related Documentation**
- [Configuring MAC Limiting](#)
  - [Monitoring Port Security on page 5381](#)
  - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
  - [Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks](#)
  - [Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks](#)

## Verifying That MAC Move Limiting Is Working Correctly

**Purpose** Verify that MAC move limiting is working on the switch.

**Action** Display the MAC addresses in the Ethernet switching table when MAC move limiting has been configured for a VLAN. The following sample shows the results after two of the hosts on **ge-0/0/2** sent packets after the MAC addresses for those hosts had moved to other interfaces more than five times in 1 second. The VLAN, **employee-vlan**, was set to a MAC move limit of 5 with the action **drop**:

```
user@switch> show ethernet-switching table
```

```
Ethernet-switching table: 7 entries, 4 learned
```

| VLAN          | MAC address       | Type  | Age | Interfaces |
|---------------|-------------------|-------|-----|------------|
| employee-vlan | 00:05:85:3A:82:77 | Learn | 0   | ge-0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:79 | Learn | 0   | ge-0/0/1.0 |
| employee-vlan | 00:05:85:3A:82:80 | Learn | 0   | ge-0/0/2.0 |
| employee-vlan | 00:05:85:3A:82:81 | Learn | 0   | ge-0/0/2.0 |
| employee-vlan | *                 | Flood | -   | ge-0/0/2.0 |
| employee-vlan | *                 | Flood | -   | ge-0/0/2.0 |

**Meaning** The last two lines of the sample output show that MAC addresses for two hosts on **ge-0/0/2** were not learned, because the hosts had been moved back and forth from the original interfaces more than five times in 1 second.

- Related Documentation**
- [Configuring MAC Move Limiting \(CLI Procedure\)](#)
  - [Configuring MAC Move Limiting \(J-Web Procedure\)](#)
  - [Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces \(CLI Procedure\)](#)
  - [Example: Configuring Basic Port Security Features](#)
  - [Monitoring Port Security on page 5381](#)

## Verifying That the Port Error Disable Setting Is Working Correctly

**Purpose** Verify that the port error disable setting is working as expected for MAC limited and storm control interfaces.

**Action** Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
```

| Interface    | State | VLAN members | Blocking                |
|--------------|-------|--------------|-------------------------|
| xe-2:0/0/0.0 | up    | T1122        | unblocked               |
| xe-2:0/0/1.0 | down  | default      | MAC limit exceeded      |
| xe-2:0/0/2.0 | down  | default      | Storm control in effect |
| xe-2:0/0/3.0 | down  | default      | unblocked               |
| xe-2:0/0/4.0 | down  | default      | unblocked               |
| xe-2:0/0/5.0 | down  | default      | unblocked               |
| xe-2:0/0/6.0 | down  | default      | unblocked               |

**Meaning** For interfaces disabled by port security features, the sample output from the **show ethernet-switching interfaces** command specifies the reasons that the interfaces are disabled:

- **MAC limit exceeded**—The interface is temporarily disabled because of a *mac-limit* error. The disabled interface is automatically restored to service when the *disable-timeout* (*Port Error Disable*) expires.
- **MAC move limit exceeded**—The interface is temporarily disabled because of a *mac-move-limit* error. The disabled interface is automatically restored to service when the *disable-timeout* expires.
- **Storm control in effect**—The interface is temporarily disabled because of a *storm-control* error. The disabled interface is automatically restored to service when the *disable-timeout* (*Port Error Disable*) expires.

- Related Documentation**
- [Understanding MAC Limiting and MAC Move Limiting for Port Security on page 5266](#)
  - [port-error-disable on page 5341](#)
  - [Configuring Autorecovery for MAC Limited or Storm Control Interfaces \(CLI Procedure\)](#)

## Verifying Unicast RPF Status

**Purpose** Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

**Action** Use one of the **show interfaces *interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The following example displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge-1/0/10 extensive
```

Physical interface: ge-1/0/10, Enabled, Physical link is Down  
 Interface index: 139, SNMP ifIndex: 58, Generation: 140  
 Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,  
 Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,  
 Auto-negotiation: Enabled, Remote fault: Online  
 Device flags : Present Running  
 Interface flags: Hardware-Down SNMP-Traps Internal: 0x0  
 Link flags : None  
 CoS queues : 8 supported, 8 maximum usable queues  
 Hold-times : Up 0 ms, Down 0 ms  
 Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab

```

Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort      0              0              0
  1 assured-forw      0              0              0
  5 expedited-fo      0              0              0
  7 network-cont      0              0              0

Active alarms : LINK
Active defects : LINK
MAC statistics:
  Total octets      Receive      Transmit
  Total packets      0            0
  Unicast packets    0            0
  Broadcast packets  0            0
  Multicast packets  0            0
  CRC/Align errors   0            0
  FIFO errors        0            0
  MAC control frames 0            0
  MAC pause frames   0            0
  Oversized frames   0
  Jabber frames      0
  Fragment frames    0
  VLAN tagged frames 0
  Code violations     0
Filter statistics:
  Input packet count      0
  Input packet rejects    0
  Input DA rejects        0
  Input SA rejects        0
  Output packet count      0
  Output packet pad count  0
  Output packet error count 0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

```



```

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
  Protocol inet, Generation: 144, Route table: 0
Flags: uRPF
Addresses, Flags: Is-Preferred Is-Primary

```

**Meaning** The `show interfaces ge-1/0/10 extensive` command (and the `show interfaces ge-1/0/10 detail` command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200, EX4200, and EX4300 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

- Related Documentation**
- *show interfaces xe-*
  - *Example: Configuring Unicast RPF on an EX Series Switch*
  - [Configuring Unicast RPF \(CLI Procedure\) on page 5303](#)
  - [Disabling Unicast RPF \(CLI Procedure\) on page 5304](#)
  - *Troubleshooting Unicast RPF*

## Verifying That a Trusted DHCP Server Is Working Correctly

**Purpose** Verify that a DHCP trusted server is working on the switch. See what happens when the DHCP server is trusted and then untrusted.

**Action** Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

DHCP Snooping Information:

| MAC Address       | IP Address | Lease | Type    | VLAN          | Interface  |
|-------------------|------------|-------|---------|---------------|------------|
| 00:05:85:3A:82:77 | 192.0.2.17 | 600   | dynamic | employee-vlan | ge-0/0/1.0 |
| 00:05:85:3A:82:79 | 192.0.2.18 | 653   | dynamic | employee-vlan | ge-0/0/1.0 |
| 00:05:85:3A:82:80 | 192.0.2.19 | 720   | dynamic | employee-vlan | ge-0/0/2.0 |
| 00:05:85:3A:82:81 | 192.0.2.20 | 932   | dynamic | employee-vlan | ge-0/0/2.0 |
| 00:05:85:3A:82:83 | 192.0.2.21 | 1230  | dynamic | employee-vlan | ge-0/0/2.0 |
| 00:05:85:27:32:88 | 192.0.2.22 | 3200  | dynamic | employee-vlan | ge-0/0/2.0 |

**Meaning** When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

- Related Documentation**
- *Enabling a Trusted DHCP Server (CLI Procedure)*
  - *Enabling a Trusted Port for DHCP*
  - *Enabling a Trusted DHCP Server (J-Web Procedure)*
  - *Example: Configuring Basic Port Security Features*
  - *Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks*
  - [Monitoring Port Security on page 5381](#)
  - *Troubleshooting Port Security*

## Verifying That Three-Color Policers Are Operational

**Purpose** Verify that three-color policers in firewall filter configurations are working properly.

**Action** Use the following operational mode commands to verify that a three-color policer is working properly:

- `show class-of-service forwarding-table classifiers`
- `show interfaces interface-name extensive`
- `show interfaces queue interface-name`

**Related Documentation**

- [Overview of Policers on page 5241](#)
- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)

## Verifying That Two-Color Policers Are Operational

**Purpose** Verify that two-color policers in firewall filter configurations are working properly.

**Action** Use the `show firewall policer` operational mode command to verify that the policers are working properly:

```
user@switch> show firewall policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-limit-tcp-icmp
Policers:
Name                               Packets
icmp-connection-policer            10
tcp-connection-policer             539
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

**Meaning** The `show firewall policer` command displays the names of all firewall filters and policers that are configured. For each policer that is specified in a filter configuration, the output field shows the current packet count for all packets that exceed the specified rate limits.

**Related Documentation**

- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)
- [Configuring Firewall Filters on page 5290](#)
- [Monitoring Firewall Filter Traffic on page 5379](#)

## Monitoring Commands

- `clear arp inspection statistics`
- `clear dhcp snooping binding`
- `clear ethernet-switching port-error`
- `clear firewall`
- `show arp inspection statistics`
- `show dhcp snooping binding`
- `show firewall`

- [show firewall policer](#)
- [show interfaces filters](#)

## clear arp inspection statistics

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | clear arp inspection statistics<br><interface <i>interface</i> >  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Clear ARP inspection statistics.  |
| <b>Options</b>                  | <b>none</b> —Clears ARP statistics on all interfaces.<br><br><b>interface <i>interface-names</i></b> —(Optional) Clear ARP statistics on one or more interfaces.  |
| <b>Required Privilege Level</b> | clear   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show arp inspection statistics on page 5399</a></li> <li>• <i>Example: Configuring Basic Port Security Features</i></li> <li>• <a href="#">Verifying That DAI Is Working Correctly on page 5383</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">clear arp inspection statistics on page 5395</a>  |
| <b>Output Fields</b>            | This command produces no output.  |

## Sample Output

### clear arp inspection statistics

```
user@switch> clear arp inspection statistics
```

## clear dhcp snooping binding

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>clear dhcp snooping binding</code><br><code>&lt;mac (all   <i>mac-address</i>)&gt;</code><br><code>&lt;vlan (all   <i>vlan-name</i>)&gt;</code><br><code>&lt;vlan (all   <i>vlan-name</i>) mac (all   <i>mac-address</i>)&gt;</code>                         |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Clear the DHCP snooping database information.  |
| <b>Options</b>                  | <b>mac (all   <i>mac-address</i>)</b> —(Optional) Clear DHCP snooping information for the specified MAC address or all MAC addresses.<br><br><b>vlan (all   <i>vlan-name</i>)</b> —(Optional) Clear DHCP snooping information for the specified VLAN or all VLANs. |
| <b>Required Privilege Level</b> | clear  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Example: Configuring Basic Port Security Features</i></li><li>• <a href="#">show dhcp snooping binding on page 5400</a></li></ul>   |
| <b>List of Sample Output</b>    | <a href="#">clear dhcp snooping binding on page 5396</a>   |
| <b>Output Fields</b>            | This command produces no output.   |

### Sample Output

#### clear dhcp snooping binding

```
user@switch> clear dhcp snooping binding
```

---

## clear ethernet-switching port-error

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | clear ethernet-switching port-error<br><interface <i>interface-name</i> >   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch or from the specified interface, and restore the interfaces or the specified interface to service.   |
| <b>Options</b>                  | <b>none</b> —Clear all MAC limiting, MAC move limiting, and storm control errors from all the Ethernet switching interfaces on the switch and restore the interfaces to service.<br><br><b>interface <i>interface-name</i></b> —(Optional) Clear all MAC limiting, MAC move limiting, and storm control errors from the specified interface and restore the interface to service.           |
| <b>Required Privilege Level</b> | clear   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring MAC Limiting</i></li><li>• <i>Example: Configuring Storm Control to Prevent Network Outages</i></li><li>• <i>Configuring Port Security (CLI Procedure)</i></li><li>• <a href="#">port-error-disable on page 5341</a></li><li>• <i>Configuring Autorecovery for MAC Limited or Storm Control Interfaces (CLI Procedure)</i></li></ul> |
| <b>Output Fields</b>            | This command produces no output.  |

## clear firewall

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>clear firewall (all   counter <i>counter-name</i>   filter <i>filter-name</i>)</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Clear statistics provided by firewall filters.</p> <p>When you clear the counters of a filter, this not only impacts the counters shown by the CLI, but also the ones tracked by SNMP 2.</p>   |
| <b>Options</b>                  | <p><b>all</b>—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.</p> <p><b>counter <i>counter-name</i></b>—Clear the packet and byte counts for the specified firewall filter counter.</p> <p><b>filter <i>filter-name</i></b>—Clear the packet and byte counts for the specified firewall filter.</p> |
| <b>Required Privilege Level</b> | clear   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Verifying That Firewall Filters Are Operational on page 5382</a></li><li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 5393</a></li><li>• <a href="#">Overview of Firewall Filters on page 5209</a></li><li>• <a href="#">Overview of Policers on page 5241</a></li></ul>                     |

## Sample Output

### clear firewall all

```
user@switch> clear firewall all
```

### clear firewall counter

```
user@switch> clear firewall counter port-filter-counter
```

### clear firewall filter

```
user@switch> clear firewall filter ingress-port-filter
```



## show arp inspection statistics

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | show arp inspection statistics   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Display ARP inspection statistics.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear arp inspection statistics on page 5395</a></li> <li>• <i>Example: Configuring Basic Port Security Features</i></li> <li>• <a href="#">Verifying That DAI Is Working Correctly on page 5383</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show arp inspection statistics on page 5399</a>  |
| <b>Output Fields</b>            | <a href="#">Table 440 on page 5399</a> lists the output fields for the <b>show arp inspection statistics</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 440: show arp inspection statistics Output Fields**

| Field Name            | Field Description  | Level of Output |
|-----------------------|--|-----------------|
| Interface             | Interface on which ARP inspection has been applied.          | All levels      |
| Packets received      | Total number of packets total that underwent ARP inspection. | All levels      |
| ARP inspection pass   | Total number of packets that passed ARP inspection.          | All levels      |
| ARP inspection failed | Total number of packets that failed ARP inspection.          | All levels      |

## Sample Output

### show arp inspection statistics

```
user@switch> show arp inspection statistics
```

| Interface | Packets received | ARP inspection pass | ARP inspection failed |
|-----------|------------------|---------------------|-----------------------|
| -----     | -----            | -----               | -----                 |
| ge-0/0/0  | 0                | 0                   | 0                     |
| ge-0/0/1  | 0                | 0                   | 0                     |
| ge-0/0/2  | 0                | 0                   | 0                     |
| ge-0/0/3  | 0                | 0                   | 0                     |
| ge-0/0/4  | 0                | 0                   | 0                     |
| ge-0/0/5  | 0                | 0                   | 0                     |
| ge-0/0/6  | 0                | 0                   | 0                     |
| ge-0/0/7  | 703              | 701                 | 2                     |

## show dhcp snooping binding

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show dhcp snooping binding</b><br><b>&lt;interface <i>interface-name</i>&gt;</b><br><b>&lt;vlan <i>vlan-name</i>&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.0 for EX Series switches.<br>Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Display the DHCP snooping database information.  |
| <b>Options</b>                  | <b>interface <i>interface-name</i></b> —(Optional) Display the DHCP snooping database information for an interface.<br><br><b>vlan <i>vlan-name</i></b> —(Optional) Display the DHCP snooping database information for a VLAN.                       |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>clear dhcp snooping binding</i></li> <li>• <i>Example: Configuring Basic Port Security Features</i></li> <li>• <a href="#">Verifying That DHCP Snooping Is Working Correctly on page 5384</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show dhcp snooping binding on page 5400</a>  |
| <b>Output Fields</b>            | <a href="#">Table 441 on page 5400</a> lists the output fields for the <b>show dhcp snooping binding</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 441: show dhcp snooping binding Output Fields

| Field Name  | Field Description   | Level of Output |
|-------------|---|-----------------|
| MAC Address | MAC address of the network device; bound to the IP address. | All levels      |
| IP Address  | IP address of the network device; bound to the MAC address. | All levels      |
| Lease       | Lease granted to the IP address.                            | All levels      |
| Type        | How the MAC address was acquired.                           | All levels      |
| VLAN        | VLAN name of the network device whose MAC address is shown. | All levels      |
| Interface   | Interface address (port).                                   | All levels      |

## Sample Output

### show dhcp snooping binding

```
user@switch> show dhcp snooping binding
```

## DHCP Snooping Information:

| MAC Address       | IP Address | Lease | Type    | VLAN  | Interface   |
|-------------------|------------|-------|---------|-------|-------------|
| 00:00:01:00:00:03 | 192.0.2.0  | 640   | dynamic | guest | ge-0/0/12.0 |
| 00:00:01:00:00:04 | 192.0.2.1  | 720   | dynamic | guest | ge-0/0/12.0 |
| 00:00:01:00:00:05 | 192.0.2.5  | 800   | dynamic | guest | ge-0/0/13.0 |

## show firewall

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>show firewall   &lt;counter <i>counter-name</i>&gt;   &lt;filter <i>filter-name</i>&gt;   &lt;log &lt;detail   interface <i>interface-name</i>&gt;&gt;   &lt;terse&gt;</pre>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Display statistics about configured firewall filters.   |
| <b>Options</b>                  | <p><b>counter <i>counter-name</i></b>—(Optional) Display statistics about a particular firewall filter counter.</p> <p><b>filter <i>filter-name</i></b>—(Optional) Display statistics about a particular firewall filter.</p> <p><b>log</b>—(Optional) Display log entries for all firewall filter activity.</p> <p><b>terse</b>—(Optional) Display firewall filter names only.</p> |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying That Firewall Filters Are Operational on page 5382</a></li> <li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 5393</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5209</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul>                        |
| <b>List of Sample Output</b>    | <p><a href="#">show firewall on page 5403</a></p> <p><a href="#">show firewall filter <i>filter-name</i> on page 5404</a></p> <p><a href="#">show firewall counter <i>counter-name</i> on page 5404</a></p> <p><a href="#">show firewall log on page 5404</a></p> <p><a href="#">show firewall log detail on page 5404</a></p>  |
| <b>Output Fields</b>            | <a href="#">Table 442 on page 5402</a> lists the output fields for the <b>show firewall</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 442: show firewall Output Fields**

| Field Name | Field Description   | Level of Output |
|------------|---|-----------------|
| Filter     | Name of the filter that is configured at the <b>[edit firewall family <i>family-name</i> filter]</b> hierarchy level. | All levels      |

Table 442: show firewall Output Fields (*continued*)

| Field Name           | Field Description  | Level of Output |
|----------------------|--|-----------------|
| <b>Counters</b>      | Display filter counter information: <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of a filter counter that has been configured with the <b>count</b> firewall filter action modifier.</li> <li>• <b>Bytes</b>—Number of bytes that match the filter term where the <b>count</b> action modifier was specified.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term where the <b>count</b> action modifier was specified.</li> </ul> | All levels      |
| <b>Policers</b>      | Display policer information: <ul style="list-style-type: none"> <li>• <b>Name</b>—Name of the policer that is configured at the <b>[edit firewall policer]</b> hierarchy level.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term where the <b>policer</b> action modifier was specified. This is the number of packets that exceeded the rate limits that the policer specifies.</li> </ul>  | All levels      |
| <b>Action</b>        | Filter action: <ul style="list-style-type: none"> <li>• <b>A</b>—Accept</li> <li>• <b>D</b>—Discard</li> </ul>   | All levels      |
| <b>Interface</b>     | Interface on which the firewall filter is applied.   | All levels      |
| <b>Protocol</b>      | Name of the packet protocol.   | All levels      |
| <b>Packet Length</b> | Length of the packet.  | All levels      |
| <b>Src Addr</b>      | Source address of the packet.  | All levels      |
| <b>Dest Addr</b>     | Destination address of the packet.   | All levels      |

## Sample Output

### show firewall

```

user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                               Bytes      Packets
counter-employee-web                0           0
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes      Packets
icmp-counter                        560        10
Policers:
Name                               Packets
icmp-connection-policer            10
tcp-connection-policer              0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest

```

**show firewall filter filter-name**

```

user@switch> show firewall filter ingress-port-limit-tcp-icmp
Filter: ingress-port-limit-tcp-icmp
Counters:
Name                               Bytes          Packets
icmp-counter                       560            10
Policers:
Name                               Packets
icmp-connection-policer           10
tcp-connection-policer            0

```

**show firewall counter counter-name**

```

user@switch> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                               Bytes          Packets
icmp-counter                       560            10

```

**show firewall log**

```

user@switch> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
08:00:53  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:52  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:51  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:50  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:49  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:48  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4
08:00:47  pfe      R    ge-1/0/6.0  ICMP      192.168.3.5
192.168.3.4

```

**show firewall log detail**

```

user@switch> show firewall log detail
Log :

Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of

```

```
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2010-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
```

## show firewall policer

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>show firewall policer</code><br><code>&lt;policer-name&gt;</code>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Display statistics about configured policers.  |
| <b>Options</b>                  | <p><b>none</b>—Display the count of policed packets for all configured policers.</p> <p><b>policer-name</b>—(Optional) Display the count of policed packets for the specified policer.</p>   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Verifying That Firewall Filters Are Operational on page 5382</a></li> <li>• <a href="#">Verifying That Two-Color Policers Are Operational on page 5393</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5209</a></li> <li>• <a href="#">Overview of Policers on page 5241</a></li> </ul> |
| <b>List of Sample Output</b>    | <p><a href="#">show firewall policer on page 5406</a></p> <p><a href="#">show firewall policer policer-name on page 5407</a></p>   |
| <b>Output Fields</b>            | Table 443 on page 5406 lists the output fields for the <b>show firewall policer</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 443: show firewall policer Output Fields**

| Field Name      | Field Description   | Level of Output |
|-----------------|---|-----------------|
| <b>Filter</b>   | Name of the filter that is configured at the <code>[edit firewall family family-name filter]</code> hierarchy level.  | All levels      |
| <b>Policers</b> | Display policer information: <ul style="list-style-type: none"> <li>• <b>Filter</b>—Name of filter that specifies the <b>policer</b> action modifier.</li> <li>• <b>Name</b>—Name of policer.</li> <li>• <b>Packets</b>—Number of packets that matched the filter term in which the <b>policer</b> action modifier is specified. This is the number of packets that exceed the rate limits that the policer specifies.</li> </ul> | All levels      |

## Sample Output

### show firewall policer

```

user@switch> show firewall policer
Filter: egress-vlan-filter
Filter: ingress-port-filter
Policers:
Name                                     Packets

```



|                                  |   |
|----------------------------------|---|
| icmp-connection-policer          | 0 |
| tcp-connection-policer           | 0 |
| Filter: ingress-vlan-rogue-block |   |

#### show firewall policer policer-name

```
user@switch> show firewall policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                               Packets
tcp-connection-policer             0
```

## show interfaces filters

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>show interfaces filters</code><br><code>&lt;interface-name&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Display firewall filters that are configured on each interface in a switch.   |
| <b>Options</b>                  | <b>none</b> —Display firewall filter information about all interfaces.<br><br><b>interface-name</b> —(Optional) Display firewall filter information about a particular interface. |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><a href="#">show firewall on page 5402</a></li> </ul>  |
| <b>List of Sample Output</b>    | <a href="#">show interfaces filters on page 5408</a><br><a href="#">show interfaces filters interface-name on page 5409</a>   |
| <b>Output Fields</b>            | Table 444 on page 5408 lists the output fields for the <b>show interfaces filters</b> command. Output fields are listed in the approximate order in which they appear.            |

Table 444: show interfaces filters Output Fields

| Field Name           | Field Description  | Level of Output |
|----------------------|--|-----------------|
| <b>Interface</b>     | Name of the physical interface.  | All levels      |
| <b>Admin</b>         | Interface state: <b>up</b> or <b>down</b> .  | All levels      |
| <b>Link</b>          | Link state: <b>up</b> or <b>down</b> .   | All levels      |
| <b>Proto</b>         | Protocol that is configured on the interface.  | All levels      |
| <b>Input Filter</b>  | Name of the firewall filter to be evaluated when packets are received on the interface.    | All levels      |
| <b>Output Filter</b> | Name of the firewall filter to be evaluated when packets are transmitted on the interface. | All levels      |

## Sample Output

### show interfaces filters

```

user@switch> show interfaces filters
Interface      Admin Link Proto Input Filter      Output Filter
ge-0/0/6       up   up   eth-switch ingress-port-limit-tcp-icmp
ge-0/0/6.0     up   up   eth-switch ingress-port-limit-tcp-icmp
ge-0/0/7       up   down
ge-0/0/8       up   down

```

|             |    |      |
|-------------|----|------|
| ge-0/0/9    | up | down |
| ge-0/0/10   | up | down |
| ge-0/0/10.0 | up | down |

#### show interfaces filters interface-name

```
user@switch> show interfaces filters ge-0/0/6
```

| Interface  | Admin | Link | Proto      | Input Filter                | Output Filter |
|------------|-------|------|------------|-----------------------------|---------------|
| ge-0/0/6   | up    | up   |            |                             |               |
| ge-0/0/6.0 | up    | up   | eth-switch | ingress-port-limit-tcp-icmp |               |



## CHAPTER 63

# Troubleshooting

- [Troubleshooting Procedures on page 5411](#)

## Troubleshooting Procedures

---

- [Troubleshooting Firewall Filter Configuration on page 5411](#)
- [Troubleshooting Policer Configuration on page 5418](#)

## Troubleshooting Firewall Filter Configuration

Use the following information to troubleshoot your firewall filter configuration.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 5411](#)
- [Filter Counts Previously Dropped Packet on page 5413](#)
- [Matching Packets Not Counted on page 5414](#)
- [Counter Reset When Editing Filter on page 5414](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 5414](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 5415](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 5415](#)
- [Egress Firewall Filters with Private VLANs on page 5415](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 5416](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 5416](#)
- [Invalid Statistics for Policer on page 5416](#)
- [Policers can Limit Egress Filters on page 5416](#)

### Firewall Filter Configuration Returns a No Space Available in TCAM Message

**Problem**    **Description:** When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

```
No space available in tcam.  
Rules for filter filter-name will not be installed.
```

A switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of space available

in the TCAM table. The filter is not applied, but the commit operation for the firewall filter configuration is completed in the CLI module.

**Solution** When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter ingress-vlan-rogue-block
user@switch# delete vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# delete vlans employee-vlan filter input ingress-vlan-rogue-block
```

2. Commit the changes:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space. For example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block ...
```

4. Apply (bind) the new firewall filter to a port, VLAN, or Layer 3 interface. For example:

```
[edit]
user@switch# set vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

5. Commit the changes:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the binding of the original filter—for example:

```
[edit]
user@switch# set vlans employee-vlan description "smaller filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

```
[edit]
user@switch# commit
```



**NOTE:** The original filter is not deleted and is still available in the configuration.

### Filter Counts Previously Dropped Packet

- Problem**    **Description:** If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:
- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
  - You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

**Solution**    This is expected behavior.

### Matching Packets Not Counted

---

**Problem** **Description:** If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet.  
For example:

- You configure an egress port filter with a counter for interface xe-0/0/1.
- You configure an egress VLAN filter with a counter for the **adminVLAN**, and interface xe-0/0/1 is a member of that VLAN.
- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

**Solution** This is expected behavior.

### Counter Reset When Editing Filter

---

**Problem** **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

### Cannot Include loss-priority and policer Actions in Same Term

---

**Problem** **Description:** You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: "cannot support policer action if loss-priority is configured."

**Solution** This is expected behavior.



### Cannot Egress Filter Certain Traffic Originating on QFX Switch

**Problem** **Description:** On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

**Solution** This is expected behavior.

### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

**Problem** **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same Ethertype.

### Egress Firewall Filters with Private VLANs

**Problem** **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

#### Egress Filtering of L2PT Traffic Not Supported

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

---

#### Cannot Drop BGP Packets in Certain Circumstances

**Problem** **Description:** BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

**Solution** This is expected behavior.

---

#### Invalid Statistics for Policer

**Problem** **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

---

#### Policers can Limit Egress Filters

**Problem** **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume

two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

**Related  
Documentation**

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5546](#)
- [Configuring Firewall Filters on page 5290](#)
- [Verifying That Firewall Filters Are Operational on page 5382](#)

## Troubleshooting Policer Configuration

- [Incomplete Count of Packet Drops on page 5418](#)
- [Counter Reset When Editing Filter on page 5418](#)
- [Invalid Statistics for Policer on page 5418](#)
- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 5419](#)
- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 5420](#)
- [Policers Can Limit Egress Filters on page 5420](#)

---

### Incomplete Count of Packet Drops

**Problem**    **Description:** Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

**Solution**    This is expected behavior.

---

### Counter Reset When Editing Filter

**Problem**    **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution**    This is expected behavior.

---

### Invalid Statistics for Policer

**Problem**    **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

### Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem** **Description:** If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate might be twice the configured rate, depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/12 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

**Solution** This is expected behavior.

### Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

---

**Problem** **Description:** You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

**Solution** To prevent this unexpected behavior, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5236](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

### Policers Can Limit Egress Filters

---

**Problem** **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters. Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.

- Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.





## PART 18

# Services

- [Overview on page 5425](#)
- [Configuration on page 5433](#)
- [Administration on page 5497](#)
- [Troubleshooting on page 5501](#)



## CHAPTER 64

# Overview

- [Port Mirroring on page 5425](#)
- [DHCP Relay on page 5431](#)

## Port Mirroring

---

- [Understanding Port Mirroring on page 5425](#)
- [Understanding Layer 3 Logical Interfaces on page 5430](#)

## Understanding Port Mirroring

- [Port Mirroring Overview on page 5425](#)
- [Port Mirroring Instance Types on page 5426](#)
- [Port-Mirroring Terminology on page 5426](#)
- [Port Mirroring and STP on page 5428](#)
- [Port Mirroring Constraints and Limitations on page 5428](#)

### Port Mirroring Overview

---

Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring or to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is needed for traffic analysis on a switch because a switch normally sends packets only to the port to which the destination device is connected. You configure port mirroring on the switch to send copies of unicast traffic to a local interface or a VLAN and run an analyzer application on a device connected to the interface or VLAN. You configure port mirroring by using the **analyzer** statement.

Keep performance in mind when configuring port mirroring. For example, If you mirror traffic from multiple ports, the mirrored traffic may exceed the capacity of the output interface. We recommend that you limit the amount of copied traffic by selecting specific interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter to send specific traffic to a port mirroring instance. Mirroring only the necessary packets reduces the possibility of a performance impact.

You can use port mirroring to copy any of the following:

- All packets entering or exiting an interface (in any combination)—For example, you can send copies of the packets entering some interfaces and the packets exiting other interfaces to the same local interface or VLAN. If you configure port mirroring to copy packets exiting an interface, traffic that originates on that switch or Node device (in a QFabric system) is not copied when it egresses. Only switched traffic is copied on egress. (See the limitation on egress mirroring below.)
- All packets entering a VLAN—You cannot use port mirroring to copy packets exiting a VLAN.
- Firewall-filtered sample—Sample of packets entering a port or VLAN. Configure a firewall filter to select certain packets for mirroring.



**NOTE:** Firewall filters are not supported on egress ports; therefore, you cannot specify policy-based sampling of packets exiting an interface.

### Port Mirroring Instance Types

To configure port mirroring, you configure an instance of one of the following types:

- Analyzer instance: You must specify the input and output for the instance. This instance type is useful for ensuring that all traffic transiting an interface or VLAN is mirrored and sent to the analyzer device.
- Port-mirroring instance: You do not specify an input for this instance type. Instead, you create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored. When you use a port-mirroring instance, you can direct traffic to it in the following ways:
  - Specify the name of the port-mirroring instance in the firewall filter using the **port-mirror-instance *instance-name*** action. You should use this approach if there are multiple port-mirroring instances defined.
  - Configure the filter to send the mirrored packets to the output interface defined in the instance using the **port-mirror** action. You can use this approach if there is only one port-mirroring instance defined.

### Port-Mirroring Terminology

Table 445 on page 5426 lists the terms used in the documentation about port mirroring and provides definitions.

**Table 445: Port Mirroring Terms and Definitions**

| Term              | Description  |
|-------------------|--|
| Analyzer instance | Port-mirroring configuration that includes a name, source interfaces or source VLAN, and a destination for mirrored packets (either a local access interface or a VLAN). |

Table 445: Port Mirroring Terms and Definitions (*continued*)

|   |  |
|---|--|
| Port mirroring instance   | A port-mirroring configuration that does not specify an input.. A firewall filter must be used to send traffic to the port mirror. Use the action <b>port-mirror-instance <i>instance-name</i></b> in the firewall filter configuration to send packets to the port mirror.  |
| Output interface (also known as monitor interface)              | <p>Access interface to which packet copies are sent and to which a device running an analyzer application is connected.</p> <p>The following limitations apply to an output interface:</p> <ul style="list-style-type: none"> <li>• Cannot also be a source port.</li> <li>• Cannot be used for switching.</li> <li>• Cannot be an aggregated Ethernet interface (LAG).</li> <li>• Does not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP).</li> <li>• Loses any existing VLAN associations when you configure it as an analyzer output interface.</li> </ul> <p>If the capacity of the output interface is insufficient to handle the traffic from the source ports, overflow packets are dropped.</p>  |
| Output IP address   | <p>IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)</p> <ul style="list-style-type: none"> <li>• An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</li> <li>• If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).</li> </ul> |
| Output VLAN (also known as monitor or analyzer VLAN)            | <p>VLAN to which copies are sent and to which a device running an analyzer application is connected. The analyzer VLAN can span multiple switches.</p> <p>The following limitations apply to an output VLAN:</p> <ul style="list-style-type: none"> <li>• Cannot be a private VLAN or VLAN range.</li> <li>• Cannot be shared by multiple <b>analyzer</b> statements.</li> <li>• An output VLAN interface cannot be a member of any other VLAN.</li> <li>• An output VLAN interface cannot be an aggregated Ethernet interface (LAG).</li> <li>• On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.</li> </ul>   |
| Input interface (also known as mirrored or monitored interface) | Interface that provides traffic to be mirrored. This traffic can be entering or exiting the interface. (Ingress or egress traffic can be mirrored.) An input interface cannot also be an output interface for an analyzer.   |
| Monitoring station  | Computer running an analyzer application.  |
| Local port mirroring  | Port-mirroring configuration in which the mirrored packets are sent to an interface on the same switch.  |
| Remote port mirroring   | Flooding mirrored packets to an analyzer VLAN that you create to receive mirror traffic or sending the mirrored packets to a remote IP address. (You cannot send mirrored packets to a remote IP address on a QFabric system.)   |

Table 445: Port Mirroring Terms and Definitions (*continued*)

|                        |  |
|------------------------|--|
| Policy-based mirroring | Mirroring of packets that match the match a firewall filter term. The action <b>analyzer analyzer-name</b> is used in the firewall filter to send the packets to the analyzer. |
|------------------------|--|

### Port Mirroring and STP

The behavior of STP in a port-mirroring configuration depends on the version of Junos OS you are using:

- Junos OS 13.2X50, Junos OS 13.2X51-D25 or earlier, Junos OS 13.2X52: If you enable STP, port mirroring might not work because STP might block the mirrored packets.
- Junos OS 13.2X51-D30, Junos OS 14.1X53: STP is disabled for mirrored traffic. You must ensure that your topology prevents loops for this traffic.

### Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 5428](#)
- [Remote Port Mirroring Only on page 5430](#)

#### Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**

- **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### ***Remote Port Mirroring Only***

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

#### **Related Documentation**

- *Configuring Port Mirroring*
- *Example: Configuring Port Mirroring for Local Analysis*
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 5438](#)
- [Troubleshooting Port Mirroring on page 5501](#)

## **Understanding Layer 3 Logical Interfaces**

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks switch to a Layer 2 switch. Only one physical connection is required between the switches.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series and EX4600 switches support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.



- Related Documentation**
- [Interfaces Overview on page 2389](#)
  - [Configuring a Layer 3 Logical Interface on page 2593](#)
  - [Configuring DHCP and BOOTP Relay](#)
  - [Junos OS Network Interfaces Library for Routing Devices](#)

## DHCP Relay

---

- [DHCP and BOOTP Relay Overview on page 5431](#)

### DHCP and BOOTP Relay Overview

You can configure a Juniper Networks switch to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that if the switch receives a broadcast DHCP or BOOTP request from a locally attached host (client), it relays the message to a specified DHCP or BOOTP server. You should configure the switch to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.



**NOTE:** Because DHCP and BOOTP messages are broadcast and are not directed to a specific server, switch, or router, Juniper switches cannot function as both a DHCP server and a DHCP/BOOTP relay agent at the same time. The Junos operating system (Junos OS) generates a commit error if both options are configured at the same time, and the commit operation does not succeed until one of the options is removed.

- Related Documentation**
- [Configuring DHCP and BOOTP Relay](#)
  - [bootp](#)



# Configuration

- [Configuration Examples on page 5433](#)
- [Configuration Tasks on page 5446](#)
- [Configuration Statements for Port Mirroring on page 5450](#)
- [Configuration Statements for Encryption on page 5464](#)
- [Configuration Statements for DHCP on page 5484](#)

## Configuration Examples

---

- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 5438](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 5443](#)

### Examples: Configuring Port Mirroring for Local Analysis

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies to a local interface for local monitoring.



**NOTE:** This example uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Port Mirroring for Local Analysis*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

This example describes how to configure port mirroring to copy traffic sent by employee computers to a switch to an access interface on the same switch.

- [Requirements on page 5434](#)
- [Overview and Topology on page 5434](#)
- [Example: Mirroring All Employee Traffic for Local Analysis on page 5434](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 5435](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2
- A switch

## Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering interfaces on the switch to an access interface on the same switch. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

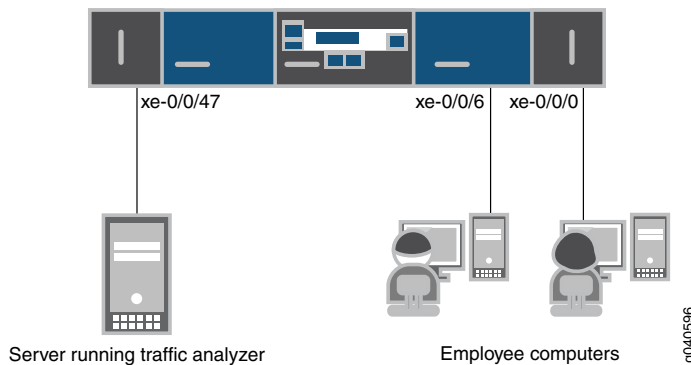
In this example, **xe-0/0/0** and **xe-0/0/6** serve as connections for employee computers. Interface **xe-0/0/47** is connected to a device running an analyzer application.



**NOTE:** Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

Figure 189 on page 5434 shows the network topology for this example.

Figure 189: Network Topology for Local Port Mirroring Example



### Example: Mirroring All Employee Traffic for Local Analysis

To configure port mirroring for all traffic sent by employee computers for local analysis, perform the tasks explained in this section.

#### CLI Quick Configuration

To quickly configure local port mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into a switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching
set interfaces xe-0/0/6 unit 0 family ethernet-switching
set interfaces xe-0/0/47 unit 0 family ethernet-switching
set forwarding-options analyzer employee-monitor input ingress interface xe-0/0/0.0
```

```
set forwarding-options analyzer employee-monitor input ingress interface xe-0/0/6.0
set forwarding-options analyzer employee-monitor output interface xe-0/0/47.0
```

### Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the output interface:

1. Configure the interfaces connected to employee computers as input interfaces for the port-mirror analyzer **employee-monitor**:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/6.0
```

2. Configure the output analyzer interface for the **employee-monitor** analyzer. This will be the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor output interface xe-0/0/47.0
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show forwarding-options analyzer
employee-monitor {
  input {
    ingress {
      interface xe-0/0/0.0;
      interface xe-0/0/6.0;
    }
  }
  output {
    interface {
      xe-0/0/47.0;
    }
  }
}
```

### Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 5435](#)
- [Overview on page 5435](#)
- [Configuring on page 5436](#)
- [Verification on page 5438](#)

#### Requirements

This example uses the following hardware and software components:

- One switch
- Junos 13.2X51

#### Overview

Rather than mirror all traffic, it is usually desirable to mirror only certain traffic. This is a more-efficient use of your bandwidth and hardware and might be necessary because

constraints on these assets. To select specific traffic for mirroring, you use a firewall filter to match the desired traffic and direct it to a port-mirroring instance. The port-mirroring instance then copies the packets and sends them to the output VLAN, interface, or IP address.

### **Configuring**

To specify that the only traffic that will be mirrored is traffic sent by employees to the Web, perform the tasks explained in this section. To select this traffic for mirroring, you use a firewall filter to specify this traffic and direct it to a port-mirroring instance.

#### **CLI Quick Configuration**

To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:

```
[edit]
set forwarding-options port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor
set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

#### **Step-by-Step Procedure**

To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the output interface:

```
[edit interfaces]
user@switch# set xe-0/0/47 unit 0 family ethernet-switching
```

2. Configure the **employee-web-monitor** output interface. (Configure only the output—the input comes from the filter.)

```
[edit forwarding-options]
user@switch# set port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
```

3. Configure a firewall filter called **watch-employee** that includes a term to match traffic sent to the Web and send it to the port-mirroring instance **employee-web-monitor**. Traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the instance:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-corp from destination-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
```

- ```

user@switch# set filter watch-employee term employee-to-web then port-mirror-instance
employee-web-monitor

```
4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers):
- ```

[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee

```

**Results** Check the results of the configuration:

```

[edit]
user@switch# show
forwarding-options {
  port-mirroring {
    instance {
      employee-web-monitor {
        output {
          interface xe-0/0/47.0;
        }
      }
    }
  }
}
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-web {
      from {
        destination-port 80;
      }
      then port-mirror-instance employee-web-monitor;
    }
  }
}
...
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  xe-0/0/6 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}

```

### Verification

#### Verifying That the Analyzer Has Been Correctly Created

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```
user@switch> show forwarding-options analyzer
  Port mirror name           : employee-monitor
  Mirror rate                 : 1
  Maximum packet length      : 0
  State                       : up
  Ingress monitored interfaces : xe-0/0/0.0
  Ingress monitored interfaces : xe-0/0/6.0
  Output interface           : xe-0/0/47.0
```

**Meaning** This output shows that the port-mirroring instance **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (0 indicates the entire packet), the state of the configuration (is up indicates that the instance is mirroring the traffic entering the xe-0/0/0, and xe-0/0/6 interfaces, and sending the mirrored traffic to the xe-0/0/47 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the instance will not be programmed for mirroring.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 5438](#)

### Example: Configuring Port Mirroring for Remote Analysis

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies either to a local interface for local monitoring or to a VLAN for remote monitoring. This example describes how to configure port mirroring for remote analysis.





**NOTE:** This example uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Example: Configuring Port Mirroring for Remote Analysis*. For ELS details, see “Getting Started with Enhanced Layer 2 Software” on page 43.

- Requirements on page 5439
- Overview and Topology on page 5439
- Mirroring All Employee Traffic for Remote Analysis on page 5439
- Mirroring Employee-to-Web Traffic for Remote Analysis on page 5440
- Verification on page 5442

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2 for the QFX Series
- A switch

### Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to an analyzer VLAN so that you can perform analysis using a remote device. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

In this example:

- Interfaces **ge-0/0/0** and **ge-0/0/1** are Layer 2 interfaces that connect to employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects to another switch.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.



**NOTE:** In addition to performing the configuration steps described here, you must also configure the analyzer VLAN (**remote-analyzer** in this example) on the other switches that are used to connect the source switch (the one in this configuration) to the one that the monitoring station is connected to.

### Mirroring All Employee Traffic for Remote Analysis

#### CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your

network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor output vlan remote-analyzer
```

#### Step-by-Step Procedure

To configure basic remote port mirroring:

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):  

```
[edit vlans]
user@switch# set vlans remote-analyzer vlan-id 999
```
2. Configure the interface connected to another switch for trunk mode and associate it with the **remote-analyzer** VLAN:  

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```
3. Configure the **employee-monitor** analyzer:  

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```
4. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
forwarding-options {
  analyzer employee-monitor {
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
    output {
      vlan {
        remote-analyzer;
      }
    }
  }
}
```

---

#### Mirroring Employee-to-Web Traffic for Remote Analysis

#### CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your

network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer employee-web-monitor loss-priority high output vlan 999
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then analyzer
employee-web-monitor
set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

### Step-by-Step Procedure

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):  

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```
2. Configure an interface to associate it with the **remote-analyzer** VLAN:  

```
[edit interfaces]
user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```
3. Configure the **employee-web-monitor** analyzer. (Configure only the output—the input comes from the filter.)  

```
[edit forwarding-options]
user@switch# set forwarding-options analyzer employee-web-monitor output vlan 999
```
4. Configure a firewall filter called **watch-employee** to match traffic sent to the Web and send it to the analyzer **employee-web-monitor**:  

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then analyzer
employee-web-monitor
```
5. Apply the firewall filter to the appropriate interfaces as an ingress filter:  

```
[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```
6. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  ...
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members remote-analyzer;
        }
      }
    }
  }
}
```

```
}
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}
...
firewall {
  family ethernet-switching {
    ...
    filter watch-employee {
      term employee-to-web {
        from {
          destination-port 80;
        }
        then analyzer employee-web-monitor;
      }
    }
  }
}
forwarding-options analyzer {
  employee-web-monitor {
    output {
      vlan {
        999;
      }
    }
  }
}
vpls {
  remote-analyzer {
    vlan-id 999;
  }
}
```

---

## Verification

### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
user@switch> show analyzer
Analyzer name           : employee-monitor
Output VLAN             : remote-analyzer
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```

**Meaning** This output shows that the **employee-monitor** analyzer is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1** and is sending the mirror traffic to the analyzer **remote-analyzer**.

- Related Documentation**
- [Understanding Port Mirroring on page 5425](#)
  - [Configuring Port Mirroring on page 5446](#)
  - [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
  - [Overview of Firewall Filters on page 5209](#)

## Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 5443](#)
- [Overview on page 5443](#)
- [Configuring on page 5443](#)
- [Verification on page 5445](#)

### Requirements

This example uses the following hardware and software components:

- One switch
- Junos 13.2X51

### Overview

Rather than mirror all traffic, it is usually desirable to mirror only certain traffic. This is a more-efficient use of your bandwidth and hardware and might be necessary because constraints on these assets. To select specific traffic for mirroring, you use a firewall filter to match the desired traffic and direct it to a port-mirroring instance. The port-mirroring instance then copies the packets and sends them to the output VLAN, interface, or IP address.

### Configuring

To specify that the only traffic that will be mirrored is traffic sent by employees to the Web, perform the tasks explained in this section. To select this traffic for mirroring, you use a firewall filter to specify this traffic and direct it to a port-mirroring instance.

**CLI Quick Configuration** To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:

```
[edit]
set forwarding-options port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor
set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

**Step-by-Step Procedure** To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the output interface:

```
[edit interfaces]
user@switch# set xe-0/0/47 unit 0 family ethernet-switching
```

2. Configure the **employee-web-monitor** output interface. (Configure only the output—the input comes from the filter.)

```
[edit forwarding-options]
user@switch# set port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
```

3. Configure a firewall filter called **watch-employee** that includes a term to match traffic sent to the Web and send it to the port-mirroring instance **employee-web-monitor**. Traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the instance:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-corp from destination-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then port-mirror-instance
employee-web-monitor
```

4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers):

```
[edit interfaces]
user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
```

```

forwarding-options {
  port-mirroring {
    instance {
      employee-web-monitor {
        output {
          interface xe-0/0/47.0;
        }
      }
    }
  }
}
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-web {
      from {
        destination-port 80;
      }
      then port-mirror-instance employee-web-monitor;
    }
  }
}
...
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  xe-0/0/6 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}

```

## Verification

### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```
user@switch> show forwarding-options analyzer
  Port mirror name           : employee-monitor
  Mirror rate                 : 1
  Maximum packet length      : 0
  State                       : up
  Ingress monitored interfaces : xe-0/0/0.0
  Ingress monitored interfaces : xe-0/0/6.0
  Output interface           : xe-0/0/47.0
```

**Meaning** This output shows that the port-mirroring instance **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (**0** indicates the entire packet), the state of the configuration (is up indicates that the instance is mirroring the traffic entering the xe-0/0/0, and xe-0/0/6 interfaces, and sending the mirrored traffic to the xe-0/0/47 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the instance will not be programmed for mirroring.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)

---

## Configuration Tasks

- [Configuring Port Mirroring on page 5446](#)
- [Configuring DHCP and BOOTP on page 5449](#)

### Configuring Port Mirroring

You use port mirroring to copy packets and send the copies to a device running an application such as a network analyzer or intrusion detection application so that you can analyze traffic without delaying it. You can mirror traffic entering or exiting a port or entering a VLAN, and you can send the copies to a local access interface or to a VLAN through a trunk interface.

We recommend that you disable port mirroring when you are not using it. To avoid creating a performance issue If you do enable port mirroring, we recommend that you select specific input interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter.



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Port Mirroring*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

---





**NOTE:** If you want to create additional analyzers without deleting an existing analyzer, first disable the existing analyzer using the `disable analyzer analyzer-name` command.



**NOTE:** You must configure port mirroring output interfaces as family `ethernet-switching`.

- [Configuring Port Mirroring for Local Analysis on page 5447](#)
- [Configuring Port Mirroring for Remote Analysis on page 5448](#)
- [Filtering the Traffic Entering an Analyzer on page 5448](#)

### Configuring Port Mirroring for Local Analysis

To mirror interface traffic to a local interface on the switch:

1. If you want to mirror traffic that is ingressing or egressing specific interfaces, choose a name for the port-mirroring configuration and configure what traffic should be mirrored by specifying the interfaces and direction of traffic:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some VLAN packets might contain incorrect VLAN IDs.



**NOTE:** If you configure mirroring for packets that egress an access interface, the original packets lose any VLAN tags when they exit the access interface, but the mirrored (copied) packets retain the VLAN tags when they are sent to the analyzer system.

2. If you want to specify that all traffic entering a VLAN should be mirrored, choose a name for the port-mirroring configuration and specify the VLAN:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input ingress vlan vlan-name
```



**NOTE:** You cannot configure port mirroring to copy traffic that egresses a VLAN.

3. Configure the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name output interface interface-name
```

## Configuring Port Mirroring for Remote Analysis

---

To mirror traffic to a VLAN for analysis at a remote location:

1. Configure a VLAN to carry the mirrored traffic:

```
[edit]
user@switch# set vlans vlan-name vlan-id number
```

2. Configure the interface that connects to another switch (the uplink interface) to trunk mode and associate it with the appropriate VLAN:

```
[edit]
user@switch# set interfaces interface-name unit 0 family ethernet-switching port-mode
trunk vlan members (vlan-name | vlan-id)
```

3. Configure the analyzer:

- a. Choose a name for the analyzer:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name
```

- b. Specify the interface to be mirrored and whether the traffic should be mirrored on ingress or egress:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```

- c. Specify the appropriate IP address or VLAN as the output (a VLAN is specified in this example:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name output vlan (vlan-name | vlan-id)
```

If you specify an IP address as the output, note the following constraints:

- The address cannot be in the same subnet as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (**inet.0** routing table).
- The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)

## Filtering the Traffic Entering an Analyzer

---

In addition to specifying which traffic to mirror by configuring an analyzer, you can also use a firewall filter to exercise more control over which packets are copied. For example, you might use a filter to specify that only traffic from certain applications be mirrored. The filter can use any of the available match conditions and must have an action of modifier of **port-mirror-instance** *instance-name*. If you use the same analyzer in multiple filters or terms, the output packets are copied only once.

When you use a firewall filter as the input to a port-mirroring instance, you send the copied traffic to a local interface or a VLAN just as you do when a firewall is not involved.

To configure port mirroring with filters:

1. Configure a port-mirroring instance for local or remote analysis. Configure only the output. For example, for local analysis enter:

```
[edit forwarding-options]
user@switch# set port-mirroring-instance instance-name output interface interface-name
```



**NOTE:** You cannot configure input to this instance.

2. Create a firewall filter using any of the available match conditions. In a **then** term, specify include the action modifier **port-mirror-instance** *instance-name*.
3. Apply the firewall filter to the interfaces or VLAN that should provide the input to the analyzer:

```
[edit]
user@switch# set interfaces interface-name unit 0 family ethernet-switching filter input
filter-name
[edit]
user@switch# set vlan (vlan-name | vlan-id) filter input filter-name
```

#### Related Documentation

- [Understanding Port Mirroring on page 5425](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 5438](#)
- [Overview of Firewall Filters on page 5209](#)

## Configuring DHCP and BOOTP

You can configure a QFX Series switch to act as a Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) server or DHCP relay agent. When a switch is a relay agent, if a locally attached host issues a DHCP or BOOTP request as a broadcast message, the switch relays the message to a specified DHCP or BOOTP server. You should configure a switch to be a DHCP and BOOTP relay agent if you have locally attached hosts and a remote DHCP or BOOTP server.



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring DHCP and BOOTP Relay*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 43.

To configure a switch to be a server, use the **dhcp-local-server** statement. To configure a switch to be a relay agent, use the **dhcp-relay** statement.

If you want to enable BOOTP support when the switch is configured to be a DHCP server, enter the following statement:

```
[edit system services dhcp-local-server]
user@switch# set overrides bootp-support
```

If you want to enable BOOTP support when the switch is configured to be a DHCP relay agent, enter the following statement:

```
[edit forwarding-options dhcp-relay]  
user@switch# set overrides bootp-support
```

---

## Configuration Statements for Port Mirroring

---

- [analyzer on page 5451](#)
- [egress on page 5452](#)
- [ethernet-switching \(Port Mirroring\) on page 5453](#)
- [family \(Port Mirroring\) on page 5454](#)
- [inet \(Port Mirroring\) on page 5455](#)
- [ingress \(Port Mirroring\) on page 5456](#)
- [input on page 5457](#)
- [instance \(Port Mirroring\) on page 5458](#)
- [interface \(Port Mirroring\) on page 5459](#)
- [ip-address \(Port Mirroring\) on page 5460](#)
- [output on page 5461](#)
- [port-mirroring on page 5462](#)
- [routing-instance \(Port Mirroring\) on page 5463](#)
- [vlan \(Port Mirroring\) on page 5464](#)

## analyzer

```
Syntax analyzer {
    name {
        input {
            egress {
                interface (all | interface-name);
            }
            ingress {
                interface (all | interface-name);
                vlan (vlan-id | vlan-name);
            }
        }
        output {
            interface interface-name;
            ip-address ip-address;
            routing-instance
            vlan (vlan-id | vlan-name);
        }
    }
}
```

**Hierarchy Level** For platforms without ELS:

[edit ethernet-switching-options]

For platforms with ELS:

[edit forwarding-options]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
 Option **output vlan** added in Junos OS Release 12.1 for the QFX Series.  
 Option **output ip-address** added in Junos OS Release 12.3 for the QFX Series for non-ELS platforms and added in 14.1X53-D10 for ELS platforms.

**Description** Configure port mirroring. You can create a total of four port-mirroring configurations on the QFX Series, subject to the following limits:

- There can be no more than two configurations that mirror ingress traffic.
- There can be no more than two configurations that mirror egress traffic.

**Default** Port mirroring is disabled, and Junos OS creates no default analyzers.

**Options** **all**—Mirror all the access interfaces. Using this option does not cause the QSFP+ or management interfaces to be mirrored.



**CAUTION:** Configuring the **all** option in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

**name**—Name of the analyzer. The name can include as many as 125 characters; must begin with a letter; and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.

The remaining statements are explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)

---

## egress

**Syntax** egress {  
    **interface** (all | *interface-name*);  
}

**Hierarchy Level** For platforms without ELS:  
  
[edit ethernet-switching-options **analyzer name input**]  
  
For platforms with ELS:  
  
[edit forwarding-options **analyzer name input**]

**Release Information** Statement introduced in Junos OS Release 11.2 for the QFX Series.

**Description** Specify interfaces for which egressing traffic is mirrored.

The statement is explained separately.



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some of the mirrored packets might contain incorrect VLAN IDs.

---

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)

---

## ethernet-switching (Port Mirroring)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>ethernet-switching;<br/>  output {<br/>    interface <i>interface-name</i> {<br/>    }<br/>    no-filter-check;<br/>  }<br/>  vlan <i>vlan-name</i> {<br/>    no-tag;<br/>  }<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit forwarding-options port-mirroring [ <i>instance name</i> ] family]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 13.2 for the QFX Series.  |
| <b>Description</b>              | Specify that the output interface for the port mirror will be configured as an <b>ethernet-switching</b> interface.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 5425</a></li><li>• <a href="#">Configuring Port Mirroring on page 5446</a></li><li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 5433</a></li></ul> |

## family (Port Mirroring)

---

**Syntax**    **family (Port Mirroring)**  
              **ethernet-switching** {  
                  **output** {  
                      **interface** *interface-name* {  
                          }  
                      no-filter-check;  
                      }  
                      **vlan** *vlan-name* {  
                          no-tag;  
                      }  
                  }  
              **inet**  
                  **output** {  
                      **ip-address** *address* {  
                          }  
                      **routing-instance** *instance-name* {  
                          **ip-address** *address* {  
                              }  
                      }  
                  }  
              }

**Hierarchy Level**    [edit forwarding-options port-mirroring [*instance name*] ]

**Release Information**    Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description**    Specify the type of interface that will be used to forward port mirrored packet to an analyzer device..

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)



## inet (Port Mirroring)

```
Syntax  inet {
        output {
            ip-address address {
            }
            routing-instance instance-name {
                ip-address address {
                }
            }
        }
    }
```

**Hierarchy Level** [edit forwarding-options port-mirroring [instance *name*] family]

**Release Information** Statement introduced in Junos OS Release 14.1X53 for the QFX Series.

**Description** Specify that the output interface will be of type **inet**. Use this statement so that you can send the mirrored packets to the IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)



**NOTE:** An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.



**NOTE:** If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)

## ingress (Port Mirroring)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>ingress {<br/>    interface (all   interface-name);<br/>    vlan (vlan-id   vlan-name);<br/>}</pre>   |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options analyzer name input]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options analyzer name input]</pre>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>Specify the interfaces or VLANs for which incoming traffic is mirrored as part of a port mirroring configuration.</p> <p>The statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 5425</a></li><li>• <a href="#">Configuring Port Mirroring on page 5446</a></li><li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 5433</a></li></ul> |

## input

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> input {   ingress {     interface (all   <i>interface-name</i>);     vlan (<i>vlan-id</i>   <i>vlan-name</i>);   }   egress {     interface (all   <i>interface-name</i>);   } } </pre>  |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <p>[edit ethernet-switching-options <a href="#">analyzer name</a>]</p> <p>For platforms with ELS:</p> <p>[edit forwarding-options <a href="#">analyzer name</a>]</p>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | <p>Define the traffic to be mirrored. The definition can be a combination of traffic entering or exiting specific ports or VLANs.</p> <p>The statements are explained separately.</p>  |
| <b>Default</b>                  | No default.  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 5425</a></li> <li>• <a href="#">Configuring Port Mirroring on page 5446</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 5433</a></li> </ul> |

## instance (Port Mirroring)

---

**Syntax**    `instance instance-name{  
              family (Port Mirroring)  
              ethernet-switching {  
                  output {  
                      interface interface-name {  
                          }  
                      no-filter-check;  
                      }  
                      vlan vlan-name {  
                          no-tag;  
                      }  
                  }  
              inet  
                  output {  
                      ip-address address {  
                          }  
                      routing-instance instance-name {  
                          ip-address address {  
                              }  
                          }  
                  }  
              }  
          }`

**Hierarchy Level**    [edit forwarding-options port-mirroring]

**Release Information**    Statement introduced in Junos OS Release 13.2 for the QFX Series.


**Description**    Specify a port-mirroring configuration (instance). You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This is useful for controlling which types of traffic should be mirrored.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**



- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 5435](#)

## interface (Port Mirroring)

|   |  |
|---|--|
| <b>Syntax</b>   | interface (all   <i>interface-name</i> );  |
| <b>Hierarchy Level</b>  | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <i>analyzer name</i> input (egress   ingress)], [edit ethernet-switching-options <i>analyzer name</i> output]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options <i>analyzer name</i> input (egress   ingress)] [edit forwarding-options <i>analyzer name</i> output] [edit forwarding-options port-mirroring[instance <i>name</i>] family ethernet-switching <i>output</i>]</pre> |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>  | Specify the interfaces for which ingressing traffic is mirrored. Specify the interface that mirrored traffic should be copied to (the output interface).   |
| <b>Options</b>  | <p>all—Apply port mirroring to all interfaces on the switch (except the output interface). Mirroring a high volume of traffic can cause performance issues, so you should generally select specific input interfaces.</p>  |
| <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>CAUTION:</b> Configuring <i>all</i> in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.</p> </div> </div> |  |
| <p><i>interface-name</i>—Apply port mirroring to the specified interface only.</p>  |  |
| <b>Required Privilege Level</b>   | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 5425</a></li> <li>• <a href="#">Configuring Port Mirroring on page 5446</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 5433</a></li> </ul>   |

## ip-address (Port Mirroring)

---

|   |   |
|---|---|
| <b>Syntax</b>   | <code>ip-address <i>ip-address</i>;</code>  |
| <b>Hierarchy Level</b>  | <code>[edit forwarding-options] <i>analyzer name</i> <i>output</i>]</code><br><code>[edit forwarding-options port-mirroring [<i>instance name</i>] family ethernet-switching <i>output interface name</i>]</code>   |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 14.1X53 for the QFX Series.  |
| <b>Description</b>  | Specify the IP address to which traffic should be mirrored (the IP address of the analyzer system). The device can be on a remote network. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.) This statement is not supported on QFabric systems. |
| <div><b>NOTE:</b> An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</div>  |   |
| <div><b>NOTE:</b> If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).</div> |   |
| <b>Required Privilege Level</b>   | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 5425</a></li><li>• <a href="#">Configuring Port Mirroring on page 5446</a></li></ul>   |

## output

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>output {   interface <i>interface-name</i>;   ip-address <i>ip-address</i>;   vlan (<i>vlan-id</i>   <i>vlan-name</i>);   routing-instance <i>instance-name</i> {     ip-address <i>address</i> { </pre>   |
| <b>Hierarchy Level</b>          | <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <i>analyzer name</i>]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options <i>analyzer name</i>] [edit forwarding-options port-mirroring [<i>instance name</i>] family ethernet-switching ]</pre> |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure the destination for mirrored traffic, either an interface on the switch (for local monitoring) or a VLAN (for remote monitoring).</p> <p>The statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 5425</a></li> <li>• <a href="#">Configuring Port Mirroring on page 5446</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 5433</a></li> </ul>        |

## port-mirroring

---

```
Syntax  port-mirroring {  
        family {  
            ethernet-switching  
            output {  
                interface interface-name {  
                }  
                no-filter-check;  
            }  
            vlan vlan-name {  
                no-tag;  
            }  
        }  
        inet  
        output {  
            ip-address address {  
            }  
            routing-instance instance-name {  
                ip-address address {  
                }  
            }  
        }  
    }  
    instance instance-name {  
        family (Port Mirroring)  
        ethernet-switching {  
            output {  
                interface interface-name {  
                }  
                no-filter-check;  
            }  
            vlan vlan-name {  
                no-tag;  
            }  
        }  
        inet  
        output {  
            ip-address address {  
            }  
            routing-instance instance-name {  
                ip-address address {  
                }  
            }  
        }  
    }  
}
```

**Hierarchy Level** [edit forwarding-options ]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Create a port-mirroring configuration.



**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 5435](#)

## routing-instance (Port Mirroring)

**Syntax** routing-instance *instance-name*;

**Hierarchy Level** [edit forwarding-options] [analyzer name output](#)  
[edit forwarding-options port-mirroring [instance *name*] family ethernet-switching [output interface name](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Configure a port mirroring instance. You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- [Configuring Port Mirroring on page 5446](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 5433](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 5435](#)

## vlan (Port Mirroring)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>vlan (<i>vlan-id</i>   <i>vlan-name</i>);</code>  |
| <b>Hierarchy Level</b>          | For platforms without ELS:<br><br>[edit ethernet-switching-options <i>analyzer name</i> input ingress],<br>[edit ethernet-switching-options <i>analyzer name</i> output]<br><br>For platforms with ELS:<br><br>[edit forwarding-options <i>analyzer name</i> input (egress   ingress)]<br>[edit forwarding-options <i>analyzer name</i> output]<br>[edit forwarding-options port-mirroring[instance <i>name</i> ] family ethernet-switching output] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Option <b>output</b> <b>vlan</b> added in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Specify that traffic entering into a VLAN should be mirrored. Configure mirrored traffic to be sent to a VLAN for remote monitoring (output).   |
| <b>Options</b>                  | <i>vlan-id</i> —Numeric VLAN identifier.<br><br><i>vlan-name</i> —Name of the VLAN.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 5425</a></li><li>• <a href="#">Configuring Port Mirroring on page 5446</a></li><li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 5433</a></li></ul>  |

## Configuration Statements for Encryption

---

- [authentication-key-chains on page 5466](#)
- [cache-size on page 5467](#)
- [cache-timeout-negative on page 5468](#)
- [ca-name on page 5468](#)
- [certificates on page 5469](#)
- [certification-authority on page 5470](#)
- [crl \(Encryption Interface\) on page 5470](#)
- [encoding on page 5471](#)
- [enrollment-retry on page 5471](#)
- [enrollment-url on page 5472](#)
- [file on page 5472](#)

- [key \(Authentication Keychain\) on page 5473](#)
- [key-chain \(Security\) on page 5474](#)
- [ldap-url on page 5475](#)
- [local on page 5476](#)
- [maximum-certificates on page 5477](#)
- [path-length on page 5477](#)
- [secret on page 5478](#)
- [security on page 5479](#)
- [ssh-known-hosts on page 5480](#)
- [start-time \(Authentication Key Transmission\) on page 5481](#)
- [traceoptions on page 5483](#)

## authentication-key-chains

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>authentication-key-chains {<br/>  key-chain <i>key-chain-name</i> {<br/>    description <i>text-string</i>;<br/>    key <i>key</i> {<br/>      algorithm (md5   hmac-sha-1);<br/>      options (basic   isis-enhanced);<br/>      secret <i>secret-data</i>;<br/>      start-time <i>yyyy-mm-dd.hh:mm:ss</i>;<br/>    }<br/>    tolerance <i>seconds</i>;<br/>  }<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit security]   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. When the <b>authentication-key-chains</b> statement is configured at the <b>[edit security]</b> hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the <b>[edit protocols]</b> hierarchy level or with the BFD protocol using the <b>bfd-liveness-detection</b> statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837</a></li></ul>   |

## cache-size

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | cache-size <i>bytes</i> ;  |
| <b>Hierarchy Level</b>     | [edit security <a href="#">certificates</a> ]  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the cache size for digital certificates.  |
| <b>Options</b>             | <b>bytes</b> —Cache size for digital certificates.<br><b>Range:</b> 64 through 4,294,967,295<br><b>Default:</b> 2 megabytes (MB)   |



**NOTE:** We recommend that you limit your cache size to 4 MB.

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Digital Certificates for an ES PIC</i></li> </ul>        |

## cache-timeout-negative

---

|                            |  |
|----------------------------|--|
| <b>Syntax</b>              | cache-timeout-negative <i>seconds</i> ;  |
| <b>Hierarchy Level</b>     | [edit security <a href="#">certificates</a> ]  |
| <b>Release Information</b> | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>         | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure a negative cache for digital certificates.  |
| <b>Options</b>             | <b>seconds</b> —Negative time to cache digital certificates, in seconds.<br><b>Range:</b> 10 through 4,294,967,295<br><b>Default:</b> 20   |



**CAUTION:** Configuring a large negative cache value can lead to a denial-of-service attack.

---

|                                 |  |
|---------------------------------|--|
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>        |

## ca-name

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | ca-name <i>ca-identity</i> ;   |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the certificate authority (CA) identity to use in the certificate request.                      |
| <b>Options</b>                  | <b>ca-identity</b> —CA identity to use in the certificate request.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>  |

## certificates

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> certificates {   cache-size bytes;   cache-timeout-negative seconds;   certification-authority ca-profile-name {     ca-name ca-identity;     crt file-name;     encoding (binary   pem);     enrollment-url url-name;     file certificate-filename;     ldap-url url-name;   }   enrollment-retry attempts;   local certificate-name {     certificate-key-string;     load-key-file URL filename;   }   maximum-certificates number;   path-length certificate-path-length; } </pre> |
| <b>Hierarchy Level</b>          | [edit security]   |
| <b>Release Information</b>      | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>   |
| <b>Description</b>              | <p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure the digital certificates for IPsec.</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li><i>Configuring Digital Certificates for an ES PIC</i></li> </ul>   |

## certification-authority

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>certification-authority <i>ca-profile-name</i> {<br/>    <i>ca-name</i> <i>ca-identity</i>;<br/>    <i>crl</i> <i>file-name</i>;<br/>    <i>encoding</i> (binary   pem);<br/>    <i>enrollment-url</i> <i>url-name</i>;<br/>    <i>file</i> <i>certificate-filename</i>;<br/>    <i>ldap-url</i> <i>url-name</i>;<br/>}</code> |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure a certificate authority profile name.<br><br>The remaining statements are explained separately.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>  |

## crl (Encryption Interface)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>crl <i>file-name</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis. |
| <b>Options</b>                  | <i>file-name</i> —Specify the file from which to read the CRL.  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>   |



## encoding

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | encoding (binary   pem);   |
| <b>Hierarchy Level</b>          | [edit security ike policy <i>ike-peer-address</i> ],<br>[edit security certificates <b>certification-authority</b> <i>ca-profile-name</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the file format used for the <b>local-certificate</b> and <b>local-key-pair</b> statements.     |
| <b>Options</b>                  | <b>binary</b> —Binary file format.<br><br><b>pem</b> —Privacy-enhanced mail (PEM), an ASCII base 64 encoded format.<br><b>Default:</b> binary  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> <li>• <i>Configuring an IKE Policy for Digital Certificates for an ES PIC</i></li> </ul>   |

## enrollment-retry

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | enrollment-retry <i>attempts</i> ;   |
| <b>Hierarchy Level</b>          | [edit security <b>certificates</b> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.  |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify how many times a router or switch can resend a digital certificate request. |
| <b>Options</b>                  | <b>attempts</b> —Number of enrollment retries.<br><b>Range:</b> 0 through 100<br><b>Default:</b> 0   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>  |

## enrollment-url

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>enrollment-url <i>url-name</i>;</code>  |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL). |
| <b>Options</b>                  | <i>url-name</i> —Certificate authority URL.   |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>   |

## file

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>file <i>certificate-filename</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Specify the file from which to read the digital certificate.  |
| <b>Options</b>                  | <i>certificate-filename</i> —File from which to read the digital certificate.  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>  |

## key (Authentication Keychain)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>key key {   algorithm (md5   hmac-sha-1);   options (basic   isis-enhanced);   secret secret-data;   start-time yyyy-mm-dd.hh:mm:ss; }</pre>   |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> ]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> |
| <b>Description</b>              | Configure the authentication element.   |
| <b>Options</b>                  | <p><b>key</b>—Each key within a keychain is identified by a unique integer value.</p> <p><b>Range:</b> 0 through 63</p> <p>The remaining statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li> <li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li> <li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837</a></li> </ul>  |

## key-chain (Security)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>keychain <i>key-chain-name</i> {<br/>  description <i>text-string</i>;<br/>  key <i>key</i> {<br/>    algorithm (md5   hmac-sha-1);<br/>    options (basic   isis-enhanced);<br/>    secret <i>secret-data</i>;<br/>    start-time <i>yyyy-mm-dd.hh:mm:ss</i>;<br/>  }<br/>  tolerance <i>seconds</i>;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> |
| <b>Description</b>              | Create the key-chain configuration for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.   |
| <b>Options</b>                  | <b><i>key-chain-name</i></b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">authentication-key-chains on page 5466</a></li><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837</a></li></ul>            |

---

## ldap-url

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <ldap-url <i>url-name</i> >;   |
| <b>Hierarchy Level</b>          | [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]   |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>(Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.       |
| <b>Options</b>                  | <i>url-name</i> —Name of the LDAP URL.   |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>  |

## local

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>local <i>certificate-name</i> {<br/>    <i>certificate-key-string</i>;<br/>    load-key-file <i>URL filename</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Import a paired X.509 private key and authentication certificate, to enable Junos XML protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.  |
| <b>Options</b>                  | <p><b><i>certificate-namecertificate-key-string</i></b>—String of alphanumeric characters that constitute the private key and certificate.</p> <p><b><i>certificate-name</i></b>—Name that uniquely identifies the certificate.</p> <p><b><i>load-key-file URL filename</i></b>—File that contains the private key and certificate. It can be one of two types of values:</p> <ul style="list-style-type: none"><li>• Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)</li><li>• URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)</li></ul> |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <i>Importing SSL Certificates for Junos XML Protocol Support</i></li></ul>   |

## maximum-certificates

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>maximum-certificates <i>number</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the maximum number of peer digital certificates to be cached.                                 |
| <b>Options</b>                  | <b><i>number</i></b> —Maximum number of peer digital certificates to be cached.<br><b>Range:</b> 64 through 4,294,967,295 peer certificates<br><b>Default:</b> 1024 peer certificates          |
| <b>Required Privilege Level</b> | system—To view this statement in the configuration.<br>system-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>  |

## path-length

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>path-length <i>certificate-path-length</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit security <a href="#">certificates</a> ]  |
| <b>Release Information</b>      | Statement introduced before Junos OS Release 7.4.<br>Statement introduced in Junos OS Release 9.0 for EX Series switches.<br>Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | (Encryption interface on M Series and T Series routers and EX Series switches only)<br>Configure the digital certificate path length.  |
| <b>Options</b>                  | <b><i>certificate-path-length</i></b> —Digital certificate path length.<br><b>Range:</b> 2 through 15 certificates<br><b>Default:</b> 15 certificates  |
| <b>Required Privilege Level</b> | admin—To view this statement in the configuration.<br>admin-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>  |

## secret

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>secret <i>secret-data</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> |
| <b>Description</b>              | Specify a password in encrypted text or plain text format. The secret password always appears in encrypted format.  |
| <b>Options</b>                  | <b><i>secret-data</i></b> —Password to use; it can include spaces if the character string is enclosed in quotation marks.   |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Example: Configuring BFD Authentication for Static Routes on page 2931</a></li><li>• <a href="#">Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837</a></li></ul>  |



## security

```
Syntax  security {
    authentication-key-chains {
        key-chain key-chain-name {
            key key {
                secret secret-data;
                start-time yyyy-mm-dd.hh:mm:ss;
            }
        }
    }
    certificates {
        cache-size bytes;
        cache-timeout-negative seconds;
        certification-authority ca-profile-name {
            ca-name ca-identity;
            crl file-name;
            encoding (binary | pem);
            enrollment-url url-name;
            file certificate-filename;
            ldap-url url-name;
        }
        enrollment-retry attempts;
        local certificate-filename {
            certificate-key-string;
            load-key-file key-file-name;
        }
        maximum-certificates number;
        path-length certificate-path-length;
    }
    ssh-known-hosts {
        host {
            fetch-from-server host-name;
            load-key-file file-name;
        }
    }
    traceoptions {
        file filename <files number> <size size>;
        flag flag;
        level level;
        no-remote-trace
    }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure security services. Most of the configuration statements do not have default values. If you do not specify an identifier for a statement that does not have a default value, you cannot commit the configuration.

Required Privilege  
Level

Related  
Documentation

## ssh-known-hosts

---


|                          |   |
|--------------------------|---|
| Syntax                   | <pre>ssh-known-hosts {<br/>  host <i>host-name</i> {<br/>    fetch-from-server <i>host-name</i>;<br/>    load-key-file <i>file-name</i>;<br/>  }<br/>}</pre>  |
| Hierarchy Level          | [edit security ssh-known-hosts]   |
| Release Information      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| Description              | Configure SSH support for known hosts and for administering SSH host key updates.   |
| Options                  | <p><b>host <i>host-name</i></b>—Hostname of the SSH known host entry. This option has the following suboptions:</p> <ul style="list-style-type: none"><li>• <b>fetch-from-server <i>host-name</i></b>—Retrieve SSH public host key information from a specified server.</li><li>• <b>load-key-file <i>filename</i></b>—Import SSH host key information from the <code>/var/tmp/ssh-known-hosts</code> file.</li></ul> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>  |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">Understanding Security Features on the QFabric System</a></li><li>• <a href="#">Configuring SSH Host Keys for Secure Copying of Data on page 1359</a></li></ul>   |

## start-time (Authentication Key Transmission)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>start-time (now   yyyy-mm-dd.hh:mm:ss);</code>   |
| <b>Hierarchy Level</b>          | [edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p>  |
| <b>Description</b>              | <p>Specify a start time for key transmission. You do not need to specify an end time for the key. If a new key is present with a new start time, the keychain rolls over to the new one. The start time must be unique within the keychain.</p>  |
| <b>Options</b>                  | <p><b>now</b>—Start time as the current year, month, day, hour, minute, and second.</p> <p><b>daydays</b>—Start time as the specified number of days after the current day. For example, if the current day is the 12th and you configure <b>start-time 2day</b>, the start time will be on the 14th, exactly two days after the configuration is entered.</p> <p><b>hourhours</b>—Start time as the specified number of hours after the current hour. For example, if the current hour is 9:00 and you configure <b>start-time 3hour</b>, the start time will be in 12:00, exactly three hours after the configuration is entered.</p> <p><b>minuteminutes</b>—Start time as the specified number of minutes after the current minute. For example, if the current minute is 27 minutes after the hour and you configure <b>start-time 5min</b>, the start time will be in 32 minutes after the hour, exactly five minutes after the configuration is entered.</p> <p><b>monthmonths</b>—Start time as the specified number of months after the current month. For example, if the current month is March and you configure <b>start-time 4month</b>, the start time will be in July, exactly four months after the configuration is entered.</p> <p><b>secondseconds</b>—Start time as the specified number of seconds after the current second. For example, if the current second is 10:20:40 and you configure <b>start-time 10seconds</b>, the start time will be 10:20:50, exactly 10 seconds after the configuration is entered.</p> <p><b>yearyears</b>—Start time as the specified number of years after the current year. For example, if the current year is 2011 and you configure <b>start-time 1year</b>, the start time will be in 2012, exactly one year after the configuration is entered.</p> <p><b>yyyy-mm-dd.hh:mm:ss</b>—Start time in UTC (Coordinated Universal Time). The start time must be unique within the keychain.</p> |
| <b>Required Privilege Level</b> | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>   |

- Related Documentation**
- *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*
  - [Example: Configuring BFD Authentication for Static Routes on page 2931](#)
  - [Example: Configuring BFD Authentication for Static Routes on page 2931](#)
  - [Example: Configuring Hitless Authentication Key Rollover for IS-IS on page 3837](#)

## traceoptions

|   |  |
|---|--|
| <b>Syntax</b>   | <pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt;;     flag all;     flag certificates;     flag database;     flag general;     flag ike;     flag parse;     flag policy-manager;     flag routing-socket;     flag timer;     level     no-remote-trace } </pre>   |
| <b>Hierarchy Level</b>  | <p>[edit security],<br/>[edit services ipsec-vpn]</p> <p>Trace options can be configured at either the <b>[edit security]</b> or the <b>[edit services ipsec-vpn]</b> hierarchy level, but not at both levels.</p>   |
| <b>Release Information</b>  | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>  |
| <b>Description</b>  | <p>Configure security trace options.</p> <p>To specify more than one trace option, include multiple <b>flag</b> statements. Trace option output is recorded in the <b>/var/log/kmd</b> file.</p>   |
| <div style="display: flex; align-items: center;">  <p><b>NOTE:</b> The <b>traceoptions</b> statement is not supported on QFabric systems.</p> </div> |  |
| <b>Options</b>  | <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file (for example, <b>kmd</b>) reaches its maximum size, it is renamed <b>kmd.0</b>, then <b>kmd.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 0 files</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, <b>kmd</b>) reaches this size, it is renamed, <b>kmd.0</b>, then <b>kmd.1</b> and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Default:</b> 1024 KB</p> |

**flag *flag***—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

- **all**—Trace all security events.
- **certificates**—Trace certificate events.
- **database**—Trace database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **timer**—Trace internal timer events.

**level *level***—(Optional) Set traceoptions level.

- **all**—match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

**no-remote-trace**—(Optional) Disable remote tracing

|                           |   |
|---------------------------|---|
| <b>Required Privilege</b> | admin—To view the configuration.                          |
| <b>Level</b>              | admin-control—To add this statement to the configuration. |

|                              |   |
|------------------------------|---|
| <b>Related Documentation</b> | • <i>Configuring Tracing Operations for Security Services</i> |
|------------------------------|---|

---

## Configuration Statements for DHCP

---

- [dhcp-local-server on page 5485](#)
- [dhcp-relay on page 5490](#)

## dhcp-local-server

```
Syntax  dhcp-local-server {
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dhcpv6 {
            authentication {
                ...
            }
            group group-name {
                authentication {
                    ...
                }
                interface interface-name {
                    exclude;
                    liveness-detection {
                        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                        method {
                            bfd {
                                version (0 | 1 | automatic);
                                minimum-interval milliseconds;
                                minimum-receive-interval milliseconds;
                                multiplier number;
                                no-adaptation;
                                transmit-interval {
                                    minimum-interval milliseconds;
                                    threshold milliseconds;
                                }
                                detection-time {
                                    threshold milliseconds;
                                }
                            }
                            session-mode (automatic | multihop | singlehop);
                            holddown-interval milliseconds;
                        }
                    }
                }
            }
            overrides {
                interface-client-limit number;
                multi-address-embedded-option-response;
                process-inform {
                    pool pool-name;
                }
            }
        }
    }
```

```
    }
    rapid-commit;
  }
  service-profile dynamic-profile-name;
  trace;
  upto upto-interface-name;
}
overrides {
  delegated-pool;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
}
overrides {
  delegated-pool;
  interface-client-limit number;
  multi-address-embedded-option-response;
  process-inform {
    pool pool-name;
  }
  rapid-commit;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
```




```

        radius-disconnect;
    }
}
route-suppression;
service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    authentication {
        ...
    }
    dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
        primary-profile-name>;
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                    minimum-receive-interval milliseconds;
                    multiplier number;
                    no-adaptation;
                    transmit-interval {
                        minimum-interval milliseconds;
                        threshold milliseconds;
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}

```

```
    }
  }
  requested-ip-network-match subnet-mask;
  route-suppression;
  service-profile dynamic-profile-name;
}
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
overrides {
  client-discover-match (option60-and-option82 | incoming-interface);
  interface-client-limit number;
  process-inform {
    pool pool-name;
  }
}
pool-match-order {
  external-authority;
  ip-address-first;
  option-82;
}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}
```

|                                 |  |
|---------------------------------|--|
| <b>Hierarchy Level</b>          | <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> system services],</p> <p>[edit logical-systems <i>logical-system-name</i> system services],</p> <p>[edit routing-instances <i>routing-instance-name</i> system services],</p> <p>[edit system services]</p>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.</p> <p>The DHCP local server and the DHCP/BOOTP relay server, which are configured under the <b>[edit forwarding-options helpers]</b> hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.</p> <p>The <b>dhcpv6</b> stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.</p> <hr/> <div>  <p><b>NOTE:</b> When you configure the <b>dhcp-local-server</b> statement at the routing instance hierarchy level, you must use a routing instance type of <b>virtual-router</b>.</p> </div> <hr/> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Extended DHCP Local Server Overview</a></li> <li>• <a href="#">DHCPv6 Local Server Overview</a></li> <li>• <a href="#">Configuring a DHCP Server on Switches (CLI Procedure) on page 155</a></li> </ul>   |

## dhcp-relay

---

```
Syntax  dhcp-relay {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
    }
    dhcpv6 {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dynamic-profile profile-name {
            aggregate-clients (merge | replace);
            use-primary primary-profile-name;
        }
        group group-name {
            active-server-group server-group-name;
            authentication {
                ...
            }
            dynamic-profile profile-name {
                ...
            }
            interface interface-name {
                exclude;
                liveness-detection {
                    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                    method {
```

```

bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
}
}
overrides {
    ...
}
relay-option {
    ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
route-suppression:
service-profile dynamic-profile-name;
overrides {
    ...
}
relay-agent-interface-id {
    ...
}
relay-agent-remote-id {
    ...
}
relay-option {
    ...
}
route-suppression;
server-response-time seconds;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {

```

```
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    detection-time {
        threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
}
}
}
overrides {
    allow-snooped-clients;
    delay-authentication;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
duplicate-clients-in-subnet (incoming-interface | option-82):
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
    dynamic-profile profile-name {
        ...
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            method {
                bfd {
                    version (0 | 1 | automatic);
                    minimum-interval milliseconds;
                }
            }
        }
    }
}
```

```

        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        detection-time {
            threshold milliseconds;
        }
        session-mode(automatic | multihop | singlehop);
        holddown-interval milliseconds;
    }
}
}
overrides {
    ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    ...
}
relay-option {
    ...
}
relay-option-82 {
    ...
}
route-suppression:
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode(automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
}
overrides {

```

```
allow-snooped-clients;
always-write-giaddr;
always-write-option-82;
client-discover-match (option60-and-option82 | incoming-interface);
delay-authentication;
disable-relay;
interface-client-limit number;
layer2-unicast-replies;
no-allow-snooped-clients;
no-bind-on-request;
proxy-mode;
replace-ip-source-with;
send-release-on-delete;
trust-option-82;
}
relay-option {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group group-name;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
route-suppression:
server-response-time seconds;
service-profile dynamic-profile-name;
}
```



|                                 |  |
|---------------------------------|--|
| <b>Hierarchy Level</b>          | <p>[edit forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options]</p>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch and enable the router (or switch) to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.</p> <p>DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.</p> <p>The extended DHCP and DHCPv6 relay agent options configured with the <b>dhcp-relay</b> and <b>dhcpv6</b> statements are incompatible with the DHCP/BOOTP relay agent options configured with the <b>bootp</b> statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router (or switch) at the same time.</p> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Extended DHCP Relay Agent Overview</i></li> <li>• <i>DHCPv6 Relay Agent Overview</i></li> <li>• <i>DHCP Relay Proxy Overview</i></li> <li>• <i>Using External AAA Authentication Services with DHCP</i></li> </ul>   |



## CHAPTER 66

# Administration

- [Monitoring Commands for Port Mirroring on page 5497](#)

### Monitoring Commands for Port Mirroring

---

- [show analyzer](#)

## show analyzer

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show analyzer</b> < <i>analyzer-name</i> >  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Display information about port mirroring.  |
| <b>Options</b>                  | <i>analyzer-name</i> —(Optional) Displays the status of a specific analyzer (port-mirroring configuration).  |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Layer 2 Port Mirroring Overview</i></li> <li>• <a href="#">Port Mirroring Constraints and Limitations on page 5428</a></li> <li>• <i>Example: Configuring Port Mirroring for Local Analysis</i></li> <li>• <i>Example: Configuring Port Mirroring for Remote Analysis</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show analyzer on page 5498</a>   |
| <b>Output Fields</b>            | <a href="#">Table 446 on page 5498</a> describes the output fields for the <b>show analyzer</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 446: show analyzer Output Fields**

| Field Name                   | Field Description   |
|------------------------------|---|
| Analyzer name                | Name of the analyzer.   |
| Output interface             | Local interface to which mirror packets are sent. If you configure an output interface, you cannot also configure an output VLAN. |
| Output VLAN                  | VLAN to which mirror packets are sent. If you configure an output VLAN, you cannot also configure an output interface.            |
| Egress monitored interfaces  | Interfaces for which egress traffic is mirrored.  |
| Ingress monitored interfaces | Interfaces for which ingress traffic is mirrored.   |
| Ingress monitored VLANs      | VLANs for which ingress traffic is mirrored.  |

## Sample Output

### show analyzer

```

user@switch> show analyzer
Analyzer name      : employee-monitor
Output interface   : ge-0/0/10.0
Output VLAN        : remote-analyzer

```

```
Egress monitored interfaces : ge-0/0/7.0  
Ingress monitored interfaces : ge-0/0/8.0  
Ingress monitored interfaces : ge-0/0/9.0
```



## CHAPTER 67

# Troubleshooting

- [Troubleshooting Procedures on page 5501](#)

## Troubleshooting Procedures

---

- [Troubleshooting Port Mirroring on page 5501](#)

## Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 5501](#)
- [Egress Port Mirroring with VLAN Translation on page 5503](#)
- [Egress Port Mirroring with Private VLANs on page 5503](#)

### Port Mirroring Constraints and Limitations

---

- [Local and Remote Port Mirroring on page 5501](#)
- [Remote Port Mirroring Only on page 5503](#)

#### ***Local and Remote Port Mirroring***

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have



incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### ***Remote Port Mirroring Only***

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

### **Egress Port Mirroring with VLAN Translation**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

### **Egress Port Mirroring with Private VLANs**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

**Related Documentation**

- [Understanding Port Mirroring on page 5425](#)
- *Example: Configuring Port Mirroring for Local Analysis*
- *Example: Configuring Port Mirroring for Remote Analysis*

## PART 19

# Storage

- [Overview on page 5507](#)
- [Configuration on page 5595](#)
- [Administration on page 5707](#)
- [Troubleshooting on page 5771](#)



## CHAPTER 68

# Overview

- [Software Features Overview on page 5507](#)
- [FCoE and FIP Snooping on page 5513](#)
- [DCBX on page 5580](#)
- [Learn About Technology on page 5593](#)

### Software Features Overview

---

- [Overview of Fibre Channel on page 5508](#)
- [Overview of FIP on page 5513](#)

## Overview of Fibre Channel

Fibre Channel (FC) is a high-speed network technology that interconnects network elements and allows them to communicate with one another. The International Committee for Information Technology Standards (INCITS) T11 Technical Committee sets FC standards.

FC networks provide high-performance characteristics such as lossless transport combined with flexible network topology. FC is primarily used in storage area networks (SANs) because it provides reliable, lossless, in-order frame transport between initiators and targets. FC components include initiators, targets, and FC-capable switches that interconnect FC devices and may also interconnect FC devices with Fibre Channel over Ethernet (FCoE) devices. Initiators originate I/O commands. Targets receive I/O commands. For example, a server can initiate an I/O request to a storage device target.

The Juniper Networks QFX3500 Switch has native FC ports as well as Ethernet access ports, and can function as an FCoE-FC gateway or as an FCoE transit switch. All other QFX Series switches and EX4600 switches have Ethernet access ports and can function as an FCoE transit switch.

FCoE transports native FC frames over an Ethernet network by encapsulating the unmodified frames in Ethernet. It also provides protocol extensions to discover FCoE devices through the Ethernet network. FCoE requires that the Ethernet network support data center bridging (DCB) extensions that ensure lossless transport and allow the Layer 2 Ethernet domain to meet the requirements of FC transport.

The FCoE-FC gateway functionality is a licensed feature on the QFX Series that is available only on QFX3500 switches. As an FCoE-FC gateway, the switch connects FCoE devices on an Ethernet network to a SAN FC switch.

You do not need a license to use the switch as an FCoE transit switch. As an FCoE transit switch, the switch:

- Is a Layer 2 data center bridging (DCB) switch that can transport FCoE frames.
- Implements FCoE Initialization Protocol (FIP) snooping.
- Connects multiple FCoE endpoints to the FC network.



**NOTE:** Standalone switches support FCoE. Virtual Chassis (VC) and mixed-mode Virtual Chassis Fabric (VCF) configurations do not support FCoE. Pure QFX5100 switch VCFs (consisting of only QFX5100 switches) support FCoE.

---

This topic describes:

- [Fibre Channel Transport Protocol on page 5509](#)
- [How FC Works on the Switch on page 5509](#)

- [Supported FC Features and Functions on page 5512](#)
- [Lossless Transport Support on page 5512](#)

### Fibre Channel Transport Protocol

The Fibre Channel Protocol is a transport protocol that consists of five layers as shown in [Table 447 on page 5509](#):

**Table 447: Fibre Channel Protocol Layers**

| FC Protocol Layer | Description                                |
|-------------------|--|
| FC-0              | Physical (cabling, connectors, and so on)  |
| FC-1              | Data link layer                            |
| FC-2              | Network layer (defines the main protocols) |
| FC-3              | Common services                            |
| FC-4              | Protocol mapping                           |

The FC protocol layers are generally split into three groups:

- FC-0 and FC-1 are the physical layers.
- FC-2 is the protocol layer, similar to OSI Layer 3.
- FC-3 and FC-4 are the services layers.

The FCoE-FC gateway operates the physical layers and the protocol layer, and provides FIP and service redirection at the services layer.

### How FC Works on the Switch

The switch connects devices that support FC and Ethernet (such as FCoE servers on an Ethernet network) to an FC SAN, thus converging the Ethernet and FC networks on a single physical network infrastructure. The switch provides the class-of-service (CoS) features needed to handle the different types of traffic appropriately.

To converge FC and Ethernet networks, you can configure the switch as an:

- [FCoE-FC Gateway on page 5509](#)
- [FCoE Transit Switch on page 5510](#)
- [FCoE VLANs on page 5510](#)

#### **FCoE-FC Gateway**

When the switch functions as an FCoE-FC gateway, the switch aggregates FCoE traffic and performs the encapsulation and de-encapsulation of native FC frames in Ethernet as it transports the frames between FCoE devices in the Ethernet network and the FC switch. In effect, the switch translates Ethernet to FC and FC to Ethernet.

The gateway receives FC frames encapsulated in Ethernet from FCoE devices through an FCoE VLAN interface composed of one or more 10-Gigabit Ethernet interfaces. The gateway removes the Ethernet encapsulation from the FC frames, and then sends the native FC frames to the FC switch through a native FC interface.

The gateway receives native FC frames from the FC switch on the gateway's native FC interfaces. The gateway encapsulates the native FC frames in Ethernet, and then sends the encapsulated frames to the appropriate FCoE device through the FCoE VLAN interface.

To FCoE devices, the gateway behaves like an FC switch and can present multiple virtual F\_Ports (VF\_Ports) on a single interface. To an FC switch, the gateway behaves like an FC node that is doing N\_Port ID virtualization (NPIV).

### ***FCoE Transit Switch***

When the switch functions as an FCoE transit switch, it forwards traffic (including FCoE traffic) based on Layer 2 media access control (MAC) forwarding and is a normal DCB-enabled Layer 2 switch that also performs FIP snooping. The switch aggregates FCoE traffic and passes it through to an FCF. The switch does not remove the Ethernet encapsulation from the FC frames, but it does preserve the class of service (CoS) required to transport FC frames.

The switch inspects (snoops) FIP information in order to create filters that permit only valid FCoE traffic to flow through the switch between FCoE devices and the FCF. The switch does not use native FC ports because the FC frames are encapsulated in Ethernet when they flow between the FCoE devices and the FCF. Virtual point-to-point links between each FCoE device and the FCF pass transparently through the switch, so the switch is not seen as a terminating point or an intermediate point by FCoE devices or by the FCF.

### ***FCoE VLANs***

All FCoE traffic must travel in a VLAN dedicated to transporting only FCoE traffic. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE or FIP traffic must travel in a different VLAN.



**NOTE:** The same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.

---





**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



**NOTE:** IGMP snooping is enabled by default on all VLANs in all software versions before Junos OS R13.2. Disable IGMP snooping on FCoE VLANs if you are using software that is older than 13.2.

You can configure more than one FCoE VLAN, but any given virtual link must be in only one FCoE VLAN.



**NOTE:** All 10-Gigabit Ethernet interfaces that connect to FCoE devices must have a native VLAN configured in order to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.



**BEST PRACTICE:** Only FCoE traffic is permitted on the FCoE VLAN. A native VLAN might need to carry untagged traffic of different types and protocols. Therefore, it is a good practice to keep the native VLAN separate from FCoE VLANs.

## Supported FC Features and Functions

---

The following features and functionality are supported:

- As an FCoE-FC gateway:
  - DCB, including Data Center Bridging Capability Exchange protocol (DCBX), priority-based flow control (PFC), enhanced transmission service (ETS), and 10-Gigabit Ethernet interfaces
  - FCoE Initialization Protocol (FIP)
  - Proxy for FCoE devices when communicating with FC switches and acts as a proxy for FC switches when communicating with FCoE devices
  - Up to 12 native FC interfaces per QFX3500 switch (each interface can be configured as a 2-Gigabit, 4-Gigabit, or 8-Gigabit Ethernet interface)
- As an FCoE transit switch:
  - DCB functions
  - FIP snooping
  - Transparent Layer 2 MAC forwarding of FCoE frames

## Lossless Transport Support

---

Up to six lossless forwarding classes are supported. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from a standalone switch or QFabric system Node device to other devices cannot exceed 300 meters.
- The internal cable length from a QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.

### Related Documentation

- *Understanding Fibre Channel*
- *Understanding an FCoE-FC Gateway*
- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Understanding FCoE on page 5518](#)
- [Understanding DCB Features and Requirements on page 5515](#)
- [Overview of FIP on page 5513](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- *Understanding Interfaces on an FCoE-FC Gateway*

- *Understanding FCoE LAGs*
- [Understanding Fibre Channel Terminology on page 5569](#)

## Overview of FIP

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) is a Layer 2 protocol that establishes and maintains Fibre Channel (FC) virtual links between pairs of FCoE devices such as server FCoE Nodes (ENodes) and FC switches. FIP can also establish and maintain virtual links between FCoE devices and an FCoE-FC gateway (such as the QFX3500 switch), where the gateway acts on behalf of the FC switch.

FIP enables FCoE devices to discover one another and to initialize and maintain virtual links over a physical Ethernet network. This allows FCoE devices in the Ethernet network to access storage devices in the FC storage area network (SAN).

FIP solves the problem presented by the FC requirement for point-to-point connections (FC does not permit point-to-multipoint connections) by creating a unique virtual link for each connection between an ENode VN\_Port and an FC switch VF\_Port. Multiple virtual links can use a single physical link and virtual links can traverse Ethernet transit (passthrough) switches while appearing to be direct point-to-point connections to the FC switch.

FIP has its own EtherType (0x8914) to distinguish its traffic from payload-carrying FCoE traffic and other Ethernet traffic. FIP operations occur on a per-VLAN basis.

For more details about FIP, see the Technical Committee T11 organization document *Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00* available at <http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf>.

### Related Documentation

- [Overview of Fibre Channel on page 5508](#)
- *Understanding Fibre Channel*
- *Understanding FIP Functions*
- *Understanding FIP Implementation on an FCoE-FC Gateway*
- *Understanding FIP Parameters on an FCoE-FC Gateway*
- *Understanding Fibre Channel Virtual Links*
- [Understanding FCoE on page 5518](#)
- *Understanding an FCoE-FC Gateway*
- *Configuring FIP on an FCoE-FC Gateway*
- [Understanding Fibre Channel Terminology on page 5569](#)

## FCoE and FIP Snooping

- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding FCoE on page 5518](#)

- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Understanding FCoE and FIP Session High Availability on page 5528](#)
- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 5530](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)
- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5546](#)
- [Understanding MC-LAGs on an FCoE Transit Switch on page 5555](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding Fibre Channel Terminology on page 5569](#)

## Understanding DCB Features and Requirements

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.



Video: [What is Data Center Bridging?](#)

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series and EX4600 switch support the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 5515](#)
- [ETS on page 5516](#)
- [DCBX on page 5517](#)

### Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.

This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 5516](#)
- [Buffer Management on page 5516](#)
- [Physical Interfaces on page 5516](#)

### **PFC**

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

### **Buffer Management**

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

### **Physical Interfaces**

The switch supports 10-Gbps, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps (or faster) Ethernet interfaces.

### **ETS**

---

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).
- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.

- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows. We recommend that you always configure a shaping rate to limit the amount of bandwidth a strict-high priority flow can consume by including the [shaping-rate](#) statement in the [\[edit class-of-service schedulers\]](#) hierarchy on the strict-high priority scheduler. This prevents a strict-high priority from starving other queues on the port.

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

## DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

### Related Documentation

- [Overview of Fibre Channel on page 5508](#)
- [Understanding FCoE on page 5518](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding DCBX on page 5580](#)
- [Understanding Fibre Channel Terminology on page 5569](#)

- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)

## Understanding FCoE

Fibre Channel over Ethernet (FCoE) is a method of supporting converged Fibre Channel (FC) and Ethernet traffic on a data center bridging (DCB) network. FCoE encapsulates unmodified FC frames in Ethernet to transport the FC frames over a physical Ethernet network. The T11 Technical Committee, which is the International Committee for Information Technology Standards (INCITS) committee responsible for FC interfaces, developed the FCoE standard to provide a method for transporting FC frames over a DCB network. The T11 document *Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00* at <http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf> provides details about the FCoE version 1 standard.



**NOTE:** The switch does not support T11 Annex F *FCoE Pre-FIP Virtual Link Instantiation Protocol*.

To the Ethernet network, an FCoE frame is the same as any other Ethernet frame because the Ethernet encapsulation provides the header information needed to forward the frames. However, to achieve the lossless behavior that FC transport requires, the Ethernet network must conform to DCB standards.

DCB standards create an environment over which FCoE can transport native FC traffic encapsulated in Ethernet while preserving the mandatory class of service (CoS) and other characteristics that FC traffic requires.

Supporting FCoE in a DCB network requires that the FCoE devices in the Ethernet network and the FC switches at the edge of the SAN network handle both Ethernet and native FC traffic. To handle Ethernet traffic, an FC switch does one of two things:

- Incorporates FCoE interfaces.
- Uses an FCoE-FC gateway such as a QFX3500 switch to de-encapsulate FCoE traffic from FCoE devices into native FC and to encapsulate native FC traffic from the FC switch into FCoE and forward it to FCoE devices through the Ethernet network.



**NOTE:** Standalone switches support FCoE. Virtual Chassis (VC) and mixed-mode Virtual Chassis Fabric (VCF) configurations do not support FCoE. Pure QFX5100 switch VCFs (consisting of only QFX5100 switches) support FCoE.

FCoE concepts include:

- [FCoE Devices on page 5519](#)
- [FCoE Frames on page 5520](#)



- [Virtual Links on page 5521](#)
- [FCoE VLANs on page 5521](#)

### FCoE Devices

Each FCoE device has a converged network adapter (CNA) that combines the functions of an FC host bus adapter (HBA) and a lossless Ethernet network interface card (NIC) with 10-Gbps Ethernet ports. The portion of the CNA that handles FCoE traffic is called an FCoE Node (ENode). An ENode combines FCoE termination functions and the client part of the FC stack on the CNA.

ENodes present virtual FC interfaces to FC switches in the form of virtual N\_Ports (VN\_Ports). A VN\_Port is an endpoint in a virtual point-to-point connection called a virtual link. The other endpoint of the virtual link is an FC switch (or FCF) port. A VN\_Port emulates a native FC N\_Port and performs similar functions: handling the creation, detection, and flow of messages to and from the FC switch. A single ENode can host multiple VN\_Ports. Each VN\_Port has a separate, unique virtual link with a FC switch.

ENodes contain at least one lossless Ethernet media access controller (MAC). Each Ethernet MAC is paired with an FCoE controller. The lossless Ethernet MAC is a full-duplex Ethernet MAC that implements Ethernet extensions to avoid frame loss due to congestion and supports frames of at least 2500 bytes. The FCoE controller instantiates and terminates VN\_Port instances dynamically as they are needed for FCoE sessions. Each VN\_Port instance has a unique virtual link to an FC switch.



**NOTE:** A *session* is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.

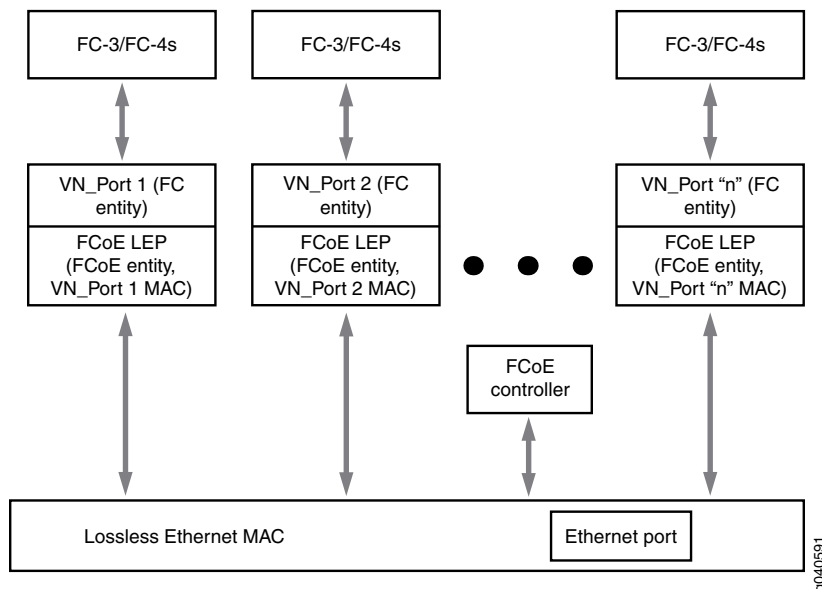
ENodes also contain one FCoE link end point (LEP) for each VN\_Port connection. An FCoE LEP is a virtual FC interface mapped onto the physical Ethernet interface.

An FCoE LEP:

- Transmits and receives FCoE frames on the virtual link.
- Handles FC frame encapsulation for traffic going from the server to the FC switch.
- Performs frame de-encapsulation of traffic received from the FC switch.

[Figure 190 on page 5520](#) shows a block diagram of the major ENode components.

Figure 190: ENode Components



### FCoE Frames

The FCoE protocol specification replaces the FC0 and FC1 layers of the FC stack with Ethernet, but retains the FC frame header. Retaining the FC frame header enables the FC frame to pass directly to a native FC SAN after de-encapsulation. The FCoE header carries the FC start of file (SOF) bits and end of file (EOF) bits in an encoded format. FCoE supports two frame types, control frames and data frames. FCoE Initialization Protocol (FIP) carries all of the discovery and fabric login frames.

FIP control frames handle FCoE device discovery, initializing communication, and maintaining communication. They do not carry a data payload. FIP has its own EtherType (0x8914) to distinguish FIP traffic from FCoE traffic and other Ethernet traffic. To establish communication, the ENode uses the globally unique MAC address assigned to it by the CNA manufacturer.

After FIP establishes a connection between FCoE devices, the FCoE data frames handle the transport of the FC frames encapsulated in Ethernet. FCoE also has its own EtherType (0x8906) to distinguish FCoE frames from other Ethernet traffic and ensure the in-order frame handling that FC requires. FCoE frames include:

- 2112 bytes FC payload
- 24 bytes FC header
- 14 bytes standard Ethernet header
- 14 bytes FCoE header
- 8 bytes cyclic redundancy check (CRC) plus EOF
- 4 bytes VLAN header
- 4 bytes frame check sequence (FCS)

The payload, headers, and checks add up to 2180 bytes. Therefore, interfaces that carry FCoE traffic should have a configured maximum transmission unit (MTU) of 2180 or larger. An MTU size of 2180 bytes is the minimum size; some network administrators prefer an MTU of 2240 or 2500 bytes.

### Virtual Links

Native FC uses point-to-point physical links between FC devices. In FCoE, virtual links replace the physical links. A virtual link emulates a point-to-point link between two FCoE device endpoints, such as a server VN\_Port and an FC switch (or FCF) VF\_Port.

Each FCoE interface can support multiple virtual links. The MAC addresses of the FCoE endpoints (the VN\_Port and the VF\_Port) uniquely identify each virtual link and allow traffic for multiple virtual links to share the same physical link while maintaining data separation and security.

A virtual link exists in one FCoE VLAN and cannot belong to more than one VLAN. Although the FC switch and the FCoE device detect a virtual link as a point-to-point connection, virtual links do not need to be direct connections between a VF\_Port and a VN\_Port. A virtual link can traverse one or more transit switches, also known as passthrough switches. A transit switch can transparently aggregate virtual links while still appearing and functioning as a point-to-point connection to the FCoE devices. However, a virtual link must remain within a single Layer 2 domain.

### FCoE VLANs

All FCoE traffic must travel in a VLAN dedicated to transporting only FCoE traffic. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE or FIP traffic must travel in a different VLAN.



**NOTE:** On a standalone switch or QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



**NOTE:** IGMP snooping is enabled by default on all VLANs in all software versions before Junos OS R13.2. Disable IGMP snooping on FCoE VLANs if you are using software that is older than 13.2.

You can configure more than one FCoE VLAN, but any given virtual link must be in only one FCoE VLAN.



**NOTE:** All 10-Gigabit Ethernet interfaces that connect to FCoE devices must have a native VLAN configured in order to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.

On switches that use the Enhanced Layer 2 Software (ELS) CLI, it is not sufficient only to configure the native VLAN on the interface, the interface must also be configured as a member of the native VLAN. (This is because the ELS CLI does not support tagged-access interface mode, so interfaces that are members of FCoE VLANs must use trunk mode, and trunk port interfaces must be explicitly included as members of a native VLAN.)

In addition, the VLAN ID must match the native VLAN ID that you configure on the physical interface. For example, to configure a native VLAN with an ID of 20 on interface xe-0/0/15 that is a member of an FCoE VLAN, you must include both of the following statements in the configuration:

1. Configure the native VLAN on the interface:

```
user@switch# set interfaces xe-0/0/15 native-vlan-id 20
```

(The equivalent configuration statement on a non-ELS device switch would be `set interfaces xe-0/0/15 unit 0 family ethernet-switching native-vlan-id 20`.)

2. Configure the port as a member of the native VLAN (this step is not required on switches that do not use the ELS software):

```
user@switch# set interfaces xe-0/0/15 unit 0 family ethernet-switching vlan members 20
```



**BEST PRACTICE:** Only FCoE traffic is permitted on the FCoE VLAN. A native VLAN might need to carry untagged traffic of different types and protocols. Therefore, it is a good practice to keep the native VLAN separate from FCoE VLANs.

#### Related Documentation

- [Overview of Fibre Channel on page 5508](#)
- [Understanding Fibre Channel](#)
- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Understanding FCoE LAGs](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding Fibre Channel Terminology on page 5569](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)
- [Configuring an FCoE LAG](#)

- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)

## Understanding FCoE Transit Switch Functionality

You can use the switch as a Fibre Channel over Ethernet (FCoE) transit switch. An FCoE transit switch is a Layer 2 data center bridging (DCB) switch that can transport FCoE frames and implements FCoE Initialization Protocol (FIP) snooping. A DCB switch transports both FCoE and Ethernet LAN traffic over the same network infrastructure while preserving the class of service (CoS) that Fibre Channel (FC) traffic requires.

An FCoE transit switch does not encapsulate or de-encapsulate FC frames in Ethernet. It is an access switch that transports FC frames that have already been encapsulated in Ethernet between FCoE initiators such as servers and a storage area network (SAN) FC switch that supports both Ethernet and native FC traffic on its interfaces. The transit switch acts as a passthrough switch and is transparent to the FC switch, which detects each connection to an FCoE device as a direct point-to-point link.

When a switch acts as a transit switch, the VLANs you configure for FCoE traffic can use any of the switch ports because the traffic in both directions is standard Ethernet traffic, not native FC traffic.



.....

**NOTE:** The Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets. It is a good practice to keep the native VLAN separate from the VLANs that carry FCoE traffic. FCoE VLANs should carry only FCoE traffic, but other types of untagged traffic might use the native VLAN.

Switches and QFabric system Node devices that use the original CLI (not the Enhanced Layer 2 (ELS) software) only require that you configure the native VLAN on the FCoE interfaces that belong to the FCoE VLAN by including the `[set interfaces interface-name unit unit family ethernet-switching native-vlan-id native-vlan-id]` statement in the configuration.

QFX5100 and EX4600 switches use ELS software and require that you include two statements in the configuration to configure a native VLAN on FCoE interfaces. Include the `[set interfaces interface-name native-vlan-id vlan-id]` statement in the configuration to configure the native VLAN on the interface, and also include the `[set interfaces interface-name unit unit family ethernet-switching native-vlan-id vlan-id]` statement in the configuration to configure the port as a member of the native VLAN.

.....

FCoE traffic should use a VLAN dedicated only to FCoE traffic. Do not mix FCoE traffic with standard Ethernet traffic on a VLAN on the switch.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



**NOTE:** IGMP snooping is enabled by default on all VLANs in all software versions before Junos OS R13.2. Disable IGMP snooping on FCoE VLANs if you are using software that is older than 13.2.



**NOTE:** On a QFX3500 switch or on a QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode. (Only QFX3500 switches can be configured in FCoE-FC gateway mode.) If you configure both a transit switch and an FCoE-FC gateway on the same QFX3500 switch or QFabric system Node device, configure different FCoE VLANs for the transit switch and the FCoE-FC gateway.

Transit switch architecture differs from FCoE-FC gateway architecture. As an FCoE-FC gateway, the system transports traffic to the FC SAN as native FC frames, and the VLAN must use an FCoE VLAN interface and native FC interfaces to transport that traffic. As a transit switch, the system forwards Ethernet traffic, and requires DCB configuration for lossless transport of that traffic and FIP snooping at FCoE device access ports, but not the FCoE-FC gateway features necessary for transporting FC traffic.

With the exception of Virtual Chassis (VC) and mixed-mode Virtual Chassis Fabric (VCF) configurations, switches support the DCB standards for ensuring lossless transport and low latency, and provide 10-Gbps ports for FCoE traffic. VCF configurations that use only

QFX5100 switches support DCB standards. For lossless transport to function correctly, you must use priority-based flow control (PFC, described in IEEE 802.1Qbb) to create bandwidth reservations and ensure proper CoS for FCoE traffic.

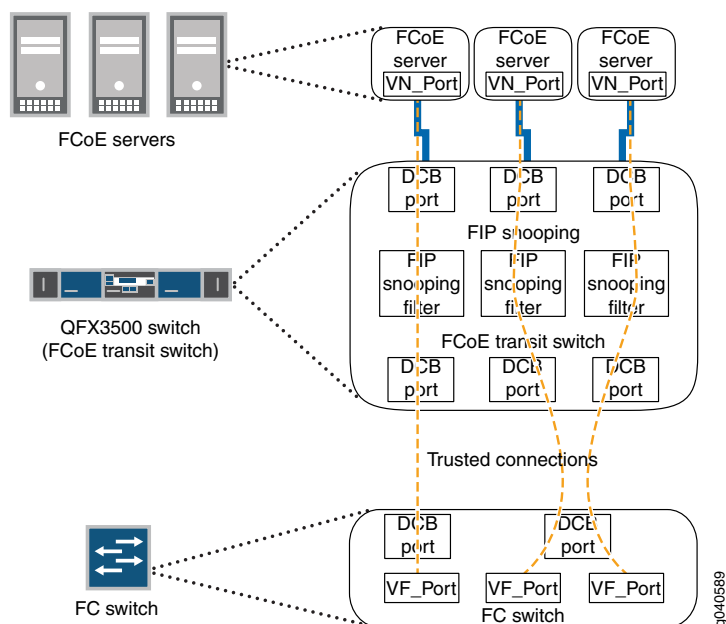
FIP snooping adds security by filtering access so that only traffic from servers that have successfully logged in to the FC network passes through the transit switch and reaches the FC network. The Technical Committee T11 organization specifications describe two types of FIP snooping:

- The FC-BB-5 specification describes VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping, which provides security for communication between FCoE device VN\_Ports on the Ethernet network and FCF or FC switch VF\_Ports.
- The FC-BB-6 specification describes VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping, which provides security for communication between FCoE device VN\_Ports on the Ethernet network.

To accommodate the larger size of Ethernet-encapsulated frames, FCoE interfaces should be configured with a maximum transmission unit (MTU) size of at least 2180 bytes.

The transit switch transparently connects FCoE-capable devices such as servers in an Ethernet LAN to an FC switch or to a gateway switch (hereafter referred to as the FC switch), as shown in [Figure 191 on page 5526](#). The transit switch acts as a transparent DCB access layer between FCoE servers and the FC switch.

**Figure 191: FCoE Transit Switch Connecting FCoE Devices to an FC Switch**



The transit switch performs FIP snooping at the ports connected to the FCoE devices. For VN2VF\_Port FIP snooping, at the SAN edge, the FC switch must be able to convert the FCoE traffic to native FC traffic. (VN2VN\_Port FIP snooping switches traffic between



VN\_Ports directly through the transit switch, without going through the FC switch, so no conversion of FCoE traffic to native FC traffic is needed.)

Encapsulated FCoE traffic flows through the transit switch to the FCoE ports on the FC switch. The FC switch removes the Ethernet encapsulation from the FCoE frames to restore the native FC frames. Native FC traffic travels out native FC ports to storage devices in the FC SAN.

Native FC traffic from storage devices flows to the FC switch FC ports, and the FC switch encapsulates that traffic in Ethernet as FCoE traffic. The FCoE traffic flows through the transit switch to the appropriate FCoE device.



**NOTE:** The FC switch and FC fabric apply appropriate zoning checks on traffic to and from each ENode and provide FC services (for example, name server, fabric login server, or event server).



**NOTE:** VN\_Port to VN\_Port FIP snooping is supported to allow FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch. An FCoE VLAN can support either VN2VF\_Port FIP snooping (FC-BB-5) or VN2VN\_Port FIP snooping (FC-BB-6), but not both. The same switch can have multiple FCoE VLANs configured, some FCoE VLANs for VN2VF FIP snooping traffic and others for VN2VN FIP snooping traffic.

#### Related Documentation

- [Overview of Fibre Channel on page 5508](#)
- [Understanding DCB Features and Requirements on page 5515](#)
- [\*Understanding an FCoE-FC Gateway\*](#)
- [Understanding FCoE on page 5518](#)
- [\*Understanding FCoE LAGs\*](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)
- [Understanding Fibre Channel Terminology on page 5569](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)
- [\*Disabling Enhanced FIP Snooping Scaling\*](#)
- [\*Configuring an FCoE LAG\*](#)

## Understanding FCoE and FIP Session High Availability

High availability features maintain storage network sessions when a system process is terminated and during certain types of upgrades:

- [High Availability for Fibre Channel Process Termination \(FCoE-FC Gateway Mode\) on page 5528](#)
- [High Availability for FIP Snooping on page 5528](#)
- [Nonstop Software Upgrade \(QFabric Systems\) on page 5529](#)

### High Availability for Fibre Channel Process Termination (FCoE-FC Gateway Mode)

In FCoE-FC gateway mode, the QFX3500 switch provides high availability to restore the FCoE sessions running on the switch in case the Fibre Channel (FC) process is terminated. A session is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric, not an end-to-end server-to-storage session.

The switch stores FCoE session data in a persistent storage module. If the FC process terminates, the switch restores the existing FCoE sessions on the same interfaces that they were on before the FC process terminated. Data traffic for existing sessions is not affected during session restoration.

For a brief time, the system does not process control traffic because of the FC process restart and session restoration. During this brief time, no new FCoE sessions can be established, and no existing sessions can log out.



.....  
**NOTE:** During the restoration process, if the FC process does not receive an *interface up* notification from a particular interface within a certain time, the switch times out the restore operation and discards the data on that interface. The previously existing FCoE sessions on that interface are not restored, and the ENodes must log in again.  
.....



.....  
**NOTE:** An FC process restart and session restoration resets the Fibre Channel statistics.  
.....

If the FC process terminates repeatedly, the operating system disables the process until you manually restart it. To restart the FC process manually, issue the **restart fibre-channel** command.

### High Availability for FIP Snooping

You can configure the system to perform FIP snooping on Ethernet interfaces that are connected to FCoE devices that have ENodes. The high availability function restores running FIP snooping sessions in case the Ethernet switching process is terminated.

The Ethernet switching process stores the FIP snooping state in a persistent storage module. If the Ethernet switching process terminates, the switch restores the existing

FIP snooping sessions on the same interfaces that they were on before the Ethernet switching process terminated. The high availability features preserve:

- Logged in ENodes
- Discovered FCFs
- Existing sessions
- Existing FIP snooping filters

The complete restoration process, including reconciling all valid states, takes a maximum of 8 seconds. During the restoration process, the switch can learn a new FCF or a new FC switch, and new ENodes can log in to the FC network. However, FDISC messages from an ENode that is already logged in to the network might be dropped if the ENode has not yet been restored.

When the Ethernet switching process terminates ungracefully, the FIP keepalive timer is reset to the normal initial value, not the value at the time of the Ethernet switching process termination.

In the event of an Ethernet switching process termination, ENodes remain logged in, and existing sessions are not interrupted.



**NOTE:** An Ethernet switching process restart and session restoration resets the FIP snooping statistics.

### Nonstop Software Upgrade (QFabric Systems)

On QFabric system Node groups that have more than one Node device, nonstop software upgrade (NSSU) enables you to upgrade the Node devices with minimal packet loss and maximum uptime. NSSU automates software upgrades on the QFabric system components in an orderly and consistent manner to maximize system uptime.

The system upgrades components with redundant architectures, such as redundant server Node groups and network Node groups that have two or more members, in stages. While the system upgrades one component, the redundant component continues to function.

For example, while one member of a redundant server Node group is upgraded, the other member continues to forward traffic. When the first Node group member completes the upgrade, it comes online while the system upgrades the second member.

NSSU provides high availability for the lossless traffic forwarding required to support storage networks. If your system design includes redundancy (redundant Node devices in Node groups, LAGs, and so on) so that an alternate traffic path is available, when you upgrade a Node device, traffic is not impacted.

In fully redundant topologies, NSSU preserves FIP session, FIP snooping filter, VN2VF\_Port session, and VN2VN\_Port session information and prevents traffic loss in most cases.

An exception is that Node devices that are directly connected to ENodes experience momentary traffic loss when the Node device reboots.

**Related  
Documentation**

- *Understanding an FCoE-FC Gateway*
- [Understanding FCoE on page 5518](#)
- *Understanding Nonstop Software Upgrade for QFabric Systems*
- *Performing a Nonstop Software Upgrade on the QFabric System*

## Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches

The originator exchange identifier (OxID) field is one of several fields that the switch can use in its hash function computation for FCoE traffic load balancing over multiple outgoing links in an Ethernet link aggregation group (LAG) on ports that face an FCoE forwarder (FCF). The originator of an exchange between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) uses the OxID field as an identifier for that exchange. The originator also uses the OxID field to track the progress of the series of sequences that comprise the exchange.

When FCoE traffic traverses a LAG that faces an FCF, it can take multiple different links between the source and destination endpoints. The idea is to distribute the FCoE traffic across the FCF-facing LAG links, thus balancing the link load. The switch creates a hash value from some of the packet header fields, and uses the hash value to assign each packet to one of the LAG links. The switch always uses five packet header fields to compute the hash value:

- Source ID (SID)
- Destination ID (DID)
- Fabric ID (FID)
- Source Port ID (SPID)
- Source Module ID (SMID)

In addition, the OxID field is included by default in the FCoE load-balancing hash computation. However, if you do not want to use the OxID field in the FCoE load-balancing hash computation, you can remove it from the computation by using the **set forwarding-options hash-key family fcoe oxid disable** command.

Including the OxID field in the load-balancing hash computation allows different exchanges between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) to take different paths across the network, thus improving the aggregate network throughput.

However, if the paths between different sets of FC endpoints have common links, congestion on one set of FC endpoints can affect the other set of endpoints. Such congestion can happen if the FCoE traffic on the two sets of endpoints uses the same priority (IEEE 802.1p code point). It is common for networks to use priority 3 (IEEE 802.1p code point 011) for FCoE traffic. However, you can assign different IEEE priorities to

different lossless FCoE flows as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5837](#) to further separate the traffic flows.

- Related Documentation**
- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on QFabric Systems](#)
  - [Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 5656](#)

## Understanding VN\_Port to VF\_Port FIP Snooping on an FCoE Transit Switch

Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping is a security mechanism that is designed to prevent unauthorized access and data transmission to a Fibre Channel (FC) network. It works by filtering traffic to permit only servers that have logged in to an FC network to access that network.

You explicitly enable VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (FC-BB-5) on FCoE VLANs when the switch is an FCoE transit switch at the access edge that connects FCoE devices on the Ethernet network to FC switches or gateways at the FC storage area network (SAN) edge. The transit switch applies FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VF\_Port FIP snooping. An FCoE transit switch is a data center bridging (DCB) switch with FIP snooping capability.

An FCoE device that has a converged network adapter (CNA) uses the FIP process to log in to the FC network as an FCoE Node (ENode). The login process establishes a dedicated virtual link between a virtual N\_Port (VN\_Port) on the ENode and a virtual F\_Port (VF\_Port) on the FC switch. This dedicated virtual link emulates a point-to-point connection. The emulated connection is called a virtual link.

Virtual links pass transparently through the transit switch. The ENode VN\_Port and the FC switch VF\_Port do not detect the transit switch, and virtual links appear to be direct point-to-point links.

The switch applies VN2VF\_Port FIP snooping firewall filters at the FCoE-network facing ports associated with the FCoE VLANs on which you enable VN2VF\_Port FIP snooping. FIP snooping provides security for virtual links by creating firewall filters based on information gathered (snooped) about FC devices during FIP transactions.

The switch also supports VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping (FC-BB-6) to allow FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch, as described in [“Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch” on page 5539](#).



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping (FC-BB-5) or VN2VN\_Port FIP snooping (FC-BB-6), but not both. The same switch can have multiple FCoE VLANs configured, some for VN2VF\_Port FIP snooping traffic and others for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port snooping VLANs, VN2VF\_Port FIP snooping traffic is dropped.

When you enable VN2VF\_Port FIP snooping on an FCoE VLAN, the system snoops VN\_Port to VF\_Port packets and enforces security only on VN2VF\_Port virtual links.

When you enable VN2VN\_Port FIP snooping on an FCoE VLAN, the system snoops VN\_Port to VN\_Port packets and enforces security only on VN2VN\_Port virtual links.

---

This topic describes:

- [FC Network Security on page 5532](#)
- [VN2VF\\_Port FIP Snooping Functions on page 5533](#)
- [FIP Snooping Firewall Filters on page 5534](#)
- [FIP Snooping Session Scalability on page 5534](#)
- [VN2VF\\_Port FIP Snooping Implementation on page 5534](#)
- [T11 VN2VF\\_Port FIP Snooping Specification on page 5538](#)

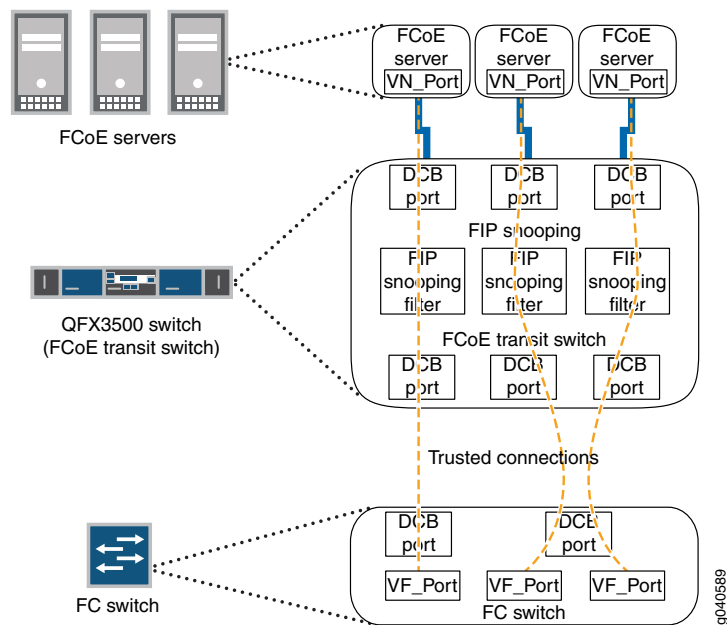
## **FC Network Security**

---

In traditional FC networks, the FC switch is usually a trusted entity, and server ENodes connect directly to its VF\_Ports. After an ENode gains access to the network through the fabric login (FLOGI) process, the FC switch enforces zoning configurations, ensures that the ENode uses valid addresses, monitors the connection, and performs other security functions to prevent unauthorized access.

However, FCoE exposes FC frames to Ethernet networks, which do not have the same level of security as native FC networks. VN2VF\_Port FIP snooping firewall filters emulate the native FC network security functions by preventing unauthorized access to the FC switch through the transit switch and by ensuring the security of the virtual link between each ENode and the FC switch, as shown in [Figure 192 on page 5533](#). VN2VF\_Port FIP snooping also prevents man-in-the-middle attacks.

Figure 192: FCoE Transit Switch Performs VN2VF\_Port FIP Snooping



The transit switch performs VN2VF\_Port FIP snooping at the ports connected to the FCoE devices. At the SAN edge, the FC switch must be able to convert the FCoE traffic to native FC traffic.

### VN2VF\_Port FIP Snooping Functions

When VN2VF\_Port FIP snooping is enabled, the transit switch sets and applies filters to block all FCoE traffic by default. The transit switch monitors FIP logs, solicitations, and advertisements that pass through it and gathers information about the ENode address and the address of the port on the FC switch. The transit switch uses the information to construct firewall filters that permit access only to logged-in ENodes. All other traffic on the VLAN is denied.

For example, when an ENode on an FCoE VLAN performs a successful login to an FC switch port, the transit switch snoops the FIP information and constructs a firewall filter that provides access for the ENode to that port on the FC switch.

The firewall filters enable FCoE frames to pass through the transit switch only on a virtual link established between an FCoE device ENode VN\_Port and the FC switch VF\_Port to which it has logged in. The firewall filters ensure that ENodes can only connect to the FC switches they have successfully logged in to and that only valid FCoE traffic along valid paths is transmitted. VN2VF\_Port FIP snooping maintains the filters by tracking FCoE sessions (ENode to FCF sessions).

### FIP Snooping Firewall Filters

---

The effect of the firewall filters is to protect the FCoE ports. VN2VF\_Port FIP snooping performs the following actions and checks to ensure that FCoE traffic is valid:

- Denies ENodes that use the FC switch media access control (MAC) address as the source address.
- Enables ENodes to transmit FIP and FCoE frames to the FC switch address.
- Ensures that the FCoE source address the FC switch assigns or accepts is only used for FCoE traffic.
- Ensures that FCoE frames are only addressed to the accepting FC switch.

### FIP Snooping Session Scalability

---

Enhanced FIP snooping session scaling, which supports up to 2,500 sessions, is enabled by default. On QFabric systems, if you want to disable enhanced FIP snooping scaling (which reduces the number of supported sessions to 376 sessions), you can do so as described in *Disabling Enhanced FIP Snooping Scaling*.

By default, up to 2500 total FIP snooping sessions are supported on an interface, an FCoE-FC gateway fabric (only supported on QFX3500 switches configured as standalone switches or as QFabric system Node devices), a switch, a QFabric Node device, or a QFabric Node group. For example, you can:

- Place all 2500 sessions on one FCoE interface.
- Split the 2500 sessions among multiple FCoE interfaces on one FCoE VLAN.
- Split the 2500 sessions among multiple FCoE interfaces on multiple FCoE VLANs.
- Split the 2500 sessions among the FCoE interfaces on multiple gateway FC fabrics on a switch.
- Split the 2500 sessions among the FCoE interfaces on multiple gateway FC fabrics on multiple Node devices in a QFabric Node group.

Regardless of how you allocate the sessions among interfaces and local FC fabrics on a switch or on a QFabric system Node device or Node group, the combined FIP session limit is a maximum of 2500 sessions.



**NOTE:** The total number of sessions the system can support is the combined number of VN2VF\_Port sessions and VN2VN\_Port sessions. If VN2VN\_Port sessions are active, the total number of available VN2VF\_Port sessions is reduced.

---

### VN2VF\_Port FIP Snooping Implementation

---

You enable VN2VF\_Port FIP snooping on a per-VLAN basis on VLANs that carry FCoE traffic. The switch snoops FIP frames at the ports associated with FCoE VLANs enabled



for VN2VF\_Port FIP snooping. The switch then installs the resulting firewall filters on the ports to ensure that all VN2VF\_Port FIP snooping occurs on the switch network edge.

VN2VF\_Port FIP snooping FCoE VLANs must meet the following criteria:

- An FCoE VLAN should be dedicated to FCoE traffic only.
- An FCoE VLAN cannot support both VN2VF\_Port FIP snooping and VN2VN\_Port FIP snooping simultaneously. You must configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic.



**NOTE:** Changing an FCoE VLAN from VN2VF\_Port FIP snooping mode to VN2VN\_Port snooping mode terminates the existing virtual links on the VLAN. The transit switch removes the existing FIP snooping filters, creates the new FIP snooping filters, and applies them to the FIP snooping ports. If you downgrade the software to Junos OS Release 12.1 or earlier, VLANs configured for VN2VN\_Port FIP snooping revert to VN2VF\_Port FIP snooping VLANs.

- For systems that use software that does not support Enhanced Layer 2 Software (ELS) CLI, configure all access ports that belong to an FCoE VLAN (ports connected to a converged network adapter [CNA] in an FCoE device) in **tagged-access** port mode. Access ports associated with an FCoE VLAN should not be configured as access ports or trunk ports on these platforms, although trunk port configuration is supported.

However, on switches that use the ELS CLI, configure access ports that belong to an FCoE VLAN in **trunk** interface mode.

- All ports connected to an FC switch (or FCoE forwarder) must be configured in **trunk** port mode. Ports connected to an FC switch must be configured as trusted ports.
- FIP traffic uses the native VLAN (FIP VLAN discovery and notification frames are exchanged as untagged packets).
- All FCoE VLAN traffic must be tagged and cannot belong to the native VLAN.
- FCoE VLAN traffic cannot be untagged or priority-tagged.

When you enable VN2VF\_Port FIP snooping, the switch inspects FIP frames.

The VN2VF\_Port FIP snooping implementation includes these considerations:

- [ENode-Facing Interfaces on page 5535](#)
- [Network-Facing Interfaces on page 5537](#)
- [FC-MAP on page 5537](#)

### ***ENode-Facing Interfaces***

When the interfaces that belong to an FCoE VLAN connect directly to FCoE devices (there is no other transit switch between the FCoE devices and the switch), we recommend that you enable VN2VF\_Port FIP snooping on all FCoE VLANs that connect VN\_Ports to VF\_Ports. Enabling FIP snooping ensures secure connections between server ENodes and FC switches. (Enabling VN2VN\_Port FIP snooping ensures secure connections on

FCoE VLANs that connect VN\_Ports to other VN\_Ports). FIP snooping should always be enabled at the access edge.

Systems that run Enhanced Layer 2 Software (ELS) support a slightly different configuration on ENode-facing interfaces than systems that do not run ELS. This section describes:

- [Non-ELS Port Mode for FCoE Interfaces on page 5536](#)
- [ELS Interface Mode for FCoE Interfaces on page 5536](#)
- [Trusted and Untrusted FCoE Interfaces on page 5536](#)

#### ***Non-ELS Port Mode for FCoE Interfaces***

The interfaces that belong to FCoE VLANs (interfaces that connect to CNAs in FCoE devices) on systems that do not support ELS should be configured in **tagged-access** port mode. After you enable VN2VF\_Port FIP snooping on an FCoE VLAN, the transit switch denies FCoE traffic from any ENode on that VLAN until the ENode performs a valid fabric login with an FC switch.

The **tagged-access** port mode was not available in Junos OS Release 11.3 and prior releases. In Release 11.3 and earlier, **trunk** port mode was used for Ethernet interfaces that connected to FCoE access devices. Because **tagged-access** mode is now available, using **trunk** mode for interfaces connected to FCoE CNAs is not recommended.

If an existing configuration uses **trunk** mode for ports connected to FCoE CNAs, you can change the port mode to **tagged-access** without disrupting traffic. Although we recommend changing the port mode of these ports from **trunk** to **tagged-access** as a best practice, it is not mandatory. New configurations should use **tagged-access** mode for interfaces that connect to FCoE devices.

#### ***ELS Interface Mode for FCoE Interfaces***

The interfaces that belong to FCoE VLANs (interfaces that connect to CNAs in FCoE devices) on systems that support ELS should be configured in **trunk** interface mode. After you enable VN2VF\_Port FIP snooping on an FCoE VLAN, the transit switch denies FCoE traffic from any ENode on that VLAN until the ENode performs a valid fabric login with an FC switch.

#### ***Trusted and Untrusted FCoE Interfaces***

Do not configure ENode-facing interfaces as FCoE trusted interfaces when VN2VF\_Port FIP snooping is enabled on those interfaces. If you enable VN2VF\_Port FIP snooping on an FCoE VLAN and you configure ENode-facing interfaces that are members of the FIP snooping VLAN as **fcoe-trusted**, then FCoE devices might not be able to log in to the FC network.

Changing ports from untrusted to trusted removes any existing VN2VF\_Port FIP snooping filters from the ports and terminates the existing sessions. Changing the fabric ports from trusted to untrusted forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN\_Ports log in again, the switch can build the appropriate VN2VF\_Port FIP snooping filters.

### Network-Facing Interfaces

When the switch acts as an FCoE transit switch, you must configure any interface that is connected to a switch as an FCoE trusted interface in **trunk** port mode and as a 10-Gigabit Ethernet interface.

Switch-facing Ethernet interfaces have the following requirements and behaviors:

- You must explicitly configure switch-facing trunk ports on an FCoE transit switch as FCoE trusted interfaces.
- After you configure an FC switch-facing trunk port as a trusted interface, the FCoE transit switch always processes FC switch frames because they come from a source on a trusted interface.
- All ports in an FCoE VLAN must be configured as tagged access or trunk ports.

### FC-MAP

When the switch acts as an FCoE transit switch and you enable VN2VF\_Port FIP snooping on an FCoE VLAN, you can optionally specify a 24-bit FCoE mapped address prefix (FC-MAP) value. On a given VLAN, the transit switch learns only those FC switches that have a matching FC-MAP value. If the transit switch FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, the transit switch does not discover the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. An FCoE VLAN can have one and only one FC-MAP value.

The FC-MAP value is a MAC address prefix unique to an FC switch in the FC SAN fabric that the FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). The FC switch combines the FC-MAP value with a unique 24-bit FCID value for the ENode VN\_Port during the login process. This creates a 48-bit identifier that is unique to the fabric. The FC switch assigns this 48-bit value to the ENode VN\_Port as its MAC address and unique identifier for the session. Each VN\_Port session the ENode establishes with the FC switch receives a unique FCID from the FC switch, so an FCoE device can host multiple virtual links (one for each VN\_Port) to an FC switch, each with a 48-bit MAC address that is unique to the fabric.

The VN2VF\_Port FIP snooping filter compares the configured FC-MAP value with the FC-MAP value in the header of frames coming from the ENode VN\_Port. If the values do not match, the transit switch denies access.



**NOTE:** Changing the FC-MAP value causes all logins to be dropped and forces ENodes to log in again.



NOTE: Do not configure static MAC addresses with the FC-MAP value as a prefix (the first 24 bits of the MAC address). If you configure a static MAC address that uses the FC-MAP value as a prefix, the system deletes the static MAC address automatically after you enable FIP snooping. The static MAC address configuration is not restored even if you disable FIP snooping later. (The system considers a static MAC address with the FC-MAP value as the prefix to be a misconfiguration.) Do not use a MAC address with the FC-MAP value as the prefix for any traffic other than the FIP snooping traffic when the switch is acting as a transit switch.

---

### T11 VN2VF\_Port FIP Snooping Specification

---

For more details about VN2VF\_Port FIP snooping, see <http://www.t11.org/ftp/t11/pub/fc/bb-5/08-264v3.pdf> for the Technical Committee T11 organization document *Increasing FCoE Robustness using FIP Snooping*.

#### Related Documentation

- [Overview of Fibre Channel on page 5508](#)
- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Understanding an FCoE-FC Gateway](#)
- [Overview of FIP on page 5513](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)
- [Understanding FCoE LAGs](#)
- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5546](#)
- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)
- [Disabling VN2VF\\_Port FIP Snooping on an FCoE-FC Gateway Switch Interface](#)
- [Disabling Enhanced FIP Snooping Scaling](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)
- [Configuring an FCoE LAG](#)
- [Understanding Fibre Channel Terminology on page 5569](#)

## Understanding VN\_Port to VN\_Port FIP Snooping on an FCoE Transit Switch

VN\_Port to VN\_Port (VN2VN\_Port) Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping (FC-BB-6) on an FCoE transit switch is conceptually similar to VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (FC-BB-5) on an FCoE transit switch. An FCoE transit switch is a data center bridging (DCB) switch with FIP snooping capability. VN2VN\_Port FIP snooping provides security in the form of filters. The filters help prevent unauthorized access and data transmission on a bridge that connects ENodes on the Ethernet network.

The main difference between VN2VN\_Port FIP snooping and VN2VF\_Port FIP snooping is that you use VN2VN\_Port FIP snooping when the FCoE devices reside on the Ethernet network, so there is no need to forward traffic between FCoE devices to the Fibre Channel (FC) network, and you use VN2VF\_Port FIP snooping when FCoE devices on the Ethernet network need to access targets on the FC network, so FCoE traffic must be forwarded to the FC network. See “[Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch](#)” on page 5531 for information about VN2VF\_Port FIP snooping.

You enable VN2VN\_Port FIP snooping on the FCoE VLAN that transports the VN2VN traffic. The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

A key benefit of VN2VN\_Port FIP snooping is that it enables FCoE initiators and targets to communicate directly through the switch without going through an FCoE forwarder (FCF) or an FC switch. The transit switch does not differentiate between initiators and targets because the transit switch sees both VN\_Ports as FIP virtual link end points. Direct VN2VN\_Port communication requires secure access (FIP snooping filters) because ENodes are not trusted entities.

This topic describes:

- [VN2VN\\_Port FIP Snooping and FIP Snooping Virtual Links](#) on page 5539
- [VN2VN\\_Port Communication Modes](#) on page 5540
- [Network Security](#) on page 5541
- [VN2VN\\_Port FIP Snooping Functions](#) on page 5541
- [Scalability](#) on page 5541
- [VN2VN\\_Port FIP Snooping Implementation](#) on page 5541
- [ENode-Facing Interfaces](#) on page 5542
- [Network-Facing Interfaces \(Connecting to Another Transit Switch\)](#) on page 5543
- [Beacon Period \(VN2VN\\_Port FIP Snooping Link Maintenance\)](#) on page 5544
- [QFabric System Differences in VN2VN\\_Port FIP Snooping Traffic Handling](#) on page 5544

### VN2VN\_Port FIP Snooping and FIP Snooping Virtual Links

FIP snooping under the T11 FC-BB-5 specification requires that an FC switch or an FCF be in the path between two VN\_Ports when they communicate. Introduced in the T11 FC-BB-6 specification (see <http://www.t11.org/ftp/t11/pub/fc/bb-6/10-019v3.pdf>), VN2VN\_Port FIP snooping allows the FCoE transit switch to connect two VN\_Ports to

each other directly, without going through an FC switch or an FCF, provided that the ENodes have logged in to the FC network.

In VN2VF\_Port FIP snooping, when an ENode logs in to the FC network, the FCoE transit switch snoops the FIP communication between the ENode and the FC switch. In VN2VN\_Port FIP snooping mode, the transit switch creates filters on the switch access ports to control VN\_Port access to other VN\_Ports on the Ethernet network. The VN2VN\_Port FIP snooping filters allow the switch to establish a dedicated virtual link that emulates a point-to-point connection between two VN\_Ports, through the switch.

Virtual links pass transparently through the transit switch. The VN\_Ports do not detect the transit switch, and virtual links appear to be direct point-to-point links.

You explicitly enable VN2VN\_Port FIP snooping on FCoE VLANs when the switch or QFabric system Node device is an FCoE transit switch connecting FCoE devices on the Ethernet network to each other and to FC switches or gateways at the FC storage area network (SAN) edge.



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port FIP snooping VLANs, VN\_Port to VF\_Port traffic is dropped.

When you enable FIP snooping, the system snoops VN2VF\_Port packets and enforces security only on VN\_Port to VF\_Port virtual links. When you enable VN2VN\_Port FIP snooping, the system snoops VN\_Port to VN\_Port FIP packets and enforces security only on VN\_Port to VN\_Port virtual links.

The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN\_Port FIP snooping. VN2VN\_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

---

### VN2VN\_Port Communication Modes

---

The transit switch supports two VN2VN\_Port communication modes:

- Point-to-point mode
- Multipoint mode

In point-to-point mode, two ENodes are connected to the network and form a single VN\_Port to VN\_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.

In multipoint mode, multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN\_Ports. This is analogous to loop mode in traditional FC networks.

The VN2VN\_Port communication mode is not configured; it is determined by the number of ENodes connected to the network.

### Network Security

In traditional FC networks, the FC switch is usually a trusted entity and the server ENodes are untrusted entities. The ENodes connect directly to the FC switch VF\_Ports. After an ENode gains access to the network through the fabric login (FLOGI) process, the FC switch enforces zoning configurations, ensures that the ENode uses valid addresses, monitors the connection, and performs other security functions to prevent unauthorized access.

However, FCoE exposes FC frames to Ethernet networks, which do not have the same level of security as native FC networks. VN2VN\_Port FIP snooping filters emulate the native FC network security functions by preventing unauthorized access and by ensuring the security of the virtual link between ENode VN\_Ports. The transit switch performs VN2VN\_Port FIP snooping at the ports connected to the FCoE VN\_Port devices.

### VN2VN\_Port FIP Snooping Functions

When you enable VN2VN\_Port FIP snooping, the transit switch sets and applies filters to block all FCoE traffic on the VLAN by default. The transit switch monitors FIP logins, solicitations, and advertisements that pass through it and gathers information about the ENode address. The transit switch uses the information to construct filters that permit access only to logged-in ENodes. All other traffic on the VLAN is denied.

The filters enable FCoE frames to pass through the transit switch only on a virtual link established between two VN\_Ports. The filters ensure that ENodes can only connect to other ENodes if they have successfully logged in to each other, and that only valid FCoE traffic along valid paths is transmitted. VN2VN\_Port FIP snooping maintains the filters by tracking VN\_Port to VN\_Port sessions.

### Scalability

Because ENodes are untrusted and the system needs to apply filters to untrusted FIP snooping interfaces, the total number of combined VN2VN\_Port FIP snooping sessions per switch is 376 sessions (ENode to ENode sessions) on untrusted interfaces. On interfaces that are configured as trusted interfaces, no FIP snooping filters are applied.



**NOTE:** The total number of sessions the system can support is the combined number of VN2VF\_Port sessions and VN2VN\_Port sessions. If VN2VF\_Port sessions are active, the total number of available VN2VN\_Port sessions is reduced.

### VN2VN\_Port FIP Snooping Implementation

You enable VN2VN\_Port FIP snooping on a per-VLAN basis on VLANs that carry FCoE traffic. The switch snoops FIP frames at the ports associated with FCoE VLANs enabled for VN2VN\_Port FIP snooping. The switch then installs the resulting filters on the ENode-facing ports to ensure that all FIP snooping occurs on the switch network edge.

VN2VN\_Port FIP snooping FCoE VLANs must meet the following criteria:

- An FCoE VLAN should be dedicated to FCoE traffic only.
- An FCoE VLAN cannot support both VN2VF\_Port FIP snooping (FC-BB-5) and VN2VN\_Port FIP snooping (FC-BB-6) simultaneously. You must configure separate FCoE VLANs for FIP snooping traffic and for VN2VN\_Port FIP snooping traffic.



**NOTE:** Changing an FCoE VLAN from VN2VF\_Port FIP snooping mode to VN2VN\_Port FIP snooping mode terminates the existing virtual links on the VLAN. The transit switch removes the existing FIP snooping filters, creates the new FIP snooping filters, and applies them to the FIP snooping ports. If you downgrade the software to Junos OS Release 12.1 or earlier, VLANs configured for VN2VN\_Port FIP snooping revert to VN2VF\_Port FIP snooping VLANs.

- For switches that do not run Enhanced Layer 2 Software (ELS), as a best practice, you should configure all access ports that belong to an FCoE VLAN (ports connected to a converged network adapter [CNA] in an FCoE device) in **tagged-access** port mode. However, access and trunk port modes are also supported. For switches that use ELS, configure access ports that belong to an FCoE VLAN in **trunk** interface mode.
- Access ports should be configured as untrusted ports.
- All ports connected to another transit switch must be configured in **trunk** port mode.
- FIP traffic uses the native VLAN.
- You can enable VN2VN\_Port FIP snooping on a native VLAN.

---

### ENode-Facing Interfaces

When the interfaces that belong to an FCoE VLAN connect directly to FCoE devices (there is no other transit switch between the FCoE devices and the switch), we recommend that you either enable VN2VN\_Port FIP snooping on all FCoE VLANs to ensure secure connections between VN\_Ports, or enable VN2VF\_Port FIP snooping on FCoE VLANs that connect ENodes to an FC switch. FIP snooping should always be enabled at the access edge.

Systems that run Enhanced Layer 2 Software (ELS) support a slightly different configuration on ENode-facing interfaces than systems that do not run ELS. This section describes:

- [Non-ELS Port Mode for FCoE Interfaces on page 5542](#)
- [ELS Interface Mode for FCoE Interfaces on page 5543](#)
- [Trusted and Untrusted FCoE Interfaces on page 5543](#)

#### **Non-ELS Port Mode for FCoE Interfaces**

The interfaces that belong to FCoE VLANs (interfaces that connect to CNAs in FCoE devices) should be configured in **tagged-access** port mode, unless your CNA does not support tagged VN2VN traffic. After you enable VN2VN\_Port FIP snooping on an FCoE



VLAN, the transit switch denies FCoE traffic from any ENode on that VLAN until the ENode performs a valid fabric login (FIP FLOGI) with another ENode.

The **tagged-access** port mode was not available in Junos OS Release 11.3 and prior releases. In Release 11.3 and earlier, **trunk** port mode was used for Ethernet interfaces that connected to FCoE access devices. Because **tagged-access** mode is now available, using **trunk** mode for interfaces connected to FCoE CNAs is not recommended.

If an existing configuration uses **trunk** mode for ports connected to FCoE CNAs, you can change the port mode to **tagged-access** without disrupting traffic. Although we recommend changing the port mode of these ports from **trunk** to **tagged-access** as a best practice, it is not mandatory. New configurations should use **tagged-access** mode for interfaces that connect to FCoE devices.

### ***ELS Interface Mode for FCoE Interfaces***

The interfaces that belong to FCoE VLANs (interfaces that connect to CNAs in FCoE devices) on systems that support ELS should be configured in **trunk** interface mode. After you enable VN2VF\_Port FIP snooping on an FCoE VLAN, the transit switch denies FCoE traffic from any ENode on that VLAN until the ENode performs a valid fabric login with an FC switch.

### ***Trusted and Untrusted FCoE Interfaces***

Do not configure ENode-facing interfaces as FCoE trusted interfaces when VN2VF\_Port FIP snooping is enabled on those interfaces. If you enable VN2VF\_Port FIP snooping on an FCoE VLAN and you configure ENode-facing interfaces that are members of the FIP snooping VLAN as **fcoe-trusted**, then FCoE devices might not be able to log in to the FC network.

Changing ports from untrusted to trusted removes any existing VN2VF\_Port FIP snooping filters from the ports and terminates the existing sessions. Changing the fabric ports from trusted to untrusted forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN\_Ports log in again, the switch can build the appropriate VN2VF\_Port FIP snooping filters.

### ***Network-Facing Interfaces (Connecting to Another Transit Switch)***

Configure any interface that is connected to another transit switch (not to an ENode) as an FCoE trusted interface, in **trunk** port mode, and as a 10-Gigabit Ethernet interface.

Network-facing Ethernet interfaces have the following requirements and behaviors:

- You must explicitly configure network-facing trunk ports on an FCoE transit switch as FCoE trusted interfaces.
- After you configure a network-facing trunk port as a trusted interface, the FCoE transit switch always processes frames from the connected switch because they come from a source on a trusted interface.
- As a best practice, configure ports in an FCoE VLAN as tagged access ports, but access and trunk port modes are also supported to accommodate whatever types of VN2VN traffic your CNA supports.

### Beacon Period (VN2VN\_Port FIP Snooping Link Maintenance)

---

The transit switch needs to maintain the virtual links between VN\_Ports, and needs to know when sessions begin and end, and when to install and remove the FIP snooping filters. FIP snooping uses a FIP keepalive advertisement to accomplish this task. VN2VN\_Port FIP snooping does not exchange FIP keepalive timer information. Instead, you configure a *beacon period*, which performs the same function as a keepalive timer.

The beacon period is the time interval between messages which verify that the connection is still valid and that the device at the other end of the virtual link is still reachable. You set the beacon period value for each FCoE VLAN that you configure to do VN2VN\_Port FIP snooping.



**NOTE:** Explicitly set the beacon period when you configure VN2VN\_Port FIP snooping. VN\_Ports do not automatically send beacons.

---

ENodes transmit periodic multicast N\_Port\_ID beacons to the ALL-VN2VN-ENode-MACs address. The transmission period varies by a random delay of between 0 ms and 100 ms to avoid synchronized bursts of multicast traffic on the network.

If the transit switch does not receive a beacon message from an ENode within 2.5 times the configured beacon period, the transit switch considers the virtual link to be down and terminates the virtual link to that ENode.

### QFabric System Differences in VN2VN\_Port FIP Snooping Traffic Handling

---

Configuring VN2VN\_Port FIP snooping on a QFabric system is the same as configuring VN2VN\_Port FIP snooping on a standalone switch. However, there are internal differences in the way a QFabric system handles VN2VN\_Port FIP snooping traffic compared to the way a standalone switch handles VN2VN\_Port FIP snooping traffic. The internal differences are transparent. Whether you configure VN2VN\_Port FIP snooping on a QFabric system or on a standalone switch, the proper FIP snooping filters and forwarding information are installed on each device.


On standalone switches, the VN2VN\_Port FIP snooping traffic does not cross a fabric (Interconnect device). VN2VN\_Port traffic enters and exits ports on a single switch, so the ingress port and the egress port have access to the same *local* forwarding and FIP snooping databases.

However, on a QFabric system, VN2VN\_Port FIP snooping traffic might enter on the ingress port of one Node device, traverse the Interconnect device fabric, and exit on the egress port of a different Node device. In this case, the QFabric system must ensure that the FIP snooping database and forwarding information for the VN2VN\_Port traffic is installed correctly on both of the Node devices so that traffic is correctly filtered and forwarded.

For example, [Figure 193 on page 5545](#) shows that VN2VN\_Port traffic from FCoE host ENode E1 enters the QFabric system at Node device ND1, traverses the Interconnect device fabric, and then exits from Node device ND2 before arriving at FCoE host ENode E2. Similarly,

VN2VN\_Port traffic from FCoE host ENode E2 enters the QFabric system at Node device ND2, traverses the Interconnect device fabric, and then exits from Node device ND1 before arriving at FCoE host ENode E1.

**Figure 193: VN2VN\_Port Traffic Across a QFabric Interconnect Device**



When the QFabric system receives a FLOGI ACC from either ENode E1 or ENode E2, the QFabric system creates and installs the correct VN2VN\_Port FIP snooping filters on both Node devices, and updates the forwarding tables accordingly.

In addition, the QFabric system must also ensure that the VN2VN\_Port FIP snooping session statistics are correctly counted. Even though a session is running on each of the two Node devices, the QFabric system counts the complete VN2VN\_Port connection as one session because the two Node devices belong to the same session. This ensures that VN2VN\_Port sessions that traverse the Interconnect device fabric are counted as one unique session, not as two separate sessions.

#### Related Documentation

- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Overview of FIP on page 5513](#)
- [Understanding Fibre Channel Terminology on page 5569](#)
- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5546](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\)](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\)](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\)](#)
- [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)

## Understanding FIP Snooping, FBF, and MVR Filter Scalability

The VLAN filter processor (VFP) ternary content addressable memory (TCAM) stores the VLAN filter configuration for three filter types:

- Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping—FIP snooping filters prevent an FCoE device from gaining unauthorized access to a Fibre Channel (FC) storage device or to another FCoE device. VN2VF\_Port FIP snooping filters prevent an FCoE device from gaining unauthorized access to devices on an FC network. VN2VN\_Port FIP snooping filters prevent an FCoE device from gaining unauthorized access to another FCoE device directly through the standalone switch or QFabric system, without traversing the FC network.

The VFP TCAM stores the VN2VF\_Port and VN2VN\_Port FIP snooping filters that the switch automatically creates when you enable FIP snooping on a VLAN that carries FCoE traffic. See [“Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch” on page 5531](#) and [“Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch” on page 5539](#) for more information.

- Filter-based forwarding (FBF)—FBF enables you to use firewall filters to direct packets to virtual routing instances. The switch then forwards the matching packets based on the configuration of the routing instances. The VFP TCAM stores the terms you configure for FBF filters. See [“Understanding Filter-Based Forwarding” on page 5212](#) for more information.
- Multicast VLAN registration (MVR)—MVR enables you to configure a multicast source VLAN (MVLAN) that is shared across a Layer 2 network. An MVLAN distributes IPTV multicast streams across different VLANs without having to create a separate multicast stream for each VLAN, and without compromising the security and separation of traffic in the different VLANs. The VFP TCAM stores the MVR rules you configure for MVLANS. See [“Understanding Multicast VLAN Registration” on page 4773](#) for more information.

FIP snooping filters, FBF filters, and MVR rules share the VFP TCAM memory space. In most use cases, the VFP TCAM memory is sufficient to store filter terms and information for all three applications.

- [VFP TCAM Architecture and Allocation on page 5546](#)
- [VFP TCAM Entry Consumption on page 5547](#)
- [Rejected Filter Configurations \(No Available VFP TCAM Space\) on page 5550](#)
- [VFP TCAM Allocation and Consumption \(Scaling\) Examples on page 5551](#)
- [Filter Configuration Recommendations on page 5553](#)

---

### VFP TCAM Architecture and Allocation

When packets arrive at an ingress interface, the VFP TCAM is the first TCAM in the packet pipeline. The VFP TCAM stores a total of 1024 entries. The 1024 entries are partitioned into four equal *slices* of 256 entries.

The VFP TCAM allocates entries to three filter types (FIP snooping filters, FBF filter terms, and MVR rules) in 256-entry slices. The VFP TCAM dynamically allocates the minimum number of memory slices required to store the filters for a particular filter type, as needed.

The TCAM does not allocate partial slices to a filter type, and slices cannot be shared among filter types. At any given time, each slice contains entries for one and only one filter type.

For example, if you configure one MVR rule, the system allocates a whole slice to MVR rules, even if the MVR rule consumes only one TCAM entry. The remaining 256 entries in the slice allocated to MVR rules can store subsequently configured MVR rules, but not FIP snooping or FBF filters. Similarly, if FIP snooping filters consume 50 entries of a 256-entry slice, the remaining 206 entries in the FIP snooping slice are available only to store more FIP snooping filters, not to store FBF filter terms or MVR rules.

The VFP TCAM allocates slices to a filter type only if there is at least one configured filter or rule for that filter type. If no filters exist for a filter type, then the VFP TCAM does not allocate a slice to that filter type.



**NOTE:** The VFP TCAM rejects partial filters. For example, if an FBF filter contains six terms, but there is only space in the TCAM for four of those terms, the whole filter is not committed.

Each filter type can use from zero slices to all four slices of VFP TCAM space. However, if one filter type uses three slices, then only one slice remains, so only one other filter type can use the remaining slice. In that situation, if you configure filters for all three filter types, the last filter type that you configure receives no TCAM space for its filter entries. Filters that receive no TCAM entry space are not implemented.

### VFP TCAM Entry Consumption

FIP snooping filters, FBF filters, and MVR rules consume VFP TCAM entry space in different ways:

- [FIP Snooping Filter VFP TCAM Consumption on page 5547](#)
- [FBF Filter VFP TCAM Consumption on page 5548](#)
- [MVR Filter VFP TCAM Consumption on page 5549](#)
- [VFP TCAM Consumption Summary Table on page 5549](#)

#### *FIP Snooping Filter VFP TCAM Consumption*

VN2VF\_Port FIP snooping filters consume VFP TCAM entry space differently than VN2VN\_Port FIP snooping filters:

- [VN2VF\\_Port FIP Snooping Filter VFP TCAM Consumption on page 5548](#)
- [VN2VN\\_Port FIP Snooping Filter VFP TCAM Consumption on page 5548](#)



**NOTE:** One FCoE VLAN cannot support both VN2VF\_Port traffic and VN2VN\_Port traffic. Configure separate FCoE VLANs for VN2VF\_Port traffic and for VN2VN\_Port traffic.

### ***VN2VF\_Port FIP Snooping Filter VFP TCAM Consumption***

The switch uses an algorithm that allows one 256-entry slice of the VFP TCAM to store the maximum possible number of VN2VF\_Port FIP snooping filters (2500 filters). VN2VF\_Port FIP snooping filters never consume more than one slice of the VFP TCAM.

Regardless of whether there is one VN2VF\_Port FIP snooping session or there are 2500 VN2VF\_Port FIP snooping sessions, VN2VF\_Port FIP snooping filters consume one slice of the VFP TCAM. (If there are no VN2VF\_Port or VN2VN\_Port FIP snooping sessions, the TCAM does not allocate a slice for FIP snooping filters.)

### ***VN2VN\_Port FIP Snooping Filter VFP TCAM Consumption***

VN2VN\_Port FIP snooping filters consume one VFP TCAM entry for each VN2VN\_Port session. The maximum number of VN2VN\_Port FIP snooping sessions is 376 sessions per switch. (If you configure an interface that carries VN2VN\_Port FIP snooping traffic as a trusted interface, the switch does not apply filters on the trusted interface.)

Because the switch can have up to 376 VN2VN\_Port sessions running simultaneously, with each session consuming one entry, VN2VN\_Port FIP snooping filters consume VFP TCAM space as follows:

- 1–256 filters consume one slice
- 257–376 filters consume two slices

### ***FBF Filter VFP TCAM Consumption***

Each FBF filter term is double-wide, so each FBF filter term consumes two entries in the VFP TCAM. One 256-entry slice can contain up to 128 FBF filter terms. FBF filters consume VFP TCAM space as follows:

- 1–128 entries consume one slice
- 129–256 entries consume two slices
- 257–384 entries consume three slices
- 385–512 entries consume four slices



**NOTE:** In practice, FBF filters can consume only three slices of the VFP TCAM because FBF filters are also stored simultaneously in the ingress filter processor (IFP) TCAM, and the IFP TCAM can store only 384 FBF filter terms (768 entries, or 3 TCAM slices).

For example, if you configure FBF filters that contain 200 terms, then the FBF filters require 400 VFP TCAM entries and consume 2 slices.

FBF filter entries are simultaneously stored in the VFP TCAM and the IFP TCAM. The IFP TCAM can only contain up to 768 entries—256 fewer entries (1 slice) than the VFP TCAM. As with the VFP TCAM, FBF filters consume two IFP TCAM entries per filter term. In addition to FBF filter terms, the IFP TCAM stores filter entries for firewall filters.



**CAUTION:** There must be enough space in the VFP TCAM *and* the IFP TCAM for the FBF filter entries. If both TCAMs do not have enough space for the FBF filters, the switch rejects the portion of the configuration that it cannot store and sends a syslog message to notify you.

For example, if you configure FBF filters that have 400 terms, even though the VFP TCAM has enough space to store the resulting 800 entries, the switch rejects a portion of the configuration because the IFP TCAM can store a maximum of only 768 entries. If the IFP TCAM stores no other filter entries, the switch rejects 32 FBF filter entries.

In another example, if you configure firewall filters that have a total of 200 terms, which consume 200 entries in the IFP TCAM, and you then configure FBF filters that have a total of 300 terms, the switch rejects a portion of the configuration because the FBF filters require 600 entries. Combined with the 200 entries required for the firewall filters, the total number of 800 entries exceeds the maximum of 768 entries that the IFP TCAM can store. In this case, the switch accepts the first 768 entries and rejects the rest of the filter entries. The switch installs the filter entries in the order that they are committed; the rejected entries are the last entries the switch attempts to commit after the TCAM space is exhausted.

The IFP TCAM limit of 768 entries means that the true maximum number of FBF filter terms is 384 terms, even though the VFP TCAM can store up to 512 FBF terms.

#### ***MVR Filter VFP TCAM Consumption***

Each MVR rule consumes one entry in the VFP TCAM, so MVR rules consume VFP TCAM space as follows:

- 1–256 rules consume one slice
- 257–512 rules consume two slices
- 513–758 rules consume three slices
- 759–1024 rules consume four slices

#### ***VFP TCAM Consumption Summary Table***

Table 448 on page 5550 summarizes VFP TCAM consumption.



**NOTE:** FBF filters are simultaneously stored in the VFP TCAM and in the IFP TCAM. Due to the IFP TCAM limit of 768 entries (384 FBF filters), which is 256 entries fewer than the VFP TCAM, the effective VFP TCAM consumption limit for FBF filters is lower than the total amount of VFP TCAM entry space, even when no other filters consume VFP TCAM space.

Table 448: VFP TCAM Entry Consumption Summary

| Filter Type                     | VFP TCAM Entry Consumption         | Maximum VFP TCAM Slices Consumed             | Other Limitations                        |
|---------------------------------|------------------------------------|--|--|
| VN2VF_Port FIP snooping filters | Never consumes more than one slice | One slice (regardless of number of sessions) | 2500 session maximum                     |
| VN2VN_Port FIP snooping filters | One entry per session              | Two  | 376 session maximum                      |
| FBF filters                     | Two entries per filter             | Three (due to IFP TCAM limitation)           | 384 filters (due to IFP TCAM limitation) |
| MVR rules                       | One entry per rule                 | Four   | 1024 rule maximum                        |

#### Rejected Filter Configurations (No Available VFP TCAM Space)

If there is not enough space available in the VFP TCAM to store the FIP snooping filters, the configured FBF filters, and the MVR rules, the switch rejects only the portion of the configuration that it cannot store. Any portion of the filter configuration that the TCAM can store, is stored. In most cases, even if the switch rejects part of the configuration, part of the configuration is also stored.

If the switch rejects any portion of a configuration, the switch sends a syslog message to notify you of the failure. The switch does not generate a commit error, and the rejected portion of the configuration remains on the switch, even though the rejected configuration does not function. (The accepted portions of the configuration function as expected.) The syslog message shows you the filter configuration that the switch rejected.

We strongly recommend that you always delete rejected filter configurations from the switch. It is important to delete rejected filter configurations because:

- Even though the rejected configuration remains on the switch, it does not function.
- After a reboot, there is no guarantee that the same filters will be rejected. The previously rejected filters might be accepted, and other filters that had previously been accepted might be rejected. Therefore, the functioning filter configuration could be changed inadvertently and unexpectedly.
- Even if a VFP TCAM slice becomes available, the switch does not automatically allocate the available slice to the rejected configuration. To use the available slice, you must delete and reconfigure the rejected configuration.

For example, you configure FBF filters and MVR rules on a switch, and that switch also transports FCoE traffic with VN2VF\_Port FIP snooping (never consumes more than one slice) enabled on FCoE access interfaces. After you commit the configuration, you check the syslog. You find that the VN2VF\_Port FIP snooping and FBF filters consume all four slices of the VFP TCAM, and the MVR configuration was rejected. Instead of deleting the MVR configuration, you leave it on the switch. Subsequently, all VN2VF\_Port FIP snooping sessions end, the FIP snooping filters time out and are removed from the VFP TCAM, so the slice that was allocated to VN2VF\_Port FIP snooping filters becomes free. However, the MVR rules do *not* automatically receive the free slice.



To force the switch to allocate the free slice to the MVR rules, you should delete the MVR rules from the configuration and then reconfigure the MVR rules. When you commit the new configuration, check the syslog messages to ensure that the MVR rule configuration was accepted.

In this example, you could also choose to free a VFP TCAM slice for MVR rule storage by deleting some of the FBF filters. To do this, you delete both the unneeded FBF filters and the MVR rule configuration. Then you reconfigure the MVR rules, and check the syslog to ensure that the configuration was successful.

### VFP TCAM Allocation and Consumption (Scaling) Examples

The following examples illustrate how FIP snooping entries, FBF filter entries, and MVR rule entries consume VFP TCAM slices:

- [Example 1: Three Filter Types Consume Three Slices on page 5551](#)
- [Example 2: Three Filter Types Consume Four Slices on page 5551](#)
- [Example 3: Two Filter Types Consume Four Slices on page 5552](#)
- [Example 4: Three Filter Types Oversubscribe the VFP TCAM on page 5552](#)

#### **Example 1: Three Filter Types Consume Three Slices**

Filters and rules are configured in the following sequence:

- 100 VN2VN\_Port FIP snooping filters (1 slice)
- 2 MVR rules (1 slice, 2 entries)
- 60 FBF filter terms (1 slice, 120 entries)

One slice remains free. The slice allocated to VN2VN\_Port FIP snooping filters can store 156 more filters before another slice is required. The slice allocated to MVR rules can store 254 more rules before another slice is required. The slice allocated to FBF filters can store 68 more filter terms (136 entries) before another slice is required. Providing that the IFP TCAM has space for the FBF filter terms, the switch accepts this configuration and rejects no filters.

#### **Example 2: Three Filter Types Consume Four Slices**

Filters and rules are configured in the following sequence:

- 2000 VN2VF\_Port FIP snooping filters (always 1 slice)
- 18 MVR rules (1 slice, 18 entries)
- 150 FBF filter terms (2 slices, 300 entries)

All four slices are allocated to filter types. The slice allocated to MVR rules can store 238 more rules before it is full. The slice allocated to FBF filters can store 106 more filter terms (212 entries) before it is full. Providing that the IFP TCAM has space for the FBF filter terms, the switch accepts this configuration and rejects no filters.



**NOTE:** If you configure more MVR rules or FBF filters than entry space remaining in the slices, the switch rejects those rules and filters because no slice is available. The switch installs filters in the order that they were configured, so if filters are rejected, the filters configured last are rejected.

---

**Example 3: Two Filter Types Consume Four Slices**

Filters and rules are configured in the following sequence:

- 50 VN2VF\_Port FIP snooping filters (always 1 slice)
- 300 FBF filter terms (3 slices, 600 entries)

All four slices are allocated to filter types. No slices are available for MVR rules. The third slice allocated to FBF filters can store 84 more filter terms (168 entries) before it consumes all of its entry space. Providing that the IFP TCAM has space for the FBF filter terms, the switch accepts this configuration and rejects no filters.



**NOTE:** If you configure MVR rules or if you configure more than 84 more FBF filters, the switch rejects those rules and filters because no slice is available for the MVR rules, and the FBF filter slice has entry space for only 84 more filter terms.

---

**Example 4: Three Filter Types Oversubscribe the VFP TCAM**

Filters and rules are configured in the following sequence:

- 1750 VN2VF\_Port FIP snooping filters (always 1 slice)
- 10 MVR rules (1 slice, 10 entries)
- 275 FBF filter terms (2 slices, 512 accepted entries, 38 rejected entries)

All four slices are allocated to filter types. The slice allocated to MVR rules can store 246 more rules before it is full, but the number of FBF filter terms exceeds the amount of available VFP TCAM storage space. (The 275 FBF filter terms consume 550 VFP TCAM entries. However, there are only two available slices, for a total of 512 available entry spaces, so only 256 FBF filter terms can be stored, leaving 19 rejected FBF filter terms.)

The switch accepts the VN2VF\_Port FIP snooping filters, the MVR rules, and 256 FBF filter terms. The switch retains the excess FBF filters in the configuration, but does not install those filters in the VFP TCAM. In this case, you delete the rejected FBF filter terms from the configuration. Alternatively, you could delete the MVR rules from the configuration to free a slice of the TCAM, and then delete and reconfigure the rejected FBF filters so that the system allocates the freed slice to the FBF filters.



**NOTE:** The sequence of configuration makes a difference; if there is not enough VFP TCAM space for a given filter type, the switch installs the filters that fit in the order they are configured. For example, if you configure the FBF filters before you configure the MVR rules, the VFP TCAM allocates one slice to FIP snooping filters, three slices to FBF filters (assuming the IFP TCAM has available space), and no slices to MVR rules, because all four slices are allocated before the switch attempts to install the MVR rules in the VFP TCAM.

### Filter Configuration Recommendations

To utilize the VFP TCAM space most efficiently:

- [Configure and Maintain the Fewest Number of Filters Needed on page 5553](#)
- [Always Delete Rejected Filter Configurations on page 5554](#)

#### ***Configure and Maintain the Fewest Number of Filters Needed***

To conserve VFP TCAM entry space, and because FBF filter storage also depends on the availability of IFP TCAM space, we recommend that you configure as few FBF filters and MVR rules as is practical to serve your network needs. The more filters you configure, the greater the possibility of exceeding TCAM storage capacity.

Several factors determine VFP TCAM consumption:

- **Type of filters configured**—Different filter types consume different amounts of VFP TCAM space. VN2VF\_Port FIP snooping filters never consume more than one slice. MVR rules and VN2VN\_Port FIP snooping filters consume entries in a slice at a rate of one entry per MVR rule or VN2VN\_Port session. FBF filter terms consume entries in a slice at a rate of two entries per FBF filter term.
- **Number of filters configured**—Although the number of filters does not affect the number of slices allocated to the VN2VF\_Port FIP snooping filter type (it is always one slice for one or more VN2VF\_Port FIP snooping filters and no slice for no FIP snooping filters), the number of VN2VN\_Port FIP snooping filters, MVR rules, and FBF filter terms that you configure determine how many VFP TCAM slices are required for each filter type.

For example, if you configure 257 MVR rules, the MVR rule entries consume 2 slices. One slice stores 256 MVR rules (entries), and one slice stores 1 MVR rule (entry). In this case, if you can eliminate one MVR rule, you can free a slice to allocate to other filter types.

- **Sequence of filter configuration**—If you configure too many filters for the VFP TCAM to store, the last filters you configure are not stored in the TCAM.

Always check the syslog after you configure FBF filters or MVR rules to ensure that the configuration was not rejected. If you enable FIP snooping on access ports, check the syslog to ensure that the configuration was not rejected due to lack of VFP TCAM space.

If you check the syslog and a filter configuration has been rejected, delete the filters that were rejected from the configuration.



**TIP:** If you no longer need an FBF filter or an MVR rule, delete it from the configuration to conserve VFP TCAM space. Enable VN2VF\_Port or VN2VN\_Port FIP snooping on access ports only if the switch port is directly connected to FCoE devices. (FIP snooping should be performed at the access edge. FIP snooping should not be performed on traffic that has already been snooped and filtered at the access edge. If another switch that is physically between the transit switch (or QFabric system) and the FCoE devices already performs FIP snooping, you do not have to enable FIP snooping on the transit switch or QFabric system, but you can.)

---

### ***Always Delete Rejected Filter Configurations***

The switch does not return a commit error if it rejects any portion of a configuration. Instead, the switch sends a syslog message to report the rejected portion of the configuration. The rejected portion of the configuration remains on the switch, but does not function.

After you configure FBF filters or MVR rules, or enable FIP snooping, check the syslog messages to ensure that the switch accepted the configuration. If the switch rejected any portion of the configuration, delete that portion of the configuration. (You do not need to delete the portion of the configuration that was accepted, unless you want to reconfigure those filters or rules.)



**CAUTION:** If you do not delete rejected filter configurations, and if you reboot the system, you cannot predict which filters the system installs after the reboot. For example, a switch with the following configuration has more configured filters than the VFP TCAM can support:

- VN2VF\_Port FIP snooping sessions (always consumes one slice)
- 20 MVR rules (consume one slice)
- 300 FBF filters (attempt to consume three slices, but because only two slices are available, 256 filters consume two slices, and the remaining 44 filters are rejected)

If you do not delete the 44 rejected FBF filters, then if the switch reboots, the 44 FBF filters that were rejected might be accepted, and 44 different FBF filters might be rejected. This unpredictable behavior is the reason that you should check the syslog messages after you configure filters, and if any filters were rejected, you should always delete the rejected filters from the configuration.

---

### **Related Documentation**

- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)
- [Understanding Filter-Based Forwarding on page 5212](#)
- [Understanding Multicast VLAN Registration on page 4773](#)

- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)
- *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)*
- *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)*
- *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)*
- [Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on page 5279](#)
- [Configuring Multicast VLAN Registration \(CLI Procedure\)](#)

## Understanding MC-LAGs on an FCoE Transit Switch

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because MC-LAGs do not carry forwarding class and IEEE 802.1p priority information.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as passthrough transit switch ports.

Standalone switches support MC-LAGs. QFabric system Node devices do not support MC-LAGs. Virtual Chassis (VC) and mixed-mode Virtual Chassis Fabric (VCF) configurations do not support FCoE. Only pure QFX5100 VCFs (consisting of only QFX5100 switches) support FCoE.

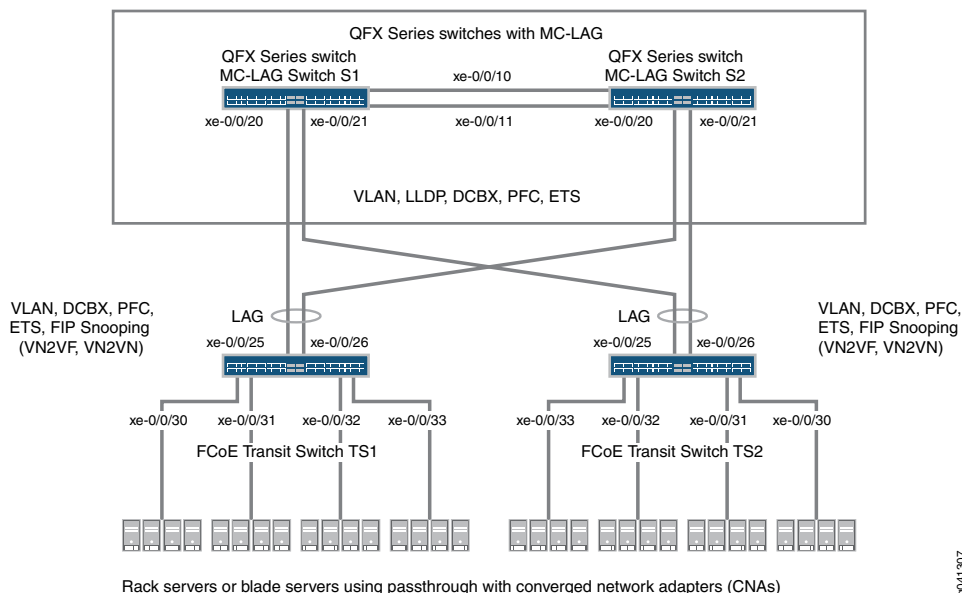
This topic describes:

- [Supported Topology on page 5555](#)
- [FIP Snooping and FCoE Trusted Ports on page 5557](#)
- [CoS and Data Center Bridging \(DCB\) on page 5558](#)

### Supported Topology

Switches that are not directly connected to FCoE hosts and that act as passthrough transit switches support MC-LAGs for FCoE traffic in an *inverted-U* network topology. [Figure 194 on page 5556](#) shows an inverted-U topology using QFX3500 switches.

Figure 194: Supported Topology for an MC-LAG on an FCoE Transit Switch



The following rules and guidelines apply to MC-LAGs when used for FCoE traffic. The rules and guidelines help ensure the proper handling and lossless transport characteristics required for FCoE traffic:

- The two switches that form the MC-LAG (Switches S1 and S2) cannot use ports that are part of an FCoE-FC gateway fabric. The MC-LAG switch ports must be passthrough transit switch ports (used as part of an intermediate transit switch that is not directly connected to FCoE hosts).
- MC-LAG Switches S1 and S2 cannot be directly connected to the FCoE hosts.
- The two switches that serve as access devices for FCoE hosts (FCoE Transit Switches TS1 and TS2) use standard LAGs to connect to MC-LAG Switches S1 and S2. FCoE Transit Switches TS1 and TS2 can be standalone switches or they can be Node devices in a QFabric system.
- Transit Switches TS1 and TS2 must use transit switch ports for the FCoE hosts and for the standard LAGs to MC-LAG Switches S1 and S2.
- Enable FIP snooping on the FCoE VLAN on Transit Switches TS1 and TS2. You can configure either VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping or VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping, depending on whether the FCoE hosts need to access targets in the FC SAN (VN2VF\_Port FIP snooping) or targets in the Ethernet network (VN2VN\_Port FIP snooping).

FIP snooping should be performed at the access edge and is not supported on MC-LAG switches. Do not enable FIP snooping on MC-LAG Switches S1 and S2. (Do not enable FIP snooping on the MC-LAG ports that connect Switches S1 and S2 to Switches TS1 and TS2 or on the LAG ports that connect Switch S1 to S2.)

- The CoS configuration must be consistent on the MC-LAG switches. Because MC-LAGs carry no forwarding class or priority information, each MC-LAG switch needs to have

the same CoS configuration to support lossless transport. (On each MC-LAG switch, the name, egress queue, and CoS provisioning of each forwarding class must be the same, and the priority-based flow control (PFC) configuration must be the same.)

### ***Transit Switches (Server Access)***

The role of FCoE Transit Switches TS1 and TS2 is to connect FCoE hosts in a multihomed fashion to the MC-LAG switches. In essence, Transit Switches TS1 and TS2 act as access switches for the FCoE hosts. (FCoE hosts are directly connected to Transit Switches TS1 and TS2.)

The transit switch configuration depends on whether you want to do VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, and whether the transit switches also have ports configured as part of an FCoE-FC gateway virtual fabric. Ports that a QFX3500 switch uses in an FCoE-FC gateway virtual fabric cannot be included in the transit switch LAG connection to the MC-LAG switches. (Ports cannot belong to both a transit switch and an FCoE-FC gateway; you must use different ports for each mode of operation.)

### ***MC-LAG Switches (FCoE Aggregation)***

The role of MC-LAG Switches S1 and S2 is to provide redundant, load-balanced connections between FCoE transit switches. In essence, MC-LAG Switches S1 and S2 act as aggregation switches. FCoE hosts are not directly connected to the MC-LAG switches.

The MC-LAG switch configuration is the same regardless of which type of FIP snooping that FCoE Transit Switches TS1 and TS2 perform.

### **FIP Snooping and FCoE Trusted Ports**

To maintain secure access, enable VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping at the transit switch access ports connected directly to the FCoE hosts. FIP snooping should be performed at the access edge of the network to prevent unauthorized access. For example, in [Figure 194 on page 5556](#), you enable FIP snooping on the FCoE VLANs on Transit Switches TS1 and TS2 that include the access ports connected to the FCoE hosts.

Do not enable FIP snooping on the switches used to create the MC-LAG. For example, in [Figure 194 on page 5556](#), you would not enable FIP snooping on the FCoE VLANs on Switches S1 and S2.

Configure links between switches as FCoE trusted ports to reduce FIP snooping overhead and ensure that the system performs FIP snooping only at the access edge. In the sample topology, configure the Transit Switch TS1 and TS2 LAG ports connected to the MC-LAG switches as FCoE trusted ports, configure the Switch S1 and S2 MC-LAG ports connected to Switches TS1 and TS2 as FCoE trusted ports, and configure the ports in the LAG that connects Switches S1 to S2 as FCoE trusted ports.

## CoS and Data Center Bridging (DCB)

---

The MC-LAG links do not carry forwarding class or priority information. The following CoS properties must have the same configuration on each MC-LAG switch or on each MC-LAG interface to support lossless transport:

- FCoE forwarding class name—For example, the forwarding class for FCoE traffic could use the default **fcoe** forwarding class on both MC-LAG switches.
- FCoE output queue—For example, the **fcoe** forwarding class could be mapped to queue 3 on both MC-LAG switches (queue 3 is the default mapping for the **fcoe** forwarding class).
- Classifier—The forwarding class for FCoE traffic must be mapped to the same IEEE 802.1p code point on each member interface of the MC-LAG on both MC-LAG switches. For example, the FCoE forwarding class **fcoe** could be mapped to IEEE 802.1p code point **011** (code point **011** is the default mapping for the **fcoe** forwarding class).
- Priority-based flow control (PFC)—PFC must be enabled on the FCoE code point on each MC-LAG switch and applied to each MC-LAG interface using a congestion notification profile.

You must also configure enhanced transmission selection (ETS) on the MC-LAG interfaces to provide sufficient scheduling resources (bandwidth, priority) for lossless transport. The ETS configuration can be different on each MC-LAG switch, as long as enough resources are scheduled to support lossless transport for the expected FCoE traffic.

LLDP and DCBX must be enabled on each MC-LAG member interface (LLDP and DCBX are enabled by default on all interfaces).



**NOTE:** As with all other FCoE configurations, FCoE traffic requires a dedicated VLAN that carries only FCoE traffic, and IGMP snooping must be disabled on the FCoE VLAN.

---

### Related Documentation

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5995](#)
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)



## Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

Two methods of peer-to-peer flow control are supported:

- IEEE 802.3X Ethernet PAUSE
- IEEE 802.1Qbb priority-based flow control (PFC)

Ethernet PAUSE and PFC are link-level flow control mechanisms.



**NOTE:** For end-to-end congestion control, see [“Understanding CoS Explicit Congestion Notification” on page 5926](#).

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority is mapped to a 3-bit IEEE 802.1p CoS code point flag in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).



**Video:** [Why Use PFC in a Data Center Network?](#)

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on specified types of traffic (for example, Fibre Channel over Ethernet traffic).



**NOTE:** Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error. Ethernet PAUSE and PFC are mutually exclusive configurations on an interface.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

- [Ethernet PAUSE on page 5560](#)

- [PFC on page 5564](#)
- [Lossless Transport Support Summary on page 5567](#)

## Ethernet PAUSE

---

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

*Symmetric flow control* means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

*Asymmetric flow control* allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, the PFC configuration overrides Ethernet PAUSE flow control.) Both symmetric and asymmetric flow control are supported.

- [Symmetric Flow Control on page 5561](#)
- [Asymmetric Flow Control on page 5561](#)

### ***Symmetric Flow Control***

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

### ***Asymmetric Flow Control***

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 449 on page 5561](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

**Table 449: Asymmetric Ethernet PAUSE Flow Control Configuration**

| Receive (Rx) Buffer | Transmit (Tx) Buffer | Configured Flow Control State   |
|---------------------|----------------------|---|
| On                  | Off                  | Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).         |
| Off                 | On                   | Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.) |
| On                  | On                   | Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.   |
| Off                 | Off                  | Ethernet PAUSE flow control is disabled.  |

The configured flow control is the Ethernet PAUSE state configured on the interface.

On 1-Gigabit Ethernet interfaces, autonegotiation of Ethernet PAUSE with the connected peer is supported. (Autonegotiation on 10-Gigabit Ethernet interfaces is not supported.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected peer using a combination of the Ethernet PAUSE and ASM\_DIR bits, as described in [Table 450 on page 5562](#):

**Table 450: Flow Control State Advertised to the Connected Peer (Autonegotiation)**

| Rx Buffer State | Tx Buffer State | PAUSE Bit | ASM_DIR Bit | Description  |
|-----------------|-----------------|-----------|-------------|--|
| Off             | Off             | 0         | 0           | The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.   |
| On              | On              | 1         | 0           | The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).  |
| On              | Off             | 0         | 1           | The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).   |
| Off             | On              | 1         | 1           | The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.) |

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control behavior (resolution) between them, as shown in [Table 451 on page 5563](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series or EX4600 switch and on the connected peer (also known as the link partner). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations

on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



**NOTE:** In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

**Table 451: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces**

| Local Interface (QFX Series or EX4600 Switch) |             | Peer Interface |             | Local Resolution  | Peer Resolution   |
|---|-------------|----------------|-------------|---|---|
| PAUSE Bit                                     | ASM_DIR Bit | PAUSE Bit      | ASM_DIR Bit |   |   |
| 0   | 0           | Don't care     | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0   | 1           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0   | 1           | 1              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0   | 1           | 1              | 1           | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive | Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive |
| 1   | 0           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1   | 0           | 1              | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |
| 1   | 1           | 0              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1   | 1           | 0              | 1           | Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive |
| 1   | 1           | Don't care     | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |



**NOTE:** For your convenience, [Table 451 on page 5563](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

## PFC

---

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- Input—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- Output—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic with the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as Fibre Channel over Ethernet (FCoE), LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of link-level flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)

- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 452 on page 5565](#):

**Table 452: Default PFC Priority to Queue and Forwarding Class Mapping**

| IEEE 802.1p Priority (Code Point) | Queue | Forwarding Class |
|-----------------------------------|-------|------------------|
| 0 (000)                           | 0     | best-effort      |
| 1 (001)                           | 1     | best-effort      |
| 2 (010)                           | 2     | best-effort      |
| 3 (011)                           | 3     | fcoe             |
| 4 (100)                           | 4     | no-loss          |
| 5 (101)                           | 5     | best-effort      |
| 6 (110)                           | 6     | network-control  |
| 7 (111)                           | 7     | network-control  |

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



**NOTE:** By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default forwarding class configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 802.1p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5837](#) and [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#).

You enable PFC on a priority by:

1. Specifying the IEEE 802.1p code point to pause in the input stanza of a CNP
2. Applying the CNP to the ingress interfaces on which you want to pause the traffic



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
  1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
  2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
  3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

Although unicast traffic and multdestination (multicast, broadcast, and destination lookup fail) traffic must use different classifiers, you can map a unicast queue (queue 0 through 7) and a multdestination queue (queue 8, 9, 10, or 11) to the same PFC priority so that both unicast and multicast traffic use that priority. Do not map multdestination traffic to lossless priorities. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.





**NOTE:** You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone switches, you can create two CNPs with an explicitly configured output stanza.

When a switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 5837 describes configuring output flow control.

### Lossless Transport Support Summary

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



**NOTE:** Junos OS Release 12.2 introduced changes to the way the switch handles lossless forwarding classes (including the default fcoe and no-loss forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the fcoe and no-loss forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the fcoe or the no-loss forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the fcoe or the no-loss forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit fcoe and no-loss forwarding class configuration before you upgrade to Junos OS Release 12.2.

See *Overview of CoS Changes Introduced in Junos OS Release 12.2* for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the fcoe and no-loss forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new no-loss packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5837](#) provides detailed information about the explicit configuration of lossless priorities and about the default configuration of lossless priorities, including the input and output stanzas of the CNP.



**NOTE:** PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

**Related Documentation**

- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)

- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding CoS Explicit Congestion Notification on page 5926](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)

## Understanding Fibre Channel Terminology

To understand the Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) capabilities of the switches, you should become familiar with the terms defined in [Table 453 on page 5569](#).

**Table 453: Fibre Channel Terms**

| Term                    | Definition  |
|-------------------------|---|
| addressing mode         | <p>Format for the locally unique MAC address the FC switch assigns to FCoE devices for FCoE transactions after FIP establishes a connection between an FCoE device and the FC switch. The two addressing modes are <i>fabric-provided MAC address (FPMA)</i> and <i>server-provided MAC address (SPMA)</i>. Only FPMA is supported.</p> <p>During FLOGI or FDISC, the ENode advertises the addressing modes it supports. If the FC switch supports an addressing mode that the ENode uses, the virtual link can be established, and the devices can communicate.</p> <p>See also <i>fabric-provided MAC address (FPMA)</i> and <i>server-provided MAC address (SPMA)</i>.</p> |
| ALL-ENode-MACs          | <p>Well-known multicast MAC address to which all FCoE ENodes listen. FCFs send multicast <i>FIP discovery advertisement</i> messages and <i>FIP keepalive</i> messages to the ALL-ENode-MACs address so that ENodes can discover and maintain connections to FCFs. The hexadecimal format of the address is <b>01:10:18:01:00:01</b>.</p> <p>See also <i>well-known address (WKA)</i>.</p>  |
| ALL-FCF-MACs            | <p>Well-known multicast MAC address to which all FCFs listen. ENodes send multicast <i>FIP discovery solicitation</i> messages to the ALL-FCF-MACs address to find out which FCFs can accept a login. The hexadecimal format of the address is <b>01:10:18:01:00:02</b>.</p> <p>See also <i>well-known address (WKA)</i>.</p>   |
| congestion notification | See <i>quantized congestion notification (QCN)</i> .  |

Table 453: Fibre Channel Terms (*continued*)

| Term   | Definition  |
|--|---|
| converged network adapter (CNA)                          | <p>Physical adapter that combines the functions of a Fibre Channel <i>host bus adapter (HBA)</i> to process FCoE frames and a <i>lossless Ethernet network interface card (NIC)</i> to process non-FCoE Ethernet frames. CNAs have one or more Ethernet ports. CNAs encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel.</p> <p>See also <i>host bus adapter (HBA)</i>.</p>   |
| data center bridging (DCB)                               | <p>Set of IEEE specifications that enhance Ethernet to allow it to support converged Ethernet (LAN) and Fibre Channel (SAN) traffic on one Ethernet network. DCB features include <i>priority-based flow control (PFC)</i>, <i>enhanced transmission selection (ETS)</i>, <i>Data Center Bridging Capability Exchange protocol (DCBX)</i>, <i>quantized congestion notification (QCN)</i>, and full-duplex 10-Gigabit Ethernet ports.</p> <p>See also <i>priority-based flow control (PFC)</i>, <i>Ethernet PAUSE</i>, <i>enhanced transmission selection (ETS)</i>, <i>Data Center Bridging Capability Exchange protocol (DCBX)</i>, and <i>quantized congestion notification (QCN)</i>.</p> |
| expansion port (E_Port)                                  | <p>An expansion port in an FC switch/FCF that connects the FC switch/FCF to the E_Port of another FC switch/FCF to form an <i>Interswitch Link (ISL)</i> in a common FC fabric.</p>   |
| Data Center Bridging Capability Exchange protocol (DCBX) | <p>Discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1AB)</p> <p>See also <i>data center bridging (DCB)</i>.</p>  |
| enhanced transmission selection (ETS)                    | <p>Mechanism that provides finer granularity of bandwidth management within a link.</p> <p>See also <i>data center bridging (DCB)</i>.</p>  |
| ENode  | <p>See <i>FCoE Node (ENode)</i></p>   |
| ENode MAC  | <p><i>Lossless Ethernet MAC</i> paired with an <i>FCoE controller</i> in an ENode.</p> <p>See also <i>FCoE node (ENode)</i>.</p>  |
| ENode MAC address  | <p>Globally unique address assigned to the CNA by the manufacturer and used to identify the node for FIP transactions.</p>  |

Table 453: Fibre Channel Terms (*continued*)

| Term                     | Definition   |
|--------------------------|--|
| Ethernet PAUSE           | <p>As defined in IEEE 802.3X, a flow control mechanism that temporarily stops the transmission of Ethernet frames on a link for a specified period. A receiving element sends an Ethernet PAUSE frame when a sender transmits data faster than the receiver can accept it. Ethernet PAUSE affects the entire link, not just an individual flow. An Ethernet PAUSE frame temporarily stops all traffic transmission on the link and allows the receiver's input buffer to empty sufficiently to restart traffic on the link. Ethernet PAUSE messages are sent to the previous hop and do not automatically propagate to the source of the congestion.</p> <p>See also <i>priority-based flow control (PFC)</i>.</p>   |
| fabric                   | Interconnection of network nodes using one or more network switches that function as a network single logical entity.  |
| fabric discovery (FDISC) | <p>Subsequent logins from the same ENode for different users, applications, or virtual machines after an ENode performs an initial FLOGI to log in to a switch.</p> <p>FC and FIP FDISC messages serve the same function in FC and FCoE networks, respectively. N_Ports send FC FDISC messages to the FC switch and VN_Ports send FIP FDISC messages to the FCF.</p> <p>After an N_Port acquires its initial N_Port ID through the FC FLOGI process, it can acquire additional N_Port IDs by sending an FC FDISC with a new worldwide port name and a source ID of 0x000000. The new port name and blank source ID tell the FC switch to assign a new N_Port ID to the N_Port. The different N_Port IDs allow multiple virtual machines or users on the N_Port to have separate, secure virtual links on the same physical N_Port. These additional ports are also referred to as VN_Ports.</p> <p>FIP FDISC works the same way, except the VN_Port logs in using a FIP FLOGI message.</p> <p>See also <i>fabric login (FLOGI)</i> and <i>N_Port ID</i>.</p> |
| fabric login (FLOGI)     | <p>Creation of a logical connection to the FC switch and establishment of a node's operating environment.</p> <p>For FC devices, an N_Port logs in to the FC network by sending an FC FLOGI message to the F_Port of an FC switch.</p> <p>For FCoE devices, a VN_Port logs in to the FC network by sending a FIP FLOGI message to the VF_Port of an FC switch.</p>   |
| fabric port (F_Port)     | <p>FC port on an FC switch or an FCF that connects point-to-point to an FC node port (N_Port) on an FC host (server or storage device). An F_Port provides access to fabric services for FC devices.</p> <p>F_Ports are intermediate ports in a connection between FC device end-point N_Ports. For example, a connection between an FC host server and an FC storage device through an FC switch looks like this: FC server N_Port to FC switch ingress F_Port to FC switch egress F_Port to FC storage device N_Port.</p> <p>See also <i>node port (N_Port)</i>.</p>   |

Table 453: Fibre Channel Terms (*continued*)

| Term                                | Definition   |
|-------------------------------------|--|
| fabric-provided MAC address (FPMA)  | <p>MAC address that an FCF assigns to a single ENode MAC through the FLOGI or FDISC process that is unique to the local fabric. The FPMA uniquely identifies a single VN_Port at that ENode MAC in FCoE transactions with the FCF.</p> <p>Because an ENode can have more than one ENode MAC, an FCF can assign multiple FPMAs to an ENode, one FPMA per ENode MAC.</p> <p>An FPMA is a 48-bit value that consists of two 24-bit values, the N_Port ID and the FC-MAP value. The N_Port ID uniquely identifies the VN_Port and the FC-MAP value identifies the FCF.</p> <p>See also <i>FCoE node (ENode)</i>, <i>N_Port ID</i>, and <i>FCoE mapped address prefix (FC-MAP)</i>.</p> |
| FCF-MAC                             | Lossless Ethernet MAC paired with an FCoE controller in an FCF. The FCF-MAC enables the FCF to handle FCoE traffic.  |
| FCoE controller                     | <p>Instantiates and terminates VN_Port and VF_Port instances on an ENode. An ENode can have more than one FCoE controller. Each FCoE controller is paired with a lossless Ethernet MAC on the ENode.</p> <p>See also <i>lossless Ethernet MAC</i>.</p>   |
| FC forwarder (FCF)                  | Alternative term and acronym to refer to an FC switch that has all physical Fibre Channel ports and the necessary set of services as defined in the T11 Organization <i>Fibre Channel Switched Fabric</i> (FC-SW) standards.   |
| FCoE forwarder (FCF)                | Defined by the <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification available at <a href="http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf">http://www.t11.org/ftp/t11/pub/fc/bb-5/09-056v5.pdf</a> as a device that has the necessary set of services as defined in FC-SW and the FCoE capabilities to act as an FCoE-based FC switch.   |
| FCoE Initialization Protocol (FIP)  | <p>Layer 2 protocol for endpoint discovery, fabric login, and fabric association. FIP enables FCoE devices and FC switches to discover one another. Through FIP, FCoE nodes can log in to an FC switch, access the SAN FC fabric, and communicate with target FC devices. FIP messages also maintain the connection between the FCoE initiator and the FCF.</p> <p>FIP has its own EtherType (0x8914) to distinguish its traffic from payload-carrying FCoE traffic and other Ethernet traffic.</p>  |
| FCoE link endpoint (LEP)            | Virtual FC interface mapped onto a physical Ethernet interface to handle FC frame encapsulation and de-encapsulation and transmission and reception of FC frames encapsulated in Ethernet through a single virtual link.   |
| FCoE mapped address prefix (FC-MAP) | <p>24-bit value that identifies the FC switch and is half of the 48-bit FPMA MAC address. The FC-MAP value can be configured on the FC switch and has a default value of 0EFC00h. The FC-MAP value was originally called the Fibre Channel Organizationally Unique Identifier (FC-OUI).</p> <p>See also <i>fabric-provided MAC address (FPMA)</i>.</p>   |

Table 453: Fibre Channel Terms (*continued*)

| Term  | Definition  |
|---|---|
| FCoE node (ENode)   | <p>Fibre Channel node that has one or more lossless Ethernet MACs, each paired with an <i>FCoE Controller</i> in order to transmit FCoE frames. An ENode combines FCoE termination functions and the FC stack on a CNA. ENodes present virtual FC interfaces to FC switches or FCFs in the form of VN_Ports, which can establish FCoE virtual links with FC switch/FCF VF_Ports. ENodes perform FCoE related functions in a <i>converged network adapter (CNA)</i>.</p> <p>See also <i>converged network adapter (CNA)</i>.</p>   |
| FCoE-FC gateway   | A form of N_Port virtualizer in which the node-facing ports are FCoE ports and the FC switch-facing ports are FC ports.   |
| FCoE-FCoE gateway   | A form of N_Port virtualizer in which the node-facing ports are FCoE ports and the FC switch-facing ports are FCoE ports.   |
| FC-FC gateway   | A form of N_Port virtualizer in which the node-facing ports are FC ports and the FC switch-facing ports are FC ports.   |
| FCoE transit switch (also known as a FIP snooping bridge) | <p>Switch with a minimum set of features designed to support FCoE Layer 2 forwarding and FCoE security. The switch can also have optional additional features.</p> <p>Minimum feature support is:</p> <ul style="list-style-type: none"> <li>• Priority-based flow control (PFC)</li> <li>• Enhanced transmission selection (ETS)</li> <li>• Data Center Bridging Capability Exchange Protocol (DCBX), including the FCoE application TLV</li> <li>• FIP snooping (minimum support is FIP automated filter programming at the ENode edge)</li> </ul> <p>Additional FIP snooping capabilities can include learning the virtual FC connection paths (VN2VF, VN2VN, or VE2VE) and monitoring the FIP keepalive mechanisms. Other optional capabilities can also enhance FCoE within the standards. FIP snooping is typically configurable on a per-VLAN basis.</p> <p>A transit switch has an FC stack even though it is not an FC switch or an FCF.</p> |
| FCoE VLAN   | VLAN dedicated to carrying only FCoE traffic. FCoE traffic must travel in a VLAN. Only FCoE interfaces should be members of an FCoE VLAN. Ethernet traffic that is not FCoE traffic must travel in a different VLAN.  |
| Fibre Channel   | High-speed network technology used for storage area networks (SANs).  |
| Fibre Channel fabric                                      | <p>Network of Fibre Channel devices that allows communication among devices, device name lookup, security, and redundancy.</p> <p>Also a local fabric on a QFX3500 switch with FCoE interfaces connected to FCoE devices on the Ethernet network and native FC interfaces connected to an FC switch in a SAN.</p>   |

Table 453: Fibre Channel Terms (*continued*)

| Term                               | Definition   |
|------------------------------------|--|
| Fibre Channel ID (FCID)            | <p>24-bit value the FC switch assigns to the N_Port or VN_Port as a unique identifier within the local FC network. The FCID consists of an 8-bit domain value, an 8-bit area value, and an 8-bit port value. The FCID is sometimes called an N_Port ID.</p> <p>See also <i>N_Port ID</i>.</p>  |
| Fibre Channel over Ethernet (FCoE) | <p>Standard for transporting FC frames over Ethernet networks. FCoE encapsulates Fibre Channel frames in Ethernet so that the same high-speed Ethernet physical infrastructure can transport both data and storage traffic while preserving the lossless CoS that FC requires. FCoE has its own EtherType (0x8906) to differentiate it from other Ethernet traffic.</p> <p>FCoE runs on a DCB network. FCoE servers connect to a switch that supports both FCoE and native FC protocols. This allows FCoE servers on the Ethernet network to access FC storage devices in the SAN fabric on one converged network.</p> <p>See also <i>data center bridging (DCB)</i>.</p>  |
| Fibre Channel services             | Functions required for establishing FC network connectivity among devices and for managing devices on the FC network, such as login servers, domain managers, name servers, and zone servers.  |
| FC stack                           | <p>FC or FCoE protocol capability implemented on a device to support the FC or FCoE functionality. Having an FC stack does not imply consuming a domain ID.</p> <p>Each FC or FCoE enabled server or storage device has an FC stack. Similarly, an FC or FCoE switch, an FCF, an FCoE-FC gateway, and an FCoE transit switch have FC stacks.</p>   |
| Fibre Channel switch               | Network switch that implements the Fibre Channel protocol.   |
| FIP discovery advertisement        | <p>Multicast or unicast message that the FC switch (or FCF) transmits to ENodes to advertise the switch's presence on the network so that ENodes can discover the switch and request to log in to the FC fabric.</p> <p>The FC switch periodically sends multicast FIP discovery advertisements to the ALL-ENode-MACs address, a well-known address to which all ENodes listen. The multicast messages advertise the FC switch to all ENodes on the VLAN and serve as keepalive messages to maintain connectivity between the FC switch and ENodes.</p> <p>When an ENode sends a FIP discovery solicitation message to the FC switch, the FC switch responds with a unicast FIP discovery advertisement to that ENode.</p> |



Table 453: Fibre Channel Terms (*continued*)

| Term                       | Definition  |
|----------------------------|---|
| FIP discovery solicitation | <p>Multicast or unicast message that an ENode transmits to FC switches (or FCFs) to find compatible switches in the network.</p> <p>When an ENode initializes, it sends a multicast FIP discovery solicitation to the ALL-FCF-MACs address, a well-known address to which all FC switches and FCFs listen. Compatible switches reply with a unicast FIP discovery advertisement.</p> <p>The ENode compiles a list of compatible switches, selects a switch, and logs in to that switch.</p>   |
| FIP keepalive              | Periodic multicast FIP discovery advertisement sent from the FC switch or FCF to all ENodes to maintain connectivity.   |
| FIP snooping               | <p>For VN_Port to VF_Port (VN2VF) paths (Technical Committee T11 BB-FC-5 specification), FIP snooping is a security feature enabled for FCoE VLANs on an Ethernet switch that connects ENodes to FC switches or FCFs. FIP snooping inspects data in FIP frames and uses that data to create firewall filters. The filters permit only traffic from sources that perform a successful FLOGI to the FC switch. All other traffic on the VLAN is denied. FIP snooping filters are installed on the ports in the FCoE VLAN.</p> <p>For VN_Port to VN_Port (VN2VN) paths (Technical Committee T11 BB-FC-6 specification), the FIP snooping security feature filters access between VN_Ports in a similar manner to VN2VF_Port FIP snooping.</p> <p>FIP snooping can also apply similarly to VE_Port to VE_Port (VE2VE) paths.</p> <p>FIP snooping can also snoop to provide additional visibility of FCoE Layer 2 operation.</p> <p>See also <i>FCoE node (ENode)</i>.</p> |
| FIP snooping bridge        | See <i>FCoE transit switch</i> and <i>FIP snooping</i> .  |
| host bus adapter (HBA)     | Physical mechanism that connects a host system to other FC network and storage devices. HBAs have a unique worldwide node name (WWNN) for the HBA node, which all of the ports on the HBA share, and each port on an HBA has a unique worldwide port name (WWPN).   |
| initiator                  | System component that originates an I/O command over an I/O bus or network. An FCoE server sending a request to an FC storage device is an example of an initiator.   |

Table 453: Fibre Channel Terms (*continued*)

| Term                      | Definition  |
|---------------------------|---|
| iSCSI transit switch      | <p>Layer 2 Ethernet switch with a minimum set of best-practice Ethernet features to support iSCSI, along with optional enhancements. Minimum feature support is:</p> <ul style="list-style-type: none"> <li>• IEEE 802.3X asymmetric and symmetric flow control on ports not running in DCB mode</li> <li>• Priority-based flow control (PFC)</li> <li>• Enhanced transmission selection (ETS)</li> <li>• Data Center Bridging Capability Exchange Protocol (DCBX), including the iSCSI application TLV</li> </ul> <p>Other capabilities such as Internet storage name service (iSNS) are optional.</p> |
| interswitch link (ISL)    | <p>Link between the <i>E_Ports</i> of two FC switches in a common FC fabric. When two FCoE-based FC switches are connected together, there is a virtual ISL through Layer 2.</p>  |
| logout (LOGO)             | <p>For FC devices, an N_Port logs out from the FC network by sending an FC LOGO message to the F_Port of an FC switch. The switch can also send a LOGO message to an N_Port to terminate its connection.</p> <p>For FCoE devices, a VN_Port logs out from the FC network by sending a FIP LOGO message to the VF_Port of an FC switch. The switch can also send a LOGO message to a VN_Port to terminate its connection.</p>  |
| lossless Ethernet MAC     | <p>Full-duplex Ethernet MAC that implements Ethernet extensions to avoid Ethernet frame loss due to congestion and supports at least 2.5-KB jumbo frames. Each lossless Ethernet MAC combines with an FCoE Controller to perform FCoE termination functions on an ENode.</p> <p>See also <i>priority-based flow control (PFC)</i>, <i>quantized congestion notification (QCN)</i>, <i>FCoE controller</i>, and <i>FCoE node (ENode)</i>.</p>  |
| lossless Ethernet network | <p>Ethernet network composed of only full-duplex links and lossless Ethernet MACs and with CoS and flow control to prevent dropping of frames.</p>  |
| lossless transport        | <p>In DCB networks, the ability to switch FCoE frames over an Ethernet network without dropping any frames. Lossless transport uses mechanisms such as priority-based flow control and quantized congestion notification to control traffic flows and avoid congestion.</p>   |
| N_Port ID                 | <p>See <i>Fibre Channel ID (FCID)</i>.</p>  |

Table 453: Fibre Channel Terms (*continued*)

| Term                            | Definition   |
|---------------------------------|--|
| N_Port ID virtualizer           | <p>Presents itself as an FC or FCoE switch to external devices, but connects to an actual FC or FCoE switch in the other direction to provide the FC-SW services.</p> <p>An N_Port ID virtualizer logs in to the actual FC or FCoE switch in the same way as a normal node device and uses the NPIV mechanism to proxy incoming FLOGIs to FDISCs on the actual FC or FCoE switch.</p> <p>An N_Port ID virtualizer has an FC stack even though it is not an FC switch or an FCF.</p> <p>The acronym <i>NPV</i> is commonly used for N_Port ID virtualizer even though the acronym is not defined in the standards.</p>  |
| N_Port ID Virtualization (NPIV) | <p>NPIV enables a physical N_Port to acquire multiple N_Port IDs. Each N_Port ID maps to a different application (such as a virtual machine) or to a different user. This allows you to associate one F_Port with many N_Port IDs and create multiple discrete, secure virtual links over one physical point-to-point connection.</p> <p>NPIV increases resource and bandwidth utilization and allows the implementation of access control, zoning, and port security on a per-application or per-user basis.</p> <p>After an N_Port performs a FLOGI and receives its first N_Port ID, it can request more N_Port IDs by sending FDISC messages.</p> <p>See also <i>fabric login (FLOGI)</i>, <i>fabric discovery (FDISC)</i>, and <i>virtual link</i>.</p> |
| node port (N_Port)              | <p>N_Ports can be in two modes:</p> <ul style="list-style-type: none"> <li>• Fabric N_Port—Node port that is an FC host or storage device end port in a point-to-point link between the device and the F_Port of an FC switch. The point-to-point link can be virtual or physical.</li> <li>• Point-to-point N_Port—Node port that connects to another N_Port. The switch does not support this configuration.</li> </ul> <p>N_Ports handle creation, detection, and flow of messages to and from the connected devices.</p>   |
| node worldwide name (NWWN)      | <p>WWN that is unique worldwide and is assigned to an FC node. An NWWN is valid for multiple ports that are on that node (this identifies the ports as network interfaces of a particular node).</p>   |
| port mode                       | <p>Role that the port plays in the FC fabric (endpoint device, FC switch connection to endpoint devices, interswitch link).</p> <p>See also <i>node port (N_Port)</i>, <i>virtual node port (VN_Port)</i>, <i>proxy node port (NP_Port)</i>, <i>fabric port (F_Port)</i>, and <i>virtual fabric port (VF_Port)</i>.</p>  |
| port worldwide name (PWWN)      | <p>WWN that is unique worldwide and is assigned to an FC port.</p>   |

Table 453: Fibre Channel Terms (*continued*)

| Term                                    | Definition  |
|---|---|
| priority-based flow control (PFC)       | <p>Link-level flow control mechanism defined by IEEE 802.1Qbb that allows independent flow control for each class of service (as defined in the 3-bit CoS field of the Ethernet header by IEEE 802.1Q tags) to ensure that no frame loss from congestion occurs in DCB networks.</p> <p>PFC is an enhancement of the Ethernet PAUSE mechanism, but PFC controls classes of flows, whereas Ethernet PAUSE indiscriminately pauses all of the traffic on a link. With PFC, a receiving device can signal a transmitting device to pause transmission based on traffic class.</p> <p>PFC provides application-specific bandwidth reservations so you can ensure that time-critical protocols and applications such as FCoE receive the priority necessary to prevent frame loss. PFC allows the same physical link to carry FCoE traffic and provide lossless service while also carrying loss-tolerant Ethernet traffic.</p> <p>See also <i>Ethernet PAUSE</i>.</p> |
| proxy gateway mode                      | Connects FCoE initiators to FC switches in a converged Ethernet and Fibre Channel network and acts as an intermediary for these devices. The FCoE-FC gateway represents and acts for the FCoE initiators in transactions from the FCoE initiators destined for an FC switch, including converting FIP and FCoE frames to FC frames. The gateway represents and acts for an FC switch in transactions from the FC switch destined for an FCoE initiator, including converting FC frames to FIP frames and encapsulating FC frames in Ethernet.   |
| proxy node port (NP_Port)               | N_Port on the QFX3500 switch that performs proxy functions when it is configured as an FCoE-FC gateway. The NP_Port acts as a proxy for the FCoE device VN_Ports in transactions with the FC switch.  |
| quantized congestion notification (QCN) | Mechanism defined by IEEE 802.1Qau that manages network congestion within a Layer 2 domain. When a queue reaches a configured threshold, QCN throttles traffic at the source of the congestion by transmitting messages that propagate back to the source and temporarily stop the source from transmitting. When the queue crosses the threshold that indicates the congestion has dissipated, QCN sends a message to allow the source to resume transmitting frames.  |
| session                                 | Fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. Session does not refer to end-to-end server-to-storage sessions.   |
| server-provided MAC address (SPMA)      | <p>MAC address that an ENode assigns to one of its ENode MACs and is not assigned to any other ENode MAC in the same FCoE VLAN. An SPMA can be associated with more than one VN_Port at that ENode MAC.</p> <p>The switch does not support SPMA.</p> <p>See also <i>ENode MAC</i> and <i>fabric-provided MAC address (FPMA)</i>.</p>  |
| storage area network (SAN)              | Network whose primary purpose is the transfer of data between computer systems and storage devices. This term is most commonly used in the context of any network that supports block storage, usually iSCSI, FC, and FCoE networks.  |

Table 453: Fibre Channel Terms (*continued*)

| Term                          | Definition   |
|-------------------------------|--|
| target                        | System component that receives an I/O command. An FC storage device that receives a request from a server is an example of a target.   |
| VE_Port                       | Virtual ports created to form a connection (an <i>interswitch link</i> ) between two FCoE-based FC switches as part of a common FC fabric.   |
| VE2VE (VE_Port to VE_Port)    | The <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification capability of FCFs to connect to each other as a single FCoE FC SAN.  |
| VN2VF (VN_Port to VF_Port)    | The <i>Fibre Channel Backbone - 5 (FC-BB-5) Rev 2.00</i> specification capability of an ENode to connect to an FCF or to an FCoE-enabled FC SAN.   |
| VN2VN (VN_Port to VN_Port)    | The <i>Fibre Channel Backbone - 6 (FC-BB-6)</i> specification capability of an ENode to connect directly over Layer 2 to another ENode without the need of any FC-related services. This capability is most often used in small-scale FCoE SANs.   |
| virtual fabric port (VF_Port) | <p>Data-forwarding component that emulates an F_Port. A VF_Port is dynamically instantiated on successful completion of a FIP FLOGI exchange and connects to one or more VN_Ports. The term <i>virtual</i> indicates the use of a non-FC link such as an FCoE link.</p> <p>See also <i>fabric port (F_Port)</i>.</p>   |
| virtual link                  | <p>Logical link connecting two FCoE Link End Points (LEPs) over a lossless Ethernet network, for example, the link between a VF_Port and a VN_Port. The MAC addresses of the two LEPs identifies a virtual link.</p> <p>See also <i>FCoE link end point (LEP)</i> and <i>lossless Ethernet network</i>.</p>  |
| virtual node port (VN_Port)   | <p>Data-forwarding component that emulates an N_Port. With FCoE, a VN_Port is dynamically instantiated on successful completion of a FIP FLOGI exchange and connects to one or more VF_Ports. The term <i>virtual</i> indicates the use of a non-FC link such as an FCoE link.</p> <p>VN_Port is also used for the virtual N_Ports created in both FC and FCoE when additional NPIV-based logins occur over a previously created N_Port-to-VN_Port or N_Port-to-VF_Port connection.</p> <p>See also <i>node port (N_Port)</i>.</p> |
| well-known address (WKA)      | Address identifier used to access a service provided by an FC fabric. The service can be distributed in many elements throughout a fabric, or it can be centralized in one element. A WKA is always accessible, regardless of zoning. An example of a WKA is the <i>ALL-FCF-MACs</i> address to which all FCFs listen.   |
| worldwide name (WWN)          | 64-bit identifier that is similar to a MAC address except that it is not used for forwarding. It uniquely identifies an FC device. The WWN is derived from the IEEE organizationally unique identifier (OUI) and vendor-supplied information. A WWN is unique worldwide.   |
| worldwide node name (WWNN)    | See <i>node worldwide name (NWWN)</i> .  |

Table 453: Fibre Channel Terms (*continued*)

| Term                       | Definition                              |
|----------------------------|---|
| worldwide port name (WWPN) | See <i>port worldwide name (PWWN)</i> . |

**Related  
Documentation**

- [Overview of Fibre Channel on page 5508](#)
- [Understanding QFabric System Terminology](#)

## DCBX

- [Understanding DCBX on page 5580](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)

### Understanding DCBX

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.



Video: [What is DCBX Protocol?](#)

This topic describes:

- [DCBX Basics on page 5580](#)
- [DCBX Modes and Support on page 5581](#)
- [DCBX Attribute Types on page 5584](#)
- [DCBX Application Protocol TLV Exchange on page 5585](#)
- [DCBX and PFC on page 5586](#)
- [DCBX and ETS on page 5587](#)

#### DCBX Basics

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



**NOTE:** When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

## DCBX Modes and Support

This section describes DCBX support:

- [DCBX Modes \(Versions\) on page 5581](#)
- [Autonegotiation on page 5583](#)
- [CNA Support for DCBX Modes on page 5584](#)
- [Interface Support for DCBX on page 5584](#)

### ***DCBX Modes (Versions)***

The two most common DCBX modes are supported:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



**NOTE:** The switch does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The switch drops LLDP frames that contain pre-CEE DCBX TLVs.

Table 454 on page 5582 summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:

**Table 454: Summary of Differences Between IEEE DCBX and DCBX Version 1.01**

| Characteristic   | IEEE DCBX  | DCBX Version 1.01  |
|--|--|--|
| OUI  | 0x0080c2   | 0x001b21   |
| Frame Format   | Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs   | Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer.   |
| Symmetric/asymmetric configuration with peer                                     | Asymmetric or symmetric  | Symmetric only   |
| Differences in the <b>show dcbx interface interface-name</b> operational command | <ul style="list-style-type: none"> <li>• Synchronization information is not shown because symmetric configuration is not required.</li> <li>• Operational state information is not shown because the operational states do not have to be symmetric.</li> <li>• TLV type is shown because unique TLVs are sent for each DCBX attribute.</li> <li>• ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs.</li> </ul> | <ul style="list-style-type: none"> <li>• Synchronization information is shown because symmetric configuration is required.</li> <li>• Operational state information is shown because the operational states do have to be symmetric.</li> <li>• TLV type is not shown because one TLV is used for all attribute information.</li> <li>• Recommendation TLV is not sent (DCBX Version 1.01 uses the “willing” bit to determine whether or not an interface uses the peer interface configuration).</li> </ul> |



For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—  
<http://www.ieee802.org/1/files/public/docs2008/az-wadkar-dcbx-capability-exchange-discovery-protocol-1108-v1.01.pdf>
- For IEEE DCBX—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>



**NOTE:** As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.
- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



**NOTE:** On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

### ***Autonegotiation***

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on standalone switches compared to QFabric systems:

- Standalone switches—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.
- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



**NOTE:** If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

---

### ***CNA Support for DCBX Modes***

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on switch interfaces depends on the DCBX features that the CNAs in your network support.

### ***Interface Support for DCBX***

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

---

### **DCBX Attribute Types**

DCBX has three attribute types:

- Informational—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- Asymmetric—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- Symmetric—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 5585](#)
- [Symmetric Attributes on page 5585](#)

### ***Asymmetric Attributes***

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing,” the configuration on the interface cannot be overridden by the peer interface configuration.
- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

### ***Symmetric Attributes***

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

### ***DCBX Application Protocol TLV Exchange***

DCBX advertises the switch’s capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 5585](#)
- [FCoE Application Protocol TLV Exchange on page 5586](#)
- [Disabling Application Protocol TLV Exchange on page 5586](#)

### ***Application Protocol TLV Exchange***

For all applications, DCBX advertises the application’s state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application’s TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

### ***FCoE Application Protocol TLV Exchange***

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map



**NOTE:** If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

---

If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.

---



**NOTE:** If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

---

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE, DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

### ***Disabling Application Protocol TLV Exchange***

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

### ***DCBX and PFC***

---

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device

connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.

You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

### DCBX and ETS

---

This section describes:

- [Default DCBX ETS Advertisement on page 5587](#)
- [ETS Advertisement and Peer Configuration on page 5587](#)
- [ETS Recommendation TLV on page 5588](#)

#### ***Default DCBX ETS Advertisement***

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

#### ***ETS Advertisement and Peer Configuration***

DCBX does not control the switch’s ETS (hierarchical scheduling) operational state. If the connected peer is configured as “willing,” DCBX pushes the switch’s ETS configuration to the switch’s peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority (essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

### ***ETS Recommendation TLV***

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.



**NOTE:** You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

---

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

#### **Related Documentation**

- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Understanding DCB Features and Requirements on page 5515](#)

- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding FCoE on page 5518](#)
- [Configuring the DCBX Mode on page 5668](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)

## Understanding DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 5590](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case (see [“Application Maps” on page 5590](#)) and only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic.

This topic describes:

- [Applications on page 5590](#)
- [Application Maps on page 5590](#)
- [Classifying and Prioritizing Application Traffic on page 5591](#)

- [Enabling Interfaces to Exchange Application Protocol Information on page 5592](#)
- [Disabling DCBX Application Protocol Exchange on page 5592](#)

## Applications

---

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise, except FCoE if FCoE is the only application that you want the interface to advertise.



**NOTE:** If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

---

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

## Application Maps

---

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.



To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

### Classifying and Prioritizing Application Traffic

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application

receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.

### Enabling Interfaces to Exchange Application Protocol Information

---

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



**NOTE:** You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

---

### Disabling DCBX Application Protocol Exchange

---

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

#### Related Documentation

- [Understanding DCBX on page 5580](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)

- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)

## Learn About Technology

---

- [Data Center Technology Overview Videos on page 5593](#)

### Data Center Technology Overview Videos

Juniper Information Experience (iX) videos provide brief, high-level overviews of data center technologies and concepts. Each video runs approximately one-and-a-half to two minutes in length. This document contains SDN-related videos and links to conceptual documents that contain other data center technology videos:

- [Learn About Video: Why Do We Need an IP Fabric? on page 5593](#)
- [Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric? on page 5593](#)
- [Learn About Video: Why Use an Overlay Network in a Data Center? on page 5593](#)
- [Conceptual Documents That Contain Technology Overview Videos on page 5594](#)

#### Learn About Video: Why Do We Need an IP Fabric?

---

The video *Why Do We Need an IP Fabric?* presents a brief overview of IP Fabric use cases.



Video: [Why Do We Need an IP Fabric?](#)

#### Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?

---

The video *What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?* presents a brief overview of the arguments for using Border Gateway Protocol (BGP) as the data center IP fabric control plane protocol.



Video: [What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?](#)

#### Learn About Video: Why Use an Overlay Network in a Data Center?

---

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of data center overlay networks.



Video: [Why Use an Overlay Network in a Data Center?](#)

### Conceptual Documents That Contain Technology Overview Videos

The following conceptual documents include brief video overviews of the technology:

- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding DCBX on page 5580](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 5950](#)
- [Virtual Chassis Fabric Overview on page 7033](#)
- [“Understanding In-Service Software Upgrade \(ISSU\)” on page 25 and “In-Service Software Upgrade \(ISSU\) System Requirements” on page 13 \(same video\)](#)

# Configuration

- [Configuration Examples on page 5595](#)
- [Configuration Examples on page 5614](#)
- [FCoE and FIP Snooping Configuration Tasks on page 5655](#)
- [DCBX Configuration Tasks on page 5667](#)
- [Configuration Statements on page 5675](#)

## Configuration Examples

---

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)

### Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The switch handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the switch exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE

priority (the FCoE IEEE 802.1p code point). The default priority mapping for the FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).

- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 5596](#)
- [Overview on page 5596](#)
- [Configuration on page 5600](#)
- [Verification on page 5601](#)

## Requirements

---

This example uses the following hardware and software components:

- Juniper Networks QFX Series device
- Junos OS Release 12.1 or later for the QFX Series

## Overview

---

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



**NOTE:** DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 5669](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.

- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 455 on page 5597](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 456 on page 5597](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

**Table 455: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

**Table 456: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |

**Table 456: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier) (continued)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

**Topology**

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



**NOTE:** You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 457 on page 5598 shows the configuration components for this example.

**Table 457: Components of DCBX Application Protocol Exchange Configuration Topology**

| Component                   | Settings   |
|-----------------------------|--|
| Hardware                    | QFX Series device  |
| LLDP                        | Enabled by default on Ethernet interfaces  |
| DCBX                        | Enabled by default on Ethernet interfaces  |
| iSCSI application (Layer 4) | Application name— <b>iscsi</b><br>protocol— <b>TCP</b><br>destination-port— <b>3260</b><br>code-points— <b>111</b> |
| PTP application (Layer 2)   | Application name— <b>ptp</b><br>ether-type— <b>0x88F7</b><br>code-points— <b>001, 101</b>                          |



**Table 457: Components of DCBX Application Protocol Exchange Configuration Topology (*continued*)**

| Component   | Settings   |
|---|--|
| FCoE application (Layer 2)  | <p>Application name—<b>fcoe</b></p> <p>ether-type—<b>0x8906</b></p> <p>code-points—<b>011</b></p> <p><b>NOTE:</b> You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map.</p>   |
| Application maps  | <p><b>dcbx-iscsi-fcoe-app-map</b>—Maps the iSCSI and FCoE applications to IEEE 802.1p code points</p> <p><b>dcbx-iscsi-ptp-app-map</b>—Maps iSCSI and PTP applications to IEEE 802.1p code points</p>  |
| Interfaces  | <p><b>xe-0/0/10</b>—Configured to exchange FCoE and iSCSI application TLVs (uses application map <b>dcbx-iscsi-fcoe-app-map</b>, carries FCoE traffic, and has PFC enabled on the FCoE priority)</p> <p><b>xe-0/0/11</b>—Configured to exchange iSCSI and PTP application TLVs (uses application map <b>dcbx-iscsi-ptp-app-map</b>)</p>  |
| PFC congestion notification profile for FCoE application exchange   | <p><b>fcoe-cnp:</b></p> <ul style="list-style-type: none"> <li>Code point—<b>011</b></li> <li>Interface—<b>xe-0/0/10</b></li> </ul>  |
| Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point) | <p><b>fcoe-iscsi-cl1:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>fcoe</b> forwarding class to the IEEE 802.1p code point used for the FCoE application (<b>011</b>) and a loss priority of <b>high</b></li> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (<b>111</b>) and a loss priority of <b>high</b></li> <li>Applied to interface <b>xe-0/0/10</b></li> </ul> <p><b>iscsi-ptp-cl2:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (<b>111</b>) and a loss priority of <b>low</b></li> <li>Maps the <b>best-effort</b> forwarding class to the IEEE 802.1p code points used for the PTP application (<b>001</b> and <b>101</b>) and a loss priority of <b>low</b></li> <li>Applied to interface <b>xe-0/0/11</b></li> </ul> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

## Configuration

### CLI Quick Configuration

To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ptp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ptp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

### Configuring DCBX Application Protocol TLV Exchange

### Step-by-Step Procedure

To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.  
  

```
[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88F7
```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.  
  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.  
  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
```

4. Apply the iSCSI and FCoE application map to interface **xe-0/0/10**, and apply the iSCSI and PTP application map to interface **xe-0/0/11**.

```
[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
```

5. Create the congestion notification profile to enable PFC on the FCoE code point (**011**), and apply the congestion notification profile to interface **xe-0/0/10**.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
```

6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority
high code-points 011
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class network-control
loss-priority high code-points 111
```

7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi-ntp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
```

8. Apply the classifiers to the appropriate interfaces.

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ntp-cl2
```

## Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 5601](#)
- [Verifying the Application Map Configuration on page 5602](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 5602](#)
- [Verifying the PFC Configuration on page 5603](#)
- [Verifying the Classifier Configuration on page 5604](#)

### Verifying the Application Configuration

**Purpose** Verify that DCBX applications have been configured.

**Action** List the applications by using the configuration mode command **show applications**:

```
user@switch# show applications
```

```
application iSCSI {
  protocol tcp;
  destination-port 3260;
}

application fcoe {
  ether-type 0x8906;
}

application ptp {
  ether-type 0x88F7;
}
```

**Meaning** The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

#### *Verifying the Application Map Configuration*

**Purpose** Verify that the application maps have been configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
  application iSCSI code-points 111;
  application FCoE code-points 011;
}

dcbx-iscsi-ptp-app-map {
  application iSCSI code-points 111;
  application PTP code-points [001 101];
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ptp-app-map**.

The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point **111**, and the FCoE application, which is mapped to IEEE 802.1p code point **011**.

The application map **dcbx-iscsi-ptp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point **111**, and the PTP application, which is mapped to IEEE 802.1p code points **001** and **101**.

#### *Verifying DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application maps have been applied to the correct interfaces.

**Action** List the application maps by using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
    application-map dcbx-iscsi-ptp-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ptp-app-map**.

### *Verifying the PFC Configuration*

**Purpose** Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

**Action** Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
    input {
        ieee-802.1 {
            code-point 011 {
                pfc;
            }
        }
    }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
    congestion-notification-profile fcoe-cnp;
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

**Meaning** The **show class-of-service congestion-notification-profile** configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile **fcoe-cnp** has been configured and has enabled PFC on the IEEE 802.1p code point **011** (the default FCoE code point).

The **show class-of-service interfaces** configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile **fcoe-cnp**, which enables PFC on the FCoE code point, is applied to interface **xe-0/0/10**.

### *Verifying the Classifier Configuration*

**Purpose** Verify that the classifiers have been configured and applied to the correct interfaces.

**Action** Display the classifier configuration by using the configuration mode command **show class-of-service**:

```
user@switch# show class-of-service
classifiers {
  ieee-802.1 fcoe-iscsi-cl1 {
    import default;
    forwarding-class network-control {
      loss-priority high code-points 111;
    }
    forwarding-class fcoe {
      loss-priority high code-points 011;
    }
  }
  ieee-802.1 iscsi-ptp-cl2 {
    import default;
    forwarding-class network-control {
      loss-priority low code-points 111;
    }
    forwarding-class best-effort {
      loss-priority low code-points [ 001 101 ];
    }
  }
}
interfaces {
  xe-0/0/10 {
    congestion-notification-profile fcoe-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-iscsi-cl1;
      }
    }
  }
  xe-0/0/11 {
    unit 0 {
      classifiers {
        ieee-802.1 iscsi-ptp-cl2;
      }
    }
  }
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

**Meaning** The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ptp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ptp-cl2** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ptp-cl2** is mapped to interface **xe-0/0/11.0**.

**Related Documentation**

- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [show dcbx on page 5723](#)
- [show dcbx neighbors on page 5724](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Using DCBX Protocol to Lower Costs](#)

## Example: Configuring CoS PFC for FCoE Traffic

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS flag in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

To configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic.
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.
- Configure the bandwidth scheduling for the FCoE forwarding class output queue.
- Create a forwarding class set (priority group) that includes the FCoE forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).
- Configure the bandwidth scheduling for the FCoE priority group.
- Apply the scheduling to the egress interfaces.



.....  
**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding classes for the lossless fcoe forwarding class. If you explicitly configure default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

.....  
This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 5607](#)
- [Overview on page 5607](#)
- [Configuration on page 5608](#)
- [Verification on page 5612](#)



## Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to:

- Assign FCoE traffic to the FCoE priority at the ingress.
- Create and apply CoS for the FCoE traffic using ETS (hierarchical port scheduling).
- Apply PFC to the FCoE traffic.
- Apply the configuration to ingress and egress interfaces.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Each interface in this example is configured as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and port scheduling are applied to all of the interfaces.

## Topology

Table 458 on page 5607 shows the configuration components for this example.

**Table 458: Components of the PFC for FCoE Traffic Configuration Topology**

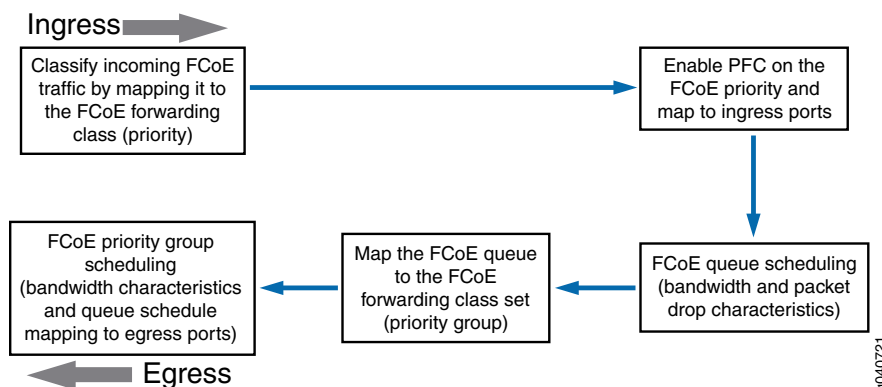
| Component   | Settings  |
|---|---|
| Hardware  | QFX3500 switch  |
| Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point) | Code point <b>011</b> to forwarding class <b>fcoe</b> and loss priority <b>low</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b> |
| PFC congestion notification profile   | <b>fcoe-cnp:</b><br>Code point <b>011</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>  |
| FCoE queue scheduler  | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>   |

Table 458: Components of the PFC for FCoE Traffic Configuration Topology (*continued*)

| Component                                  | Settings   |
|--|--|
| Forwarding class-to-scheduler mapping      | Scheduler map <b>fcoe-map</b> :<br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b>                           |
| Forwarding class set (FCoE priority group) | <b>fcoe-pg</b> :<br>Forwarding class <b>fcoe</b><br>Egress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b> |
| Traffic control profile                    | <b>fcoe-tcp</b> :<br>Scheduler map <b>fcoe-map</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b>       |

Figure 195 on page 5608 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 195: PFC for FCoE Traffic Configuration Components Block Diagram



### Configuration

#### CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```

[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100

```

```

set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp

```

### Step-by-Step Procedure

To configure the FCoE forwarding class (priority), ingress classifier, output queue scheduling, forwarding class set (priority group) and its output port scheduling, PFC application, and interfaces to set up PFC for FCoE traffic:

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:

```

[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority
low code-points 011

```

2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc

```

3. Apply the PFC configuration to the ingress interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp

```

4. Assign the classifier to the ingress interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier

```

5. Configure output scheduling for the FCoE queue:

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```

6. Map the FCoE forwarding class to the FCoE scheduler:

```

[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched

```

7. Configure the forwarding class set for the FCoE traffic:

```

[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe

```

8. Define the traffic control profile for the FCoE forwarding class set:

```

[edit class-of-service]

```

```
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

9. Apply the FCoE forwarding class set and traffic control profile to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

### Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class):

```
user@switch> show configuration class-of-service
```

```
classifiers {
  ieee-802.1 fcoe-classifier {
    forwarding-class fcoe {
      loss-priority low code-points 011;
    }
  }
}
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 3000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
}
interfaces {
  xe-0/0/31 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
```

```

        output-traffic-control-profile fcoe-tcp;
    }
}
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}
xe-0/0/32 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    unit 0 {
        classifiers {
            ieee-802.1 fcoe-classifier;
        }
    }
}
xe-0/0/33 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    unit 0 {
        classifiers {
            ieee-802.1 fcoe-classifier;
        }
    }
}
xe-0/0/34 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    unit 0 {
        classifiers {
            ieee-802.1 fcoe-classifier;
        }
    }
}
}
scheduler-maps {
    fcoe-map {
        forwarding-class fcoe scheduler fcoe-sched;
    }
}
schedulers {
    fcoe-sched {

```

```

        transmit-rate 3000000000;
        shaping-rate percent 100;
        priority low;
    }
}

```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5612](#)
- [Verifying the Ingress Interface PFC Configuration on page 5613](#)

### *Verifying That Priority-Based Flow Control Has Been Enabled*

**Purpose** Verify that PFC is enabled on the FCoE queue to enable lossless transport.

**Action** List the congestion notification profiles using the operational mode command `show class-of-service congestion-notification`:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: fcoe-cnp, Index: 51697
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      |                     |
| 010      | 1                   |
| 011      | 2                   |
| 100      | 3                   |
| 101      | 4                   |
| 110      | 5                   |
| 111      | 6                   |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point 011 for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### *Verifying the Ingress Interface PFC Configuration*

**Purpose** Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.

**Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}
```

**Meaning** The **show configuration class-of-service interfaces** commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Configuration Examples

---

- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5614](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5636](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5641](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5647](#)

### Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX Series switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).



**NOTE:** This example uses the Junos OS Enhanced Layer 2 Software (ELS) configuration style for QFX Series switches. If your switch runs software that does not support ELS, see [“Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG” on page 5995](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

You can use an MC-LAG to provide a redundant aggregation layer for Fiber Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX Series switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.

Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.





**NOTE:** This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two QFX Series switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the QFX Series switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see [“Example: Configuring Multichassis Link Aggregation” on page 2471](#). However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

QFX3500 and QFX3600 Virtual Chassis switches do not support FCoE.

This topic describes:

- [Requirements on page 5615](#)
- [Overview on page 5615](#)
- [Configuration on page 5619](#)
- [Verification on page 5628](#)

## Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX5100 Switches running the ELS CLI that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX5100 Switches running the ELS CLI that provide FCoE server access in transit switch mode and that connect to the MC-LAG switches.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 13.2 or later for the QFX Series.

## Overview

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX5100 switches that form the MC-LAG, including priority-based flow control (PFC) and enhanced transmission selection (ETS; hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group).



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

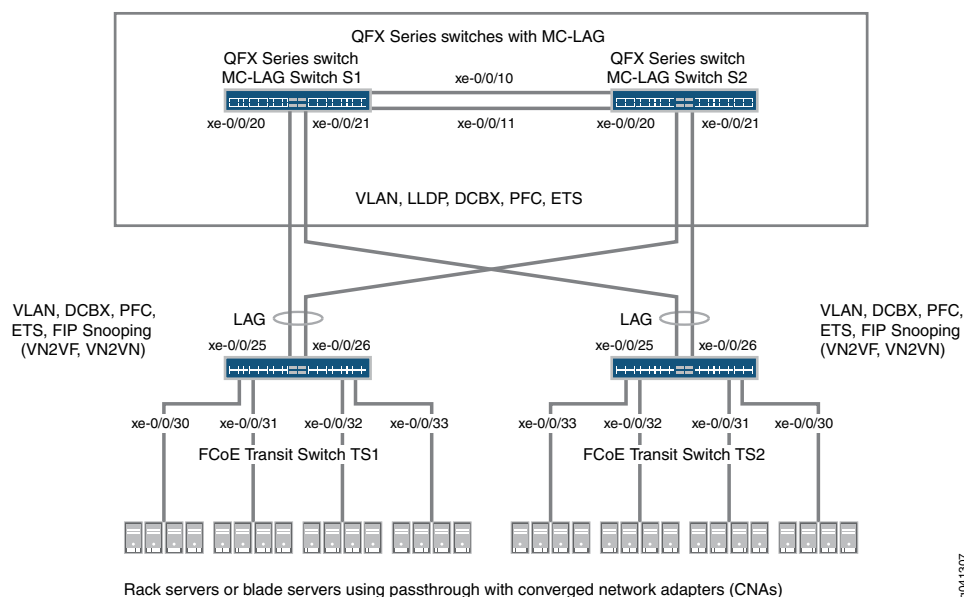


**NOTE:** Do not enable IGMP snooping on the FCoE VLAN. (IGMP snooping is enabled on the default VLAN by default, but is disabled by default on all other VLANs.)

### Topology

QFX5100 switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 196 on page 5616](#).

**Figure 196: Supported Topology for an MC-LAG on an FCoE Transit Switch**



[Table 459 on page 5616](#) shows the configuration components for this example.

**Table 459: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology**

| Component  | Settings   |
|--|--|
| Hardware   | Four QFX5100 switches running the ELS CLI (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access). |
| Forwarding class (all switches)  | Default <b>fcoe</b> forwarding class.  |
| Classifier (forwarding class mapping of incoming traffic to IEEE priority) | Default IEEE 802.1p trusted classifier on all FCoE interfaces.   |

**Table 459: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)**

| Component  | Settings   |
|--|--|
| LAGs and MC-LAG  | <p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S1 to Switch S2. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S2 to Switch S1. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p><b>NOTE:</b> Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>trunk</b> interface mode, with an MTU of <b>2180</b>.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> interface mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>trunk</b> interface mode, with an MTU of <b>2180</b>.</p> |
| FCoE queue scheduler (all switches)                      | <p><b>fcoe-sched:</b><br/> Minimum bandwidth <b>3g</b><br/> Maximum bandwidth <b>100%</b><br/> Priority <b>low</b></p>   |
| Forwarding class-to-scheduler mapping (all switches)     | <p>Scheduler map <b>fcoe-map</b>:<br/> Forwarding class <b>fcoe</b><br/> Scheduler <b>fcoe-sched</b></p>   |
| Forwarding class set (FCoE priority group, all switches) | <p><b>fcoe-pg:</b><br/> Forwarding class <b>fcoe</b></p> <p>Egress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>  |
| Traffic control profile (all switches)                   | <p><b>fcoe-tcp:</b><br/> Scheduler map <b>fcoe-map</b><br/> Minimum bandwidth <b>3g</b><br/> Maximum bandwidth <b>100%</b></p>   |

Table 459: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (*continued*)

| Component  | Settings  |
|--|---|
| PFC congestion notification profile (all switches) | <b>fcoe-cnp:</b><br>Code point <b>011</b><br><br>Ingress interfaces: <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>   |
| FCoE VLAN name and tag ID                          | Name— <b>fcoe_vlan</b><br>ID— <b>100</b><br><br>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.   |
| FIP snooping                                       | Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.<br><br>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration. |



**NOTE:** This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode interfaces if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is queue **3**.
- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk interface mode. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (**011**) to the FCoE forwarding class. If you choose to configure the BA classifier instead of using the default classifier, you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.

- Configure a congestion notification profile that enables PFC on the FCoE code point (code point 011 in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure enhanced transmission selection (ETS, also known as hierarchical scheduling) on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic.
- Apply the ETS scheduling to the interfaces.
- Configure the interface mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. ([“Example: Configuring Multichassis Link Aggregation” on page 2471](#) describes how to configure MC-LAGs.)
- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. ([“Configuring Link Aggregation” on page 2593](#) describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

## Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [Configuring MC-LAG Switches S1 and S2 on page 5621](#)
- [Configuring FCoE Transit Switches TS1 and TS2 on page 5623](#)
- [Results on page 5625](#)

### CLI Quick Configuration

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the **[edit]** hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

## Switch S1 and Switch S2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
set vlans fcoe_vlan forwarding-options fip-security interface ae0 fcoe-trusted
set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the **[edit]** hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

## Switch TS1 and Switch TS2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
```

```

set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching interface-mode trunk vlan members fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
set vlans fcoe_vlan forwarding-options fip-security examine-vn2v2 beacon-period 90000

```

### Configuring MC-LAG Switches S1 and S2

#### Step-by-Step Procedure

To configure CoS resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point 011, so you do not configure them):

1. Configure output scheduling for the FCoE queue:

```

[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100

```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```

[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched

```

3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:

```

[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe

```

4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:

```

[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100

```

5. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:

```

[edit class-of-service]
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp

```

6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
```
7. Apply the PFC configuration to the LAG and MC-LAG interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```
8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):  

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
9. Add the member interfaces to the LAG between the two MC-LAG switches:  

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```
10. Add the member interfaces to the MC-LAG:  

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```
11. Configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):  

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```
12. Set the MTU to **2180** for the LAG and MC-LAG interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:  

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```
13. Set the LAG and MC-LAG interfaces as FCoE trusted ports. Ports that connect to other switches should be trusted and should not perform FIP snooping:  

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae0 fcoe-trusted
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```



### Configuring FCoE Transit Switches TS1 and TS2

**Step-by-Step Procedure** The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:  

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):  

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:  

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
5. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```
6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
7. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```

```
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
```

8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

9. Add the member interfaces to the LAG:

```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```

10. On the LAG (**ae1**), configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```

11. On the FCoE access interfaces (**xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**), configure the interface mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching interface-mode trunk
vlan members fcoe_vlan
```

12. Set the MTU to **2180** for the LAG and FCoE access interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

13. Set the LAG interface as an FCoE trusted port. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security interface ae1 fcoe-trusted
```



**NOTE:** Access ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are not configured as FCoE trusted ports. The access ports remain in the default state as untrusted ports because they connect directly to FCoE devices and must perform FIP snooping to ensure network security.

14. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses VN2VN\_Port FIP snooping; the example is equally valid if you use VN2VF\_Port FIP snooping):

```
[edit]
user@switch# set vlans fcoe_vlan forwarding-options fip-security examine-vn2vn
beacon-period 90000
```

### Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are the same):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 30000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
}
interfaces {
  ae0 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  ae1 {
```

```
forwarding-class-set {
  fcoe-pg {
    output-traffic-control-profile fcoe-tcp;
  }
}
congestion-notification-profile fcoe-cnp;
}
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
```



**NOTE:** The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see [“Example: Configuring Multichassis Link Aggregation” on page 2471](#).

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are the same):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 3000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
```

```

    }
  }
  interfaces {
    xe-0/0/30 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
    xe-0/0/31 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
    xe-0/0/32 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
    xe-0/0/33 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
    ae1 {
      forwarding-class-set {
        fcoe-pg {
          output-traffic-control-profile fcoe-tcp;
        }
      }
      congestion-notification-profile fcoe-cnp;
    }
  }
  scheduler-maps {
    fcoe-map {
      forwarding-class fcoe scheduler fcoe-sched;
    }
  }
  schedulers {
    fcoe-sched {
      transmit-rate 3000000000;
      shaping-rate percent 100;
      priority low;
    }
  }
}

```



**NOTE:** The forwarding class and classifier configurations are not shown because the `show` command does not display default portions of the configuration.

## Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default **fcoe** forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 5628](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created on page 5629](#)
- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created on page 5629](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5630](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 5631](#)
- [Verifying That the Interfaces Are Correctly Configured on page 5632](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 5635](#)
- [Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 on page 5635](#)

### *Verifying That the Output Queue Schedulers Have Been Created*

**Purpose** Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

**Action** List the scheduler map using the operational mode command `show class-of-service scheduler-map fcoe-map`:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
  Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
  Buffer Limit: none, Priority: low
  Excess Priority: unspecified
  Shaping rate: 100 percent,
  drop-profile-map-set-type: mark
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      <default-drop-profile>
    Medium high   any       1      <default-drop-profile>
    High          any       1      <default-drop-profile>
```

**Meaning** The `show class-of-service scheduler-map fcoe-map` command lists the properties of the scheduler map `fcoe-map`. The command output includes:

- The name of the scheduler map (`fcoe-map`)
- The name of the scheduler (`fcoe-sched`)
- The forwarding classes mapped to the scheduler (`fcoe`)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

***Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created***

**Purpose** Verify that the traffic control profile `fcoe-tcp` has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

**Action** List the FCoE traffic control profile properties using the operational mode command `show class-of-service traffic-control-profile fcoe-tcp`:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
  Shaping rate: 100 percent
  Scheduler map: fcoe-map
  Guaranteed rate: 3000000000
```

**Meaning** The `show class-of-service traffic-control-profile fcoe-tcp` command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (`fcoe-tcp`)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (`fcoe-map`)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

***Verifying That the Forwarding Class Set (Priority Group) Has Been Created***

**Purpose** Verify that the FCoE priority group has been created and that the `fcoe` priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.

**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
  Forwarding class          Index
  fcoe                      1
```

**Meaning** The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

#### *Verifying That Priority-Based Flow Control Has Been Enabled*

**Purpose** Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

**Action** List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
  Priority    PFC      MRU
  000        Disabled
  001        Disabled
  010        Disabled
  011        Enabled   2500
  100        Disabled
  101        Disabled
  110        Disabled
  111        Disabled
Type: Output
  Priority    Flow-Control-Queues
  000
  001        0
  010        1
  011        2
  100        3
  101        4
  110        5
  111        6
  111        7
```

**Meaning** The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point 011 (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.



The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### ***Verifying That the Interface Class of Service Configuration Has Been Created***

**Purpose** Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

**Action** List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}

ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
}
```

```
        congestion-notification-profile fcoe-cnp;
    }
    xe-0/0/33 {
        forwarding-class-set {
            fcoe-pg {
                output-traffic-control-profile fcoe-tcp;
            }
        }
        congestion-notification-profile fcoe-cnp;
    }
    ae1 {
        forwarding-class-set {
            fcoe-pg {
                output-traffic-control-profile fcoe-tcp;
            }
        }
        congestion-notification-profile fcoe-cnp;
    }
}
```

**Meaning** The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



**NOTE:** Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

---

### ***Verifying That the Interfaces Are Correctly Configured***

**Purpose** Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

**Action** List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
```

```

xe-0/0/10 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/11 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/20 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/21 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}

```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```

user@switch> show configuration interfaces
xe-0/0/25 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/26 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/30 {
    mtu 2180;
    unit 0 {

```

```
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/31 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/32 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/33 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
```

**Meaning** The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The interface mode (**trunk** mode both for interfaces that connect two switches and for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe\_vlan**)

***Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces***

**Purpose** Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the port security configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration vlans fcoe\_vlan forwarding-options fip-security**:

```
user@switch> show configuration vlans fcoe_vlan forwarding-options fip-security
interface ae1.0 {
    fcoe-trusted;
}
examine-vn2vn {
    beacon-period 90000;
}
```

**Meaning** The **show configuration vlans fcoe\_vlan forwarding-options fip-security** command lists VLAN FIP security information, including whether a port member of the VLAN is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- VN2VN\_Port FIP snooping is enabled (**examine-vn2vn**) on the FCoE VLAN and the beacon period is set to 90000 milliseconds. On Transit Switches TS1 and TS2, all interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Transit Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

***Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2***

**Purpose** Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan,      Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
...
```



**NOTE:** The output has been truncated to show only the relevant information.

---

**Meaning** The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe\_vlan**
- The FIP snooping mode is VN2VN\_Port FIP snooping (**VN2VN Snooping**)

**Related Documentation**

- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
- [Configuring Link Aggregation on page 2593](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Understanding MC-LAGs on an FCoE Transit Switch on page 5555](#)

### Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

This example shows how to configure VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping when the hosts are directly connected to the same FCoE transit switch.



**NOTE:** This example uses the Junos OS Enhanced Layer 2 Software (ELS) configuration style for QFX Series switches. If your switch runs software that does not support ELS, see *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 43.

---

VN2VN\_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN\_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN\_Port FIP snooping is conceptually similar to VN2VN\_Port FIP snooping between VN\_Ports and VF\_Ports, but VN2VN\_Port FIP snooping does not require traffic between VN\_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN\_Port communicates transparently through the transit switch on a virtual link that emulates a direct connection to the VN\_Port at the other end of the virtual link.

To configure VN2VN\_Port FIP snooping when the hosts are directly connected to the same FCoE transit switch, you must follow these configuration rules:

- VN2VN\_Port traffic must use a dedicated FCoE VLAN, and all ENodes that communicate using VN2VN\_Port FIP snooping must use that FCoE VLAN. You cannot mix VN2VN\_Port FIP snooping traffic with VN2VF\_Port FIP snooping traffic in the same FCoE VLAN.



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port FIP snooping VLANs, VN\_Port to VF\_Port (FIP snooping) traffic is dropped.

- ENode-facing ports must be set in **trunk** interface mode.
- ENode-facing ports must be untrusted ports.
- Network-facing (switch-facing) ports must be set in **trunk** interface mode.
- Network-facing ports must be FCoE trusted ports.
- Explicitly configure the beacon period. The beacon period is essentially a keepalive timer for virtual link maintenance.

When you enable VN2VF\_Port FIP snooping, the system snoops VN\_Port to VF\_Port packets and enforces security only on VN\_Port to VF\_Port virtual links. When you enable VN2VN\_Port FIP snooping, the system snoops VN\_Port to VN\_Port packets and enforces security only on VN\_Port to VN\_Port virtual links.

The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

This example describes how to configure VN2VN\_Port FIP snooping when the FCoE hosts are directly connected to the same transit switch:

- [Requirements on page 5637](#)
- [Overview on page 5638](#)
- [Configuration on page 5639](#)
- [Verification on page 5639](#)

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX5100 Switch running the ELS CLI and used as a transit switch

- Junos OS Release 13.2 or later for the QFX Series
- Two FCoE hosts that have ENodes

Overview

This example shows you how to:

- Set the correct interface mode on the transit switch.
- Configure the interfaces to use the dedicated FCoE VLAN for VN2VN\_Port FIP snooping.
- Configure the dedicated FCoE VLAN for VN2VN\_Port FIP snooping traffic.
- Enable VN2VN\_Port FIP snooping on the FCoE VLAN and configure the beacon period.

Topology


Table 460 on page 5638 shows the configuration components for this example.

Table 460: Components of the VN2VN\_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)

| Component                           | Settings  |
|-------------------------------------|---|
| Hardware                            | QFX5100 switch running the ELS CLI (FCoE transit switch TS1)<br><br>Two FCoE hosts that have ENodes (ENode1 and ENode2, respectively)   |
| Interfaces and interface mode       | <ul style="list-style-type: none"><li>• Interface <b>xe-0/0/20</b>, interface mode <b>trunk</b>, connects directly to the FCoE host with ENode1.</li><li>• Interface <b>xe-0/0/21</b>, interface mode <b>trunk</b>, connects directly to the FCoE host with ENode2.</li></ul> |
| Interface VLAN membership           | Both interfaces use VLAN <b>vlan200</b> .   |
| VN2VN_Port FIP snooping VLAN        | VLAN name— <b>vlan200</b><br>VLAN ID—200  |
| FIP snooping mode and beacon period | Set <b>examine-vn2vn</b> (VN2VN_Port FIP snooping)<br>Beacon period—90000 ms  |

Figure 197 on page 5638 shows the network topology for this example.

Figure 197: VN2VN\_Port FIP Snooping (FCoE Hosts Connected to Same Transit Switch) Topology





## Configuration

**CLI Quick Configuration** To quickly configure VN2VN\_Port FIP snooping for FCoE hosts connected directly to the same transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

### *Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)*

**Step-by-Step Procedure** To configure interface mode, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN\_Port traffic, configure the VLAN, set the beacon period, and enable VN2VN\_Port FIP snooping:

1. Configure the modes of the interfaces that connect directly to the FCoE host ENodes:

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces connected to the ENodes are members of the dedicated VN2VN\_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the FCoE VLAN dedicated to VN2VN\_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

4. Enable VN2VN\_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

## Verification

To verify that the VN2VN\_Port FIP snooping configuration has been created and is operating properly, perform these tasks:

- [Verifying That VN2VN\\_Port FIP Snooping is Enabled on the FCoE VLAN on page 5639](#)

### *Verifying That VN2VN\_Port FIP Snooping is Enabled on the FCoE VLAN*

**Purpose** Verify that VN2VN\_Port FIP snooping is enabled on the correct VLAN (**vlan200**), the beacon period is set to **90000** milliseconds, and the correct interfaces (**xe-0/0/20** and **xe-0/0/21**) are members of the VLAN.

**Action** List the FIP snooping information using the operational mode command **show fip snooping detail**.

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/20
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/21
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
```

**Meaning** The **show fip snooping detail** command lists all of the transit switch information about VN2VN\_Port FIP snooping and VN2VF\_Port FIP snooping. The command shows that:

- The VLAN is **vlan200**.
- The mode is FIP snooping mode **VN2VN**, for VN2VN\_Port FIP snooping. (If the Mode field shows **VN2VF**, then the FIP snooping mode is VN2VF\_Port FIP snooping.)
- The beacon period is **90000**.
- The interfaces for the ENodes are **xe-0/0/20** and **xe-0/0/21**.

In addition, this useful command shows information about the ENodes and the VN2VN\_Port sessions.

- Related Documentation**
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5641](#)
  - [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5647](#)
  - [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)
  - [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)

## Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)

This example shows how to configure VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping when the hosts are directly connected to different FCoE transit switches, and the transit switches are directly connected to each other.



**NOTE:** This example uses the Junos OS Enhanced Layer 2 Software (ELS) configuration style for QFX Series switches. If your switch runs software that does not support ELS, see *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)*. For ELS details, see “[Getting Started with Enhanced Layer 2 Software](#)” on page 43.

VN2VN\_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN\_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN\_Port FIP snooping is conceptually similar to VN2VF\_Port FIP snooping between VN\_Ports and VF\_Ports, but VN2VN\_Port FIP snooping does not require traffic between VN\_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN\_Port communicates transparently through one or more transit switches on a virtual link that emulates a direct connection to the VN\_Port at the other end of the virtual link.

To configure VN2VN\_Port FIP snooping when the hosts are directly connected to different FCoE transit switches, and the transit switches are directly connected to each other, you must follow these configuration rules:

- VN2VN\_Port traffic must use a dedicated FCoE VLAN, and all ENodes that communicate using VN2VN\_Port FIP snooping must use that FCoE VLAN. The FCoE VLAN must be configured on each transit switch. You cannot mix VN2VN\_Port FIP snooping traffic with VN2VF\_Port FIP snooping traffic in the same FCoE VLAN.



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port FIP snooping VLANs, VN2VF\_Port traffic is dropped.

- ENode-facing ports must be set in **trunk** interface mode.
- ENode-facing ports must be untrusted ports.
- Network-facing (switch-facing) ports must be set in **trunk** interface mode.

- Network-facing ports must be FCoE trusted ports.
- Explicitly configure the beacon period. The beacon period is essentially a keepalive timer for virtual link maintenance.

When you enable VN2VF\_Port FIP snooping, the system snoops VN\_Port to VF\_Port packets and enforces security only on VN\_Port to VF\_Port virtual links. When you enable VN2VN\_Port FIP snooping, the system snoops VN\_Port to VN\_Port packets and enforces security only on VN\_Port to VN\_Port virtual links.

The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

This example describes how to configure VN2VN\_Port FIP snooping when the FCoE hosts are directly connected to different transit switches, and the transit switches are directly connected to each other:

- [Requirements on page 5642](#)
- [Overview on page 5642](#)
- [Configuration on page 5643](#)
- [Verification on page 5645](#)

---

## Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX5100 Switches running the ELS CLI and used as transit switches
- Junos OS Release 13.2 or later for the QFX Series
- Two FCoE hosts that have ENodes

---

## Overview

This example shows you how to:

- Set the correct interface mode on the transit switch.
- Configure the interfaces to use the dedicated FCoE VLAN for VN2VN\_Port FIP snooping.
- Configure the network-facing interfaces as FCoE trusted interfaces.
- Configure the dedicated FCoE VLAN for VN2VN\_Port FIP snooping traffic.
- Enable VN2VN\_Port FIP snooping on the FCoE VLAN and configure the beacon period.

## Topology


[Table 461 on page 5643](#) shows the configuration components for this example.

**Table 461: Components of the VN2VN\_Port FIP Snooping Configuration Topology (FCoE Hosts Directly Connected to Different FCoE Transit Switches)**

| Component                           | Settings   |
|-------------------------------------|--|
| Hardware                            | Two QFX5100 switches running the ELS CLI (FCoE transit switch TS1 and FCoE transit switch TS2)<br><br>Two FCoE hosts that have ENodes (ENode1 and ENode2, respectively)  |
| Interfaces and interface mode       | <ul style="list-style-type: none"> <li>Interface <b>xe-0/0/20</b>, interface mode <b>trunk</b>, connects directly from transit switch TS1 to the FCoE host with ENode1.</li> <li>Interface <b>xe-0/0/21</b>, interface mode <b>trunk</b>, connects directly from transit switch TS1 to transit switch TS2.</li> <li>Interface <b>xe-0/0/31</b>, interface mode <b>trunk</b>, connects directly from transit switch TS2 to transit switch TS1.</li> <li>Interface <b>xe-0/0/30</b>, interface mode <b>trunk</b>, connects directly from transit switch TS2 to the FCoE host with ENode2.</li> </ul> |
| Interface VLAN membership           | The interfaces on both transit switches use VLAN <b>vlan200</b> .  |
| VN2VN_Port FIP snooping VLAN        | VLAN name (both transit switches)— <b>vlan200</b><br>VLAN ID—200   |
| FIP snooping mode and beacon period | Set <b>examine-vn2vn</b> (VN2VN_Port FIP snooping)<br>Beacon period—90000 ms   |

Figure 198 on page 5643 shows the network topology for this example.

**Figure 198: VN2VN\_Port FIP Snooping (FCoE Hosts Connected to Different Transit Switches) Topology**



### Configuration

To configure VN2VN\_Port FIP snooping for VN\_Ports that are directly connected to different transit switches (and the transit switches are directly connected to each other), perform these tasks:

- [Configuring VN2VN\\_Port FIP Snooping on FCoE Transit Switch TS1 on page 5644](#)
- [Configuring VN2VN\\_Port FIP Snooping on FCoE Transit Switch TS2 on page 5645](#)

### CLI Quick Configuration

The configuration for each FCoE transit switch is shown separately.

To quickly configure VN2VN\_Port FIP snooping for FCoE hosts connected directly to different transit switches, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS1:

## FCoE Transit Switch TS1

```
set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set vlans vlan200 forwarding-options fip-security interface xe-0/0/21 fcoe-trusted
set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

To quickly configure VN2VN\_Port FIP snooping for FCoE hosts connected directly to different transit switches, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS2:

## FCoE Transit Switch TS2

```
set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set vlans vlan200 forwarding-options fip-security interface xe-0/0/31 fcoe-trusted
set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

### *Configuring VN2VN\_Port FIP Snooping on FCoE Transit Switch TS1*

#### Step-by-Step Procedure

To configure interface mode, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN\_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN\_Port FIP snooping:

1. Configure the modes of the interfaces that connect directly to the FCoE host with ENode1 (**xe-0/0/20**) and to FCoE transit switch TS2 (**xe-0/0/21**):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN\_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the FCoE VLAN dedicated to VN2VN\_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

4. Configure the network-facing port (**xe-0/0/21**) as an FCoE trusted port:

```
user@switch# set vlans vlan200 forwarding-options fip-security interface xe-0/0/21 fcoe-trusted
```

5. Enable VN2VN\_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

### Configuring VN2VN\_Port FIP Snooping on FCoE Transit Switch TS2

- Step-by-Step Procedure** To configure interface mode, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN\_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN\_Port FIP snooping:
1. Configure the modes of the interfaces that connect directly to the FCoE host with ENode2 (**xe-0/0/30**) and to FCoE transit switch TS1 (**xe-0/0/31**):  

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk
```
  2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN\_Port VLAN (**vlan200**):  

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
```
  3. Configure the FCoE VLAN dedicated to VN2VN\_Port FIP snooping:  

```
user@switch# set vlans vlan200 vlan-id 200
```
  4. Configure the network-facing port (**xe-0/0/31**) as an FCoE trusted port:  

```
user@switch# set vlans vlan200 forwarding-options fip-security interface xe-0/0/31 fcoe-trusted
```
  5. Enable VN2VN\_Port FIP snooping on the VLAN and configure the beacon period:  

```
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

### Verification

To verify that the VN2VN\_Port FIP snooping configuration has been created and is operating properly on both switches, perform these tasks:

- [Verifying That VN2VN\\_Port FIP Snooping is Enabled on the FCoE VLAN \(Transit Switches TS1 and TS2\) on page 5645](#)

#### *Verifying That VN2VN\_Port FIP Snooping is Enabled on the FCoE VLAN (Transit Switches TS1 and TS2)*

- Purpose** Verify that VN2VN\_Port FIP snooping is enabled on the correct VLAN (**vlan200**), the beacon period is set to **90000** milliseconds, and that the correct interfaces (**xe-0/0/20** and **xe-0/0/21** on TS1, and **xe-0/0/30** and **xe-0/0/31** on TS2) are members of the VLAN.

**Action** List the FIP snooping information on transit switch TS1 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/20
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fc:00:01:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0b:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/21
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fc:00:01:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0a:01
```

List the FIP snooping information on transit switch TS2 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/30
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/31
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
```

**Meaning** The **show fip snooping detail** command lists all of the transit switch information about VN2VN\_Port FIP snooping and VN2VF\_Port FIP snooping on each transit switch. The command shows that:

- The VLAN is **vlan200**.
- The mode is FIP snooping mode **VN2VN**, for VN2VN\_Port FIP snooping. (If the Mode field shows **VN2VF**, then the FIP snooping mode is VN2VF\_Port FIP snooping.)



- The beacon period is **90000**.
- The interfaces connected to the ENodes are **xe-0/0/20** and **xe-0/0/21** on transit switch TS1, and **xe-0/0/30** and **xe-0/0/31** on transit switch TS2. Because the transit switches are transparent passthrough switches, the network-facing trunk ports “see” the FCoE host ENodes at the far end of the VN2VN\_Port virtual link.

In addition, this useful command shows information about the ENodes and the VN2VN\_Port sessions.

#### Related Documentation

- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5636](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5647](#)
- [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)

### Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)

This example shows how to configure VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping when the hosts are indirectly connected through an aggregation layer FCoE transit switch. Each FCoE host ENode is directly connected to an FCoE transit switch, but the FCoE transit switches are not directly connected to each other. The FCoE transit switches are both connected to a third FCoE transit switch that acts as an aggregation layer switch.



**NOTE:** This example uses the Junos OS Enhanced Layer 2 Software (ELS) configuration style for QFX Series switches. If your switch runs software that does not support ELS, see [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\)](#). For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).

VN2VN\_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN\_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN\_Port FIP snooping is conceptually similar to VN2VN\_Port FIP snooping between VN\_Ports and VF\_Ports, but VN2VN\_Port FIP snooping does not require traffic between VN\_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN\_Port communicates transparently through one or more transit switches on a virtual link that emulates a direct connection to the VN\_Port at the other end of the virtual link.

To configure VN2VN\_Port FIP snooping when the hosts are indirectly connected, you must follow these configuration rules:

- VN2VN\_Port traffic must use a dedicated FCoE VLAN, and all ENodes that communicate using VN2VN\_Port FIP snooping must use that FCoE VLAN. The FCoE VLAN must be configured on each transit switch. You cannot mix VN2VN\_Port FIP snooping traffic with VN2VF\_Port FIP snooping traffic in the same FCoE VLAN.



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port FIP snooping VLANs, VN\_Port to VF\_Port traffic is dropped.

- ENode-facing ports must be set in **trunk** interface mode.
- ENode-facing ports must be untrusted ports.
- Network-facing (switch-facing) ports must be set in **trunk** interface mode.
- Network-facing ports must be FCoE trusted ports.
- Explicitly configure the beacon period. The beacon period is essentially a keepalive timer for virtual link maintenance.

When you enable FIP snooping, the system snoops VN\_Port to VF\_Port packets and enforces security only on VN\_Port to VF\_Port virtual links. When you enable VN2VN\_Port FIP snooping, the system snoops VN\_Port to VN\_Port packets and enforces security only on VN\_Port to VN\_Port virtual links.

The transit switch applies VN2VN\_Port FIP snooping filters at the ports associated with the FCoE VLANs on which you enable VN2VN FIP snooping.

This example describes how to configure VN2VN\_Port FIP snooping when the FCoE hosts are indirectly connected across an aggregation layer FCoE transit switch:

- [Requirements on page 5648](#)
- [Overview on page 5649](#)
- [Configuration on page 5650](#)
- [Verification on page 5653](#)

### Requirements

---

This example uses the following hardware and software components:

- Three Juniper Networks QFX5100 Switches running the ELS CLI and used as transit switches
- Junos OS Release 13.2 or later for the QFX Series
- Two FCoE hosts that have ENodes

## Overview

This example shows you how to:

- Set the correct interface mode on the transit switch.
- Configure the interfaces to use the dedicated FCoE VLAN for VN2VN\_Port FIP snooping.
- Configure the network-facing interfaces as FCoE trusted interfaces.
- Configure the dedicated FCoE VLAN for VN2VN\_Port FIP snooping traffic.
- Enable VN2VN\_Port FIP snooping on the FCoE VLAN and configure the beacon period.

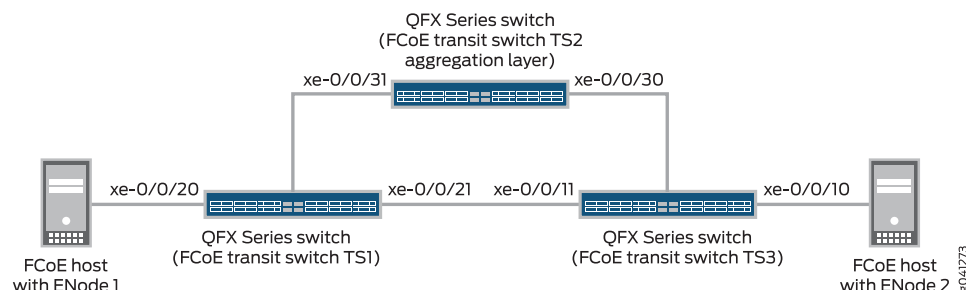
## Topology

Table 462 on page 5649 shows the configuration components for this example.

**Table 462: Components of the VN2VN\_Port FIP Snooping Configuration Topology (FCoE Hosts Indirectly Connected Across an Aggregation Layer FCoE Transit Switch)**

| Component                           | Settings   |
|-------------------------------------|--|
| Hardware                            | <p>Three QFX5100 switches running the ELS CLI, two of which are FCoE transit switches that are directly attached to the FCoE hosts (transit switches TS1 and TS2) and one of which is an aggregation layer FCoE transit switch (TS3)</p> <p>Two FCoE hosts that have ENodes (ENode1 and ENode2, respectively)</p>  |
| Interfaces and interface mode       | <ul style="list-style-type: none"> <li>• Interface <b>xe-0/0/20</b>, interface mode <b>trunk</b>, connects directly from transit switch TS1 to the FCoE host with ENode1.</li> <li>• Interface <b>xe-0/0/21</b>, interface mode <b>trunk</b>, connects directly from transit switch TS1 to aggregation layer transit switch TS2.</li> <li>• Interface <b>xe-0/0/31</b>, interface mode <b>trunk</b>, connects directly from aggregation layer transit switch TS2 to transit switch TS1.</li> <li>• Interface <b>xe-0/0/30</b>, interface mode <b>trunk</b>, connects directly from aggregation layer transit switch TS2 to transit switch TS3.</li> <li>• Interface <b>xe-0/0/11</b>, interface mode <b>trunk</b>, connects directly from transit switch TS3 to aggregation layer transit switch TS2.</li> <li>• Interface <b>xe-0/0/10</b>, interface mode <b>trunk</b>, connects directly from transit switch TS3 to the FCoE host with ENode2.</li> </ul> |
| Interface VLAN membership           | The interfaces on all three switches use VLAN <b>vlan200</b> .   |
| VN2VN_Port FIP snooping VLAN        | <p>VLAN name (all three switches)—<b>vlan200</b></p> <p>VLAN ID—200</p>  |
| FIP snooping mode and beacon period | <p>Set <b>examine-vn2vn</b> (VN2VN_Port FIP snooping)</p> <p>Beacon period—90000 ms</p>  |

Figure 199 on page 5650 shows the network topology for this example.

**Figure 199: VN2VN\_Port FIP Snooping (FCoE Hosts Indirectly Connected) Topology**

### Configuration

To configure VN2VN\_Port FIP snooping for VN\_Ports that are indirectly connected across an aggregation layer FCoE transit switch, perform these tasks:

- [Configuring VN2VN\\_Port FIP Snooping on FCoE Transit Switch TS1 on page 5651](#)
- [Configuring VN2VN\\_Port FIP Snooping on Aggregation Layer FCoE Transit Switch TS2 on page 5652](#)
- [Configuring VN2VN\\_Port FIP Snooping on FCoE Transit Switch TS3 on page 5652](#)

#### CLI Quick Configuration

The configuration for each FCoE transit switch is shown separately.

To quickly configure VN2VN\_Port FIP snooping for FCoE hosts that are indirectly connected across an aggregation layer FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS1:

#### FCoE Transit Switch TS1

```
set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set vlans vlan200 forwarding-options fip-security interface xe-0/0/21 fcoe-trusted
set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

To quickly configure VN2VN\_Port FIP snooping for FCoE hosts that are indirectly connected across an aggregation layer FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS2:

#### FCoE Transit Switch TS2

```
set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
```

```
set vlans vlan200 forwarding-options fip-security interface xe-0/0/30 fcoe-trusted
set vlans vlan200 forwarding-options fip-security interface xe-0/0/31 fcoe-trusted
set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

To quickly configure VN2VN\_Port FIP snooping for FCoE hosts that are indirectly connected across an aggregation layer FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. To configure FCoE transit switch TS3:

### FCoE Transit Switch TS3

```
set interfaces xe-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/11 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members vlan200
set vlans vlan200 vlan-id 200
set vlans vlan200 forwarding-options fip-security interface xe-0/0/11 fcoe-trusted
set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

#### *Configuring VN2VN\_Port FIP Snooping on FCoE Transit Switch TS1*

#### Step-by-Step Procedure

To configure interface mode, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN\_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN\_Port FIP snooping:

1. Configure the modes of the interfaces that connect directly to the FCoE host with ENode1 (**xe-0/0/20**) and to aggregation layer FCoE transit switch TS2 (**xe-0/0/21**):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN\_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the FCoE VLAN dedicated to VN2VN\_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

4. Configure the network-facing port (**xe-0/0/21**) as an FCoE trusted port:

```
user@switch# set vlans vlan200 forwarding-options fip-security interface xe-0/0/21 fcoe-trusted
```

5. Enable VN2VN\_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

**Configuring VN2VN\_Port FIP Snooping on Aggregation Layer FCoE Transit Switch TS2**

**Step-by-Step Procedure** To configure interface mode, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN\_Port traffic, set the network-facing ports as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN\_Port FIP snooping:

1. Configure the mode of the interfaces that connect directly to FCoE transit switches TS1 (**xe-0/0/31**) and TS3 (**xe-0/0/30**). Both interfaces are network-facing and must be configured as trunk interfaces:

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/31 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN\_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/31 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the FCoE VLAN dedicated to VN2VN\_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

4. Configure the network-facing ports (**xe-0/0/30** and **xe-0/0/31**) as FCoE trusted ports:

```
user@switch# set vlans vlan200 forwarding-options fip-security interface xe-0/0/30 fcoe-trusted
user@switch# set vlans vlan200 forwarding-options fip-security interface xe-0/0/31 fcoe-trusted
```

5. Enable VN2VN\_Port FIP snooping on the VLAN and configure the beacon period:

```
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period 90000
```

**Configuring VN2VN\_Port FIP Snooping on FCoE Transit Switch TS3**

**Step-by-Step Procedure** To configure interface mode, configure interface VLAN membership in the FCoE VLAN dedicated to VN2VN\_Port traffic, set the network-facing port as FCoE trusted, configure the VLAN, set the beacon period, and enable VN2VN\_Port FIP snooping:

1. Configure the mode of the interfaces that connect directly to the FCoE host with ENode2 (**xe-0/0/10**) and to aggregation layer FCoE transit switch TS2 (**xe-0/0/11**):

```
user@switch# set interfaces xe-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/11 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the interface VLAN membership so that the interfaces are members of the dedicated VN2VN\_Port VLAN (**vlan200**):

```
user@switch# set interfaces xe-0/0/10 unit 0 family ethernet-switching vlan members vlan200
set interfaces xe-0/0/11 unit 0 family ethernet-switching vlan members vlan200
```

3. Configure the FCoE VLAN dedicated to VN2VN\_Port FIP snooping:

```
user@switch# set vlans vlan200 vlan-id 200
```

4. Configure the network-facing port (**xe-0/0/11**) as an FCoE trusted port:  

```
user@switch# set vlans vlan200 forwarding-options fip-security interface xe-0/0/11
fcoe-trusted
```
5. Enable VN2VN\_Port FIP snooping on the VLAN and configure the beacon period:  

```
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn
beacon-period 90000
```

### Verification

To verify that the VN2VN\_Port FIP snooping configuration has been created and is operating properly on all three switches, perform these tasks:

- [Verifying That VN2VN\\_Port FIP Snooping Is Enabled on the FCoE VLAN \(All Three Transit Switches\) on page 5653](#)

#### *Verifying That VN2VN\_Port FIP Snooping Is Enabled on the FCoE VLAN (All Three Transit Switches)*

**Purpose** Verify that VN2VN\_Port FIP snooping is enabled on the correct VLAN (**vlan200**), the beacon period is set to **90000** milliseconds, and that the correct interfaces (**xe-0/0/20** and **xe-0/0/21** on TS1, **xe-0/0/30** and **xe-0/0/31** aggregation layer TS2, and **xe-0/0/10** and **xe-0/0/11** on TS3) are members of the VLAN.

**Action** List the FIP snooping information on transit switch TS1 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/20
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fc:00:01:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0b:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/21
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fc:00:01:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0a:01
```

List the FIP snooping information on aggregation layer transit switch TS2 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/30
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fc:00:01:0b:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0a:01
    Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/31
    Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fc:00:01:0a:01
      Active Sessions : 1
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fc:00:01:0b:01
```

List the FIP snooping information on transit switch TS3 using the operational mode command **show fip snooping detail**

```
user@switch> show fip snooping detail
VLAN: vlan200, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Point-to-Point
  Enode Information
```



```

Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/10
Active VN_Ports : 1
VN_Port Information
VN-Port MAC: 0e:fd:00:00:0b:01
Active Sessions : 1
Session Information
Vlink far-end VN-Port-MAC: 0e:fd:00:00:0a:01
Enode-MAC: 10:10:94:01:00:02,      Interface: xe-0/0/11
Active VN_Ports : 1
VN_Port Information
VN-Port MAC: 0e:fd:00:00:0a:01
Active Sessions : 1
Session Information
Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01

```

**Meaning** The **show fip snooping detail** command lists all of the transit switch information about VN2VN\_Port FIP snooping and VN2VF\_Port FIP snooping on each transit switch. The command shows that:

- The VLAN is **vlan200**.
- The mode is FIP snooping mode **VN2VN**, for VN2VN\_Port FIP snooping. (If the Mode field shows **VN2VF**, then the FIP snooping mode is VN2VF\_Port FIP snooping.)
- The beacon period is **90000**.
- The interfaces connected to the ENodes are **xe-0/0/20** and **xe-0/0/21** on transit switch TS1, **xe-0/0/30** and **xe-0/0/31** on aggregation layer transit switch TS2, and **xe-0/0/10** and **xe-0/0/11** on transit switch TS3. Because the transit switches are transparent passthrough switches, the network-facing trunk ports “see” the FCoE host ENodes at the far end of the VN2VN\_Port virtual link.

In addition, this useful command shows information about the ENodes and the VN2VN\_Port sessions.

- Related Documentation**
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\) on page 5636](#)
  - [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\) on page 5641](#)
  - [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)
  - [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)

## FCoE and FIP Snooping Configuration Tasks

- [Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 5656](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)

- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)
- [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)

## Enabling and Disabling CoS OxID Hash Control on Standalone Switches

The originator exchange identifier (OxID) field is one of several fields that the switch can use in its hash function computation for FCoE traffic load balancing over multiple outgoing links in an Ethernet link aggregation group (LAG) on ports that face an FCoE forwarder (FCF). You can configure whether or not the switch uses the OxID in the hash computation.

Including the OxID field in the load-balancing hash computation allows different exchanges between a pair of Fibre Channel (FC) endpoints (such as an FCoE host and an FC storage device) to take different paths across the network, thus improving the aggregate network throughput.

However, if the paths between different sets of FC endpoints have common links, congestion on one set of FC endpoints can affect the other set of endpoints. Such congestion can happen if the FCoE traffic on the two sets of endpoints uses the same priority (IEEE 802.1p code point). It is common for networks to use priority 3 (IEEE 802.1p code point 011) for FCoE traffic. However, you can assign different IEEE priorities to different lossless FCoE flows as described in "[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)" on page 5837 to further separate the traffic flows.

OxID hash control is enabled by default.

- To enable OxID hash control field for FCoE traffic load balancing:

```
[edit forwarding-options hash-key]
user@switch# set family fcoe oxid enable
```

- To disable OxID hash control field for FCoE traffic load balancing:

```
[edit forwarding-options hash-key]
user@switch# set family fcoe oxid disable
```

### Related Documentation

- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 5530](#)
- [Understanding OxID Hash Control for FCoE Traffic Load Balancing on QFabric Systems](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)

## Configuring VLANs for FCoE Traffic on an FCoE Transit Switch

When you configure a switch as a Fibre Channel over Ethernet (FCoE) transit switch, you must configure a VLAN that transports only FCoE traffic. FCoE traffic requires a dedicated VLAN and cannot share a VLAN with any other type of traffic. Because FCoE traffic is tagged traffic, the port (or interface) mode cannot be access mode, it must be either tagged-access port-mode (for switches that run the original CLI) or trunk interface-mode (for switches that run the Enhanced Layer 2 Software (ELS) CLI).

However, each interface that belongs to an FCoE VLAN must not only transport the tagged FCoE traffic, it must also transport the untagged FCoE Initialization Protocol (FIP) traffic. FIP communicates with the storage area network (SAN) Fibre Channel (FC) switch to set up the FCoE session for the FCoE client.

To transport untagged traffic on a tagged-access or trunk mode interface, the interface must have a native VLAN configured on it. Therefore, each interface that belongs to an FCoE VLAN must also have a native VLAN on it.

There are slight differences in the way you configure a native VLAN on an interface, depending on whether the switch uses the ELS CLI or the original CLI. This topic describes both methods.



**NOTE:** FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

FCoE traffic cannot use a standard LAG because traffic might be hashed to different physical LAG links on different transmissions. This breaks the (virtual) point-to-point link that Fibre Channel traffic requires. If you configure a standard LAG interface for FCoE traffic, FCoE traffic might be rejected by the FC SAN.

QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across the QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG.



**NOTE:** To configure an FCoE VLAN on a QFX3500 switch that you are using as an FCoE-FC gateway, you must also configure an FCoE VLAN interface as described in *Configuring an FCoE VLAN Interface on an FCoE-FC Gateway*. (Only the QFX3500 switch supports FCoE-FC gateway configuration.)

---

FCoE VLAN configuration includes:

- Configuring a VLAN to use as a dedicated FCoE VLAN
- Configuring the interface members of the FCoE VLAN.
- Configuring a native VLAN for FIP traffic.

This topic includes two configuration procedures, one for switches that run the original CLI, and one for switches that run the ELS CLI.

## Original CLI Configuration

To configure an FCoE VLAN on a non-ELS switch:

1. Configure a dedicated FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id
```

For example, to configure a VLAN named **fcoe\_vlan** with a VLAN ID of **100** as the FCoE VLAN:

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

2. Configure the FCoE VLAN on the interface (use **ethernet-switching** as the family and **tagged-access** as the port mode):

```
[edit interfaces]
user@switch# set interface-name unit unit family family port-mode mode vlan members
vlan-name
```

For example, to configure the interface **xe-0/0/10** as a member of the FCoE VLAN **fcoe\_vlan**:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access vlan
members fcoe_vlan
```

3. Configure the Ethernet interface membership in the FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name interface interface-name
```

For example, to assign the interface **xe-0/0/10.0** to the FCoE VLAN named **fcoe\_vlan**:

```
[edit vlans]
user@switch# set fcoe_vlan interface xe-0/0/10.0
```

4. Configure a native VLAN for the untagged FIP traffic:

```
[edit vlans]
user@switch# set native vlan-id vlan-id
```

For example, to configure the native VLAN with a VLAN ID of 1:

```
[edit vlans]
user@switch# set native vlan-id 1
```

5. Assign member interfaces to the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit unit family family native-vlan-id vlan-id
```

For example, to configure the interface **xe-0/0/10** as a member of the native VLAN with the native VLAN ID 1:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
```

## ELS CLI Configuration

To configure an FCoE VLAN on a switch running ELS:

1. Configure a dedicated FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id
```

For example, to configure a VLAN named **fcoe\_vlan** with a VLAN ID of **100** as the FCoE VLAN:

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```

2. Configure the FCoE VLAN on the interface (use **ethernet-switching** as the family and **trunk** as the interface mode):

```
[edit interfaces]
user@switch# set interface-name unit unit family family interface-mode mode vlan members
vlan-name
```

For example, to configure the interface **xe-0/0/10** as a member of the FCoE VLAN **fcoe\_vlan**:

```
[edit interfaces]
user@switch# set xe-0/0/10 unit 0 family ethernet-switching interface-mode trunk vlan
members fcoe_vlan
```

3. Configure the Ethernet interface membership in the FCoE VLAN:

```
[edit vlans]
user@switch# set vlan-name interface interface-name
```

For example, to assign the interface **xe-0/0/10.0** to the FCoE VLAN named **fcoe\_vlan**:

```
[edit vlans]
user@switch# set fcoe_vlan interface xe-0/0/10.0
```

4. Configure a native VLAN on the physical Ethernet interface for the untagged FIP traffic:

```
[edit interfaces]
user@switch# set interface-name native-vlan-id vlan-id
```

For example, to configure the native VLAN on interface **xe-0/0/10** with a VLAN ID of **1**:

```
[edit interfaces]
user@switch# set xe-0/0/10 native-vlan-id 1
```

5. Configure the Ethernet interface as a member of the native VLAN:

```
[edit interfaces]
user@switch# set interface-name unit unit family family vlan members native-vlan-id
```



**NOTE:** The *native-vlan-id* number must be the same as the native VLAN ID number that you configured on the physical Ethernet interface (see step 4).

For example, to configure the interface **xe-0/0/10** as a member of the native VLAN with the native VLAN ID **1**:

```
[edit interfaces]
```

```
user@switch# set xe-0/0/10 unit 0 family ethernet-switching vlan members 1
```

**Related Documentation**

- [Understanding FCoE on page 5518](#)
- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)
- [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)
- *Disabling Enhanced FIP Snooping Scaling*
- *Configuring an FCoE LAG*
- *Configuring an FCoE VLAN Interface on an FCoE-FC Gateway*

## Configuring VN2VF\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch

VN\_Port to VF\_Port (VN2VF\_Port) Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping uses information gathered during FIP discovery and login to create firewall filters that provide security against unauthorized access to the FC switch or FCoE forwarder (FCF) through the switch when the switch is acting as an FCoE transit switch. The firewall filters allow only FCoE devices that successfully log in to the FC fabric to access the FCF through the transit switch. VN2VF\_Port FIP snooping provides security for the point-to-point virtual links that connect host FCoE Nodes (ENodes) and FCFs in the FCoE VLAN by denying access to any device that does not successfully log in to the FCF.

VN2VF\_Port FIP snooping is disabled by default. You enable VN2VF\_Port FIP snooping on a per-VLAN basis for VLANs that carry FCoE traffic. Ensure that a VLAN that carries FCoE traffic carries only FCoE traffic, because enabling VN2VF\_Port FIP snooping denies access for all other Ethernet traffic.





**NOTE:** All of the transit switch ports are untrusted by default. If an ENode on an FCoE device logs in to an FCF before you enable VN2VF\_Port FIP snooping on the VLAN and you then enable VN2VF\_Port FIP snooping, the transit switch denies traffic from the ENode because the transit switch has not snooped (learned) the ENode state. The following process automatically logs the ENode back in to the FCF to reestablish the connection:

1. VN2VF\_Port FIP snooping is enabled on an FCoE VLAN on the switch.
2. The switch denies existing connections between servers and the FCF on the FCoE VLAN by filtering the FCoE traffic and FIP traffic, so no keepalive messages from the ENodes reach the FCF.
3. The FCF port timer for each ENode and for each VN\_Port on each ENode expires.
4. The FCF sends each ENode whose port timer has expired a Clear Virtual Links (CVL) message.
5. The CVL message causes the ENode to log in again.

Because the FCF is a trusted source, you configure interfaces that connect to the FCF as FCoE trusted interfaces. FCoE trusted interfaces do not filter traffic (FIP snooping filtering should occur only at the FCoE access edge), but VN2VF\_Port FIP snooping continues to run on trusted interfaces so that the switch learns the FCF state.



**NOTE:** Do not configure ENode-facing interfaces both with FIP snooping enabled and as trusted interfaces. FCoE VLANs with interfaces that are directly connected to FCoE hosts should be configured with FIP snooping enabled and the interfaces should *not* be trusted interfaces. Ethernet interfaces that are connected to an FCF should be configured as trusted interfaces and should not have FIP snooping enabled. Interfaces that are connected to a transit switch that is performing FIP snooping can be configured as trusted interfaces if the FCoE VLAN is not enabled for FIP snooping.

Optionally, you can specify an FC-MAP value for each FCoE VLAN. On a given FCoE VLAN, the switch learns only FCFs that have a matching FC-MAP value. The default FC-MAP value is 0EFC00h for all FC devices. (Enter hexadecimal values for FC-MAP preceded by the hexadecimal indicator "0x"—for example, 0x0EFC00.) If you change the FC-MAP value of an FCF, change the FC-MAP value for the FCoE VLAN it belongs to on the switch and on the servers you want to communicate with the FCF. An FCoE VLAN can have one and only one FC-MAP value.



**NOTE:** The default enhanced FIP snooping scaling supports 2,500 sessions. On QFabric systems, starting with Junos OS Release 13.2X52, you can disable enhanced FIP snooping scaling on a per-VLAN basis if you want to do so, but only 376 sessions are supported if you disable enhanced FIP snooping scaling.

There are differences in the way you configure FIP snooping and FCoE trusted interfaces on a switch that depend on whether the switch uses the original CLI or the Enhanced Layer 2 Software (ELS) CLI. This topic includes two configuration procedures, one for switches that run the original CLI, and one for switches that run the ELS CLI.

### Original CLI Configuration

To enable VN2VF\_Port FIP snooping:

- To enable VN2VF\_Port FIP snooping on a single VLAN and specify the optional FC-MAP value:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan vlan-name examine-fip fc-map fc-map-value
```

For example, to enable VN2VF\_Port FIP snooping on a VLAN named **san1\_vlan** and change the FC-MAP value to **0x0EFC03**:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan san1_vlan examine-fip fc-map 0x0EFC03
```



**NOTE:** Changing the FC-MAP value causes all logins to drop and forces ENodes to log in again.

- To enable VN2VF\_Port FIP snooping on all VLANs and use the default FC-MAP value:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all examine-fip
```

- To configure an interface as an FCoE trusted interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface interface-name fc-e-trusted
```

For example, to configure interface **xe-0/0/30** as an FCoE trusted interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface xe-0/0/30 fc-e-trusted
```

### ELS CLI Configuration

To enable VN2VF\_Port FIP snooping:

- To enable VN2VF\_Port FIP snooping on a VLAN and specify the optional FC-MAP value:

```
[edit]
user@switch# set vlans vlan-name forwarding-options fip-security fc-map fc-map-value
examine-vn2vf
```

For example, to enable VN2VF\_Port FIP snooping on a VLAN named **san1\_vlan** and change the FC-MAP value to **0x0EFC03**:

```
[edit]
user@switch# set vlans san1_vlan forwarding-options fip-security fc-map 0x0EFC03
examine-vn2vf
```



**NOTE:** Changing the FC-MAP value causes all logins to drop and forces ENodes to log in again.

- To configure an interface as an FCoE trusted interface:

```
[edit]
user@switch# set vlans vlan-name forwarding-options fip-security interface interface-name
fcoe-trusted
```

For example, to configure interface **xe-0/0/30** on VLAN named **san1\_vlan** as an FCoE trusted interface:

```
[edit]
user@switch# set vlans san1_vlan forwarding-options fip-security interface xe-0/0/30
fcoe-trusted
```

#### Related Documentation

- *Example: Configuring an FCoE Transit Switch*
- *Configuring an FCoE VLAN Interface on an FCoE-FC Gateway*
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)
- *Configuring an FCoE LAG*
- *Disabling Enhanced FIP Snooping Scaling*
- *Understanding FIP Snooping*
- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- *Understanding FCoE LAGs*

## Enabling VN2VN\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch

VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN\_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions.

VN2VN\_Port FIP snooping is conceptually similar to VN2VF\_Port FIP snooping between VN\_Ports and VF\_Ports, but VN2VN\_Port FIP snooping does not require traffic between VN\_Ports to traverse the Fibre Channel (FC) switch or FCoE forwarder (FCF). Instead, a VN\_Port communicates transparently through the transit switch on a virtual link that emulates a direct connection to the VN\_Port at the other end of the virtual link.

VN2VN\_Port FIP snooping is disabled by default. You enable VN2VN\_Port FIP snooping on a per-VLAN basis on VLANs that carry VN2VN\_Port FCoE traffic. Ensure that the VLAN carries only FCoE traffic between VN\_Ports, because enabling VN2VN\_Port FIP snooping denies access for all other traffic, including VN2VF\_Port FIP snooping traffic.

All ENodes that you want to communicate using VN2VN\_Port FIP snooping must use an FCoE VLAN dedicated to VN2VN\_Port traffic. You cannot mix VN2VN\_Port FIP snooping traffic with VN2VF\_Port FIP snooping traffic in the same FCoE VLAN.



**NOTE:** An FCoE VLAN can support either VN2VF\_Port FIP snooping or VN2VN\_Port FIP snooping, but not both. Configure separate FCoE VLANs for VN2VF\_Port FIP snooping traffic and for VN2VN\_Port FIP snooping traffic. On FCoE VLANs that are configured as VN2VN\_Port FIP snooping VLANs, VN2VF\_Port traffic is dropped.

The *beacon period* is conceptually similar to the FIP keepalive period (timer) for VN2VF\_Port FIP snooping virtual link maintenance. The beacon period performs virtual link maintenance for VN2VN\_Port FIP snooping. It is the time interval between messages that verify the connection is still valid and the device at the other end of the virtual link is still reachable. You set the beacon period value for each FCoE VLAN that you configure to do VN2VN\_Port FIP snooping.



**NOTE:** In addition to enabling VN2VN\_Port FIP snooping and configuring the beacon period, you must also configure a dedicated FCoE VLAN for the VN2VN\_Port traffic, and set the FCoE transit switch ports in the proper port mode and trusted or untrusted state (interfaces are untrusted by default). See the VN2VN\_Port FIP snooping configuration example topics for complete configurations of several common network topologies.

There are differences in the way you configure a native VLAN on an interface that depend on whether the switch uses the original CLI or the Enhanced Layer 2 Software (ELS) CLI. This topic includes two configuration procedures, one for switches that run the original CLI, and one for switches that run the ELS CLI.

#### Original CLI Configuration

To enable VN2VN\_Port FIP snooping and set the beacon period on an FCoE VLAN that is dedicated to VN2VN\_Port traffic:

- [edit ethernet-switching-options secure-access-port]  
user@switch# **set vlan *vlan-name* examine-fip *examine-vn2vn* beacon-period *milliseconds***

For example, to enable VN2VN\_Port FIP snooping on a VLAN named **vlan200** and set the beacon period to **90000** milliseconds:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan vlan200 examine-fip examine-vn2vn beacon-period 90000
```

#### ELS CLI Configuration

To enable VN2VN\_Port FIP snooping and set the beacon period on an FCoE VLAN that is dedicated to VN2VN\_Port traffic:

- [edit]  
user@switch# **set vlans *vlan-name* forwarding-options fip-security examine-vn2vn beacon-period *milliseconds***

For example, to enable VN2VN\_Port FIP snooping on a VLAN named **vlan200** and set the beacon period to **90000** milliseconds:

```
[edit]
user@switch# set vlans vlan200 forwarding-options fip-security examine-vn2vn beacon-period
90000
```

#### Related Documentation

- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\)](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Directly Connected to Different FCoE Transit Switches\)](#)
- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\)](#)
- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)

## DCBX Configuration Tasks

- [Configuring the DCBX Mode on page 5668](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)

## Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. Three DCBX modes are supported:

- Autonegotiation—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- IEEE DCBX—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- DCBX Version 1.01—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric system Node devices other than QFX3500 switches come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



**NOTE:** Pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00 are not supported. If an interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

---

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]  
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |  
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface **xe-0/0/21**:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

### Related Documentation

- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Understanding DCBX on page 5580](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [show dcbx neighbors on page 5724](#)

## Configuring DCBX Autonegotiation

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default forwarding class. If ETS is not enabled on an interface (no priority groups are configured), all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to

the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



**NOTE:** If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

---

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE-FC gateway, it does not send or receive FCoE Initialization Protocol (FIP) packets.
- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.



To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# set enhanced-transmission-selection no-recommendation-tlv

#### Related Documentation

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)

## Disabling the ETS Recommendation TLV

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



**NOTE:** Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# **set enhanced-transmission-selection no-recommendation-tlv**

### Related Documentation

- [Configuring the DCBX Mode on page 5668](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Understanding DCBX on page 5580](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

## Defining an Application for DCBX Application Protocol TLV Exchange

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

#### Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [show dcbx neighbors on page 5724](#)

## Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [ aliases ] [ bit-patterns ]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[ 001 101 ]
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5724](#)

## Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

```
[edit protocols]
```

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5724](#)

## Configuration Statements

- [application \(Application Maps\) on page 5677](#)
- [application \(Applications\) on page 5678](#)
- [application-map on page 5679](#)
- [application-maps on page 5680](#)
- [applications \(Applications\) on page 5681](#)
- [applications \(DCBX\) on page 5682](#)

- [beacon-period on page 5683](#)
- [code-points \(Application Maps\) on page 5684](#)
- [dcbx on page 5685](#)
- [dcbx-version on page 5686](#)
- [destination-port \(Applications\) on page 5687](#)
- [disable \(DCBX\) on page 5688](#)
- [enhanced-transmission-selection on page 5689](#)
- [ether-type on page 5690](#)
- [examine-vn2vf on page 5691](#)
- [examine-vn2vn on page 5692](#)
- [family fcoe on page 5693](#)
- [fc-map on page 5694](#)
- [fip-security on page 5696](#)
- [fcoe-trusted on page 5697](#)
- [interface \(DCBX\) on page 5698](#)
- [interface \(FIP Snooping\) on page 5699](#)
- [no-recommendation-tlv on page 5700](#)
- [oxid on page 5701](#)
- [policy-options on page 5702](#)
- [priority-flow-control on page 5703](#)
- [protocol \(Applications\) on page 5704](#)
- [recommendation-tlv on page 5705](#)

## application (Application Maps)

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>application <i>application-name</i> {<br/>    code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];<br/>}</code>   |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Add an application to an application map and define the application's code points.   |
| <b>Options</b>                  | <i>application-name</i> —Name of the application.<br><br>The remaining statement is explained separately.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## application (Applications)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>application <i>application-name</i> {<br/>    <i>destination-port</i> <i>port-value</i>;<br/>    <i>protocol</i> (tcp   udp);<br/>    <i>ether-type</i> <i>type</i>;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit applications]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Configure properties to define an application.  |
| <b>Options</b>                  | <p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |



## application-map

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>application-map <i>application-map-name</i>;</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Specify an application map to apply to an interface.  |
| <b>Options</b>                  | <i>application-map-name</i> —Name of the application map.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## application-maps

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>application-maps <i>application-map-name</i> {<br/>    application <i>application-name</i> {<br/>        code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];<br/>    }<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit policy-options]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Define an application map by specifying the applications that belong to the application map.   |
| <b>Options</b>                  | <p><i>application-map-name</i>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## applications (Applications)


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> applications {   application application-name {     destination-port port-value;     protocol (tcp   udp);     ether-type type;   } } </pre>  |
| <b>Hierarchy Level</b>          | [edit]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>   |
| <b>Description</b>              | Define applications that DCBX advertises.   |
| <b>Options</b>                  | The statements are explained separately.  |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## applications (DCBX)

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre>applications {<br/>    no-auto-negotiation;<br/>}</pre>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for the EX Series  |
| <b>Description</b>              | Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.  |
| <b>Options</b>                  | The remaining statements are explained separately.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li></ul> |

## beacon-period

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>beacon-period <i>milliseconds</i>;</code>  |
| <b>Hierarchy Level</b>          | Original CLI<br><br>[edit ethernet-switching options secure-access-port vlan (all   <i>vlan-name</i> ) examine-fip <a href="#">examine-vn2vn</a> ]<br><br>ELS CLI for Platforms that Support FCoE<br><br>[edit <a href="#">vlans</a> <i>vlan-name</i> <a href="#">forwarding-options</a> <a href="#">fip-security</a> ]  |
|                                 | <div>  <p><b>NOTE:</b> The <code>beacon-period</code> configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.</p> </div>   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.<br>Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.   |
| <b>Description</b>              | <p>Set the interval between periodic beacons. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements.</p> <p>The ENode sends periodic beacons every 90 seconds on behalf of the VN_Port. Each received beacon resets the session timer for the virtual link connection to the other VN_Port. If the FCF does not receive a beacon before the beacon timer expires, the VN_Port is considered as “down” and the virtual link is terminated. The beacon timer expires in 2.5 times the configured beacon timer value.</p>  |
| <b>Options</b>                  | <p><b><i>milliseconds</i></b>—Time in milliseconds between beacons.</p> <p><b>Range:</b> 250 through 90000 milliseconds</p> <p><b>Default:</b> 8000 milliseconds</p>   |
| <b>Required Privilege Level</b> | <p>storage—To view this statement in the configuration.</p> <p>storage-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <i>Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)</i></li> <li>• <i>Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)</i></li> <li>• <i>Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)</i></li> <li>• <a href="#">Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch) on page 5636</a></li> <li>• <a href="#">Example: Configuring VN2VN_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches) on page 5641</a></li> </ul> |

- [Example: Configuring VN2VN\\_Port FIP Snooping \(FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch\) on page 5647](#)

## code-points (Application Maps)

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>   |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> <b>application</b> <i>application-name</i> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Define one or more code-point aliases or bit sets for an application.  |
| <b>Options</b>                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## dcbx

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre> dcbx {   disable;   interface (interface-name   all) {     disable;     application-map application-map-name;     applications {       no-auto-negotiation;     }     enhanced-transmission-selection {       no-auto-negotiation;       no-recommendation-tlv;       recommendation-tlv {         no-auto-negotiation;       }     }     dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);     priority-flow-control {       no-auto-negotiation;     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols</a> ]  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p><b>mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>  |
| <b>Description</b>              | Configure DCBX properties.   |
| <b>Options</b>                  | The statements are explained separately.   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li> <li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li> </ul>                              |


## dcbx-version

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <code>dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);</code>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.   |
| <b>Description</b>              | <p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p> |
| <b>Default</b>                  | The default DCBX mode is autonegotiation.   |
| <b>Options</b>                  | <p><b>auto-negotiate</b>—Automatically negotiate the DCBX version with the connected peer.</p> <p><b>ieee-dcbx</b>—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p><b>dcbx-version-1.01</b>—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li><li>• <a href="#">Understanding DCBX on page 5580</a></li></ul>  |



## destination-port (Applications)

|   |   |
|---|---|
| <b>Syntax</b>   | <code>destination-port <i>port-value</i>;</code>  |
| <b>Hierarchy Level</b>  | [edit applications <b>application</b> <i>application-name</i> ]   |
| <b>Release Information</b>  | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>  | Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with <b>protocol</b> to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at <a href="http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml">http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml</a> for a list of assigned port numbers. |
| <div>  <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>. </div> |   |
| <b>Options</b>  | <i>port-value</i> —Identifier for the port.   |
| <b>Required Privilege Level</b>   | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>  | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>               |

## disable (DCBX)

---


|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | disable  |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx</a> ]<br><br>[edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.   |
| <b>Description</b>              | Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.  |
| <b>Default</b>                  | DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.<br><br>DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li><li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li><li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li></ul> |

## enhanced-transmission-selection


|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <pre> enhanced-transmission-selection {   no-auto-negotiation;   no-recommendation-tlv;   recommendation-tlv {     no-auto-negotiation;   } } </pre>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | <p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p> |
| <b>Options</b>                  | <p><b>no-auto-negotiation</b>—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p><b>no-recommendation-tlv</b>—Disable automatic negotiation of the ETS Recommendation TLV</p> <p><b>recommendation-tlv</b>—Enable automatic negotiation of ETS Recommendation TLV</p>   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> </ul>   |

## ether-type

---

|  |  |
|--|--|
| <b>Syntax</b>  | <code>ether-type <i>ether-type</i>;</code>   |
| <b>Hierarchy Level</b>   | [edit applications <b>application</b> <i>application-name</i> ]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>   | Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See <a href="http://standards.ieee.org/develop/regauth/ethertype/eth.txt">http://standards.ieee.org/develop/regauth/ethertype/eth.txt</a> for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes.                   |
| <div> <b>NOTE:</b> To create a FIP application, use the EtherType 0x8914.</div> |  |
| <b>Options</b>   | <b>type</b> —Identifier for the EtherType.   |
| <b>Required Privilege Level</b>  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li></ul> |

## examine-vn2vf

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | examine-vn2vf   |
| <b>Hierarchy Level</b>          | [edit <a href="#">vlans</a> <i>vlan-name</i> <a href="#">forwarding-options</a> <a href="#">fip-security</a> ]  |
| <b>Release Information</b>      | Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.   |
| <b>Description</b>              | <p> <b>NOTE:</b> This statement supports the Enhanced Layer 2 Software (ELS) CLI. If your switch runs the original (non-ELS) software, see <i>examine-fip</i>. For ELS details, see “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 43.</p> <p>Enable VN_Port to VF_Port (VN2VF_Port) FIP snooping on the specified VLAN. Ensure that the VLAN is a dedicated FCoE VLAN that transports only FCoE traffic.</p> <p>If the switch also performs VN_Port to VN_Port (VN2VN_Port) FIP snooping, ensure that the VN2VN_Port traffic is on a different VLAN than the VN2VF_Port traffic. You cannot mix VN2VF_Port and VN2VN_Port traffic in the same VLAN, so you must use separate VLANs for VN2VF_Port and VN2VN_Port traffic.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">examine-vn2vn on page 5692</a></li> <li>• <a href="#">Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch on page 5531</a></li> <li>• <a href="#">Understanding FCoE Transit Switch Functionality on page 5524</a></li> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> </ul>  |

## examine-vn2vn

**Syntax** `examine-vn2vn {  
    beacon-period milliseconds;  
}`

**Hierarchy Level** Original CLI

[edit ethernet-switching options secure-access-port vlan (all | *vlan-name*) examine-fip]

ELS CLI for Platforms that Support FCoE

[edit *vlangs* *vlan-name* forwarding-options fip-security]



**NOTE:** The `examine-vn2vn` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

**Release Information** Statement introduced in Junos OS Release 12.2 for the QFX Series.

Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

**Description** Enable VN\_Port to VN\_Port (VN2VN) FIP snooping on a specified VLAN. The VLAN must be a dedicated FCoE VLAN that transports only FCoE traffic. A VLAN cannot support VN2VN FIP snooping and VN\_Port to VF\_Port FIP snooping (VN2VF) simultaneously. Configure separate VLANs for VN2VN FIP snooping and VN2VF FIP snooping.

When you enable VN2VN FIP snooping on a VLAN, the VN2VF session filters are removed and the all existing VN2VF sessions are terminated.

The remaining statement is explained separately.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

- Related Documentation**
- *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch)*
  - *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches)*
  - *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch)*
  - *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to the Same FCoE Transit Switch) on page 5636*
  - *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Directly Connected to Different FCoE Transit Switches) on page 5641*
  - *Example: Configuring VN2VN\_Port FIP Snooping (FCoE Hosts Indirectly Connected Through an Aggregation Layer FCoE Transit Switch) on page 5647*

## family fcoe

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <p>QFX Series Standalone Switches</p> <pre>family fcoe {   oxid (enable   disable); }</pre> <p>QFabric Systems</p> <pre>family fcoe {   ethernet-interfaces {     node-group (node-group-name   all) {       oxid (enable   disable);     }   }   fabric-interfaces {     node-group (node-group-name   all) {       oxid (enable   disable);     }   } }</pre>  |
| <b>Hierarchy Level</b>          | [edit forwarding-options hash-key]   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p> <p>Ethernet-interfaces and fabric-interfaces statements introduced in Junos OS Release 13.2X52-D10 for the QFX Series.</p>  |
| <b>Description</b>              | Configure whether or not to use the originator exchange identifier (Oxid) field for hash control for FCoE traffic load balancing.  |
| <b>Options</b>                  | The statement is explained separately.   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 5656</a></li> <li>• <a href="#">Enabling and Disabling CoS OxID Hash Control on QFabric Systems</a></li> <li>• <a href="#">Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 5530</a></li> <li>• <a href="#">Understanding OxID Hash Control for FCoE Traffic Load Balancing on QFabric Systems</a></li> </ul> |

## fc-map

---

**Syntax** `fc-map fc-map-value;`

**Hierarchy Level** Original CLI

[edit ethernet-switching options secure-access-port vlan (all | *vlan-name*) examine-fip]

ELS CLI for Platforms that Support FCoE

[edit **v**lans *vlan-name* **f**orwarding-**o**ptions **f**ip-**s**ecurity]



**NOTE:** The `fc-map` configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.

QFX Series that Support FCoE-FC Gateway Configuration

[edit fc-fabrics *fc-fabric-name* protocols fip]

**Release Information** Statement introduced in Junos OS Release 10.4 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

**Description** Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).

You can configure the FC-MAP value or use the default value. The default FC-MAP value is different for VN\_Port to VF\_Port (VN2VF\_Port) FIP snooping (0x0EFC00) than for VN\_Port to VN\_Port (VN2VN\_Port) FIP snooping.

The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN\_Port MAC address.

When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.



**NOTE:** Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.

**Options** `fc-map-value`—FC-MAP value, hexadecimal value preceded by “0x”.



**Range:** 0x0EFC00 through 0x0EFCFF

**Default:** 0x0EFC00 for VN2VF\_Port FIP snooping 0x0EFD00 for VN2VN\_Port FIP snooping

**Required Privilege** routing—To view this statement in the configuration.  
**Level** routing-control—To add this statement to the configuration.

**Related Documentation**

- *examine-fip*
- [show fip snooping on page 5746](#)
- *Example: Configuring an FCoE Transit Switch*
- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)

## fip-security

---

**Syntax**    `fip-security {  
              examine-vn2vf;  
              examine-vn2vn {  
                  beacon-period milliseconds;  
              }  
              fc-map fc-map-value;  
              interface interface-name {  
                  (fcoe-trusted | no-fcoe-trusted);  
              }  
          }`

**Hierarchy Level**    [edit **vlan**s *vlan-name* **forwarding-options**]

**Release Information**    Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.

**Description**



**NOTE:** This statement supports the Enhanced Layer 2 Software (ELS) CLI. If your switch runs the original (non-ELS) software, see *examine-fip*. For ELS details, see [“Getting Started with Enhanced Layer 2 Software” on page 43](#).


Configure FIP snooping and FCoE interface properties.

**Required Privilege Level**    routing—To view this statement in the configuration.  
                                  routing-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding VN\\_Port to VF\\_Port FIP Snooping on an FCoE Transit Switch on page 5531](#)
- [Understanding VN\\_Port to VN\\_Port FIP Snooping on an FCoE Transit Switch on page 5539](#)
- [Understanding FCoE Transit Switch Functionality on page 5524](#)
- [Configuring VN2VF\\_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662](#)
- [Enabling VN2VN\\_Port FIP Snooping and Configuring the Beacon Period on an FCoE Transit Switch on page 5665](#)

## fcoe-trusted


|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | fcoe-trusted;  |
| <b>Hierarchy Level</b>          | Original CLI<br><br>[edit ethernet-switching-options secure-access-port interface <i>interface-name</i> ]<br><br>ELS CLI for Platforms that Support FCoE<br><br>[edit <b>vlangs</b> <i>vlan-name</i> <b>forwarding-options fip-security interface</b> <i>interface-name</i> ]  |
|                                 | <div>  <p><b>NOTE:</b> The <b>fcoe-trusted</b> configuration statement is in a different hierarchy on the original CLI than on the Enhanced Layer 2 Software (ELS) CLI.</p> </div>  |
|                                 | <p>QFX Series that Support FCoE-FC Gateway Configuration</p> <p>[edit fc-fabrics <i>fc-fabric-name</i> protocols fip]</p>  |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced for the FC fabric in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.</p>  |
| <b>Description</b>              | <p>Configure the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p> <p>(QFX Series FCoE-FC gateway) Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the <b>fcoe-trusted</b> configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.</p> |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> </ul>  |

## interface (DCBX)

---

|                          |  |
|--------------------------|--|
| Syntax                   | <pre>interface (<i>interface-name</i>   all) {<br/>  disable;<br/>  application-map <i>application-map-name</i>;<br/>  applications {<br/>    no-auto-negotiation;<br/>  }<br/>  enhanced-transmission-selection {<br/>    no-auto-negotiation;<br/>    no-recommendation-tlv;<br/>    recommendation-tlv {<br/>      no-auto-negotiation;<br/>    }<br/>  }<br/>  dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);<br/>  priority-flow-control {<br/>    no-auto-negotiation;<br/>  }<br/>}</pre>                                   |
| Hierarchy Level          | [edit <a href="#">protocols dcbx</a> ]   |
| Release Information      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p><b>Mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>  |
| Description              | Configure DCBX properties on an interface.   |
| Options                  | <p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>   |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>   |
| Related Documentation    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li><li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## interface (FIP Snooping)

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | interface <i>interface-name</i> {<br>(fcoe-trusted  no-fcoe-trusted);<br>}  |
| <b>Hierarchy Level</b>          | [edit <i>vlan</i> <i>vlan-name</i> forwarding-options fip-security]   |
| <b>Release Information</b>      | Statement introduced for the ELS CLI in Junos OS Release 13.2 for the QFX Series.   |
| <b>Description</b>              | <div>  <p><b>NOTE:</b> This statement supports the Enhanced Layer 2 Software (ELS) CLI. If your switch runs the original (non-ELS) software, see <i>interface (Secure Access Port)</i> for how to specify an interface to configure as FCoE trusted or FCoE untrusted. For ELS details, see “<a href="#">Getting Started with Enhanced Layer 2 Software</a>” on page 43.</p> </div> <p>Specify an interface to set as FCoE trusted or as FCoE untrusted. Configure interfaces that connect to other switches as trusted interfaces. Configure interfaces that connect directly to FCoE devices as untrusted interfaces and enabled FIP snooping on the untrusted interfaces to prevent unauthorized access to the storage network.</p> |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Understanding FCoE Transit Switch Functionality on page 5524</a></li> </ul>  |

## no-recommendation-tlv

---

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | no-recommendation-tlv;  |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.   |
| <b>Description</b>              | Disable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li></ul>  |

## oxid

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | oxid (enable   disable)  |
| <b>Hierarchy Level</b>          | <p>QFX Series Standalone Switches</p> <p>[edit forwarding-options hash-key family fcoe]</p> <p>QFabric Systems</p> <p>[edit forwarding-options hash-key family fcoe ethernet-interfaces node-group (node-group-name   all) {}]</p> <p>[edit forwarding-options hash-key family fcoe fabric-interfaces node-group (node-group-name   all) {}]</p>   |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X52-D10 for the QFabric System.</p>   |
| <b>Description</b>              | Enable or disable whether the switch uses the originator exchange identifier (OxID) field for hash control for FCoE traffic load balancing.  |
| <b>Default</b>                  | OxID hash control is enabled by default.   |
| <b>Options</b>                  | <b>oxid (enable   disable)</b> —Enable or disable whether the switch uses the OxID hash control field for FCoE traffic load balancing.   |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Enabling and Disabling CoS OxID Hash Control on Standalone Switches on page 5656</a></li> <li>• <a href="#">Enabling and Disabling CoS OxID Hash Control on QFabric Systems</a></li> <li>• <a href="#">Understanding OxID Hash Control for FCoE Traffic Load Balancing on Standalone Switches on page 5530</a></li> <li>• <a href="#">Understanding OxID Hash Control for FCoE Traffic Load Balancing on QFabric Systems</a></li> </ul> |

## policy-options

---

**Syntax**    `policy-options`

```
  application-maps application-map-name {
    application application-name {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name match-type <actions>;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
    }
  }
```

**Hierarchy Level**    [edit]

**Release Information**    Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.1 for the EX Series.

**Description**    Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level**    storage—To view this statement in the configuration.  
storage-control—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)




## priority-flow-control

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>priority-flow-control {<br/>    no-auto-negotiation;<br/>}</code>  |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.   |
| <b>Description</b>              | Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.   |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable automatic negotiation of PFC.  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5606</a></li> <li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li> <li>• <a href="#">Understanding Priority-Based Flow Control</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> </ul> |

## protocol (Applications)

---

|  |   |
|--|---|
| <b>Syntax</b>  | <code>protocol (tcp   udp);</code>  |
| <b>Hierarchy Level</b>   | [edit applications <a href="#">application</a> <i>application-name</i> ]  |
| <b>Release Information</b>   | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>   | Networking protocol type, which combines with <b>destination-port</b> to identify an application type.  |
| <div> <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number 3260.</div> |   |
| <b>Options</b>   | <code>tcp</code> —Transmission Control Protocol<br><br><code>udp</code> —User Datagram Protocol   |
| <b>Required Privilege Level</b>  | <code>interface</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.   |
| <b>Related Documentation</b>   | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

---

## recommendation-tlv

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <pre>recommendation-tlv {<br/>    no-auto-negotiation;<br/>}</pre>   |
| <b>Hierarchy Level</b>          | [edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.  |
| <b>Description</b>              | Enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.   |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable sending of the ETS recommendation TLV.   |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li></ul>   |



## CHAPTER 70

# Administration

- [Operational Commands on page 5707](#)

### Operational Commands

---

- [clear fip snooping enode](#)
- [clear fip snooping statistics](#)
- [clear fip snooping vlan](#)
- [clear fip vlan-discovery statistics](#)
- [restart](#)
- [show dcbx](#)
- [show dcbx neighbors](#)
- [show fip snooping](#)
- [show fip snooping enode](#)
- [show fip snooping fcf](#)
- [show fip snooping interface](#)
- [show fip snooping statistics](#)
- [show fip snooping vlan](#)
- [show fip vlan-discovery](#)

## clear fip snooping enode

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>clear fip snooping enode <i>enode-mac</i></b><br><b>&lt;vlan <i>vlan-name</i>&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Clear FIP snooping information for the specified FCoE Node (ENode) or (optionally) only on a specified VLAN. This operation deletes the ENode state from the switch database and from the FIP snooping firewall filters, which causes the ENode to lose its connection to the FCoE forwarder (FCF) and to log in to the FCF again. |
| <b>Options</b>                  | <b><i>enode-mac</i></b> —MAC address of the ENode.<br><br><b>vlan <i>vlan-name</i></b> —(Optional) Name of the VLAN.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show fip snooping enode on page 5751</a></li></ul>   |
| <b>List of Sample Output</b>    | <a href="#">clear fip snooping enode enode-mac on page 5708</a>  |

### Sample Output

#### clear fip snooping enode enode-mac

```
user@switch> clear fip snooping enode 00:10:94:00:00:02
```

---

## clear fip snooping statistics

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>clear fip snooping statistics</code><br><code>&lt;vlan <i>vlan-name</i>&gt;</code>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series. |
| <b>Description</b>              | Clear FIP snooping statistics globally or on a specified VLAN.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show fip snooping statistics on page 5761</a></li></ul>                            |
| <b>List of Sample Output</b>    | <a href="#">clear fip snooping statistics on page 5709</a>   |

### Sample Output

#### clear fip snooping statistics

```
user@switch> clear fip snooping statistics
```

## clear fip snooping vlan

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <code>clear fip snooping vlan <i>vlan-name</i></code>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Clear FIP snooping information for the specified VLAN. This operation deletes all ENode and FCF information for the VLAN from the switch database and causes the ENodes to lose their connections to the FCFs. After clearing a VLAN, the switch relearns all of the FCFs and ENodes on the VLAN, and the ENodes must log in to the FCF again. |
| <b>Options</b>                  | <i>vlan-name</i> —Name of the VLAN.  |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show fip snooping vlan on page 5764</a></li></ul>  |
| <b>List of Sample Output</b>    | <a href="#">clear fip snooping vlan vlan-name on page 5710</a>   |

### Sample Output

#### clear fip snooping vlan vlan-name

```
user@switch> clear fip snooping vlan fcoevlan1
```



## clear fip vlan-discovery statistics

---

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | clear fip vlan-discovery statistics  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Clear FIP VLAN discovery statistics.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show fip vlan-discovery on page 5768</a></li></ul> |
| <b>List of Sample Output</b>    | <a href="#">clear fip vlan-discovery statistics on page 5711</a>                                       |

### Sample Output

#### clear fip vlan-discovery statistics

```
user@switch> clear fip vlan-discovery statistics
```

## restart

### List of Syntax [Syntax on page 5712](#)

[Syntax \(ACX Series Routers\) on page 5712](#)  
[Syntax \(EX Series Switches\) on page 5712](#)  
[Syntax \(Routing Matrix\) on page 5713](#)  
[Syntax \(J Series Routing Platform\) on page 5713](#)  
[Syntax \(TX Matrix Routers\) on page 5713](#)  
[Syntax \(TX Matrix Plus Routers\) on page 5713](#)  
[Syntax \(MX Series Routers\) on page 5713](#)  
[Syntax \(J Series Routers\) on page 5714](#)  
[Syntax \(QFX Series\) on page 5714](#)

### Syntax restart

```

<adaptive-services | ancpd-service | application-identification | audit-process |
  auto-configuration | captive-portal-content-delivery | ce-l2tp-service | chassis-control |
  class-of-service | clksyncd-service | database-replication | datapath-trace-service
  | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | iccp-service | idp-policy | immediately
  | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service | link-management
  | local-policy-decision-function | mac-validation | mib-process | mobile-ip | mountd-service
  | mpls-traceroute | mspd | multicast-snooping | named-service | nfsd-service |
  packet-triggered-subscribers | peer-selection-service | pgcp-service | pgm |
  pic-services-logging | pki-service | ppp | ppp-service | pppoe |
  protected-system-domain-service | redundancy-interface-process | remote-operations |
  root-system-domain-service | routing <logical-system logical-system-name> | sampling
  | sbc-configuration-process | sdk-service | service-deployment | services | services pgcp
  gateway gateway-name | snmp | soft | static-subscribers | statistics-service |
  subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control |
  vrrp | web-management>
<gracefully | immediately | soft>

```

### Syntax (ACX Series Routers)

```

restart
<adaptive-services | audit-process | auto-configuration | autoinstallation | chassis-control |
  class-of-service | clksyncd-service | database-replication | dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  | ethernet-link-fault-management | event-processing | firewall
  | general-authentication-service | gracefully | immediately | interface-control |
  ipsec-key-management | l2-learning | lacp | link-management | mib-process | mobile-ip |
  mountd-service | mpls-traceroute | mspd | named-service | nfsd-service | pgm | pki-service
  | ppp | pppoe | redundancy-interface-process | remote-operations | routing | sampling |
  sdk-service | secure-neighbor-discovery | service-deployment | services | snmp | soft
  | statistics-service | subscriber-management | subscriber-management-helper | tunnel-oamd
  | vrrp>

```

### Syntax (EX Series Switches)

```

restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management

```

|   |  |
|---|--|
|   | lldpd-service   mib-process   mountd-service   multicast-snooping   pgm  <br>redundancy-interface-process   remote-operations   routing   secure-neighbor-discovery<br>  service-deployment   sflow-service   snmp   vrrp   web-management>  |
| <b>Syntax (Routing Matrix)</b>            | restart<br><adaptive-services   audit-process   chassis-control   class-of-service   disk-monitoring  <br>dynamic-flow-capture   ecc-error-logging   event-processing   firewall   interface-control<br>  ipsec-key-management   kernel-replication   l2-learning   l2tp-service   lacp  <br>link-management   mib-process   pgm   pic-services-logging   ppp   pppoe  <br>redundancy-interface-process   remote-operations   routing <logical-system<br><i>logical-system-name</i> >   sampling   service-deployment   snmp><br><all   all-lcc   lcc <i>number</i> ><br><gracefully   immediately   soft>   |
| <b>Syntax (J Series Routing Platform)</b> | restart<br><adaptive-services   audit-process   chassis-control   class-of-service   dhcp   dialer-services<br>  dlsr   event-processing   firewall   interface-control   ipsec-key-management  <br>isdn-signaling   l2-learning   l2tp-service   mib-process   network-access-service   pgm  <br>ppp   pppoe   remote-operations   routing <logical-system <i>logical-system-name</i> >   sampling<br>  service-deployment   snmp   usb-control   web-management><br><gracefully   immediately   soft>  |
| <b>Syntax (TX Matrix Routers)</b>         | restart<br><adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service  <br>diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging  <br>event-processing   firewall   interface-control   ipsec-key-management   kernel-replication<br>  l2-learning   l2tp-service   lacp   link-management   mib-process   pgm   pic-services-logging<br>  ppp   pppoe   redundancy-interface-process   remote-operations   routing <logical-system<br><i>logical-system-name</i> >   sampling   service-deployment   snmp   statistics-service><br><all-chassis   all-lcc   lcc <i>number</i>   scc><br><gracefully   immediately   soft>   |
| <b>Syntax (TX Matrix Plus Routers)</b>    | restart<br><adaptive-services   audit-process   chassis-control   class-of-service   dhcp-service  <br>diameter-service   disk-monitoring   dynamic-flow-capture   ecc-error-logging  <br>event-processing   firewall   interface-control   ipsec-key-management   kernel-replication<br>  l2-learning   l2tp-service   lacp   link-management   mib-process   pgm  <br>pic-services-logging   ppp   pppoe   redundancy-interface-process   remote-operations  <br>routing <logical-system <i>logical-system-name</i> >   sampling   service-deployment   snmp  <br>statistics-service><br><all-chassis   all-lcc   all-sfc   lcc <i>number</i>   sfc <i>number</i> ><br><gracefully   immediately   soft>   |
| <b>Syntax (MX Series Routers)</b>         | restart<br><adaptive-services   ancpd-service   application-identification   audit-process  <br>auto-configuration   captive-portal-content-delivery   ce-l2tp-service   chassis-control  <br>class-of-service   clksyncd-service   database-replication   datapath-trace-service<br>  dhcp-service   diameter-service   disk-monitoring   dynamic-flow-capture  <br>ecc-error-logging   ethernet-connectivity-fault-management<br>  ethernet-link-fault-management   event-processing   firewall  <br>general-authentication-service   gracefully   iccp-service   idp-policy   immediately<br>  interface-control   ipsec-key-management   kernel-replication   l2-learning   l2cpd-service<br>  l2tp-service   l2tp-universal-edge   lacp   license-service   link-management<br>  local-policy-decision-function   mac-validation   mib-process   mobile-ip   mountd-service<br>  mpls-traceroute   mspd   multicast-snooping   named-service   nfsd-service |

```

packet-triggered-subscribers |peer-selection-service | pgcp-service | pgm |
pic-services-logging | pki-service | ppp | ppp-service | pppoe |
protected-system-domain-service | redundancy-interface-process | remote-operations
|root-system-domain-service | routing |routing <logical-system logical-system-name> |
sampling | sbc-configuration-process | sdk-service |service-deployment |services | services
pgcp gateway gateway-name |snmp |soft |static-subscribers |statistics-service|
subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>

```

#### Syntax (J Series Routers)

```

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp | dhcp-service
| dialer-services | diameter-service | dlsr | event-processing | firewall | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
network-access-service | pgm | ppp | pppoe | remote-operations | routing <logical-system
logical-system-name> | sampling | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

```

#### Syntax (QFX Series)

```

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
diameter-service | dlsr | ethernet-connectivity | event-processing | fibre-channel | firewall
| general-authentication-service | igmp-host-services | interface-control |
ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
redundancy-interface-process | remote-operations |logical-system-name> | routing |
sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
web-management>
<gracefully | immediately | soft>

```

#### Release Information

Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 12.2 for ACX Series routers.  
 Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsr** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

**Description** Restart a Junos OS process.



**CAUTION:** Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

**Options** **none**—Same as **gracefully**.

**adaptive-services**—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

**all-chassis**—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

**all-lcc**—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

**all-members**—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

**all-sfc**—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

**ancpd-service**—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

**application-identification**—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

**audit-process**—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

**auto-configuration**—(Optional) Restart the Interface Auto-Configuration process.

**autoinstallation**—(EX Series switches only) (Optional) Restart the autoinstallation process.

**captive-portal-content-delivery**—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

**ce-l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

**chassis-control**—(Optional) Restart the chassis management process.

**class-of-service**—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

**clksyncd-service**—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

**database-replication**—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

**dapath-trace-service**—(Optional) Restart the packet path tracing process.

**dhcp**—(J Series routers and EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

**dhcp-service**—(Optional) Restart the Dynamic Host Configuration Protocol process.

**dialer-services**—(J Series routers and EX Series switches only) (Optional) Restart the ISDN dial-out process.

**diameter-service**—(Optional) Restart the diameter process.

**disk-monitoring**—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

**dlsw**—(J Series routers and QFX Series only) (Optional) Restart the data link switching (DLSw) service.

**dot1x-protocol**—(EX Series switches only) (Optional) Restart the port-based network access control process.

**dynamic-flow-capture**—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

**ecc-error-logging**—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

**ethernet-connectivity-fault-management**—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

**ethernet-link-fault-management**—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

**ethernet-switching**—(EX Series switches only) (Optional) Restart the Ethernet switching process.

**event-processing**—(Optional) Restart the event process (eventd).

**fibre-channel**—(QFX Series only) (Optional) Restart the Fibre Channel process.

**firewall**—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

**general-authentication-service**—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

**gracefully**—(Optional) Restart the software process.

**iccp-service**—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

**idp-policy**—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

**immediately**—(Optional) Immediately restart the software process.

**interface-control**—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

**ipsec-key-management**—(Optional) Restart the IPsec key management process.

**isdn-signaling**—(J Series routers and QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

**kernel-replication**—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

**l2-learning**—(Optional) Restart the Layer 2 address flooding and learning process.

**l2cpd-service**—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

**l2tp-service**—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

**l2tp-universal-edge**—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

**lACP**—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG,

and to enable the transmission and reception processes for the link to function in an orderly manner.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**license-service**—(EX Series switches only) (Optional) Restart the feature license management process.

**link-management**—(TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

**lldpd-service**—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

**local**—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

**local-policy-decision-function**—(Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

**mac-validation**—(Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

**member *member-id***—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**mib-process**—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

**mobile-ip**—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.



**mountd-service**—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

**mpls-traceroute**—(Optional) Restart the MPLS Periodic Traceroute process.

**mspd**—(Optional) Restart the Multiservice process.

**multicast-snooping**—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

**named-service**—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

**network-access-service**—(J Series routers and QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

**nfsd-service**—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

**packet-triggered-subscribers**—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

**peer-selection-service**—(Optional) Restart the Peer Selection Service process.

**pgcp-service**—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

**pgm**—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

**pic-services-logging**—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

**pki-service**—(Optional) Restart the PKI Service process.

**ppp**—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

**ppp-service**—(Optional) Restart the Universal Edge PPP process, which is the encapsulation protocol process for transporting IP traffic across Universal Edge routers.

**pppoe**—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

**protected-system-domain-service**—(Optional) Restart the Protected System Domain (PSD) process.

**redundancy-interface-process**—(Optional) Restart the ASP redundancy process.

**remote-operations**—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

**root-system-domain-service**—(Optional) Restart the Root System Domain (RSD) service.

**routing**—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

**routing <logical-system *logical-system-name*>**—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

**sampling**—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

**sbc-configuration-process**—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

**scc**—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

**sdk-service**—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

**secure-neighbor-discovery**—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

**sfc number**—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace **number** with **0**.

**service-deployment**—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

**services**—(Optional) Restart a service.

**services pgcp gateway *gateway-name***—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

**sflow-service**—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

**snmp**—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

**soft**—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

**static-subscribers**—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

**statistics-service**—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

**subscriber-management**—(Optional) Restart the Subscriber Management process.

**subscriber-management-helper**—(Optional) Restart the Subscriber Management Helper process.

**tunnel-oamd**—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 PDUs across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

**usb-control**—(J Series routers and MX Series routers only) (Optional) Restart the USB control process.

**vrrp**—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

**web-management**—(J Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

**Required Privilege Level** reset

**Related Documentation** [• Overview of Junos OS CLI Operational Mode Commands on page 58](#)

**List of Sample Output** [restart interfaces on page 5721](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```



## show dcbx

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | show dcbx   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.3 for the QFX Series.   |
| <b>Description</b>              | List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.   |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> </ul> |
| <b>Output Fields</b>            | <a href="#">Table 463 on page 5723</a> lists the output fields for the <b>show dcbx</b> command. Output fields are listed in the approximate order in which they appear.      |

**Table 463: show dcbx output fields**

| Field Name | Field Description   |
|------------|---|
| DCBX       | Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"> <li>• Enabled—DCBX is enabled on the switch or on the specified interface</li> <li>• Disabled—DCBX is disabled on the switch or on the specified interface</li> </ul> |
| Interface  | Name of the interface   |

## Sample Output

### show dcbx

```

user@switch> show dcbx
DCBX                : Enabled
Interface           DCBX
xe-0/0/9.0          enabled
xe-0/0/32.0         enabled
xe-0/0/36.0         enabled

```

## show dcbx neighbors

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show dcbx neighbors</b><br><b>&lt;interface interface-name&gt;</b><br><b>&lt;terse&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 11.1 for the QFX Series.<br>Command introduced in Junos OS Release 11.3 for EX Series switches.   |
| <b>Description</b>              | Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.  |
| <b>Options</b>                  | <b>none</b> —Display information about all DCBX neighbor interfaces.<br><br><b>interface-name</b> —(Optional) Display information for the specified interface.<br><br><b>terse</b> —Display the specified level of output.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> <li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li> <li>• <a href="#">dcbx on page 5685</a></li> </ul>                   |
| <b>List of Sample Output</b>    | <a href="#">show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 5737</a><br><a href="#">show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 5739</a><br><a href="#">show dcbx neighbors terse (QFX Series) on page 5741</a><br><a href="#">show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 5741</a><br><a href="#">show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 5742</a><br><a href="#">show dcbx neighbors (EX4500 Switch: Includes ETS) on page 5743</a> |
| <b>Output Fields</b>            | <a href="#">Table 464 on page 5724</a> lists the output fields for the <b>show dcbx neighbors</b> command. Output fields are listed in the approximate order in which they appear.   |

**Table 464: show dcbx neighbors Output Fields**

| Field Name | Field Description      |
|------------|------------------------|
| Interface  | Name of the interface. |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name             | Field Description  |
|------------------------|--|
| Parent Interface       | Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.  |
| Active-application-map | Name of the application map applied to the interface.  |
| Protocol-Mode          | <p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> <li>IEEE DCBX Version—The interface uses IEEE DCBX mode.</li> <li>DCBX Version 1.01—The interface uses DCBX version 1.01.</li> </ul> <p><b>NOTE:</b> On interfaces that use the IEEE DCBX mode, the <b>show dcbx neighbors interface <i>interface-name</i></b> operational command does not include application, PFC, or ETS operational state in the output.</p>  |
| Protocol-State         | <p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li><b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> <li><b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> </ul> |
| Local-Advertisement    | <p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>  |
| Operational version    | Version of the DCBX standard used.   |
| sequence-number        | <p>Number of state change messages sent to the peer.</p> <p>If the interface <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p> <p>If the interface <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p>   |
| acknowledge-id         | <p>Number of acknowledge messages received from the peer.</p> <p>If the <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p> <p>If the <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p>  |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description   |
|----------------------------|---|
| <b>Peer-Advertisement</b>  | (DCBX Version 1.01 only)<br><br>Status of advertisements that the peer sends to the local interface.  |
| <b>Operational version</b> | Version of the DCBX standard used.  |
| <b>sequence-number</b>     | <p>Number of state change messages the peer sent to the local interface.</p> <p>If this number matches the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized.</p> <p>If this number does not match the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.</p>            |
| <b>acknowledge-id</b>      | <p>Number of acknowledge messages the peer has received from the local interface.</p> <p>If this number matches the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization.</p> <p>If this number does not match the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.</p> |



Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                        | Field Description  |
|-----------------------------------|--|
| <b>Feature: PFC</b>               | Priority-based flow control (PFC) feature DCBX state information.  |
| <b>Protocol-State</b>             | (DCBX Version 1.01 only)<br><br>DCBX protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—PFC autonegotiation is disabled.</li> </ul> |
| <b>Operational State</b>          | (DCBX Version 1.01 only)<br><br>Operational state of the feature: <b>enabled</b> or <b>disabled</b> .  |
| <b>Local-Advertisement</b>        | Status of advertisements that the local interface sends to the peer.   |
| <b>Enable</b>                     | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>  |
| <b>Willing</b>                    | Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the PFC configuration from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the PFC configuration from the peer.</li> </ul>  |
| <b>Mac auth Bypass Capability</b> | (IEEE DCBX only)<br><br>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is <b>no</b> .   |
| <b>Error</b>                      | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>   |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>  |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>  |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>  |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>   |
| <b>Operational Mode</b>                               | <p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> <li>• <b>Enable</b>—PFC is enabled on the code point.</li> <li>• <b>Disable</b>—PFC is disabled on the code point.</li> </ul>   |
| <b>Peer-Advertisement</b>                             | <p>Status of advertisements that the peer sends to the local interface.</p>  |
| <b>Enable</b>   | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>   |
| <b>Willing</b>  | <p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the PFC configuration from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the PFC configuration from the local interface.</li> </ul> |
| <b>Error</b>  | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>  |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name  | Field Description  |
|---|--|
| <b>Operational State</b>                              | <p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>  |
| <b>Mac auth Bypass Capability</b>                     | <p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The connected peer supports MAC authentication bypass.</li> <li>• <b>No</b>—The connected peer does not support MAC authentication bypass.</li> </ul> |
| <b>Maximum Traffic Classes capable to support PFC</b> | <p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>   |
| <b>Code Point</b>                                     | <p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>  |
| <b>Admin Mode</b>                                     | <p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>  |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                  | Field Description  |
|-----------------------------|--|
| <b>Feature: Application</b> | State information for the DCBX application.  |
| <b>Protocol-State</b>       | <p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—The local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.</li> </ul> |
| <b>Local-Advertisement</b>  | <p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>  |
| <b>Enable</b>               | <p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>   |
| <b>Willing</b>              | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the FCoE interface state from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the FCoE interface state from the peer.</li> </ul>  |
| <b>Error</b>                | <p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. The local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. The local and peer configuration are not compatible.</li> </ul>  |
| <b>Appl-Name</b>            | Name of the application:   |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description   |
|---------------------------------------|---|
| <b>Ethernet-Type</b>                  | <p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>  |
| <b>Socket-Number</b>                  | <p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>   |
| <b>Priority-Field or Priority-Map</b> | <p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>  |
| <b>Status</b>                         | <p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> <p><b>NOTE:</b> If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p> |
| <b>Peer-Advertisement</b>             | <p>Status of advertisements that the peer sends to the local interface.</p>   |
| <b>Enable</b>                         | <p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>  |
| <b>Willing</b>                        | <p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the FCoE interface state from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the FCoE interface state from the local interface.</li> </ul>   |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                            | Field Description  |
|---------------------------------------|--|
| <b>Error</b>                          | (DCBX Version 1.01 only)<br><br>Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>   |
| <b>Appl-Name</b>                      | Name of the application: <ul style="list-style-type: none"> <li>• <b>FCoE</b>—Fibre Channel over Ethernet</li> </ul>   |
| <b>Ethernet-Type</b>                  | Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.  |
| <b>Socket-Number</b>                  | Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.   |
| <b>Priority-Field or Priority-Map</b> | Priority assigned to the application.  |
| <b>Status</b>                         | (DCBX Version 1.01 only)<br><br>Peer interface status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                 | Field Description  |
|----------------------------|--|
| <b>Feature: ETS</b>        | Enhanced Transmission Selection (ETS) DCBX state information.  |
| <b>Protocol-State</b>      | (DCBX Version 1.01 only)<br><br>ETS protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> </ul>  |
| <b>Operational State</b>   | (DCBX Version 1.01 only)<br><br>Operational state of the feature, <b>enabled</b> or <b>disabled</b> .  |
| <b>Local-Advertisement</b> | Status of advertisements that the local interface sends to the peer.   |
| <b>Enable</b>              | (DCBX Version 1.01 only)<br><br>State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>  |
| <b>TLV Type</b>            | (IEEE DCBX only)<br><br>Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Recommendation-or-Configuration</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>             | Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise <b>No</b> for this field): <ul style="list-style-type: none"> <li>• <b>Yes</b>—Local interface is willing to learn the ETS state from the peer.</li> <li>• <b>No</b>—Local interface is not willing to learn the ETS state from the peer.</li> </ul>   |
| <b>Credit Based Shaper</b> |  |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name  | Field Description   |
|---|---|
|   | (IEEE DCBX only)  |
|   | Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .  |
| <b>Error</b>  | (DCBX Version 1.01 only)<br><br>Configuration error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error. This should always be the switch ETS error state.</li> <li>• <b>Yes</b>—Error detected.</li> </ul>   |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.   |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)  |
| <b>Code Point</b>                                     | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.  |
| <b>Priority-Group</b>                                 | Class-of-service (CoS) priority group (forwarding class set) identification number.   |
| <b>Percentage B/W</b>                                 | Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.) |
| <b>Transmission Selection Algorithm</b>               | (IEEE DCBX only)<br><br>The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .   |
| <b>Peer-Advertisement</b>                             | Status of advertisements that the peer sends to the local interface.  |
| <b>Enable</b>   |   |



Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name  | Field Description   |
|---|---|
|   | (DCBX Version 1.01 only)<br><br>State that the peer advertises to the local interface: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>   |
| <b>TLV Type</b>                                       | (IEEE DCBX only)<br><br>Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Configuration/Recommendation</b>—Advertises both TLVs.</li> </ul> |
| <b>Willing</b>  | Willingness of the peer to learn the ETS state from the local interface using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—Peer is willing to learn the ETS state from the local interface.</li> <li>• <b>No</b>—Peer is not willing to learn the ETS state from the local interface.</li> </ul>   |
| <b>Credit Based Shaper</b>                            | (IEEE DCBX only)<br><br>Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .  |
| <b>Error</b>  | (DCBX Version 1.01 only)<br><br>Configuration error status of the peer: <ul style="list-style-type: none"> <li>• <b>No</b>—No error in peer ETS TLV.</li> <li>• <b>Yes</b>—Error in peer ETS TLV.</li> </ul>  |
| <b>Maximum Traffic Classes capable to support PFC</b> | (DCBX Version 1.01 only)<br><br>Largest number of traffic classes the local interface supports for PFC.   |
| <b>Maximum Traffic Classes supported</b>              | (IEEE DCBX only)<br><br>Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)  |
| <b>Code Point</b>                                     |   |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name                              | Field Description  |
|---|--|
|   | PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.   |
| <b>Priority-Group</b>                   | CoS priority group (forwarding class set) identification number.   |
| <b>Percentage B/W</b>                   | Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)  |
| <b>Transmission Selection Algorithm</b> | (IEEE DCBX only)<br><br>Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .  |
| <b>PFC</b>                              | (QFX Series, <b>terse</b> option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> <li>• Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled</li> <li>• Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled</li> <li>• Not Advt—Interface does not advertise PFC to the connected peer</li> </ul> |
| <b>ETS</b>                              | ( <b>terse</b> option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> <li>• Advt—Interface advertises ETS TLVs</li> <li>• Disabled—ETS is disabled on the interface (interface does not advertise ETS)</li> </ul>   |
| <b>ETS Rec</b>                          | ( <b>terse</b> option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> <li>• Advt—Peer interface advertises ETS TLVs</li> <li>• Not Advt—Peer interface does not advertise ETS</li> </ul> <p><b>NOTE:</b> When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>  |

Table 464: show dcbx neighbors Output Fields (*continued*)

| Field Name | Field Description  |
|------------|--|
| Version    | <p>(<b>terse</b> option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> <li>• <b>IEEE</b>—The interface uses IEEE DCBX.</li> <li>• <b>1.01</b>—The interface uses DCBX version 1.01.</li> </ul> <p>When the DCBX version used is the result of autonegotiation, the term (<b>Auto</b>) appears next to the version. For example, <b>IEEE (Auto)</b> indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p> |

## Sample Output

### show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
  Operational version: 1
  sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
  Operational version: 1
  sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode      Operational Mode
000             Disabled       Disable
001             Disabled       Disable
010             Disabled       Disable
011             Enabled        Enable
100             Enabled        Enable
101             Disabled       Disable
110             Disabled       Disable
111             Disabled       Disable

Peer-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode
000             Disabled

```

|     |          |
|-----|----------|
| 001 | Disabled |
| 010 | Disabled |
| 011 | Enabled  |
| 100 | Enabled  |
| 101 | Disabled |
| 110 | Disabled |
| 111 | Disabled |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001110     | Enabled |
| iSCSI     |               | 3260          | 10000000     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        | N/A           | 00001110     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |

|                |                |
|----------------|----------------|
| 111            | 7              |
| Priority-Group | Percentage B/W |
| 0              | 40%            |
| 1              | 5%             |

### show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

user@switch> **show dcbx neighbors interface xe-0/0/0**

Interface : xe-0/0/0.0 - Parent Interface: ae0.0

Active-application-map: app-map-1

Protocol-Mode: IEEE-DCBX Version

Feature: PFC

Local-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Peer-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application

Local-Advertisement:

| App1-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
| FCoE      | 0x8906        |               | 00001110       |
| iSCSI     |               | 3260          | 10000000       |

Peer-Advertisement:

| App1-Name | Ethernet-Type | Socket-Number | Priority-field |
|-----------|---------------|---------------|----------------|
|-----------|---------------|---------------|----------------|

|      |        |     |          |
|------|--------|-----|----------|
| FCoE | 0x8906 | N/A | 00001110 |
|------|--------|-----|----------|

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |
| 101        | 1              |
| 110        | 1              |
| 111        | 7              |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

Peer-Advertisement:

TLV Type: Recommendation

| Code Point | Priority-Group |
|------------|----------------|
| 000        | 0              |
| 001        | 7              |
| 010        | 7              |
| 011        | 7              |
| 100        | 0              |

|     |   |
|-----|---|
| 101 | 1 |
| 110 | 1 |
| 111 | 7 |

| Priority-Group | Percentage B/W |
|----------------|----------------|
| 0              | 40%            |
| 1              | 5%             |

| Priority-Group | Transmission Selection Algorithm |
|----------------|----------------------------------|
| 0              | Enhanced Transmission Selection  |
| 1              | Enhanced Transmission Selection  |

### show dcbx neighbors terse (QFX Series)

```

user@switch> show dcbx neighbors terse
Interface Parent PFC ETS ETS Version
Interface
xe-0/0/8.0 - Enabled Advt Advt IEEE (Auto)
xe-0/0/9.0 - Disabled Disabled 1.01
xe-0/0/11.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/12.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/32.0 - Enabled Advt Not Advt IEEE
xe-0/0/36.0 - Not Advt Advt Advt IEEE

```

### show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)

```

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
  Operational version: 0
  sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
  Operational version: 0
  sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 6

Code Point      Admin Mode
000             Disabled
001             Disabled
010             Disabled
011             Enabled
100             Disabled
101             Disabled
110             Disabled
111             Disabled

```

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Status  | Appl-Name | Ethernet-Type | Socket-Number | Priority-Map |
|---------|-----------|---------------|---------------|--------------|
| Enabled | FCoE      | 0x8906        |               | 00001000     |

**show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)**

user@switch&gt; show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

## Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

## Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled



## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Disabled   |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Enabled    |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

| Appl-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00001000     | Enabled |
| iscsi     |               | 3260          | 00100000     | Enabled |

**show dcbx neighbors (EX4500 Switch: Includes ETS)**

user@switch&gt; show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0  
 Protocol-State: in-sync  
 Active-application-map: map\_iscsi

## Local-Advertisement:

Operational version: 0

sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Enabled    |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Disabled   |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

| Code Point | Admin Mode |
|------------|------------|
| 000        | Enabled    |
| 001        | Disabled   |
| 010        | Disabled   |
| 011        | Disabled   |
| 100        | Enabled    |
| 101        | Disabled   |
| 110        | Disabled   |
| 111        | Disabled   |

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00000001     | Enabled |
| iscsi     |               | 3260          | 00000010     | Enabled |

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

| App1-Name | Ethernet-Type | Socket-Number | Priority-Map | Status  |
|-----------|---------------|---------------|--------------|---------|
| FCoE      | 0x8906        |               | 00010000     | Enabled |
| iscsi     |               | 3260          | 00010000     | Enabled |

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No  
Maximum Traffic Classes supported : 3

| Code Point     | Priority-Group |
|----------------|----------------|
| 000            | 7              |
| 001            | 7              |
| 010            | 7              |
| 011            | 7              |
| 100            | 7              |
| 101            | 7              |
| 110            | 7              |
| 111            | 7              |
| Priority-Group | Percentage B/W |
| 7              | 100%           |

## Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No  
Maximum Traffic Classes supported : 8

| Code Point     | Priority-Group |
|----------------|----------------|
| 000            | 0              |
| 001            | 1              |
| 010            | 0              |
| 011            | 0              |
| 100            | 2              |
| 101            | 0              |
| 110            | 0              |
| 111            | 0              |
| Priority-Group | Percentage B/W |
| 0              | 30%            |
| 1              | 40%            |
| 2              | 30%            |

## show fip snooping

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show fip snooping</b><br><b>&lt;brief   detail&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Display FIP snooping information.   |
| <b>Options</b>                  | <b>none</b> —Display FIP snooping information.<br><br><b>brief   detail</b> —(Optional) Display the specified level of output.  |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Configuring an FCoE LAG</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Example: Configuring an FCoE LAG on a Redundant Server Node Group</a></li> <li>• <a href="#">show fip snooping enode on page 5751</a></li> <li>• <a href="#">show fip snooping fcf on page 5755</a></li> <li>• <a href="#">show fip snooping interface on page 5758</a></li> <li>• <a href="#">show fip snooping statistics on page 5761</a></li> <li>• <a href="#">show fip snooping vlan on page 5764</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show fip snooping on page 5748</a><br><a href="#">show fip snooping brief (QFX Series) on page 5748</a><br><a href="#">show fip snooping detail (QFX Series Switches) on page 5749</a><br><a href="#">show fip snooping detail (QFabric System FCoE with LAG Configured) on page 5749</a><br><a href="#">show fip snooping detail (EX Series Switches) on page 5750</a>   |
| <b>Output Fields</b>            | <a href="#">Table 465 on page 5746</a> lists the output fields for the <b>show fip snooping</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 465: show fip snooping Output Fields**

| Field Name | Field Description | Level of Output |
|------------|-------------------|-----------------|
| VLAN       | Name of the VLAN. | All             |

Table 465: show fip snooping Output Fields (*continued*)

| Field Name                              | Field Description  | Level of Output |
|---|--|-----------------|
| <b>Mode</b>                             | <p>(QFX Series only)<br/>Snooping mode enabled on the FCoE VLAN:</p> <ul style="list-style-type: none"> <li>• VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port.</li> <li>• VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.</li> </ul>  | All             |
| <b>FC-MAP</b>                           | FCoE mapped address prefix of the FCoE forwarder for the VLAN.   | All             |
| <b>FCF or FCF-MAC</b>                   | MAC address of the FCF.  | All             |
| <b>Session Count or Active Sessions</b> | Current number of virtual link sessions with VN_Ports.   | All             |
| <b>VN_Port Count</b>                    | <p>(QFX Series only)<br/>Number of VN_Ports active on an ENode.</p>  | <b>brief</b>    |
| <b>Configured FKA-ADV</b>               | <p>FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p> | <b>detail</b>   |
| <b>Running FKA-ADV</b>                  | <p>Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>                                     | <b>detail</b>   |
| <b>Beacon Period</b>                    | <p>(QFX Series only)<br/>Beacon period interval in milliseconds.</p>   | <b>detail</b>   |

Table 465: show fip snooping Output Fields (*continued*)

| Field Name                       | Field Description  | Level of Output |
|----------------------------------|--|-----------------|
| <b>VN2VN Mode</b>                | (QFX Series only)<br>Mode of VN2VN_Port snooping: <ul style="list-style-type: none"> <li>Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks.</li> <li>Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.</li> </ul> | <b>detail</b>   |
| <b>ENode-MAC</b>                 | MAC address of the connected FCoE node (ENode).  | All             |
| <b>Interface</b>                 | Interface connected to the ENode.<br><br>(QFabric System only)<br>When an FCoE LAG has been configured, LAG interface connected to the ENode and LAG member interface connected to ENode.  | <b>detail</b>   |
| <b>VN-Port MAC</b>               | MAC address of a VN_Port on the ENode.   | All             |
| <b>FKA-ADV</b>                   | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.  | <b>detail</b>   |
| <b>Active VN_Ports</b>           | (QFX Series only)<br>Number of VN_Ports active on an ENode.  | <b>detail</b>   |
| <b>Vlink far-end VN-Port-MAC</b> | (QFX Series only)<br>Media access control (MAC) address of the VN_Port at the other end of the virtual link.   | <b>detail</b>   |

## Sample Output

### show fip snooping

```

user@switch> show fip snooping
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
FCF : 00:10:94:00:00:01  Session Count : 2
ENode-MAC : 00:10:94:00:00:02
VN-Port-MAC : 0E:FC:00:01:00:05
VN-Port-MAC : 0E:FC:00:01:00:01

```

### show fip snooping brief (QFX Series)

```

user@switch> show fip snooping brief
VLAN: vlan100,    Mode: VN2VF Snooping
FC-MAP: 0e:fc:00
FCF: 30:10:94:01:00:00  Session Count: 2
ENode-MAC: 10:10:94:01:00:01

```

```

VN-Port-MAC: 0e:fc:00:01:0d:01
VN-Port-MAC: 0e:fc:00:01:0e:01
VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fc:00
Enode-MAC: 10:10:94:01:00:02 VN_Port count: 1
VN-Port-MAC: 0e:fc:00:01:0a:01 Session Count: 2
Enode-MAC: 10:10:94:01:00:03 VN_Port count: 0

```

### show fip snooping detail (QFX Series Switches)

```

user@switch> show fip snooping detail
root@sw-pa02v> show fip snooping detail
VLAN: vlan100, Mode: VN2VF Snooping
FC-MAP: 0e:fc:00
FCF Information
FCF-MAC : 30:10:94:01:00:00
Active Sessions : 2
Configured FKA-ADV : 258
Running FKA-ADV : 188
Enode Information
Enode-MAC: 10:10:94:01:00:01, Interface: xe-0/0/10
Configured FKA-ADV : 258
Running FKA-ADV : 230
Session Information
VN-Port MAC: 0e:fc:00:01:0d:01, FKA-ADV : 230
VN-Port MAC: 0e:fc:00:01:0e:01, FKA-ADV : 245

VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Multi-Point
Enode Information
Enode-MAC: 10:10:94:01:00:02, Interface: xe-0/0/10
Active VN_Ports : 1
VN_Port Information
VN-Port MAC: 0e:fd:00:01:0a:01
Active Sessions : 2
Session Information
Vlink far-end VN-Port-MAC: 0e:fd:00:01:0b:01
Vlink far-end VN-Port-MAC: 0e:fd:00:01:0c:01
Enode-MAC: 10:10:94:01:00:02, Interface: xe-0/0/11
Active VN_Ports : 0

```

### show fip snooping detail (QFabric System FCoE with LAG Configured)

```

admin@qfabric> show fip snooping detail
VLAN: vlan_100, Mode: VN2VF Snooping
FC-MAP: 0e:fc:00
FCF Information
FCF-MAC : 84:18:88:d1:f5:cc
Active Sessions : 2
Configured FKA-ADV : 8000
Running FKA-ADV : 23962
Enode Information
Enode-MAC: 00:c0:dd:14:ae:6d, Interface: P4546-C:ae0 P4546-C:xe-0/0/39

Configured FKA-ADV : 8000
Running FKA-ADV : 16622
Session Information
VN-Port MAC: 0e:fc:00:6c:06:a5, FKA-ADV : 246303
Enode Information

```

```
Enode-MAC: 00:c0:dd:14:ae:6f,      Interface: P4546-C:ae0 P4546-C:xe-0/0/38

Configured FKA-ADV : 8000
Running FKA-ADV    : 16512
Session Information
VN-Port MAC: 0e:fc:00:6c:06:a4,    FKA-ADV : 238150
```

#### show fip snooping detail (EX Series Switches)

```
user@switch> show fip snooping detail
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
FCF Information
FCF-MAC           : 00:10:94:00:00:01
Active Sessions   : 2
Configured FKA-ADV : 258
Running FKA-ADV    : 244
Enode Information
Enode-MAC : 00:10:94:00:00:02      Interface : xe-0/0/1
Configured FKA-ADV : 258
Running FKA-ADV    : 248
Session Information
VN-Port MAC : 0E:FC:00:01:00:05    FKA-ADV : 264
VN-Port MAC : 0E:FC:00:01:00:01    FKA-ADV : 260
```



## show fip snooping enode

|                                 |   |  |
|---------------------------------|---|--|
| <b>Syntax</b>                   | <b>show fip snooping enode <i>enode-mac</i></b><br><b>&lt;brief   detail&gt;</b><br><b>&lt;vlan <i>vlan-name</i>&gt;</b>  |  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.  |  |
| <b>Description</b>              | Display FIP snooping FCoE node (ENode) information.   |  |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b><i>enode-mac</i></b> —Display information for the ENode specified by the MAC address.<br><br><b>vlan <i>vlan-name</i></b> —(Optional) Display FIP snooping information for the ENode on only the specified VLAN.   |  |
| <b>Required Privilege Level</b> | view  |  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">show fip snooping fcf on page 5755</a></li> <li>• <a href="#">show fip snooping interface on page 5758</a></li> <li>• <a href="#">show fip snooping statistics on page 5761</a></li> <li>• <a href="#">show fip snooping vlan on page 5764</a></li> </ul> |  |
| <b>List of Sample Output</b>    | <a href="#">show fip snooping enode on page 5753</a><br><a href="#">show fip snooping enode brief (QFX Series) on page 5753</a><br><a href="#">show fip snooping enode detail (QFX Series) on page 5753</a><br><a href="#">show fip snooping enode detail on page 5753</a>  |  |
| <b>Output Fields</b>            | <a href="#">Table 466 on page 5751</a> lists the output fields for the <b>show fip snooping enode</b> command. Output fields are listed in the approximate order in which they appear.  |  |

**Table 466: show fip snooping enode Output Fields**

| Field Name          | Field Description                 | Level of Output |
|---------------------|-----------------------------------|-----------------|
| ENode and ENode MAC | MAC address of the ENode.         | All             |
| VLAN                | Name of the VLAN.                 | All             |
| Interface           | Interface connected to the ENode. | All             |

Table 466: show fip snooping enode Output Fields (*continued*)

| Field Name                    | Field Description   | Level of Output |
|-------------------------------|---|-----------------|
| <b>Mode</b>                   | (QFX Series only)<br>Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> <li>VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port.</li> <li>VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.</li> </ul>   | All             |
| <b>VN_Port Count</b>          | (QFX Series only)<br>Number of VN_Ports active on an ENode.   | <b>brief</b>    |
| <b>Session Count</b>          | Current number of virtual link sessions with VN_Ports.  | All             |
| <b>Configured FKA-ADV</b>     | FIP keepalive interval in seconds configured on the FCoE forwarder (FCF) multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.<br><br>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module. | <b>detail</b>   |
| <b>Running FKA-ADV</b>        | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.<br><br>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.  | <b>detail</b>   |
| <b>VN-Port or VN-Port-MAC</b> | MAC address of a VN_Port on the ENode.  | All             |
| <b>FKA-ADV</b>                | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF.   | <b>detail</b>   |
| <b>FCF or FCF-MAC</b>         | MAC address of the FCF to which the VN_Port is connected.   | All             |
| <b>Beacon Period</b>          | (QFX Series only)<br>Beacon period interval in milliseconds.  | <b>detail</b>   |

Table 466: show fip snooping enode Output Fields (*continued*)

| Field Name                | Field Description  | Level of Output |
|---------------------------|--|-----------------|
| VN2VN Mode                | (QFX Series only)<br>Mode of VN2VN_Port snooping: <ul style="list-style-type: none"> <li>Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks.</li> <li>Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.</li> </ul> | detail          |
| Vlink far-end VN-Port-MAC | (QFX Series only)<br>Media access control (MAC) address of the VN_Port at the other end of the virtual link.   | detail          |

## Sample Output

### show fip snooping enode

```

user@switch> show fip snooping enode 00:10:94:00:00:02
Enode : 00:10:94:00:00:02   VLAN : vlan1   Interface : xe-0/0/1
      VN-Port-MAC          FCF-MAC
      0E:FC:00:00:00:05     00:10:94:00:00:01
      0E:FC:00:00:00:01     00:10:94:00:00:01

```

### show fip snooping enode brief (QFX Series)

```

user@switch> show fip snooping enode 10:10:94:01:00:02 brief
Enode: 10:10:94:01:00:02 ,   VLAN: vlan101,   Interface: xe-0/0/10
  Mode: VN2VF Snooping      VN_Port Count: 1
    VN_Port Information
    VN_Port Mac: 0e:fc:00:01:0a:01      Session Count: 2

```

### show fip snooping enode detail (QFX Series)

```

user@switch> show fip snooping enode 10:10:94:01:00:02 detail
Enode MAC: 10:10:94:01:00:02,   VLAN: vlan101,   Interface: xe-0/0/10
  Mode: VN2VF Snooping      VN_Port Count: 1
  Beacon_Period: 90000      VN2VN Mode: Multi-Point
    VN_Port Information
    VN_Port Mac: 0e:fc:00:01:0a:01      Session Count: 2
  Vlink far-end VN-Port-MAC: 0e:fc:00:01:0b:01
  Vlink far-end VN-Port-MAC: 0e:fc:00:01:0c:01

```

### show fip snooping enode detail

```

user@switch> show fip snooping enode 00:10:94:00:00:02 detail
Enode MAC : 00:10:94:00:00:02   VLAN : vlan1   Interface : xe-0/0/1
Configured FKA-ADV : 258      Running FKA-ADV : 213
  Session Information
  VN-Port : 0E:FC:00:00:00:05   FKA-ADV : 229   FCF : 00:10:94:00:00:01
  VN-Port : 0E:FC:00:00:00:01   FKA-ADV : 225   FCF : 00:10:94:00:00:01

```



## show fip snooping fcf

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show fip snooping fcf <i>fcf-mac</i></b><br><b>&lt;brief   detail&gt;</b><br><b>&lt;vlan <i>vlan-name</i>&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.  |
| <b>Description</b>              | Display FIP snooping FCoE forwarder (FCF) information.  |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b><i>fcf-mac</i></b> —Display information for the FCF specified by the MAC address.<br><br><b><i>vlan-name</i></b> —(Optional) Display FIP snooping information for the FCF on only the specified VLAN.  |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">show fip snooping enode on page 5751</a></li> <li>• <a href="#">show fip snooping interface on page 5758</a></li> <li>• <a href="#">show fip snooping statistics on page 5761</a></li> <li>• <a href="#">show fip snooping vlan on page 5764</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show fip snooping fcf on page 5756</a><br><a href="#">show fip snooping fcf detail on page 5756</a>   |
| <b>Output Fields</b>            | <a href="#">Table 467 on page 5755</a> lists the output fields for the <b>show fip snooping fcf</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 467: show fip snooping fcf Output Fields**

| Field Name     | Field Description                                      | Level of Output |
|----------------|--|-----------------|
| FCF or FCF-MAC | MAC address of the FCoE forwarder.                     | All             |
| VLAN           | Name of the VLAN.                                      | All             |
| Session Count  | Current number of virtual link sessions with VN_Ports. | None            |

Table 467: show fip snooping fcf Output Fields (*continued*)

| Field Name           | Field Description   | Level of Output |
|----------------------|---|-----------------|
| Configured FKA-ADV   | FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.                                  | detail          |
| Running FKA-ADV      | Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.  | detail          |
| ENode-MAC            | MAC address of the connected ENode.   | All             |
| • Interface          | Interface connected to the ENode.   | detail          |
| • Configured FKA-ADV | FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.     | detail          |
| • Running FKA-ADV    | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.   | detail          |
| • VN-Port MAC        | MAC address of a VN_Port on the ENode.  | All             |
| • FKA-ADV            | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF. | detail          |

## Sample Output

### show fip snooping fcf

```

user@switch> show fip snooping fcf 00:10:94:00:00:01
FCF : 00:10:94:00:00:01  VLAN : v1an1  Session Count : 2
  ENode-MAC : 00:10:94:00:00:02
    VN-Port-MAC : 0E:FC:00:00:00:05
    VN-Port-MAC : 0E:FC:00:00:00:01

```

### show fip snooping fcf detail

```

user@switch> show fip snooping fcf 00:10:94:00:00:01 detail
FCF-MAC : 00:10:94:00:00:01  VLAN : v1an1
Configured FKA-ADV : 258      Running FKA-ADV : 222
  ENode Information
    ENode-MAC : 00:10:94:00:00:02 Interface: xe-0/0/1
    Configured FKA-ADV : 258
    Running FKA-ADV : 226
    Session Information
      VN-Port MAC : 0E:FC:00:00:00:05  FKA-ADV : 242
      VN-Port MAC : 0E:FC:00:00:00:01  FKA-ADV : 238

```



## show fip snooping interface

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show fip snooping interface</b> <i>interface-name</i><br><brief   detail>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for the QFX Series.  |
| <b>Description</b>              | Display FIP snooping information for the specified interface.  |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —Display information for the specified interface.  |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">show fip snooping enode on page 5751</a></li> <li>• <a href="#">show fip snooping fcf on page 5755</a></li> <li>• <a href="#">show fip snooping statistics on page 5761</a></li> <li>• <a href="#">show fip snooping vlan on page 5764</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show fip snooping interface on page 5760</a><br><a href="#">show fip snooping interface detail on page 5760</a>  |
| <b>Output Fields</b>            | <a href="#">Table 468 on page 5758</a> lists the output fields for the <b>show fip snooping interface interface-name</b> command. Output fields are listed in the approximate order in which they appear.  |

Table 468: show fip snooping interface Output Fields

| Field Name                       | Field Description  | Level of Output |
|----------------------------------|--|-----------------|
| VLAN                             | Name of the VLAN.  | All             |
| FC-MAP                           | FCoE mapped address prefix of the FCoE forwarder for the VLAN. | All             |
| FCF or FCF-MAC                   | MAC address of the FCF.  | All             |
| Session Count or Active Sessions | Current number of virtual link sessions with VN_Ports.         | All             |



Table 468: show fip snooping interface Output Fields (*continued*)

| Field Name         | Field Description   | Level of Output |
|--------------------|---|-----------------|
| Configured FKA-ADV | <p>FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>                              | detail          |
| Running FKA-ADV    | <p>Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>  | detail          |
| ENode-MAC          | MAC address of the connected FCoE node (ENode).   | All             |
| Interface          | Interface connected to the ENode.   | detail          |
| Configured FKA-ADV | <p>FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p> | detail          |
| Running FKA-ADV    | <p>Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.</p> <p>For the QFX Series only, the output of this field is always 0 (zero) if the VLAN is an FCoE-FC gateway VLAN. If the VLAN is a FIP snooping VLAN (a transit switch VLAN), then the output is accurate. This is because for an FCoE-FC gateway VLAN, FIP snooping is performed internally and the keepalive advertisements are not tracked by the switch's Ethernet module.</p>   | detail          |
| VN-Port MAC        | MAC address of a VN_Port on the ENode.  | All             |

Table 468: show fip snooping interface Output Fields (*continued*)

| Field Name | Field Description   | Level of Output |
|------------|---|-----------------|
| FKA-ADV    | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF. | <b>detail</b>   |

## Sample Output

### show fip snooping interface

```

user@switch> show fip snooping interface xe-0/0/9.0
VLAN: vlan_100,    FC-MAP: 0e:fc:00
FCF: 30:10:94:01:00:00    Session Count: 1
  Enode-MAC: 10:10:94:01:00:01
    VN-Port-MAC: 0e:fc:00:01:0a:01

```

### show fip snooping interface detail

```

user@switch> show fip snooping interface xe-0/0/9.0 detail
VLAN: vlan_100, FC-MAP: 0e:fc:00
FCF Information
FCF-MAC          : 30:10:94:01:00:00
Active Sessions  : 1
Configured FKA-ADV : 368640000
Running FKA-ADV   : 0
  Enode Information
  Enode-MAC: 10:10:94:01:00:01,      Interface: xe-0/0/9
  Configured FKA-ADV : 368640000
  Running FKA-ADV    : 0
    Session Information
    VN-Port MAC: 0e:fc:00:01:0a:01,  FKA-ADV : 0

```

## show fip snooping statistics

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show fip snooping statistics</b><br><b>&lt;vlan vlan-name&gt;</b>   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Display FIP snooping statistics.   |
| <b>Options</b>                  | <b>vlan vlan-name</b> —(Optional) Display FIP snooping statistics for the specified VLAN.  |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">show fip snooping enode on page 5751</a></li> <li>• <a href="#">show fip snooping fcf on page 5755</a></li> <li>• <a href="#">show fip snooping interface on page 5758</a></li> <li>• <a href="#">show fip snooping vlan on page 5764</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show fip snooping statistics (FIP Snooping) on page 5763</a><br><a href="#">show fip snooping statistics (VN2VN_Port Snooping) on page 5763</a>  |
| <b>Output Fields</b>            | Table 469 on page 5761 lists the output fields for the <b>show fip snooping statistics</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 469: show fip snooping statistics Output Fields**

| Field Name           | Field Description   |
|----------------------|---|
| <b>VLAN</b>          | Name of the VLAN for which a set of statistics is displayed.  |
| <b>Mode</b>          | (QFX Series only)<br>Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> <li>• VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port.</li> <li>• VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.</li> </ul> |
| <b>Number of MDS</b> | Number of multicast discovery solicitation messages sent on the VLAN.   |

Table 469: show fip snooping statistics Output Fields (*continued*)

| Field Name                     | Field Description   |
|--------------------------------|---|
| Number of UDS                  | Number of unicast discovery solicitation messages sent on the VLAN.   |
| Number of FLOGI                | Number of fabric logins on the VLAN.  |
| Number of FDISC                | Number of fabric discovery logins on the VLAN.  |
| Number of LOGO                 | Number of fabric logouts on the VLAN.   |
| Number of ENode-keep-alive     | Number of ENode keepalive messages sent on the VLAN.  |
| Number of VN_Port-keep-alive   | Number of VN_Port keepalive messages sent on the VLAN.  |
| Number of MDA                  | Number of multicast discovery advertisement messages sent on the VLAN.  |
| Number of UDA                  | Number of unicast discovery advertisement messages sent on the VLAN.  |
| Number of FLOGI_ACC            | Number of fabric logins accepted on the VLAN.   |
| Number of FLOGI_RJT            | Number of fabric logins rejected on the VLAN.   |
| Number of FDISC_ACC            | Number of fabric discoveries accepted on the VLAN.  |
| Number of FDISC_RJT            | Number of fabric discoveries rejected on the VLAN.  |
| Number of LOGO_ACC             | Number of fabric logouts accepted on the VLAN.  |
| Number of LOGO_RJT             | Number of fabric logouts rejected on the VLAN.  |
| Number of CVL                  | Number of clear virtual links (CVL) actions on the VLAN.  |
| Number of VN_Port Probes Req   | (QFX Series only)<br>Number of multicast N_Port_ID probes sent to the ALL-VN2VN-ENode-MACs multicast address on the VLAN.                         |
| Number of VN_Port Claim Notif  | (QFX Series only)<br>Number of multicast N_Port_ID claim notifications sent on the VLAN.  |
| Number of VN_Port Beacons      | (QFX Series only)<br>Number of multicast beacons sent on the VLAN.  |
| Number of VN_Port Probes Reply | (QFX Series only)<br>Number of replies to N_Port_ID probes sent on the VLAN. Replies are unicast to the ENode MAC address of the probe requester. |

Table 469: show fip snooping statistics Output Fields (*continued*)

| Field Name                    | Field Description   |
|-------------------------------|---|
| Number of VN_Port Claim Reply | (QFX Series only)<br>Number of replies to N_Port_ID claim notifications sent on the VLAN. Replies are unicast to the ENode MAC address of the claim notifier. |

## Sample Output

### show fip snooping statistics (FIP Snooping)

```

user@switch> show fip snooping statistics
VLAN: fcoevlan1      Mode: VN2VF Snooping

Number of MDS:                2
Number of UDS:                2
Number of FLOGI:              2
Number of FDISC:              2
Number of LOGO:               0
Number of Enode-keep-alive: 200
Number of VNPort-keep-alive: 200

Number of MDA:                25
Number of UDA:                2
Number of FLOGI_ACC:          2
Number of FLOGI_RJT:          0
Number of FDISC_ACC:          2
Number of FDISC_RJT:          0
Number of LOGO_ACC:           0
Number of LOGO_RJT:           0
Number of CVL:                0

```

### show fip snooping statistics (VN2VN\_Port Snooping)

```

user@switch> show fip snooping statistics
VLAN: vlan101      Mode: VN2VN Snooping

Number of VN_Port Probes Req:    3
Number of VN_Port Claim Notif:  3
Number of VN_Port Beacons:      0

Number of VN_Port Probes Reply:  3
Number of VN_Port Claim Reply:   3
Number of FLOGI:                 0
Number of FLOGI_ACC:             0
Number of FLOGI_RJT:             0
Number of FDISC:                 0
Number of FDISC_ACC:             0
Number of FDISC_RJT:             0
Number of LOGO:                 0
Number of LOGO_ACC:             0
Number of LOGO_RJT:             0

```

## show fip snooping vlan

|                                 |  |
|---------------------------------|--|
| <b>Syntax</b>                   | <b>show fip snooping vlan <i>vlan-name</i></b><br><b>&lt;brief   detail&gt;</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 10.4 for EX Series switches.<br>Command introduced in Junos OS Release 11.1 for the QFX Series.   |
| <b>Description</b>              | Display FIP snooping VLAN information.   |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b><i>vlan-name</i></b> —Display information for the specified VLAN.   |
| <b>Required Privilege Level</b> | view   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring VN2VF_Port FIP Snooping and FCoE Trusted Interfaces on an FCoE Transit Switch on page 5662</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">show fip snooping on page 5746</a></li> <li>• <a href="#">show fip snooping enode on page 5751</a></li> <li>• <a href="#">show fip snooping fcf on page 5755</a></li> <li>• <a href="#">show fip snooping interface on page 5758</a></li> <li>• <a href="#">show fip snooping statistics on page 5761</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show fip snooping vlan on page 5766</a><br><a href="#">show fip snooping vlan (QFX Series, VN2VF_Port FIP Snooping) on page 5766</a><br><a href="#">show fip snooping vlan (QFX Series, VN2VN_Port FIP Snooping) on page 5766</a><br><a href="#">show fip snooping vlan detail (QFX Series, VN2VN_Port FIP Snooping) on page 5767</a><br><a href="#">show fip snooping vlan detail on page 5767</a>  |
| <b>Output Fields</b>            | <a href="#">Table 470 on page 5764</a> lists the output fields for the <b>show fip snooping vlan</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 470: show fip snooping vlan Output Fields**

| Field Name | Field Description | Level of Output |
|------------|-------------------|-----------------|
| VLAN       | Name of the VLAN. | All             |

Table 470: show fip snooping vlan Output Fields (*continued*)

| Field Name                              | Field Description  | Level of Output |
|---|--|-----------------|
| <b>Mode</b>                             | (QFX Series only)<br>Snooping mode enabled on the FCoE VLAN: <ul style="list-style-type: none"> <li>VN2VF Snooping—The FCoE VLAN is configured for FIP snooping between an ENode VN_Port and a switch VF_Port.</li> <li>VN2VN Snooping—The FCoE VLAN is configured for VN_Port to VN_Port FIP snooping between ENode VN_Ports.</li> </ul>  | All             |
| <b>VN_Port count</b>                    | (QFX Series only)<br>Number of VN_Ports active on an ENode when the mode is VN2VN_Port FIP snooping.   |                 |
| <b>FC-MAP</b>                           | FCoE mapped address prefix of the FCoE forwarder for the VLAN.   | All             |
| <b>Beacon_Period</b>                    | (QFX Series only)<br>Beacon period interval in milliseconds.   | <b>detail</b>   |
| <b>VN2VN Mode</b>                       | (QFX Series only)<br>Mode of VN2VN_Port snooping: <ul style="list-style-type: none"> <li>Multi-Point—Multiple ENodes are connected to the network and form multiple virtual links. Each virtual link is created between one pair of VN_Ports. This is analogous to the loop mode in traditional FC networks.</li> <li>Point-to-Point—Two ENodes are connected to the network and form a single VN_Port to VN_Port virtual link. This is analogous to the point-to-point FC link between an FC initiator and an FC target.</li> </ul> | <b>detail</b>   |
| <b>FCF or FCF-MAC</b>                   | MAC address of the FCF.  | All             |
| <b>Session Count or Active Sessions</b> | Current number of virtual link sessions with VN_Ports.   | All             |
| <b>Configured FKA-ADV</b>               | FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258.   | <b>detail</b>   |
| <b>Running FKA-ADV</b>                  | Runtime interval in seconds of the last FIP keepalive advertisement the FCF received. This value changes every time the FCF receives an FKA_ADV.   | <b>detail</b>   |

Table 470: show fip snooping vlan Output Fields (*continued*)

| Field Name                         | Field Description   | Level of Output |
|------------------------------------|---|-----------------|
| <b>ENode-MAC</b>                   | MAC address of the connected ENode.   | All             |
| • <b>Interface</b>                 | Interface connected to the ENode.   | <b>detail</b>   |
| • <b>Configured FKA-ADV</b>        | FIP keepalive interval in seconds configured on the FCF multiplied by three. For example, if the FKA_ADV period configured on the FCF is 86 seconds, the value of this field is 258. This value remains constant.     | <b>detail</b>   |
| • <b>Running FKA-ADV</b>           | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF. This value changes every time the ENode sends an FKA_ADV to the FCF.   | <b>detail</b>   |
| • <b>VN-Port MAC</b>               | MAC address of a VN_Port on the ENode.  | All             |
| • <b>FKA-ADV</b>                   | Runtime interval in seconds of the last FIP keepalive advertisement the ENode sent to the FCF on behalf of the VN_Port (VN_Port FKA_ADV). This value changes every time the ENode sends a VN_Port FKA_ADV to the FCF. | <b>detail</b>   |
| • <b>Active VN_Ports</b>           | (QFX Series only)<br>Number of VN_Ports active on an ENode.   | <b>detail</b>   |
| • <b>Vlink far-end VN-Port-MAC</b> | (QFX Series only)<br>Media access control (MAC) address of the VN_Port at the other end of the virtual link.  | <b>detail</b>   |

## Sample Output

### show fip snooping vlan

```

user@switch> show fip snooping vlan fcoevlan1
VLAN : fcoevlan1    FC-MAP : 0e:fc:00
FCF : 00:10:94:00:00:01  Session Count : 2
ENode-MAC : 00:10:94:00:00:02
VN-Port-MAC : 0E:FC:00:00:00:05
VN-Port-MAC : 0E:FC:00:00:00:01

```

### show fip snooping vlan (QFX Series, VN2VF\_Port FIP Snooping)

```

user@switch> show fip snooping vlan fcoevlan1
VLAN : fcoevlan1    Mode: VN2VF Snooping
FC-MAP : 0e:fc:00
FCF : 00:10:94:00:00:01  Session Count : 2
ENode-MAC : 00:10:94:00:00:02
VN-Port-MAC : 0E:FC:00:00:00:05
VN-Port-MAC : 0E:FC:00:00:00:01

```

### show fip snooping vlan (QFX Series, VN2VN\_Port FIP Snooping)

```

user@switch> show fip snooping vlan vlan101

```



```
VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
  Enode-MAC: 10:10:94:01:00:02 VN_Port count: 1
    VN-Port-MAC: 0e:fd:00:00:0a:01 Session Count: 2
  Enode-MAC: 10:10:94:01:00:03 VN_Port count: 0
```

#### show fip snooping vlan detail (QFX Series, VN2VN\_Port FIP Snooping)

```
user@switch> show fip snooping vlan vlan101 detail
VLAN: vlan101, Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
Beacon_Period: 90000
VN2VN Mode: Multi-Point
  Enode Information
    Enode-MAC: 10:10:94:01:00:02, Interface: xe-0/0/10
      Active VN_Ports : 1
    VN_Port Information
      VN-Port MAC: 0e:fd:00:00:0a:01
        Active Sessions : 2
      Session Information
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0b:01
        Vlink far-end VN-Port-MAC: 0e:fd:00:00:0c:01
      Enode-MAC: 10:10:94:01:00:02, Interface: xe-0/0/11
        Active VN_Ports : 0
```

#### show fip snooping vlan detail

```
user@switch> show fip snooping vlan fcoevlan1 detail
VLAN : fcoevlan1 FC-MAP : 0e:fc:00
FCF Information
FCF-MAC : 00:10:94:00:00:01
Active Sessions : 2
Configured FKA-ADV : 258
Running FKA-ADV : 235
  Enode Information
    Enode-MAC : 00:10:94:00:00:02 Interface : xe-0/0/1
    Configured FKA-ADV : 258
    Running FKA-ADV : 239
    Session Information
      VN-Port MAC : 0E:FC:00:00:00:05 FKA-ADV : 255
      VN-Port MAC : 0E:FC:00:00:00:01 FKA-ADV : 251
```

## show fip vlan-discovery

|                                 |   |
|---------------------------------|---|
| <b>Syntax</b>                   | <b>show fip vlan-discovery (enodes   statistics)</b>  |
| <b>Release Information</b>      | Command introduced in Junos OS Release 12.1 for the QFX Series.   |
| <b>Description</b>              | Display FCoE VLAN information from the Fibre Channel switch or FCoE forwarder (FCF).  |
| <b>Options</b>                  | <b>enodes</b> —Display VLAN discovery information for each ENode.<br><b>statistics</b> —Display VLAN discovery information statistics.  |
| <b>Required Privilege Level</b> | view  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">clear fip vlan-discovery statistics on page 5711</a></li> <li>• <i>Understanding FIP Functions</i></li> <li>• <i>Understanding FIP Implementation on an FCoE-FC Gateway</i></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show fip vlan-discovery enodes on page 5769</a><br><a href="#">show fip vlan-discovery statistics (QFX3500) on page 5769</a><br><a href="#">show fip vlan-discovery statistics (QFabric Systems) on page 5769</a>               |
| <b>Output Fields</b>            | Table 471 on page 5768 lists the output fields for the <b>show fip vlan-discovery</b> command. Output fields are listed in the approximate order in which they appear.  |

**Table 471: show fip vlan-discovery Output Fields**

| Field Name                            | Field Description   | Level of Output   |
|---------------------------------------|---|-------------------|
| <b>Enode-MAC</b>                      | Media access control (MAC) address of the ENode.  | <b>enodes</b>     |
| <b>Interface</b>                      | Name of the interface.  | <b>enodes</b>     |
| <b>Unsolicited notification count</b> | Number of unsolicited VLAN discovery notifications.   | All               |
| <b>Solicited notification count</b>   | Number of solicited VLAN discovery notifications.   | <b>statistics</b> |
| <b>Node Group Name</b>                | Displays the name of the Node group on QFabric systems.   | <b>statistics</b> |
| <b>Request count</b>                  | Number of VLAN discovery requests sent by the ENode. This number should match the <b>Solicited notification count</b> number. | <b>statistics</b> |
| <b>VLAN tags</b>                      | Tags of the FIP-enabled VLANs.  | <b>enodes</b>     |

## Sample Output

### show fip vlan-discovery enodes

```
user@switch> show fip vlan-discovery enodes
```

| Enode-MAC         | Interface  | Unsolicited<br>Notification<br>Count | Vlan Tags |
|-------------------|------------|--------------------------------------|-----------|
| 00:10:94:00:00:02 | xe-0/0/9.0 | 0                                    | 400       |

### show fip vlan-discovery statistics (QFX3500)

```
user@switch> show fip vlan-discovery statistics
```

```
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

### show fip vlan-discovery statistics (QFabric Systems)

```
user@switch> show fip vlan-discovery statistics
```

```
NW-NG-0:
```

```
-----
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

```
BBAK0399:
```

```
-----
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```

```
FCC001:
```

```
-----
Request count: 0
Solicited notification count: 0
Unsolicited notification count: 1
```



## CHAPTER 71

# Troubleshooting

- [Troubleshooting Procedures on page 5771](#)

## Troubleshooting Procedures

---

- [Troubleshooting Dropped FCoE Traffic on page 5771](#)
- [Troubleshooting Dropped FIP Traffic on page 5774](#)

## Troubleshooting Dropped FCoE Traffic

**Problem** **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

**Cause** There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.
5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.
6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



.....

**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

.....

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

**Solution** The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 5606 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

#### Related Documentation

- [show class-of-service congestion-notification on page 6305](#)
- [show class-of-service forwarding-class-set on page 6313](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Troubleshooting Dropped FIP Traffic

**Problem**    **Description:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) traffic such as FIP VLAN discovery and notification frames is dropped.

**Cause**        The interface on which the FIP traffic is dropped does not have a native VLAN configured. FIP VLAN discovery and notification messages are exchanged as untagged packets on the native VLAN. (After the FCoE session with the Fibre Channel switch is established, FCoE traffic uses the FCoE VLAN.)



**Solution** Check to ensure that every 10-Gigabit Ethernet interface that connects to an FCoE device includes a native VLAN. Configure a native VLAN on all 10-Gigabit Ethernet interfaces that connect to FCoE devices.



**NOTE:** Make sure that the native VLAN you are using is the same native VLAN that the FCoE devices use for Ethernet traffic.

The procedure for configuring a native VLAN on an interface is different on switches that use the original CLI than on switches that use the Enhanced Layer 2 Software (ELS) CLI. This topic provides the configuration procedure for each CLI.

### Configuring a Native VLAN on Switches Using the Original CLI

To configure a native VLAN on an interface:

1. Set the interface port mode to **tagged-access** if you have not already done so:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching port-mode
tagged-access
```

For example, to set the port mode to **tagged-access** for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode
tagged-access
```

2. Configure the native VLAN if it does not already exist:

```
[edit]
user@switch# set vlans vlan-name vlan-id vlan-id
```

For example, to name the native VLAN **native** and use the VLAN ID 1:

```
[edit]
user@switch# set vlans native vlan-id 1
```

3. Configure the native VLAN on the interface:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching native-vlan-id
vlan-id
```

For example, to configure a native VLAN with the VLAN ID 1 on interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
```

### Configuring a Native VLAN on Switches Using the ELS CLI

To configure a native VLAN on an interface:

1. Set the interface mode to **trunk** if you have not already done so:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching interface-mode
trunk
```

For example, to set the interface mode to **trunk** for interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching interface-mode trunk
```

2. Configure the native VLAN if it does not already exist:

```
[edit]
user@switch# set vlans vlan-name vlan-id vlan-id
```

For example, to name the native VLAN **native** and use the VLAN ID 1:

```
[edit]
user@switch# set vlans native vlan-id 1
```

3. Configure the native VLAN on the physical Ethernet interface:

```
[edit]
user@switch# set interfaces interface native-vlan-id vlan-id
```

For example, to configure a native VLAN with the VLAN ID 1 on interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 native-vlan-id 1
```

4. Configure the Ethernet interface as a member of the native VLAN:

```
[edit]
user@switch# set interfaces interface unit unit family ethernet-switching vlan members
vlan-name
```

For example, to configure an Ethernet interface as a member of a native VLAN with the VLAN ID 1 on interface **xe-0/0/6.0**:

```
[edit]
user@switch# set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members native
```

#### Related Documentation

- [interfaces on page 2686](#)
- [vlans](#)
- [Understanding FIP Functions](#)
- [Configuring VLANs for FCoE Traffic on an FCoE Transit Switch on page 5657](#)

## PART 20

# Traffic Management

- [Overview on page 5779](#)
- [Configuration on page 5965](#)
- [Administration on page 6289](#)
- [Troubleshooting on page 6451](#)



## CHAPTER 72

# Overview

- [CoS Overview on page 5779](#)
- [QFX5100 Switches Only on page 5950](#)
- [QFX3500 and QFX3600 Virtual Chassis Only on page 5952](#)
- [Virtual Chassis Fabric Only on page 5957](#)
- [Learn About Technology on page 5963](#)

## CoS Overview

---

- [Overview of Junos OS CoS for the QFX Series and EX4600 Switch on page 5781](#)
- [Overview of Policers on page 5783](#)
- [Understanding Junos CoS Components on page 5789](#)
- [Understanding CoS Packet Flow on page 5793](#)
- [CoS Inputs and Outputs Overview on page 5795](#)
- [Understanding Default CoS Settings on page 5796](#)
- [Understanding Host Inbound Traffic Classification on page 5805](#)
- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806](#)
- [Understanding CoS Code-Point Aliases on page 5808](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 5817](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5835](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)
- [Understanding Default CoS Scheduling and Classification on page 5856](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Priority Group Scheduling on page 5877](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)

- [Understanding CoS Priority Group and Queue Guaranteed Rates \(Minimum Bandwidth\) on page 5881](#)
- [Understanding CoS Priority Group Shaping and Queue Shaping \(Maximum Bandwidth\) on page 5884](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5886](#)
- [Understanding CoS Buffer Configuration on page 5891](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5916](#)
- [Understanding CoS Explicit Congestion Notification on page 5926](#)
- [Understanding DCB Features and Requirements on page 5934](#)
- [Understanding DCBX on page 5937](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5946](#)

## Overview of Junos OS CoS for the QFX Series and EX4600 Switch

When a network experiences congestion and delay, some packets must be dropped. Junos OS class of service (CoS) enables you to divide traffic into classes and set various levels of throughput and packet loss when congestion occurs. You have greater control over packet loss because you can configure rules tailored to your needs.

You can configure CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP) or IEEE 802.1p code-point bits of packets leaving an interface, thus allowing you to tailor packets for the network requirements of the remote peers.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

In designing CoS applications, you must carefully consider your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because CoS is implemented in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet forwarding and switching performance.



**NOTE:** CoS policies can be enabled or disabled on each switch interface. Also, each physical and logical interface on the switch can have associated custom CoS rules.

When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

This topic describes:

- [CoS Standards on page 5781](#)
- [How Junos CoS Works on page 5782](#)
- [Default CoS Behavior on page 5783](#)

### CoS Standards

The following RFCs define the standards for CoS capabilities:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*

- RFC 2698, *A Two Rate Three Color Marker*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

The following data center bridging (DCB) standards are also supported to provide the CoS (and other characteristics) Fibre Channel requires for transmitting storage traffic over an Ethernet network:

- IEEE 802.1Qbb, priority-based flow control (PFC)
- IEEE 802.1Qaz, enhanced transmission selection (ETS)
- IEEE 802.1AB (LLDP) extension called Data Center Bridging Capability Exchange Protocol (DCBX)

### How Junos CoS Works

---

Junos CoS works by examining traffic entering at the edge of your network. The switch classifies traffic into defined service groups to provide the special treatment of traffic across the network. For example, you can send voice traffic across certain links and data traffic across other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic to meet the policies of the targeted peer by rewriting the DSCP or IEEE 802.1 code-point bits.


To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings dictate which packets are transmitted first to the next downstream switch. Switches at the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

In [Figure 200 on page 5783](#), Switch A is receiving traffic. As each packet enters, Switch A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined on the switch. This definition allows Switch A to prioritize its resources for servicing the traffic streams it receives. Switch A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the defined traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packet according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can reclassify (rewrite) the CoS code-point bits of the packets before transmitting them.



Figure 200: Packet Flow Across the Network



### Default CoS Behavior

If you do not configure CoS settings, the software performs some CoS functions to ensure that the system forwards traffic and protocol packets with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as rewrite rules, are applied only if you explicitly associate them with an interface.

#### Related Documentation

- [Overview of Policers on page 5241](#)
- [Understanding Junos CoS Components on page 5789](#)
- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding Default CoS Settings on page 5796](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)

### Overview of Policers

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.


- [Policer Overview on page 5784](#)
- [Policer Types on page 5784](#)
- [Policer Actions on page 5785](#)
- [Policer Colors on page 5786](#)
- [Filter-Specific Policers on page 5786](#)
- [Suggested Naming Convention for Policers on page 5787](#)
- [Policer Counters on page 5787](#)
- [Policer Algorithms on page 5787](#)
- [How Many Policers are Supported? on page 5787](#)
- [Policers can Limit Egress Firewall Filters on page 5788](#)

## Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 181 on page 5242](#) illustrates this process.

**Figure 201: Flow of Tricolor Marking Policer Operation**



After you name and configure a policer, you use it by specifying it as an action in one or more firewall filters.

## Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- Two-rate three-color marker—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



**NOTE:** A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 432 on page 5243](#) for information about how metering results are applied for each of these policer types.

### Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 432 on page 5243](#) describes the policer actions.

**Table 472: Policer Actions**

| Policer                 | Marking                        | Implicit Action                  | Configurable Action |
|-------------------------|--------------------------------|----------------------------------|---------------------|
| Single-rate two-color   | Green (conforming)             | Assign low loss priority         | None                |
|                         | Red (nonconforming)            | None                             | Discard             |
| Single-rate three-color | Green (conforming)             | Assign low loss priority         | None                |
|                         | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                         | Red (above the EBS)            | Assign high loss priority        | Discard             |

Table 472: Policer Actions (*continued*)

| Policer              | Marking                        | Implicit Action                  | Configurable Action |
|----------------------|--------------------------------|----------------------------------|---------------------|
| Two-rate three-color | Green (conforming)             | Assign low loss priority         | None                |
|                      | Yellow (above the CIR and CBS) | Assign medium-high loss priority | None                |
|                      | Red (above the PIR and PBS)    | Assign high loss priority        | Discard             |



**NOTE:** If you specify a policer in an egress firewall filter, the only supported action is **discard**.

### Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

### Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5236](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

---

### Suggested Naming Convention for Policers

---

We recommend that you use the naming convention ***policertypeTCM#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named ***srTCM1-ca***. The second two-rate, color-blind three-color configured would be named ***trTCM2-cb***. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

---

### Policer Counters

---

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

---

### Policer Algorithms

---

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.

---

### How Many Policers are Supported?

---

You can configure and commit the following numbers of policers on QFX3500 and QFX3600 standalone switches and QFabric Node devices:

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

### Policers can Limit Egress Firewall Filters

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

#### **Related Documentation**

- [Understanding Color-Blind Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Blind Mode for Two-Rate Tricolor Marking on page 5249](#)
- [Understanding Color-Aware Mode for Single-Rate Tricolor Marking on page 5247](#)
- [Understanding Color-Aware Mode for Two-Rate Tricolor Marking on page 5249](#)

- [Configuring Two-Color and Three-Color Policers to Control Traffic Rates on page 5296](#)

## Understanding Junos CoS Components

This topic describes the Junos operating system (OS) class-of-service (CoS) components:

- [Code-Point Aliases on page 5789](#)
- [Policers on page 5789](#)
- [Classifiers on page 5789](#)
- [Forwarding Classes on page 5790](#)
- [Forwarding Class Sets on page 5790](#)
- [Flow Control \(Ethernet PAUSE, PFC, and ECN\) on page 5790](#)
- [WRED Profiles on page 5791](#)
- [Schedulers on page 5791](#)
- [Rewrite Rules on page 5792](#)

### Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.

### Policers

Policers limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with filters that you can associate with input interfaces.

### Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Junos OS, *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate (BA) or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value or IEEE 802.1p value.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

You can create unicast classifiers for unicast traffic and multideestination classifiers for multicast, broadcast, and destination lookup fail traffic. You cannot assign unicast traffic and multideestination traffic to the same classifier.

You can apply unicast classifiers to one or more interfaces. Multidestination classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces.

### Forwarding Classes

---

Forwarding classes group packets for transmission and CoS. You assign each packet to an output queue based on the packet's forwarding class. Forwarding classes affect the forwarding, scheduling, and rewrite marking policies applied to packets as they transit the switch.

The switch provides five default forwarding classes:

- fcoe—Fibre Channel over Ethernet traffic
- no-loss—Lossless traffic
- be—Best-effort traffic
- nc—Network control traffic
- mcast—Multicast traffic

The switch supports a total of 12 forwarding classes (8 unicast forwarding classes and 4 multicast forwarding classes), which provide flexibility in classifying traffic.

### Forwarding Class Sets

---

You can group forwarding classes (output queues) into *forwarding class sets* in order to apply CoS to groups of traffic that require similar treatment. Forwarding class sets map traffic into priority groups to support enhanced transmission selection (ETS, described in IEEE 802.1Qaz).

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. For example, you can configure different forwarding class sets to apply CoS to unicast groups of local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and configure another group for multicast traffic.

Within each forwarding class set, you can configure special CoS treatment for the traffic mapped to each individual queue. This provides the ability to configure CoS in a two-tier hierarchical manner. At the forwarding class set tier, you configure CoS for groups of traffic using a *traffic control profile*. At the queue tier, you configure CoS for individual output queues within a forwarding class set using a *scheduler* that you map to a queue (forwarding class) using a *scheduler map*.

### Flow Control (Ethernet PAUSE, PFC, and ECN)

---

Ethernet PAUSE (described in IEEE 802.3X) is a link-level flow control mechanism. During periods of network congestion, Ethernet PAUSE stops all traffic on a full-duplex Ethernet link for a period of time specified in the PAUSE message.

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is part of the IEEE data center bridging (DCB) specifications for creating a lossless Ethernet environment to transport loss-sensitive flows such as Fibre Channel over Ethernet (FCoE) traffic.



PFC is a link-level flow control mechanism similar to Ethernet PAUSE. However, Ethernet PAUSE stops all traffic on a link for a period of time. PFC decouples the pause function from the physical link and divides the traffic on the link into eight priorities (3-bit IEEE 802.1p code points). You can think of the eight priorities as eight “lanes” of traffic. You can apply pause selectively to the traffic on any priority without pausing the traffic on other priorities on the same link.

The granularity that PFC provides allows you to configure different levels of CoS for different types of traffic on the link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop methods of congestion management for IP traffic on the same link.



**NOTE:** If you transport FCoE traffic, you must enable PFC on the priority assigned to FCoE traffic (usually IEEE 802.1p code point 011 on interfaces that carry FCoE traffic).

Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

### WRED Profiles

A WRED (weighted random early detection) profile (drop profile) defines parameters that enable the network to drop packets during periods of congestion. A drop profile defines the conditions under which packets of different loss priorities drop, by determining the probability of dropping a packet for each loss priority when output queues become congested. Drop profiles essentially set a value for a level of queue fullness—when the queue fills to the level of the queue fullness value, packets drop.

You can associate different drop profiles with different loss priorities to set the probability of dropping packets. You can apply a drop profile for each loss priority to a forwarding class (output queue) by applying a drop profile to a scheduler, and then mapping the scheduler to a forwarding class using a scheduler map. When the queue mapped to the forwarding class experiences congestion, the drop profile determines the level of packet drop for traffic of each loss priority in that queue.

Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. Typically you mark packets exceeding a particular service level with a high loss priority.

### Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This

process often involves determining the sequence in which different types of packets should be transmitted.

You can define the priority (**priority**), minimum bandwidth (**transmit-rate**), maximum bandwidth (**shaping-rate**), and WRED profiles to be applied to a particular queue for packet transmission. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

### Rewrite Rules

---

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream device to classify the packet into the appropriate service group. Rewriting (marking) outbound packets is useful when the switch is at the border of a network and must change the CoS values to meet the policies of the targeted peer.



**NOTE:** Ingress firewall filters can also rewrite forwarding class and loss priority values.

---

#### Related Documentation

- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding CoS Code-Point Aliases on page 5808](#)
- [Overview of Policers on page 5241](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5835](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS Explicit Congestion Notification on page 5926](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding DCB Features and Requirements on page 5515](#)

## Understanding CoS Packet Flow

When a packet traverses a switch, the switch provides the appropriate level of service to the packet using either default class of service (CoS) settings or CoS settings that you configure. On ingress ports, the switch classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the switch applies packet scheduling and (if you have configured them) rewrite rules to re-mark packets.

You can configure CoS on Layer 2 logical interfaces, and you can configure CoS on Layer 3 physical interfaces if you have defined at least one logical interface on the Layer 3 physical interface. You cannot configure CoS on Layer 2 physical interfaces and Layer 3 logical interfaces.

For Layer 2 traffic, either use the default CoS settings or configure CoS on each logical interface. You can apply different CoS settings to different Layer 2 logical interfaces.

For Layer 3 traffic, either use the default CoS settings or configure CoS on the physical interface (not on the logical unit). The switch uses the CoS applied on the physical Layer 3 interface for all logical Layer 3 interfaces configured on the physical Layer 3 interface.

The switch applies to CoS to packets as they flow through the system:

- An interface has one or more classifiers of different types applied to it (configure this at the **[edit class-of-service interfaces]** hierarchy level). The classifier types are based on the portion of the incoming packet that the classifier examines (IEEE 802.1p code point bits or DSCP code point bits).
- When a packet enters an ingress port, the classifier assigns the packet to a forwarding class and a loss priority based on the code point bits of the packet (configure this at the **[edit class-of-service classifiers]** hierarchy level).
- The switch assigns each forwarding class to an output queue (configure this at the **[edit class-of-service forwarding-classes]** hierarchy level).
- Input (and output) policers meter traffic and can change the forwarding class and loss priority if a traffic flow exceeds its service level.
- A scheduler map is applied to each interface. When a packet exits an egress port, the scheduler map controls how it is treated (configure this at the **[edit class-of-service interfaces]** hierarchy level). A scheduler map assigns schedulers to forwarding classes (configure this at the **[edit class-of-service scheduler-maps]** hierarchy level).
- A scheduler defines how traffic is treated at the egress interface output queue (configure this at the **[edit class-of-service schedulers]** hierarchy level). You control the transmit rate, shaping rate, priority, and drop profile of each forwarding class by mapping schedulers to forwarding classes in scheduler maps, then applying scheduler maps to interfaces.
- A drop-profile defines how aggressively to drop packets that are mapped to a particular scheduler (configure this at the **[edit class-of-service drop-profiles]** hierarchy level).
- A rewrite rule takes effect as the packet leaves an interface that has a rewrite rule configured (configure this at the **[edit class-of-service rewrite-rules]** hierarchy level).

The rewrite rule writes information to the packet (for example, a rewrite rule can re-mark the code point bits of outgoing traffic) according to the forwarding class and loss priority of the packet.

Figure 202 on page 5794 is a high-level flow diagram of how packets from various sources enter switch interfaces, are classified at the ingress, and then scheduled (provided bandwidth) at the egress queues.

**Figure 202: CoS Classifier, Queues, and Scheduler**


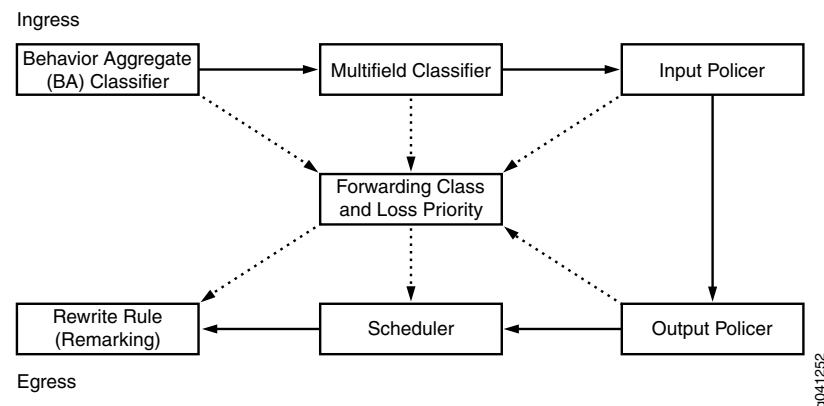


Figure 203 on page 5794 shows the packet flow through the CoS components that you can configure.

**Figure 203: Packet Flow Through Configurable CoS Components**



The middle box ("Forwarding Class and Loss Priority") represents two values that you can use on ingress and egress interfaces. The system uses these values for classifying traffic on ingress interfaces and for rewrite rule re-marking on egress interfaces. Each outer box represents a process component. The components in the top row apply to incoming packets. The components in the bottom row apply to outgoing packets.

The solid-line arrows show the direction of packet flow from ingress to egress. The dotted-line arrows show inputs and outputs or show settings and actions based on those settings.

For example, the BA classifier sets the forwarding class and loss priority of incoming packets, so the forwarding class and loss priority are outputs of the classifier and the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packets based on those settings, so the arrow points toward the scheduler.

#### Related Documentation

- [Understanding CoS Classifiers on page 5810](#)
- [Overview of Policers on page 5241](#)
- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding Default CoS Scheduling and Classification on page 5856](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Priority Group Scheduling on page 5877](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)
- [Understanding CoS Rewrite Rules on page 5914](#)

## CoS Inputs and Outputs Overview

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in [Table 473 on page 5795](#).

**Table 473: CoS Mappings—Inputs and Outputs**

| CoS Mappings   | Inputs  | Outputs   | Comments   |
|--|---|---|--|
| <a href="#">classifiers</a>                              | <a href="#">code-points</a>   | <a href="#">forwarding-class</a> ,<br><a href="#">loss-priority</a> | The map sets the forwarding class and packet loss priority (PLP) for a specific set of code points.  |
| <a href="#">drop-profile-map</a>                         | <a href="#">loss-priority</a> , <a href="#">protocol</a>            | <a href="#">drop-profile</a>  | The map sets the drop profile for a specific PLP and protocol type.  |
| <a href="#">rewrite-rules</a>                            | <a href="#">loss-priority</a> ,<br><a href="#">forwarding-class</a> | <a href="#">code-points</a>   | The map sets the code points for a specific forwarding class and PLP.  |
| <a href="#">rewrite-value (Fibre Channel Interfaces)</a> | <a href="#">forwarding-class</a>                                    | <a href="#">code-point</a>  | The map sets the code point for the forwarding class specified in the fixed classifier attached to the native Fibre Channel (NP_Port) interface. |

#### Related Documentation

- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Example: Configuring Forwarding Classes on page 6075](#)

- [Configuring CoS Drop Profile Maps on page 6164](#)
- [Defining CoS Rewrite Rules on page 6182](#)

## Understanding Default CoS Settings

If you do not configure CoS settings, Junos OS performs some CoS functions to ensure that traffic and protocol packets are forwarded with minimum delay when the network experiences congestion. Some default mappings are automatically applied to each logical interface that you configure.

You can display default CoS settings by issuing the **show class-of-service** operational mode command.

This topic describes the default configurations for the following CoS components:

- [Default Forwarding Classes and Queue Mapping on page 5796](#)
- [Default Forwarding Class Sets \(Priority Groups\) on page 5797](#)
- [Default Code-Point Aliases on page 5797](#)
- [Default Classifiers on page 5799](#)
- [Default Rewrite Rules on page 5801](#)
- [Default Drop Profile on page 5801](#)
- [Default Schedulers on page 5802](#)
- [Default Scheduler Maps on page 5804](#)
- [Default Shared Buffer Configuration on page 5804](#)

### Default Forwarding Classes and Queue Mapping

Table 474 on page 5796 shows the default mapping of the default forwarding classes to queues and packet drop attribute.

**Table 474: Default Forwarding Classes and Queue Mapping**

| Default Forwarding Class | Description  | Default Queue Mapping | Packet Drop Attribute |
|--------------------------|--|-----------------------|-----------------------|
| best-effort (be)         | Best-effort traffic class (priority 0, IEEE 802.1p code point 000)                   | 0                     | drop                  |
| fcoe                     | Guaranteed delivery for FCoE traffic (priority 3, IEEE 802.1p code point 011)        | 3                     | no-loss               |
| no-loss                  | Guaranteed delivery for TCP no-loss traffic (priority 4, IEEE 802.1p code point 100) | 4                     | no-loss               |
| network-control (nc)     | Network control traffic (priority 7, IEEE 802.1p code point 111)                     | 7                     | drop                  |

Table 474: Default Forwarding Classes and Queue Mapping (*continued*)

| Default Forwarding Class | Description              | Default Queue Mapping | Packet Drop Attribute |
|--------------------------|--------------------------|-----------------------|-----------------------|
| mcast                    | Multidestination traffic | 8                     | drop                  |

**NOTE:** You cannot configure multidestination forwarding classes as no-loss (lossless) traffic classes.

### Default Forwarding Class Sets (Priority Groups)

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set.

Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth.

The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.

### Default Code-Point Aliases

Table 475 on page 5797 shows the default mapping of code-point aliases to IEEE code points.

Table 475: Default IEEE 802.1 Code-Point Aliases

| CoS Value Types | Mapping |
|-----------------|---------|
| be              | 000     |
| be1             | 001     |
| ef              | 010     |
| ef1             | 011     |
| af11            | 100     |
| af12            | 101     |
| nc1             | 110     |
| nc2             | 111     |

Table 476 on page 5798 shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

**Table 476: Default DSCP and DCSP IPv6 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| ef              | 101110  |
| af11            | 001010  |
| af12            | 001100  |
| af13            | 001110  |
| af21            | 010010  |
| af22            | 010100  |
| af23            | 010110  |
| af31            | 011010  |
| af32            | 011100  |
| af33            | 011110  |
| af41            | 100010  |
| af42            | 100100  |
| af43            | 100110  |
| be              | 000000  |
| cs1             | 001000  |
| cs2             | 010000  |
| cs3             | 011000  |
| cs4             | 100000  |
| cs5             | 101000  |
| nc1             | 110000  |
| nc2             | 111000  |



### Default Classifiers

The switch applies default unicast IEEE 802.1, unicast DSCP, and multidestination classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

[Table 477 on page 5799](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

**Table 477: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

[Table 478 on page 5799](#) shows the default mapping of IEEE 802.1p code-point values to unicast forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

**Table 478: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |

**Table 478: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier) (continued)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

[Table 479 on page 5800](#) shows the default mapping of IEEE 802.1 code-point values to multdestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

**Table 479: Default IEEE 802.1 Multidestination Classifiers**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | mcast            | low           |
| be1 (001)  | mcast            | low           |
| ef (010)   | mcast            | low           |
| ef1 (011)  | mcast            | low           |
| af11 (100) | mcast            | low           |
| af12 (101) | mcast            | low           |
| nc1 (110)  | mcast            | low           |
| nc2 (111)  | mcast            | low           |

[Table 480 on page 5800](#) shows the default mapping of DSCP code-point values to unicast forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

**Table 480: Default DSCP IP and IPv6 Unicast Classifiers**

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| ef (101110)   | best-effort      | low           |
| af11 (001010) | best-effort      | low           |
| af12 (001100) | best-effort      | low           |
| af13 (001110) | best-effort      | low           |

Table 480: Default DSCP IP and IPv6 Unicast Classifiers (*continued*)

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| af21 (010010) | best-effort      | low           |
| af22 (010100) | best-effort      | low           |
| af23 (010110) | best-effort      | low           |
| af31 (011010) | best-effort      | low           |
| af32 (011100) | best-effort      | low           |
| af33 (011110) | best-effort      | low           |
| af41 (100010) | best-effort      | low           |
| af42 (100100) | best-effort      | low           |
| af43 (100110) | best-effort      | low           |
| be (000000)   | best-effort      | low           |
| cs1 (001000)  | best-effort      | low           |
| cs2 (010000)  | best-effort      | low           |
| cs3 (011000)  | best-effort      | low           |
| cs4 (100000)  | best-effort      | low           |
| cs5 (101000)  | best-effort      | low           |
| nc1 (110000)  | network-control  | low           |
| nc2 (111000)  | network-control  | low           |



**NOTE:** There are no default DSCP IP or IPv6 classifiers for multidestination traffic. DSCP IPv6 classifiers are not supported for multidestination traffic.

### Default Rewrite Rules

There are no default rewrite rules. If you do not explicitly configure rewrite rules, the switch does not reclassify egress traffic.

### Default Drop Profile

Table 481 on page 5802 shows the default drop profile configuration.

Table 481: Default Drop Profile

| Fill Level | Drop Probability |
|------------|------------------|
| 100        | 100              |

### Default Schedulers

Table 482 on page 5802 shows the default scheduler configuration.

Table 482: Default Schedulers

| Default Scheduler and Queue Number   | Transmit Rate (Guaranteed Minimum Bandwidth) | Shaping Rate (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|--------------------------------------|--|----------------------------------|--------------------------|----------|-------------|
| Best-effort scheduler (queue 0)      | 5%   | None                             | 5%                       | low      | 5%          |
| FCoE scheduler (queue 3)             | 35%  | None                             | 35%                      | low      | 35%         |
| No-loss scheduler (queue 4)          | 35%  | None                             | 35%                      | low      | 35%         |
| Network-control scheduler (queue 7)  | 5%   | None                             | 5%                       | low      | 5%          |
| Multidestination scheduler (queue 8) | 20%  | None                             | 20%                      | low      | 20%         |



**NOTE:** The minimum guaranteed bandwidth (transmit rate) also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue.

By default, only the five default schedulers shown in [Table 482 on page 5802](#) have traffic mapped to them. Only the queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

Default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multidestination traffic. By default, unicast traffic consists of queue 0 (**best-effort** forwarding class), queue 3 (**fcoe** forwarding class), queue 4 (**no-loss** forwarding class), and queue 7 (**network-control** forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multidestination traffic (**mcast** queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multidestination traffic receives 2-Gbps of bandwidth.



**NOTE:** Multidestination queue 11 also receives a small amount of default bandwidth from the multidestination scheduler. CPU-generated multidestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multidestination traffic to a unicast forwarding class, that traffic uses queue 11.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

You should explicitly map traffic to non-default (unconfigured) queues and create schedulers to allocate bandwidth to those queues if you want to use them to forward traffic. By default, unicast queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)



**NOTE:** All four multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

If the other unicast queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher

priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

### Default Scheduler Maps

Table 483 on page 5804 shows the default mapping of forwarding classes to schedulers.

**Table 483: Default Scheduler Maps**

| Forwarding Class | Scheduler                          |
|------------------|------------------------------------|
| best-effort      | Default BE scheduler               |
| fcoe             | Default FCoE scheduler             |
| no-loss          | No-loss scheduler                  |
| network-control  | Default network-control scheduler  |
| mcast-be         | Default multidestination scheduler |

### Default Shared Buffer Configuration

Table 484 on page 5804 and Table 485 on page 5804 show the default shared buffer allocations:

**Table 484: Default Ingress Shared Buffer Configuration**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 100%                        | 9%              | 45%                      | 46%          |

**Table 485: Default Egress Shared Buffer Configuration**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 50%             | 31%          | 19%              |

#### Related Documentation

- [Overview of Junos OS CoS for the QFX Series and EX4600 Switch on page 5781](#)
- [Understanding Junos CoS Components on page 5789](#)
- [Understanding Default CoS Scheduling and Classification on page 5856](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Understanding CoS Code-Point Aliases on page 5808](#)

- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)

## Understanding Host Inbound Traffic Classification

The destination address of traffic that enters the switch can be an external device such as another switch, a router, or a server, or the destination can be the host (the switch Routing Engine or CPU). When the destination is an external device, the DSCP and IEEE 802.1p code-point bits of incoming traffic are preserved as the traffic travels through the switch to the egress port. At the egress port, the code-point bits are either preserved when the packets are sent to the next hop or they are rewritten according to the rewrite rule attached to the egress interface.

When the destination of incoming traffic is the host, DSCP bits are preserved. However, IEEE 802.1p bits are not preserved. The IEEE 802.1p bits of traffic destined for the host are set to zero (0). This does not affect system behavior because the switch prioritizes traffic destined for the host based on the protocol type. For example, the switch gives a higher priority to BPDU traffic than to ping traffic.

### Related Documentation

- [Understanding Default CoS Scheduling and Classification on page 5856](#)
- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806](#)

## Understanding Host Routing Engine Outbound Traffic Queues and Defaults

The host Routing Engine and CPU generate outbound traffic that is transmitted using different protocols. You cannot configure a classifier to map different types of outbound traffic that the host generates to forwarding classes (queues). The traffic that the host generates is assigned to forwarding classes by default as shown in [Table 486 on page 5806](#).

If you want to separate host outbound traffic from other traffic or if you want to assign that traffic to a particular queue, you can configure a single forwarding class for all traffic that the host generates. If you configure a forwarding class for outbound host traffic, that forwarding class is used globally for all traffic generated by the host. (That is, the host outbound traffic is mapped to the selected queue on all egress interfaces.) Configuring a forwarding class for host outbound traffic does not affect transit or incoming traffic.

Whether you use the default host outbound traffic forwarding class configuration or configure a forwarding class for all host outbound traffic, the configuration applies to all Layer 2 and Layer 3 protocols and to all application-level traffic such as FTP and ping operations.

If you configure a queue for host outbound traffic, the queue must be properly configured on all interfaces.



**NOTE:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) packets generated by the CPU are always transmitted on the `fcoe` queue (queue 3), even if you configure a queue for host outbound traffic. This helps to ensure lossless behavior for FCoE traffic. QFabric systems classify FIP control packets into the same traffic class (`fcoe`) across the Interconnect device (`fabric`) and the egress Node device.

By default, traffic generated by the host is sent to the best effort queue (queue 0) or to the network control queue (queue 7). [Table 486 on page 5806](#) lists the default host traffic to output queue mapping.

**Table 486: Routing Engine Protocol Default Queue Mapping**

| Routing Engine Protocol  | Default Queue Mapping |
|--|-----------------------|
| Address Resolution Protocol (ARP) reply                          | Queue 0               |
| ARP request  | Queue 0               |
| Border Gateway Protocol (BGP)                                    | Queue 0               |
| BGP TCP Retransmission   | Queue 7               |
| Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) | Queue 3               |
| File Transfer Protocol (FTP)                                     | Queue 0               |



Table 486: Routing Engine Protocol Default Queue Mapping (*continued*)

| Routing Engine Protocol                         | Default Queue Mapping |
|---|-----------------------|
| Internet Control Message Protocol (ICMP) reply  | Queue 0               |
| ICMP request                                    | Queue 0               |
| Internet Group Management Protocol (IGMP) query | Queue 7               |
| IGMP report                                     | Queue 0               |
| Link Aggregation Control Protocol (LACP)        | Queue 7               |
| Open Shortest Path First (OSPF) hello           | Queue 7               |
| OSPF protocol data unit (PDU)                   | Queue 7               |
| OSPF link state advertisements (LSAs)           | Queue 7               |
| Protocol Independent Multicast (PIM)            | Queue 7               |
| PIM hello                                       | Queue 7               |
| Simple Network Management Protocol (SNMP)       | Queue 0               |
| Secure Shell (SSH)                              | Queue 0               |
| Telnet  | Queue 0               |
| Virtual Router Redundancy Protocol (VRRP)       | Queue 7               |
| VLAN Spanning Tree Protocol (VSTP)              | Queue 7               |
| <code>xnm-clear-text</code>                     | Queue 0               |
| <code>xnm-ssl</code>                            | Queue 0               |

- Related Documentation**
- [Understanding CoS Forwarding Classes on page 5830](#)
  - [Changing the Host Outbound Traffic Default Queue Mapping on page 6172](#)
  - [Example: Configuring Forwarding Classes on page 6075](#)

## Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs) and IEEE 802.1 bits to associate incoming packets with a particular CoS servicing level. You can assign a meaningful name or alias to the CoS values and use that alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure alias names for user-defined classifiers. If the value of an alias changes, it alters the behavior of any classifier that references it.

You can configure code-point aliases for the following type of CoS markers:

- dscp or dscp-ipv6—Handles incoming IP and IPv6 packets.
- ieee-802.1—Handles Layer 2 CoS.

This topic covers:

- [Default Code-Point Aliases on page 5808](#)

### Default Code-Point Aliases

[Table 487 on page 5808](#) shows the default mapping of code-point aliases to IEEE code points.

**Table 487: Default IEEE 802.1 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| be              | 000     |
| be1             | 001     |
| ef              | 010     |
| ef1             | 011     |
| af11            | 100     |
| af12            | 101     |
| nc1             | 110     |
| nc2             | 111     |

[Table 488 on page 5809](#) shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

**Table 488: Default DSCP and DSCP IPv6 Code-Point Aliases**

| CoS Value Types | Mapping |
|-----------------|---------|
| ef              | 101110  |
| af11            | 001010  |
| af12            | 001100  |
| af13            | 001110  |
| af21            | 010010  |
| af22            | 010100  |
| af23            | 010110  |
| af31            | 011010  |
| af32            | 011100  |
| af33            | 011110  |
| af41            | 100010  |
| af42            | 100100  |
| af43            | 100110  |
| be              | 000000  |
| cs1             | 001000  |
| cs2             | 010000  |
| cs3             | 011000  |
| cs4             | 100000  |
| cs5             | 101000  |
| nc1             | 110000  |
| nc2             | 111000  |

**Related  
Documentation**

- [Understanding Junos CoS Components on page 5789](#)

- [Defining CoS Code-Point Aliases on page 6159](#)

## Understanding CoS Classifiers

Packet classification associates incoming packets with a particular class-of-service (CoS) servicing level. Classifiers associate packets with a forwarding class and a loss priority, and assign packets to output queues based on the associated forwarding class. There are three general types of classifiers:

- Behavior aggregate (BA) classifiers—DSCP and DSCP IPv6 classify IP and IPv6 traffic, EXP classifies MPLS traffic, and IEEE 802.1p classifiers classify all other traffic. (Although this topic covers EXP classifiers, for more details about EXP classifiers, see [“Understanding CoS MPLS EXP Classifiers and Rewrite Rules” on page 4419](#). EXP classifiers are applied only on **family mpls** interfaces.)
  - Fixed classifiers—Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
  - Multifield (MF) classifiers—MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.
- [Interfaces and Output Queues on page 5810](#)
  - [Behavior Aggregate Classifiers on page 5811](#)
  - [Fixed Classifiers on Ethernet Interfaces on page 5814](#)
  - [Fixed Classifiers on Native Fibre Channel Interfaces \(NP\\_Ports\) on page 5815](#)
  - [Multifield Classifiers on page 5815](#)
  - [Packet Classification for Routed VLAN Interfaces \(RVIs\) on page 5816](#)

### Interfaces and Output Queues

---

On Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and link aggregation (LAG) interfaces, you can apply classifiers to Layer 2 logical interfaces and to Layer 3 physical interfaces if the Layer 3 physical interface has at least one defined logical interface. Classifiers applied to Layer 3 physical interfaces are used on all logical interfaces on that physical interface. [“Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces” on page 5820](#) describes the interaction between classifiers and interfaces in greater detail.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification is performed. If the two classification results conflict, the MF classification result overrides the BA classification result.

You cannot configure a fixed classifier and a BA classifier on the same interface.

You can configure both a DSCP or a DSCP IPv6 classifier and an IEEE 802.1p classifier on the same interface. IP traffic uses the DSCP or DSCP IPv6 classifier. All other traffic uses the IEEE classifier (except when you configure a global EXP classifier; in that case, MPLS traffic uses the EXP classifier providing that the interface is configured as **family**

**mpls**). You can configure only one DSCP classifier on a physical interface (either one DSCP classifier or one DSCP IPv6 classifier, but not both).

Although you can configure as many EXP classifiers as you want, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All switch interfaces that are configured as **family mpls** use the EXP classifier specified using this configuration statement to classify MPLS traffic.

You can create unicast BA classifiers for unicast traffic and multicast BA classifiers for multdestination traffic, which includes multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot assign unicast traffic and multdestination traffic to the same BA classifier.

On each interface, the switch has separate output queues for unicast traffic and for multdestination traffic:

- The switch supports 12 output queues, with 8 queues dedicated to unicast traffic and 4 queues dedicated to multdestination traffic.
- Queues 0 through 7 are unicast traffic queues. You can apply only unicast BA classifiers to unicast queues. A unicast BA classifier should contain only forwarding classes that are mapped to unicast queues.
- Queues 8 through 11 are multdestination traffic queues. You can apply only multdestination BA classifiers to multdestination queues. A multdestination BA classifier should contain only forwarding classes that are mapped to multdestination queues.

You can apply unicast classifiers to one or more interfaces. Multdestination classifiers and EXP classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces. Use the DSCP multdestination classifier for both IP and IPv6 multdestination traffic. The DSCP IPv6 classifier is not supported for multdestination traffic.

### Behavior Aggregate Classifiers

The behavior aggregate classifier maps a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

The switch supports three types of BA classifiers:

- Differentiated Services Code Point (DSCP) for IP DiffServ (IP and IPv6)
- IEEE 802.1p CoS bits
- MPLS EXP (applies only to interfaces configured as **family mpls**)

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore, core devices, which handle high traffic volumes, are normally configured to perform BA classification.

Unicast and multicast traffic cannot share the same classifier. You can map unicast traffic and multicast traffic to the same classifier CoS value, but the unicast traffic must belong to a unicast classifier and the multicast traffic must belong to a multidestination classifier.

### Default Behavior Aggregate Classification

Juniper Networks Junos OS automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. [Table 489 on page 5812](#) lists different types of interfaces and the corresponding implicit default BA classifiers.

**Table 489: Default BA Classification**

| Type of Interface  | Default BA Classification |
|--|---------------------------|
| Layer 2 interface in trunk mode or in tagged-access mode | ieee8021p-default         |
| Layer 3 interface  | dscp-default              |
| Layer 2 interface in access mode                         | ieee8021p-untrusted       |



**NOTE:** There are default BA classifiers for the best-effort, fcoe, no-loss, network-control, and mcast forwarding classes.



**NOTE:** There is no default MPLS EXP classifier. You must configure an EXP classifier and apply it globally to all interfaces that are configured as family mpls by including it in the [edit class-of-service system-defaults classifiers exp] hierarchy. On family mpls interfaces, if a fixed classifier is present on the interface, the EXP classifier overrides the fixed classifier.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier, even interfaces that have a fixed classifier.

When you explicitly associate a unicast classifier with a logical interface, you override the default unicast classifier with the explicit unicast classifier.



**NOTE:** You can apply only one classifier of each type, DSCP and IEEE 802.1p, to a Layer 2 interface. If both types of classifiers are present, DSCP classifiers take precedence over IEEE 802.1p classifiers. (If you also configure a global EXP classifier, only MPLS traffic on interfaces configured as family `mpls` uses the EXP classifier, and other traffic uses the configured or default classifier for that traffic type.)

### **Importing a Classifier**

You can use any existing classifier, including the default classifiers, as the basis for defining a new classifier. You accomplish this using the **import** statement.

The imported classifier is used as a template and is not modified. The modifications you make become part of a new classifier (and a new template) identified by the name of the new classifier. Whenever you commit a configuration that assigns a new class-name and loss-priority value to a code-point alias or set of bits, it replaces that entry in the new classifier template. As a result, you must explicitly specify every CoS value in every designation that requires modification.

### **Multidestination Classifiers**

Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. You can configure both a DSCP multidestination classifier and an IEEE multidestination classifier. IP and IPv6 traffic use the DSCP classifier, and all other traffic uses the IEEE classifier.

DSCP IPv6 multidestination classifiers are not supported, so IPv6 traffic uses the DSCP multidestination classifier.

The default multidestination classifier is the IEEE 802.1p multidestination classifier.

### **PFC Priorities**

The eight IEEE 802.1p code points correspond to the eight priorities that priority-based flow control (PFC) uses to differentiate traffic classes for lossless transport. When you map a forwarding class (which maps to an output queue) to an IEEE 802.1p CoS value, the IEEE 802.1p CoS value identifies the priority.

Although you can map a priority to any output queue (by mapping the priority to a forwarding class), we recommend that the priority and the unicast forwarding class match in a one-to-one correspondence in which priority 0 is assigned to queue 0, priority 1 is assigned to queue 1, and so on, as shown in [Table 490 on page 5813](#). A one-to-one correspondence of queue and priority numbers makes it easier to configure and maintain the mapping of forwarding classes to priorities and queues.

**Table 490: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping**

| IEEE 802.1p Code Point | PFC Priority | Unicast Output Queue | Forwarding Class and Packet Drop Attribute |
|------------------------|--------------|----------------------|--|
| 000                    | 0            | 0                    | best-effort (drop)                         |

**Table 490: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping (*continued*)**

| IEEE 802.1p Code Point | PFC Priority | Unicast Output Queue | Forwarding Class and Packet Drop Attribute |
|------------------------|--------------|----------------------|--|
| 001                    | 1            | 1                    | best-effort (drop)                         |
| 010                    | 2            | 2                    | best-effort (drop)                         |
| 011                    | 3            | 3                    | fcoe (no-loss)                             |
| 100                    | 4            | 4                    | no-loss (no-loss)                          |
| 101                    | 5            | 5                    | best-effort (drop)                         |
| 110                    | 6            | 6                    | network-control (drop)                     |
| 111                    | 7            | 7                    | network-control (drop)                     |



**NOTE:** By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires. We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.

### Fixed Classifiers on Ethernet Interfaces

Fixed classifiers map all traffic on an interface to a forwarding class and a loss priority. (As opposed to BA classifiers, which map traffic into multiple different forwarding classes based on the CoS field value in the packet header.) The forwarding class determines the output queue. Incoming traffic of all IEEE 802.1p priorities is classified into the forwarding class specified in the fixed classifier. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

You cannot configure a fixed classifier and a DSCP or IEEE 802.1p BA classifier on the same interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.





**NOTE:** Because EXP classifiers are global, you can configure both a global EXP classifier and also apply fixed classifiers on interfaces. When both the global EXP classifier and a fixed classifier are applied to an interface, MPLS traffic on interfaces configured as family mpls uses the EXP classifier and all other traffic uses the fixed classifier.

To switch from a fixed classifier to a BA classifier or to switch from a BA classifier to a fixed classifier, deactivate the existing classifier attachment on the interface, and then attach the new classifier to the interface.



**NOTE:** If you configure a fixed classifier that classifies all incoming traffic into the fcoe forwarding class (or any forwarding class designed to handle FCoE traffic), you must ensure that all traffic that enters the interface is FCoE traffic and is tagged with the FCoE IEEE 802.1p code point (priority).

### Fixed Classifiers on Native Fibre Channel Interfaces (NP\_Ports)

Applying a fixed classifier to a native Fibre Channel (FC) interface (NP\_Port) is a special case. By default, native FC interfaces classify incoming traffic from the FC SAN into the fcoe forwarding class and map the traffic to IEEE 802.1p priority 3 (code point 011). When you apply a fixed classifier to an FC interface, you also configure a priority rewrite value for the interface. The FC interface uses the priority rewrite value as the IEEE 802.1p tag value for all incoming packets instead of the default value of 3.

For example, if you specify a priority rewrite value of 5 (code point 101) for an FC interface, the interface tags all incoming traffic from the FC SAN with priority 5 and classifies the traffic into the forwarding class specified in the fixed classifier.



**NOTE:** The forwarding class specified in the fixed classifier on FC interfaces must be a lossless forwarding class.

### Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

MF classification is normally performed at the network edge because of the general lack of DiffServ code point (DSCP) support in end-user applications. On a switch at the edge of a network, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

### Packet Classification for Routed VLAN Interfaces (RVIs)

You cannot apply classifiers directly to routed VLAN interfaces (RVIs) because the members of RVIs are VLANs, not ports. However, you can apply classifiers to the VLAN port members of an RVI. You can also apply MF classifiers to RVIs.

#### **Related Documentation**

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding Default CoS Settings on page 5796](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 6162](#)

## Understanding CoS MPLS EXP Classifiers and Rewrite Rules

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. (For information about DSCP and IEEE 802.1p classifiers and general information about classifiers, see [“Understanding CoS Classifiers” on page 5810](#). For information about DSCP and IEEE 802.1p rewrite rules, see [“Understanding CoS Rewrite Rules” on page 5914](#).)

When a packet enters a customer-edge interface on the ingress provider edge (PE) switch, the switch associates the packet with a particular CoS servicing level before placing the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the classifier is translated and encoded in the MPLS header by means of the experimental (EXP) bits.

EXP classifiers map incoming MPLS packets to a forwarding class and a loss priority, and assign MPLS packets to output queues based on the forwarding class mapping. EXP classifiers are behavior aggregate (BA) classifiers.

EXP rewrite rules change (rewrite) the CoS value of the EXP bits in outgoing packets on the egress queues of the switch so that the new (rewritten) value matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** There is no default EXP classifier. There is no default EXP rewrite rule. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. If you want to rewrite the EXP bit value at the egress interface, you must configure EXP rewrite rules and apply them to logical interfaces.

EXP classifiers and rewrite rules are applied only to interfaces that are configured as **family mpls** (for example, set interfaces xe-0/0/35 unit 0 family mpls.)

This topic includes:

- [EXP Classifiers on page 5817](#)
- [EXP Rewrite Rules on page 5818](#)
- [Schedulers on page 5819](#)

### EXP Classifiers

Unlike DSCP and IEEE 802.1p BA classifiers, EXP classifiers are global to the switch and apply to all switch interfaces that are configured as **family mpls**. When you configure and apply an EXP classifier, MPLS traffic on all **family mpls** interfaces uses the EXP classifier, even on interfaces that also have a fixed classifier. If an interface has both an EXP classifier and a fixed classifier, the EXP classifier is applied to MPLS traffic and the fixed classifier is applied to all other traffic.

Also unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. If you want to classify MPLS traffic based on the EXP bits, you must explicitly configure an EXP classifier and apply it to the switch interfaces. Each EXP classifier has eight entries that correspond to the eight EXP CoS values (0 through 7, which correspond to bits 000 through 111).

You can configure as many EXP classifiers as you want. However, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All switch interfaces use the global EXP classifier to classify MPLS traffic.

Only one EXP classifier can be configured as the global EXP classifier at any time. If you want to change the global EXP classifier, delete the global EXP classifier configuration (use the **user@switch# delete class-of-service system-defaults classifiers exp** configuration statement), then configure the new global EXP classifier.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.



**NOTE:** The switch uses only the outermost label of incoming EXP packets for classification.

---



**NOTE:** MPLS packets with 802.1Q tags are not supported.

---

### EXP Rewrite Rules

---

As MPLS packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. EXP rewrite rules set the value of the EXP CoS bits within the header of the outgoing MPLS packet on **family mpls** interfaces. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes that CoS value into the packet header, replacing the old CoS value. EXP rewrite rules apply only to MPLS traffic.

EXP rewrite rules apply only to logical interfaces. You cannot apply EXP rewrite rules to physical interfaces.

There are no default EXP rewrite rules. If you want to rewrite the EXP value in MPLS packets, you must configure EXP rewrite rules and apply them to logical interfaces. If no rewrite rules are applied, all MPLS labels that are pushed have a value of zero (0). The EXP value remains unchanged on MPLS labels that are swapped.

You can configure as many EXP rewrite rules as you want, but you can only apply 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

You can apply an EXP rewrite rule to an interface that has a DSCP, DSCP IPv6, or IEEE 802.1p rewrite rule. Only MPLS traffic uses the EXP rewrite rule. MPLS traffic does not use DSCP or DSCP IPv6 rewrite rules.

If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

## Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on the switch. Default schedulers are provided only for the best-effort, fcoe, no-loss, and network-control forwarding classes. If you configure a custom forwarding class for MPLS traffic, you need to configure a scheduler to support that forwarding class and provide bandwidth to that forwarding class. See “[Understanding CoS Output Queue Schedulers](#)” on page 5868 and “[Example: Configuring Queue Schedulers](#)” on page 6081 for more information.

### Related Documentation

- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring CoS Bits for an MPLS Network on page 4478](#)

## Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p, DSCP, or MPLS EXP class of service (CoS) code point bits in the packet header. At egress interfaces, you can use rewrite rules to change (re-mark) the code point bits before the interface forwards the packets. At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p, DSCP, or MPLS EXP CoS code point bits in the packet header. At egress interfaces, rewrite rules can change (re-mark) the code point bits before the interface forwards the packets.

You can apply classifiers and rewrite rules to interfaces to control the level of CoS applied to each packet as it traverses the system and the network. This topic describes:

- [Supported Classifier and Rewrite Rule Types on page 5820](#)
- [Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration on page 5821](#)
- [Default Classifiers on page 5823](#)
- [Default Rewrite Rules on page 5824](#)
- [Classifier Precedence on page 5824](#)
- [Classifier Behavior and Limitations on page 5825](#)
- [Rewrite Rule Precedence and Behavior on page 5826](#)
- [Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration on page 5827](#)

### Supported Classifier and Rewrite Rule Types

Table 491 on page 5820 shows the supported types of classifiers and rewrite rules supports:

**Table 491: Supported Classifiers and Rewrite Rules**

| Classifier or Rewrite Rule Type  | Description   |
|--|---|
| Fixed classifier   | Classifies all ingress traffic on a physical interface into one fixed forwarding class, regardless of the CoS bits in the packet header.  |
| DSCP and DSCP IPv6 unicast classifiers                                       | Classifies IP and IPv6 traffic into forwarding classes and assigns loss priorities to the traffic.  |
| IEEE 802.1p unicast classifier   | Classifies Ethernet traffic into forwarding classes and assigns loss priorities to the traffic.   |
| MPLS EXP classifier  | Classifies MPLS traffic into forwarding classes and assigns loss priorities to the traffic on interfaces configured as <b>family mpls</b> . The system uses one global EXP classifier on all <b>family mpls</b> switch interfaces.        |
| DSCP multdestination classifier (also used for IPv6 multdestination traffic) | Classifies IP and IPv6 multicast, broadcast, and destination lookup fail (DLF) traffic into multdestination forwarding classes. Multdestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. |

Table 491: Supported Classifiers and Rewrite Rules (*continued*)

| Classifier or Rewrite Rule Type        | Description  |
|--|--|
| IEEE 802.1p multdestination classifier | Classifies Ethernet multicast, broadcast, and destination lookup fail (DLF) traffic into multdestination forwarding classes. Multdestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. |
| DSCP and DSCP IPv6 rewrite rules       | Re-marks the DSCP code points of IP and IPv6 packets before forwarding the packets.  |
| IEEE 802.1p rewrite rule               | Re-marks the IEEE 802.1p code points of Ethernet packets before forwarding the packets.  |
| MPLS EXP rewrite rule                  | Re-marks the EXP code points of MPLS packets before forwarding the packets on interfaces configured as <b>family mpls</b> .  |



**NOTE:** On native Fibre Channel (FC) interfaces (NP\_Ports) only, you can specify a rewrite value to set the IEEE 802.1p code point of incoming FC traffic when the NP\_Port encapsulates the FC packet in Ethernet before forwarding it to the FCoE network (see *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*).

DSCP, IEEE 802.1p, and MPLS EXP classifiers are behavior aggregate (BA) classifiers. Unlike DSCP and IEEE 802.1p classifiers, EXP classifiers are global and apply only to all interfaces that are configured as **family mpls**. Also unlike DSCP and IEEE 802.1p classifiers, for MPLS traffic only, EXP classifiers overwrite fixed classifiers. (An interface that has a fixed classifier uses the EXP classifier for MPLS traffic, not the fixed classifier.)

Multdestination classifiers are global and apply to all interfaces; you cannot apply a multdestination classifier to individual interfaces.

Classifying packets into forwarding classes assigns packets to the output queues associated with the forwarding classes. Classifying traffic into a forwarding class associates the CoS scheduling for the forwarding class with that traffic.



**NOTE:** In addition to BA classifiers and fixed classifiers, which classify traffic based on the CoS field in the packet header, you can use firewall filters to configure multifield (MF) classifiers. MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

### Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration

To apply a classifier to incoming traffic or a rewrite rule to outgoing traffic, you need to apply the classifier or rewrite rule to one or more interfaces. When you apply a classifier or rewrite rule to an interface, the interface uses the classifier to group incoming traffic

into forwarding classes and uses the rewrite rule to re-mark the CoS code point value of each packet before it leaves the system.

Not all interfaces types support all types of CoS configuration. This section describes:

- [Interface Types That Support Classifier and Rewrite Rule Configuration on page 5822](#)
- [Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support on page 5822](#)
- [Routed VLAN Interfaces \(RVIs\) and Integrated Routing and Bridging \(IRB\) Interfaces on page 5823](#)

### ***Interface Types That Support Classifier and Rewrite Rule Configuration***

You can apply classifiers to all Ethernet interfaces. For Layer 3 LAGs, configure BA or fixed classifiers on the LAG (ae) interface. The classifier configured on the LAG is valid on all of the LAG member interfaces.

You can apply fixed classifiers to native FC interfaces (NP\_Ports). You cannot apply other types of classifiers or rewrite rules to native FC interfaces. You can rewrite the value of the IEEE 802.1p code point of incoming FC traffic when the interface encapsulates it in Ethernet before forwarding it to the FCoE network as described in *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*.

### ***Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support***

The Ethernet ports can function as:

- Layer 2 physical interfaces (family ethernet-switching)
- Layer 2 logical interfaces (family ethernet-switching)
- Layer 3 physical interfaces (family inet/inet6)
- Layer 3 logical interfaces (family inet/inet6)

You can apply CoS classifiers and rewrite rules only to the following interfaces:

- Layer 2 logical interfaces
- Layer 3 physical interfaces if at least one logical Layer 3 interface is configured on the physical interface



**NOTE:** The CoS you configure on a Layer 3 physical interface is applied to all of the Layer 3 logical interfaces on that physical interface. This means that each Layer 3 interface uses the same classifiers and rewrite rules for all of the Layer 3 traffic on that interface.

---

You cannot apply classifiers or rewrite rules to Layer 2 physical interfaces or to Layer 3 logical interfaces. [Table 492 on page 5823](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.



Table 492: Ethernet Interface Support for Classifier and Rewrite Rule Configuration

| CoS Classifiers and Rewrite Rules | Layer 2 Physical Interfaces   | Layer 2 Logical Interfaces | Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined) | Layer 3 Logical Interfaces |
|-----------------------------------|---|----------------------------|--|----------------------------|
| Fixed classifier                  | No  | Yes                        | Yes  | No                         |
| DSCP classifier                   | No  | Yes                        | Yes  | No                         |
| DSCP IPv6 classifier              | No  | Yes                        | Yes  | No                         |
| IEEE 802.1p classifier            | No  | Yes                        | Yes  | No                         |
| EXP classifier                    | Global classifier, applies only to all switch interfaces that are configured as <b>family mpls</b> . Cannot be configured on individual interfaces. |                            |  |                            |
| DSCP rewrite rule                 | No  | Yes                        | Yes  | No                         |
| DSCP IPv6 rewrite rule            | No  | Yes                        | Yes  | No                         |
| IEEE 802.1p rewrite rule          | No  | Yes                        | Yes  | No                         |
| EXP rewrite rule                  | No  | Yes                        | Yes  | No                         |



**NOTE:** IEEE 802.1p multidestination and DSCP multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. No DSCP IPv6 multidestination classifier is supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

### ***Routed VLAN Interfaces (RVIs) and Integrated Routing and Bridging (IRB) Interfaces***

You cannot apply classifiers and rewrite rules directly to routed VLAN interfaces (RVIs) or integrated routing and bridging (IRB) interfaces because the members of RVIs and IRBs are VLANs, not ports. However, you can apply classifiers and rewrite rules to the VLAN port members of an RVI or an IRB. You can also apply MF classifiers to RVIs and IRBs.

### **Default Classifiers**

If you do not explicitly configure classifiers on an Ethernet interface, default classifiers are applied (see [“Understanding Default CoS Settings” on page 5796](#)) so that the traffic receives basic CoS treatment. The factors that determine the default classifier applied to the interface include the interface type (Layer 2 or Layer 3), the port mode (trunk, tagged-access, or access), and whether logical interfaces have been configured. The system applies a default classifier using the following rules:

- If the physical interface has at least one Layer 3 logical interface configured, it uses the default DSCP classifier.
- If the physical interface has a Layer 2 logical interface in trunk mode or tagged-access mode, it uses the default trusted classifier.
- If the physical interface has a Layer 2 logical interface in access mode, it uses the default untrusted classifier.
- If the physical interface has no logical interface configured, no default classifier is applied.
- The default multidestination classifier is the IEEE 802.1p multidestination classifier.
- There is no default MPLS EXP classifier. If you want to classify traffic using EXP bits, you must configure an EXP classifier and configure it as the global system default EXP classifier.

---

### Default Rewrite Rules

No default rewrite rules are applied to interfaces. If you want to re-mark packets at the egress interface, you must explicitly configure a rewrite rule.

---

### Classifier Precedence

You can apply multiple unicast classifiers (MF, fixed, IEEE 802.1p, DSCP, or EXP) to a physical or logical Ethernet interface to handle different types of traffic. (EXP classifiers are global and apply only to all MPLS traffic on all **family mpls** interfaces.) When you apply more than one classifier to an interface, the system uses an order of precedence to determine which classifier to use on physical and logical interfaces:

- [Unicast Classifier Precedence on Physical Ethernet Interfaces on page 5824](#)
- [Unicast Classifier Precedence on Logical Ethernet Interfaces on page 5825](#)

#### *Unicast Classifier Precedence on Physical Ethernet Interfaces*

The precedence of unicast classifiers on physical interfaces, from the highest-priority classifier to the lowest-priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the physical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface



**NOTE:** If an EXP classifier is configured, MPLS traffic uses the EXP classifier on all **family mpls** interfaces, even if an MF or fixed classifier is applied to the interface. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

---

You can apply a DSCP classifier, an IEEE 802.1p classifier, and an EXP classifier on a physical interface. When all three classifiers are on an interface, IP traffic uses the DSCP classifier, MPLS traffic uses the EXP classifier, and all other traffic uses the IEEE classifier.



**NOTE:** You cannot apply a fixed classifier and a DSCP or IEEE classifier to the same interface. If a DSCP classifier, an IEEE classifier, or both are on an interface, you cannot apply a fixed classifier to that interface unless you first delete the DSCP and IEEE classifiers. If a fixed classifier is on an interface, you cannot apply a DSCP classifier or an IEEE classifier unless you first delete the fixed classifier.

### ***Unicast Classifier Precedence on Logical Ethernet Interfaces***

The precedence of unicast classifiers on logical interfaces, from the highest priority classifier to the lowest priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the logical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface



**NOTE:** If an EXP classifier is configured, MPLS traffic uses the EXP classifier on all family mpls interfaces, even if a fixed classifier is applied to the interface. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic.

You can apply both a DSCP classifier and an IEEE 802.1p classifier on a logical interface. When both a DSCP and an IEEE classifier are on an interface, IP traffic uses the DSCP classifier, and all other traffic uses the IEEE classifier. If an MPLS EXP classifier is also applied to the interface, only MPLS traffic uses the EXP classifier.

### **Classifier Behavior and Limitations**

Consider the following behaviors and constraints when you apply classifiers to physical and logical Ethernet interfaces:

- You can configure only one DSCP classifier (IP or IPv6) on a physical interface. You cannot configure both types of DSCP classifier on one physical interface. Both IP and IPv6 traffic use whichever DSCP classifier is configured on the interface.
- When you configure a DSCP or a DSCP IPv6 classifier on a physical interface and the physical interface has at least one logical Layer 3 interface, all packets (IP, IPv6, and non-IP) use that classifier.

- An interface with both a DSCP classifier (IP or IPv6) and an IEEE 802.1p classifier uses the DSCP classifier for IP and IPv6 packets, and uses the IEEE classifier for all other packets.
- Fixed classifiers and BA classifiers (DSCP and IEEE classifiers) are not permitted simultaneously on an interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.
- When you configure an IEEE 802.1p classifier on a physical interface and a DSCP classifier is not explicitly configured on that interface, the interface uses the IEEE classifier for all types of packets. No default DSCP classifier is applied to the interface. (In this case, if you want a DSCP classifier on the interface, you must explicitly configure it.)
- The system does not apply a default classifier to a physical interface until you create a logical interface on that physical interface. If you configure a Layer 3 logical interface, the system uses the default DSCP classifier. If you configure a Layer 2 logical interface, the system uses the default IEEE 802.1p trusted classifier if the port is in trunk mode or tagged-access mode, or the default IEEE 802.1p untrusted classifier if the port is in access mode.
- MF classifiers configured on logical interfaces take precedence over BA and fixed classifiers, with the exception of the global EXP classifier, which is always used for MPLS traffic on **family mpls** interfaces. (Use firewall filters to configure MF classifiers.) When BA or fixed classifiers are present on an interface, you can still configure an MF classifier on that interface.
- There is no default EXP classifier for MPLS traffic.
- You can configure as many EXP classifiers as you want, but the switch uses only one MPLS EXP classifier as a global classifier on all **family mpls** interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All **family mpls** switch interfaces use the EXP classifier specified using this configuration statement to classify MPLS traffic, even on interfaces that have a fixed classifier. No other traffic uses the EXP classifier.

---

### Rewrite Rule Precedence and Behavior

The following rules apply on both physical and logical Ethernet interfaces for rewrite rules:

- If you configure both one DSCP (or DSCP IPv6) rewrite rule and one IEEE 802.1p rewrite rule on an interface, both rewrite rules take effect. Traffic with IP and IPv6 headers use the DSCP rewrite rule, and traffic with a VLAN tag uses the IEEE rewrite rule.
- If you do not explicitly configure a rewrite rule, there is no default rewrite rule, so the system does not apply any rewrite rule to the interface.
- You can apply a DSCP rewrite rule or a DSCP IPv6 rewrite rule to an interface, but you cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Both

IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6.

- MPLS EXP rewrite rules apply only to logical interfaces on **family mpls** interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch.
- A logical interface can use both DSCP (or DSCP IPv6) and EXP rewrite rules.
- DSCP and DSCP IPv6 rewrite rules are not applied to MPLS traffic.
- If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

### Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration

You can apply classifiers and rewrite rules only on Layer 2 logical interfaces and Layer 3 physical interfaces (if the Layer 3 physical interface has at least one defined logical interface). This section focuses on BA classifiers, but the interaction between BA classifiers and interfaces described in this section also applies to fixed classifiers and rewrite rules.



**NOTE:** Multidestination classifiers, and EXP classifiers (only on **family mpls** interfaces), are global and apply to all switch interfaces. See [“Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)” on page 6160](#) for how to configure multidestination classifiers and see [“Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)” on page 6160](#) for how to configure EXP classifiers.

There are two components to applying classifiers or rewrite rules to interfaces:

1. Setting the interface family (inet, inet6, or ethernet-switching; ethernet-switching is the default interface family) in the **[edit interfaces]** configuration hierarchy.
2. Applying a classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy.

These are separate operations that can be set and committed at different times. Because the type of classifier or rewrite rule you can apply to an interface depends on the interface family configuration, the system performs checks to ensure that the configuration is valid. The method the system uses to notify you of an invalid configuration depends on the **set** operation that causes the invalid configuration.

When applying the classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy causes an invalid configuration, the system rejects the configuration and returns a commit check error.

When setting the interface family in the **[edit interfaces]** configuration hierarchy causes an invalid configuration, the system creates a syslog error message. When you receive the error message, you need to remove the classifier or rewrite rule configuration from the logical interface and apply it to the physical interface, or remove the classifier or rewrite rule configuration from the physical interface and apply it to the logical interface. For classifiers, if you do not take action to correct the error, the system programs the default classifier for the interface family on the interface. (There are no default rewrite rules. If the commit check fails, no rewrite rule is applied to the interface.)

Two scenarios illustrate these situations:

- [Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error on page 5828](#)
- [Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error on page 5829](#)



**NOTE:** Both of these scenarios also apply to fixed classifiers and rewrite rules.

---

### ***Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error***

In Scenario 1, we set the interface family, and then specify an invalid classifier.

1. Set and commit the interface as a Layer 3 (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This commit operation succeeds.

2. Set and commit a DSCP classifier on the logical interface (this example uses a DSCP classifier named **dscp1**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This configuration is not valid, because it attempts to apply a classifier to a Layer 3 logical interface. Because the failure is caused by the class-of-service configuration and not by the interface configuration, the system rejects the commit operation and issues a commit error, not a syslog message.

Note that the commit operation succeeds if you apply the classifier to the physical Layer 3 interface as follows:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Because the logical unit is not specified, the classifier is applied to the physical Layer 3 interface in a valid configuration, and the commit check succeeds.

### Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error

In Scenario 2, we set the classifier first, then set an invalid interface type.

1. Set and commit a DSCP classifier on a Layer 3 logical interface, assuming that the interface has no existing configuration:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This commit succeeds. Because no explicit configuration existed on the interface, it is by default a Layer 2 (**family ethernet-switching**) interface. Layer 2 logical interfaces support BA classifiers, so applying the classifier is a valid configuration.

2. Set and commit the interface as a Layer 3 interface (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This configuration is not valid because it attempts to change an interface from Layer 2 (**family ethernet-switching**) to Layer 3 (**family inet**) when a classifier has already been applied to a logical interface. Layer 3 logical interfaces do not support classifiers. Because the failure is caused by the interface configuration and not by the class-of-service configuration, the system does not issue a commit error, but instead issues a syslog message.

When the system issues the syslog message, it programs the default classifier for the interface type on the interface. In this scenario, the interface has been configured as a Layer 3 interface, so the system applies the default DSCP profile to the physical Layer 3 interface.

In this scenario, to install a configured DSCP classifier, you remove the misconfigured classifier from the Layer 3 logical interface and apply it to the Layer 3 physical interface. For example:

```
[edit]
user@switch# delete class-of-service interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
user@switch# set class-of-service interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

#### Related Documentation

- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding Default CoS Settings on page 5796](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Default CoS Scheduling and Classification on page 5856](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)

- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\)](#)

## Understanding CoS Forwarding Classes

Forwarding classes group traffic and assign the traffic to output queues. Each forwarding class is mapped to an output queue. Classification identifies the output queue for each incoming packet by mapping the packet code point bits to forwarding classes. The forwarding class to queue mapping defines the output queue used for the packet.

A classifier must associate each packet with one of the following five default forwarding classes or with a user-configured forwarding class in order to assign an output queue to the packet:

- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.
- **no-loss**—Guaranteed delivery for TCP lossless traffic.
- **best-effort**—Provides best-effort delivery without a service profile. Loss priority is typically not carried in a class-of-service (CoS) value.
- **network-control**—Supports protocol control and is typically high priority.
- **mcast**—Provides no service profile for multidestination (multicast, broadcast, and destination lookup fail) packets.

The switch supports up to 12 forwarding classes, thus enabling flexible, differentiated, packet classification. For example, you can configure multiple classes of best-effort traffic such as **best-effort**, **best-effort1**, and **best-effort2**.

The switch supports up to 12 output queues: 8 output queues for unicast traffic (queues 0 through 7) and 4 output queues for multidestination traffic (queues 8 through 11). Forwarding classes mapped to unicast queues are associated with unicast traffic, and forwarding classes mapped to multidestination queues are associated with multidestination traffic. You cannot map unicast and multidestination traffic to the same queue. You cannot map a strict-high priority queue to a multidestination forwarding class (queues 8 through 11 do not support strict-high priority configuration).

- [Default Forwarding Classes on page 5831](#)
- [Forwarding Class Configuration Rules on page 5832](#)
- [Lossless Transport Support on page 5833](#)



## Default Forwarding Classes

Table 493 on page 5831 shows the four default forwarding classes defined for unicast traffic, and Table 494 on page 5832 shows the four default forwarding classes defined for multicast traffic.

If desired, you can rename the forwarding classes associated with the queues supported on your switch. Assigning a new class name to an output queue does not alter the default classification or scheduling that is applicable to that queue. CoS configurations can be quite complicated, so unless it is required by your scenario, we recommend that you not alter the default class names or queue number associations.

**Table 493: Default Forwarding Classes for Unicast Packets**

| Forwarding Class Name | Default Queue Mapping | Comments   |
|-----------------------|-----------------------|--|
| best-effort (be)      | 0                     | <p>The software does not apply any special CoS handling to packets with 000000 in the DiffServ field. This is a backward compatibility feature. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of <b>drop</b>.</p>   |
| fcoe                  | 3                     | <p>By default, the <b>fcoe</b> forwarding class is a lossless forwarding class designed to handle Fibre Channel over Ethernet (FCoE) traffic. The <b>no-loss</b> packet drop attribute is applied by default.</p> <p><b>NOTE:</b> By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the <b>fcoe</b> forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires.</p> <p>We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.</p> |
| no-loss               | 4                     | <p>By default, this is a lossless forwarding class with a packet drop attribute of <b>no-loss</b>.</p>   |
| network-control (nc)  | 7                     | <p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of <b>drop</b>.</p>   |

Table 494: Default Forwarding Classes for Multicast Packets

| Forwarding Class Name | Default Queue Mapping | Comments  |
|-----------------------|-----------------------|---|
| mcast                 | 8                     | <p>The software does not apply any special CoS handling to the multidestination packets. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of <b>drop</b>.</p> |



**NOTE:** Mirrored traffic is always sent to the queue that corresponds to the multidestination forwarding class. The switched copy of the mirrored traffic is forwarded with the priority determined by the behavior aggregate classification process.

### Forwarding Class Configuration Rules

Take the following rules into account when you configure forwarding classes:

- [Queue Assignment Rules on page 5832](#)
- [Scheduling Rules on page 5833](#)
- [Rewrite Rules on page 5833](#)

#### Queue Assignment Rules

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. The commit operation fails with a detailed message that states the total number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that appears in the default configuration is the forwarding class currently associated with that queue.
- Only unicast forwarding classes can be mapped to unicast queues (0 through 7), and only multidestination forwarding classes can be mapped to multidestination queues (8 through 11).
- Strict-high priority queues cannot be mapped to multidestination forwarding classes. (Strict-high priority traffic cannot be mapped to queues 8 through 11).
- If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

In addition, if you configure a strict-high priority queue, we recommend that you always apply a shaping rate to prevent the strict-high priority queue from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority

queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

### ***Scheduling Rules***

When you define a forwarding class that is used on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

### ***Rewrite Rules***

On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

### **Lossless Transport Support**

---

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from the switch or QFabric system Node device to other devices cannot exceed 300 meters.
- The internal cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.
- Changing any portion of a PFC configuration on a port blocks the entire port until the change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Changing the PFC configuration means any change to a congestion notification profile that is configured on a port (enabling or disabling PFC on a code point, changing the MRU or cable-length value, or specifying an output flow control queue). Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.



.....

**NOTE:** Junos OS Release 12.2 introduces changes to the way lossless forwarding classes (the `fcoe` and `no-loss` forwarding classes) are handled.

In Junos OS Release 12.1, both explicitly configuring the `fcoe` and `no-loss` forwarding classes, and using the default configuration for these forwarding classes, resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the `fcoe` or the `no-loss` forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the `fcoe` or the `no-loss` forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the explicit `fcoe` and `no-loss` forwarding class configuration before you upgrade to Junos OS Release 12.2.

See *Overview of CoS Changes Introduced in Junos OS Release 12.2* for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the `fcoe` and `no-loss` forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new `no-loss` packet drop attribute or the forwarding class is lossy.

.....

**Related  
Documentation**

- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding Junos CoS Components on page 5789](#)
- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Example: Configuring Forwarding Classes on page 6075](#)
- [Defining CoS Forwarding Classes on page 6164](#)

## Understanding CoS Forwarding Class Sets (Priority Groups)

A forwarding class set is the Junos OS configuration construct that equates to a priority group in enhanced transmission selection (ETS, described in IEEE 802.1Qaz). The switch implements ETS using a two-tier hierarchical scheduler.

A priority group is a group of queues (priorities). Mapping a forwarding class to a queue defines the traffic for that queue, so a priority equates to a queue (forwarding class). The queues in a priority group share the port bandwidth allocated to that priority group. The traffic for queues in one priority group usually share similar traffic-handling requirements.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. Only unicast forwarding classes can belong to unicast forwarding class sets. Only multicast forwarding classes can belong to the multicast forwarding class set.

If you configure a strict-high priority queue, you must observe the following rules when configuring forwarding class sets:

- You must create a separate forwarding class set for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- You cannot configure a guaranteed minimum bandwidth (guaranteed rate) for a forwarding class set that includes a strict-high priority queue. (You also cannot configure a guaranteed minimum bandwidth for a strict-high queue.)
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

You must use hierarchical scheduling to define CoS for output queues. The two-tier hierarchical scheduler defines bandwidth resources for the priority group, and then allocates those resources among the priorities that belong to the priority group.

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set. Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement.

When you explicitly configure forwarding class sets and map them to an interface, any forwarding class that you do not map to a forwarding class set receives no guaranteed

bandwidth on that interface. Forwarding classes that belong to the default forwarding class set might receive bandwidth if the other forwarding class sets are not using all of the port bandwidth. However, the amount of bandwidth forwarding classes that are not in explicitly configured forwarding class sets receive is not guaranteed. The bandwidth for the default forwarding class depends on whether extra port bandwidth is available and therefore is not deterministic.

To guarantee bandwidth for forwarding classes in a predictable manner, be sure to map all forwarding classes that you expect to carry traffic on an interface to a forwarding class set and map the forwarding class set to the interface.

**Related  
Documentation**

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Defining CoS Forwarding Class Sets on page 6166](#)

## Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

Junos OS Release 12.3 increased support for lossless priorities from two lossless forwarding classes to up to six lossless forwarding classes. Each forwarding class is mapped to an IEEE 802.1p code point (priority).



**NOTE:** Junos OS Release 13.1 introduced support for up to six lossless forwarding classes on QFabric systems. Throughout this document, features introduced on standalone switches in Junos OS Release 12.3 are introduced on QFabric systems in Junos OS Release 13.1 unless otherwise noted.

Junos OS Release 13.2 is the first QFX5100 switch release, and Junos OS 13.2X51-D25 is the first EX4600 switch release. The QFX5100 and EX4600 switches also support up to six lossless forwarding classes. However, because the QFX5100 and EX4600 switches have no native Fibre Channel (FC) interfaces, these switches do not support native FC traffic and does not support configuration as an FCoE-FC gateway. Throughout this document, features that pertain to native FC traffic and to FCoE-FC gateway configuration do not apply to QFX5100 and EX4600 switches.

Earlier Junos OS software releases supported two lossless forwarding classes, the default *fc* and *no-loss* forwarding classes, which are mapped by default to IEEE 802.1p priorities 3 (code point 011) and 4 (code point 100), respectively. Junos OS Release 12.3 also introduced a new output stanza in the congestion notification profile (CNP) to configure priority-based flow control (PFC) on output queues.



**Video:** [Why Use PFC in a Data Center Network?](#)

The default configuration is the same as the default configuration in Junos OS Release 12.2 and is backward-compatible. If you need only two (or fewer) lossless forwarding classes, use the default configuration. If you need more than two lossless forwarding classes, you can use the two default forwarding classes and configure additional lossless forwarding classes. If you do not want to use the default lossless forwarding classes, you can change them or use only the lossless forwarding classes that you explicitly configure.

- [Lossless Transport Features Introduced in Junos OS Release 12.3 on page 5838](#)
- [Default Lossless Priority Configuration on page 5838](#)
- [Configuring Lossless Priorities on page 5841](#)
- [Backward Compatibility with Junos OS Releases Earlier Than Release 12.3 on page 5854](#)
- [Configuration Rules and Recommendations on page 5855](#)

### Lossless Transport Features Introduced in Junos OS Release 12.3

---

Support for lossless transport introduced in Junos OS Release 12.3 includes:

- Configuring up to six lossless forwarding classes.
- Configuring PFC pause on output queues to program the output queues that can respond to PFC pause messages received from the connected peer. The priorities you pause on output queues must match the priorities on which you enable PFC on the corresponding ingress interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101), then you must also enable pause on priorities 3 and 5 on the corresponding ingress interfaces. Configuring flow control on the output queues and enabling PFC on the corresponding input queues allows you to pause up to six priorities (forwarding classes).
- Controlling the headroom buffer on Ethernet interfaces by configuring the maximum receive unit (MRU) size for the traffic mapped to an IEEE 802.1p priority (configured per priority) and the length of the attached cable (configured per interface). The MRU size can range up to full jumbo packet size (9216 bytes).
- Remapping (rewriting) IEEE 802.1p priorities on native Fibre Channel (FC) interfaces when the system is acting as an FCoE-FC gateway. If the Ethernet (FCoE) network uses a different IEEE 802.1p priority than priority 3 (011) for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that different priority (see *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*).

Lossless transport still requires configuring previously existing features, including enabling PFC on the lossless priorities on ingress interfaces, and configuring classifiers to classify incoming traffic into lossless forwarding classes based on the IEEE 802.1p priority tag of the packet.



**NOTE:** If you expect a large amount of lossless traffic on your network and configure multiple lossless traffic classes, ensure that you reserve enough scheduling resources (bandwidth) and lossless headroom buffer space to support the lossless flows. (“[Understanding CoS Buffer Configuration](#)” on [page 5891](#) describes how to configure buffers and provides a recommended buffer configuration for networks with larger amounts of lossless traffic.)

---

### Default Lossless Priority Configuration

---

If you do not explicitly configure forwarding classes, the system uses the default forwarding class configuration, which provides two default lossless forwarding classes (*fcoe* and *no-loss*). (If you change the forwarding class configuration, the changes apply to all traffic on that device because forwarding classes are global to a particular device.)

If you do not explicitly configure classifiers, and you do not explicitly configure flow control to pause output queues (configured in the output stanza of the CNP), the default classifier and the default output queue pause configuration are applied to all Ethernet interfaces on the switches (or Node devices). You can override the default classifier and the default



output queue pause configuration on a per-interface basis by applying an explicit configuration to an Ethernet interface. The default configuration is used on all Ethernet interfaces that do not have an explicit configuration.



**NOTE:** If you do not configure flow control on output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause in the output stanza of the CNP.

In the default configuration, only queue 3 and queue 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza of the CNP. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza of the CNP.

The default configuration is the same as the default configuration in software releases earlier than Junos OS Release 12.3, and provides the same lossless behavior:

- There are two default lossless forwarding classes (the no-loss packet drop attribute is applied automatically):  
fcoe—Mapped to output queue 3  
no-loss—Mapped to output queue 4
- The default classifier maps the fcoe forwarding class to IEEE 802.1p priority 3 (011) and the no-loss forwarding class to IEEE 802.1p priority 4 (100)
- Priority-based flow control (PFC) is enabled on Ethernet interface output queues 3 and 4 when those queues carry lossless traffic (traffic that is mapped to the fcoe and no-loss forwarding classes, respectively). In Junos OS software releases earlier than Release 12.3, output queue flow control was not user-configurable.

On native FC interfaces (NP\_Ports), default flow control is enabled on output queue 3 (IEEE 802.1p priority 3) for FCoE/FC traffic.

- PFC must be enabled explicitly on the lossless IEEE 802.1p priorities (code points) on ingress Ethernet interfaces; no default PFC configuration is applied at ingress interfaces. If you do not enable PFC on lossless priorities, those priorities might experience packet loss during periods of congestion. For example, if you want lossless FCoE traffic and you are using the default fcoe forwarding class, you use a CNP to enable PFC on priority 3 (code point 011), and apply that CNP to all ingress interfaces that carry FCoE traffic.
- On Ethernet ports, PFC buffer calculations use the following default values to determine the headroom buffer size:  
Cable length—100 meters (approximately 328 feet)  
MRU for priority 3 traffic—2500 bytes  
MRU for priority 4 traffic—9216 bytes  
Maximum transmission unit (MTU)—1522 (or the configured MTU value for the interface)



**NOTE:** If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.

- DCBX is enabled on all interfaces in autonegotiation mode, and automatically exchanges FCoE application protocol type, length, and values (TLVs) on interfaces that carry FCoE traffic. However, if you explicitly configure DCBX protocol TLV exchange for any application, then you must explicitly configure protocol TLV exchange for every application for which you want DCBX to exchange TLVs, including FCoE.

The default CoS configuration is backward-compatible with the *default* CoS configuration of software releases before Junos OS Release 12.3. If you explicitly configure lossless transport, ensure that the input and output queues corresponding to the lossless forwarding classes are explicitly configured for PFC pause.



**NOTE:** If you *explicitly* configured the lossless fcoe or no-loss forwarding classes before upgrading from a release earlier than Junos OS Release 12.3, those forwarding classes are *not* lossless after the upgrade to Junos OS Release 12.3 or later. To regain lossless behavior, you can delete the explicit configuration and use the default lossless forwarding classes, or you can use the no-loss packet drop attribute introduced in Junos OS Release 12.3 to configure the forwarding classes for lossless behavior.

Table 495 on page 5840 summarizes the default unicast forwarding classes and their mapping to output queues, IEEE 802.1p priorities, and drop attributes.

**Table 495: Mapping of Default Unicast Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute**

| Forwarding Class Name | Output Queue | Priority | Drop Attribute |
|-----------------------|--------------|----------|----------------|
| best-effort           | 0            | 0        | drop           |
| fcoe                  | 3            | 3        | no-loss        |
| no-loss               | 4            | 4        | no-loss        |
| network-control       | 7            | 7        | drop           |

There is one default multidestination forwarding class named *mcast* for multicast, broadcast, and destination lookup fail (DLF) traffic that is mapped to output queue 8 with a drop attribute of drop. (Incoming multidestination traffic on all IEEE 802.1p priorities is mapped to the mcast forwarding class by default.)

## Configuring Lossless Priorities

Configuring more than two lossless priorities (forwarding classes), or changing the default mapping of lossless forwarding classes to priorities and paused output queues, requires explicit configuration. Configuring lossless priorities includes:

- Configuring forwarding classes with the no-loss packet drop attribute
- Using a CNP to configure PFC on ingress interfaces and flow control (PFC) on egress interfaces
- Configuring a classifier to map IEEE 802.1p priorities (code points) to the correct forwarding classes (the forwarding classes for which you want lossless transport)

In addition, on Ethernet interfaces, DCBX must exchange the appropriate application protocol TLVs for the lossless traffic, and when the switch acts as an FCoE-FC gateway, you need to remap the FCoE priority on native FC interfaces if your network uses a priority other than 3 (IEEE code point 011) for FCoE traffic. This section describes:

- [Configuring Lossless Forwarding Classes \(Packet Drop Attribute\) on page 5841](#)
- [Congestion Notification Profiles \(PFC Configuration\) on page 5843](#)
- [Configuring DCBX \(Application Protocol TLV Exchange\) on page 5849](#)
- [Fate Sharing Among Traffic Classes on page 5849](#)
- [Transit Switch Configuration Versus FCoE-FC Gateway Configuration on page 5851](#)
- [Configuration Results and Commit Checks on page 5851](#)

### Configuring Lossless Forwarding Classes (Packet Drop Attribute)

Junos OS Release 12.3 introduced the *no-loss* parameter for forwarding class configuration. (Although it uses the same name, this is not the no-loss default forwarding class. It is a packet drop attribute you can specify to configure any unicast forwarding class as a lossless forwarding class.)

You can configure up to six forwarding classes (depending on system architecture and the availability of system resources) as lossless forwarding classes by including the **no-loss** drop attribute at the **[edit class-of-service forwarding-classes class forwarding-class-name queue-num queue-number]** hierarchy level.

If you use the default fcoe or no-loss forwarding classes, they include the no-loss drop attribute by default. If you explicitly configure the fcoe or no-loss forwarding classes and you want to retain their lossless behavior, you *must* include the no-loss drop attribute in the configuration.



**NOTE:** All forwarding classes mapped to the same output queue must have the same packet drop attribute. (All forwarding classes mapped to the same output queue must be either lossy or lossless. You cannot map both a lossy and a lossless forwarding class to the same queue.)

To avoid fate sharing (different flows receiving the same CoS treatment), use a one-to-one mapping of lossless forwarding classes to IEEE 802.1p code points (priorities) and queues. (Each forwarding class should be mapped to a different queue and classified into a different priority.) The classifier attached to the interface determines the forwarding class to priority mapping.

The fcoe and no-loss forwarding classes are special cases, because in the default configuration, they are configured for lossless behavior (providing that you also enable PFC on the priorities mapped to the fcoe and no-loss forwarding classes in the CNP input stanza).

[Table 496 on page 5842](#) summarizes the possible configurations of the fcoe and no-loss forwarding classes in Junos OS Release 12.3 and later, and the result of those configurations in terms of lossless traffic behavior. It is assumed that PFC, DCBX, and classifiers are properly configured.

**Table 496: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3**

| Explicit (User-Configured) or Default Forwarding Class Configuration | Packet Drop Attribute  | Result and Notes  |
|--|--|---|
| Default  | Default  | The fcoe and no-loss forwarding classes are lossless.<br><br><b>NOTE:</b> Even if you explicitly configure other forwarding classes (lossy or lossless forwarding classes), the fcoe and no-loss forwarding classes remain lossless because they are not explicitly configured.   |
| Explicit   | Not specified in the explicit forwarding class configuration                         | The fcoe and no-loss forwarding classes are lossy because they do not include the no-loss drop attribute.   |
| Explicit   | No-loss  | The fcoe and no-loss forwarding classes are lossless.   |
| Explicit, configured in Junos OS Release 12.2 or earlier             | Not specified (packet drop attribute was not available before Junos OS Release 12.3) | The fcoe and no-loss forwarding classes are lossy in Junos OS Release 12.3 and later because they do not include the no-loss drop attribute.<br><br><b>NOTE:</b> To retain lossless behavior, before you upgrade to Junos OS Release 12.3, delete the explicit configuration so that the system uses the default configuration. Alternatively, you can reconfigure the forwarding classes with the no-loss packet drop attribute after upgrading to Junos OS Release 12.3 or later. |

For all other forwarding classes, you must explicitly configure lossless transport by specifying the no-loss packet drop attribute, because the default configuration for all other forwarding classes is lossy.

### ***Congestion Notification Profiles (PFC Configuration)***

Use CNPs to configure lossless PFC characteristics on input and output interfaces.

The input stanza of a CNP enables PFC on specified IEEE 802.1p priorities (code points) and fine-tunes headroom buffer settings by configuring the maximum receive unit (MRU) value and cable length on ingress interfaces.

The output stanza of a CNP enables PFC (flow control) on output queues for specified IEEE 802.1p priorities so that the queues can respond to PFC pause messages from the connected peer on the priority of your choice. (By default, output queues 3 and 4 respond to received PFC messages when those queues carry lossless traffic in the fcoe and no-loss forwarding classes, respectively.)

To achieve lossless transport, the priority paused at the ingress interfaces must match the priority paused at the egress interfaces for a given traffic flow. For example, if you configure ingress interfaces to pause traffic tagged with IEEE 802.1p priority 5 (code point 101) and priority 5 traffic is mapped to output queue 5, then you must also configure the corresponding output interfaces to pause priority 5 on queue 5. In addition, the forwarding class mapped to queue 5 must be configured as a lossless forwarding class (using the no-loss drop attribute).



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
  1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
  2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.

3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.

- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

---

### ***Configuring Input Interface Flow Control (PFC and Headroom Buffer Calculation)***

On Ethernet interfaces, the input stanza of the CNP enables PFC on specified priorities so that the ingress interface can send a pause message to the connected peer during periods of congestion. Input CNPs also fine-tune the headroom buffers used for PFC support by allowing you to configure the MRU value and cable length (if you do not want to use the default configuration).

Headroom buffers support lossless transport by storing the traffic that arrives at an interface after the interface sends a PFC flow control message to pause incoming traffic. Until the connected peer receives the flow control message and pauses traffic, the interface continues to receive traffic and must buffer it (and the traffic that is still on the wire after the peer pauses) to prevent packet loss.

The system uses the MRU and the length of the attached physical cable to calculate buffer headroom allocation. The default configuration values are:

- MRU for priority 3 traffic—2500 bytes
- MRU for priority 4 traffic—9216 bytes
- Cable length—100 meters (approximately 328 feet)



**NOTE:** If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not explicitly configure an MRU value, the default MRU value is 2500 bytes.

---

You can fine-tune the MRU and the cable length to adjust the size of the headroom buffer on an interface. The switch has a shared global buffer pool and dynamically allocates headroom buffer space to lossless queues as needed.

A lower MRU or a shorter cable length reduces the amount of headroom buffer required on an interface and leaves more headroom buffer space for other interfaces. A higher MRU or a longer cable length increases the amount of headroom buffer space required on an interface and leaves less headroom buffer space for other interfaces.

In many cases, you can better utilize the headroom buffers by reducing the MRU value (for example, an MRU of 2180 is sufficient for most FCoE networks) and by reducing the cable length value if the physical cable is less than 100 meters long.



**NOTE:** When you configure the headroom buffers by changing the MRU or the cable length, and commit the configuration, the system performs a commit check and rejects the configuration if sufficient headroom buffer space is not available.

However, the system does not perform a commit check but instead returns a syslog error if:

- The buffers are configured on a LAG interface.
- The default classifier is used on the interface (instead of a user-configured classifier).
- The interface has not been created yet.

### ***Configuring Output Interface Flow Control (PFC)***

On Ethernet interfaces, you can use the output stanza of the CNP to configure flow control on unicast output queues and enable PFC pause response on specified IEEE 802.1p priorities. By default, output queues 3 and 4 are enabled for PFC pause on priorities 3 (IEEE 802.1p code point 011) and 4 (IEEE 802.1p code point 100). The default PFC pause response supports the default lossless forwarding class configuration, which maps the fcoe forwarding class to queue 3 and priority 3, and maps the no-loss forwarding class to queue 4 and priority 4.

Configuring PFC on output queues enables you to pause any priority on any unicast output queue on any Ethernet interface. Output flow control enables you to use more than two output queues to support lossless traffic flows (you can configure up to six lossless forwarding classes and map them to different output queues that are enabled for PFC pause). Output queue flow control also enables you to support multiple lossless forwarding classes (each mapped to a different priority and output queue) for one class of traffic.



**NOTE:** Output flow control only works when PFC is enabled in the CNP input stanza on the corresponding priorities on the interface.

For example, if the converged Ethernet network uses two different priorities for FCoE traffic (for example, priority 3 and priority 5), then you can classify those priorities into different lossless forwarding classes that are mapped to different output queues by:

1. Configuring two lossless forwarding classes for FCoE traffic, with each forwarding class mapped to a different output queue. For example, you could use the default fcoe forwarding class, which is mapped to queue 3, and you could configure a second lossless forwarding class called fcoe1 and map it to queue 5. The fcoe forwarding class is for priority 3 FCoE traffic (code point 011), and the fcoe1 forwarding class is for priority 5 (code point 101) FCoE traffic.
2. Configuring a classifier that maps each forwarding class to the desired IEEE 802.1p code point (priority). If FCoE traffic on both priorities uses one interface, the classifier

must classify both forwarding classes to the correct priorities. If FCoE traffic of different priorities uses different interfaces, the classifier configuration on each interface must map the correct priority to the corresponding lossless forwarding class.

3. Applying the classifier to the interfaces that carry FCoE traffic. The classifier determines the mapping of forwarding classes to priorities on each interface.

To configure lossless transport for these forwarding classes, you also need to:

- Enable PFC on the two priorities (3 and 5 in this example) at the ingress interfaces in the CNP input stanza.
- Configure PFC on the output queues and priorities for the forwarding classes in the CNP output stanza so that the interface can respond to pause messages received from the connected peer.



**NOTE:** When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

- Configure DCBX to exchange application protocol TLVs on both FCoE priorities.



**NOTE:** If you do not configure flow control to pause output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer, and you must explicitly enable PFC on the corresponding priorities in the CNP input stanza to achieve lossless behavior.

If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause. For example, if you explicitly configure flow control on output queue 5, the default configuration is no longer valid, and only output queue 5 is enabled for PFC pause. Output queues 3 and 4 are no longer enabled for PFC pause, so traffic using those queues no longer responds to PFC pause messages even if the corresponding forwarding class is configured with the no-loss drop attribute. To retain the pause configuration on output queues 3 and 4 and configure flow control on queue 5, you need to explicitly configure flow control on queues 3, 4, and 5.

You cannot configure flow control to pause a multidestination output queue. You can configure flow control to pause only unicast output queues.



### Output Interface Flow Control Profiles

Configuring the CNP output stanza creates an output flow control profile that tells egress ports the queues on which the Ethernet interface should respond to PFC pause messages. Although you can create an unlimited number of CNPs that contain input stanzas only, the number of CNPs that you can configure with output stanzas is limited:

- For standalone switches that are not part of a QFabric system, you can configure up to two output interface flow control profiles. (You can configure up to two CNPs with output stanzas.)
- For QFabric systems, you can configure one output interface flow control profile per Node device. (You can configure one CNP with an output stanza per Node device.)

There are a total of four output flow control profiles.

The system has a default output flow control profile that is applied to all Ethernet interfaces when the CNP attached to the interface has only an input stanza and does not include an output stanza. The default profile responds to PFC pause messages received on queue 3 (for priority 3, for the default fcoe forwarding class) and on queue 4 (for priority 4, for the default no-loss forwarding class), and is effective only if PFC is configured on those priorities in the CNP input stanza.

Additionally, the system has two internal output flow control profiles that it applies automatically to fabric (FTE) ports and to native FC interfaces (NP\_Ports). When the switch is not part of a QFabric system, the profile normally used for FTE ports is available for user configuration and provides a second user-configurable profile. (That is why standalone switches have two user-configurable output flow control profiles, but Node devices on a QFabric system have only one user-configurable output flow control profile.)

Because one output CNP can configure PFC pause response on multiple output queues (priorities), one user-configurable output CNP is usually flexible enough to specify the desired PFC response on all programmed interfaces.



**NOTE:** Each port can use one output flow control profile. You cannot apply more than one profile to one port.

Output flow control profiles can be expressed in table format. For example, [Table 497 on page 5847](#) shows the default output flow control profile that pauses priorities 3 and 4 on queues 3 and 4 (remember that PFC must also be enabled on code points 3 and 4 in the CNP input stanza in order for PFC to work):

**Table 497: Default Output Flow Control Profile**

| IEEE 802.1p Priority Specified in Received PFC Frame | Paused Output Queue |
|--|---------------------|
| 0 (000)  | —                   |
| 1 (001)  | —                   |

**Table 497: Default Output Flow Control Profile (*continued*)**

| IEEE 802.1p Priority Specified in Received PFC Frame | Paused Output Queue |
|--|---------------------|
| 2 (010)  | —                   |
| 3 (011)  | 3                   |
| 4 (100)  | 4                   |
| 5 (101)  | —                   |
| 6 (110)  | —                   |
| 7 (111)  | —                   |

Table 498 on page 5848 is an example of a user-configured output flow control profile. Using the example from the preceding section, the CNP output stanza configures flow control on output queue 5, and also explicitly configures output flow control on queues 3 and 4 for the fcoe and no-loss forwarding classes. (If you explicitly configure an output CNP, you must explicitly configure every output queue that you want to respond to PFC messages, because the user-configured profile overrides the default profile. If this example did not include queues 3 and 4, those queues would no longer respond to received PFC messages.)

**Table 498: User-Configured Output Flow Control Profile**

| IEEE 802.1p Priority Specified in Received PFC Frame | Paused Output Queue |
|--|---------------------|
| 0 (000)  | —                   |
| 1 (001)  | —                   |
| 2 (010)  | —                   |
| 3 (011)  | 3                   |
| 4 (100)  | 4                   |
| 5 (101)  | 5                   |
| 6 (110)  | —                   |
| 7 (111)  | —                   |

Remember that you must also enable PFC on code points 3, 4, and 5 in the CNP input stanza for this configuration to work. When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

### ***Configuring PFC Across Layer 3 Interfaces on QFX5100 and EX4600 Switches***

Enabling PFC on traffic flows is based on the IEEE 802.1p code point (priority) in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits). To enable PFC on traffic that crosses Layer 3 interfaces, the traffic must be classified by its IEEE 802.1p code point, not by its DSCP (or DSCP IPv6) code point.

See [“Understanding PFC Functionality Across Layer 3 Interfaces” on page 5950](#) for a conceptual overview of how to enable PFC on traffic across Layer 3 interfaces. See [“Example: Configuring PFC Across Layer 3 Interfaces” on page 6138](#) for an example of how to configure PFC on traffic that traverses Layer 3 interfaces.

### ***Configuring DCBX (Application Protocol TLV Exchange)***

For applications that require lossless transport, DCBX exchanges application protocol TLVs with the connected peer interface. By default, DCBX advertises FCoE application protocol TLVs on all interfaces that are enabled for DCBX, and by default, DCBX is enabled on all interfaces. DCBX advertises no other applications by default.

For each application (for example, iSCSI) that you want to configure for lossless transport, you must enable the interfaces which carry that application traffic to exchange DCBX protocol TLVs with the connected peer. The TLV exchange allows the peer interfaces to negotiate a compatible configuration to support the application.

If you configure DCBX to advertise any application, the default DCBX advertisement is overridden, and DCBX advertises only the configured applications. If you want an interface to advertise only the FCoE application, you do not have to configure DCBX application protocol TLV exchange; instead, you can use the default configuration.

If you want DCBX to advertise other applications, you must explicitly configure an application map and apply it to the interfaces on which you want to exchange protocol TLVs for those applications. If you want to exchange FCoE application protocol TLVs in addition to other application protocol TLVs, you must also explicitly configure the FCoE application in the application map. [“Understanding DCBX Application Protocol TLV Exchange” on page 5589](#) describes how application mapping works.



**NOTE:** Lossless transport also requires that you enable PFC on the correct priority (IEEE 802.1p code point) on the ingress interfaces using an input CNP. If the priority you pause at the ingress interfaces is not mapped to queue 3 or queue 4 (the two output queues that are enabled for PFC pause flow control by default), then you must also enable the output queues that correspond to paused input priorities to pause using the output stanza of the CNP.

### ***Fate Sharing Among Traffic Classes***

You can configure different lossless (or lossy) traffic flows to share fate—that is, to receive the same CoS treatment.

Fate sharing is not desirable for I/O convergence. Instead of independent control of the fate of each type of flow, different types of flows receive the same treatment. Fate sharing is particularly undesirable for lossless flows. If one lossless flow experiences congestion and must be paused, that affects flows that share fate with the congested flow even if the other flows are not experiencing congestion, and also can cause ingress port congestion. If your network requires that all 802.1p priorities be lossless, you can achieve that by allowing some fate sharing among the eight priorities by spreading them across up to six lossless forwarding classes.

If the number of lossless priorities is less than or equal to the number of configured lossless forwarding classes, then you can avoid fate sharing by configuring a one-to-one mapping of forwarding classes to IEEE 802.1p code points (priorities) and output queues. (Each forwarding class should be mapped to a different output queue and classified to a different priority.)

If you want to configure different traffic flows to share fate, two fate-sharing configurations are supported: mapping one forwarding class to more than one IEEE 802.1p code point (priority), and mapping two forwarding classes to the same output queue:

1. If you map one lossless forwarding class to more than one priority, the traffic tagged with each of the priorities uses the same CoS properties associated (the CoS properties associated with the forwarding class). For example, configuring a forwarding class called `fc1`, mapping it to queue 1, and mapping it to code points 101 and 110 using a classifier named `classify1` results in the traffic tagged with priorities 101 and 110 sharing fate:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 110
```

In this case, if the traffic mapped to either priority experiences congestion, both priorities are paused because they are mapped to the same forwarding class and are therefore treated similarly.

2. If you map multiple lossless forwarding classes to the same output queue, the traffic mapped to the forwarding classes uses the same output queue. This increases the amount of traffic the queue needs to buffer and forward, and can create congestion that affects all of the traffic flows that are mapped to the queue. For example, configuring two forwarding classes called `fc1` and `fc2`, mapping both forwarding classes to queue 1, and mapping the forwarding classes to code points 101 and 110 (respectively) using a classifier named `classify1` results in the traffic tagged with priorities 101 and 110 sharing fate on the same output queue:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service forwarding-classes class fc2 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc2
loss-priority low code-points 110
```

in this case, even though the two forwarding classes use different IEEE 802.1p priorities, if one forwarding class experiences congestion, it affects the other forwarding class. The reason is that if the output queue is paused because of congestion on either forwarding class, all traffic that uses that queue is paused. Since both forwarding classes are mapped to the queue, the traffic mapped to both forwarding classes is paused.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

### ***Transit Switch Configuration Versus FCoE-FC Gateway Configuration***

On a transit switch (all Ethernet ports, no native FC ports) that forwards FCoE traffic (or other traffic that requires lossless transport across the Ethernet network), the configuration of classifiers, lossless forwarding classes, DCBX, and PFC on ingress and egress interfaces to support lossless transport is as described in this document.

When the QFX3500 switch acts as an FCoE-FC gateway, the system uses native FC interfaces (NP\_Ports) to connect to the FC switch (or FCoE forwarder) at the FC network edge. You cannot apply CNPs or DCBX to native FC interfaces, only to Ethernet interfaces.

On an FCoE-FC gateway, the Ethernet interface configuration of classifiers, DCBX, and PFC is the same as the Ethernet interface configuration on a transit switch. The configuration of lossless forwarding classes is also the same.

However, supporting lossless transport on native FC interfaces requires that you rewrite the IEEE 802.1p priority value *if* your network uses any priority other than 3 (IEEE code point 011) for FCoE traffic. If your network uses priority 3 for FCoE traffic, you can and should use the default configuration on native FC interfaces.

By default, native FC interfaces tag packets with priority 3 when they encapsulate the incoming FC packets in Ethernet. If your FCoE network uses a different priority than 3 for FCoE traffic, you need to rewrite the priority value to the value that your network uses on the FC interface, classify the FCoE traffic to the correct priority on the Ethernet interfaces, and enable PFC on the correct priority on the Ethernet interfaces, as described in *Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway*.

### ***Configuration Results and Commit Checks***

Different configurations of forwarding classes and their drop attributes, classifiers, CNPs (PFC flow control), and Ethernet PAUSE (IEEE 802.3X flow control) result in different system behaviors.

[Table 499 on page 5852](#) describes the results of the possible lossless transport configurations in each case. The assumption in the *Result* column is that the system's buffer headroom calculation resulted in a successful configuration.

However, if the system calculates that there is insufficient buffer space to support the configuration, a commit check prevents you from committing the configuration on an individual Ethernet interface. For LAG interfaces, the system does not issue a commit check error but instead issues a syslog message.



**NOTE:** After you configure lossless transport for a LAG interface, be sure to check the syslog messages to confirm that the commit was successful.

**Table 499: Results of Lossless Priority Configuration**

| Classifier Configuration                               | Congestion Notification Profile Configuration                                 | Ethernet PAUSE (IEEE 802.3X) Configuration | Result  |
|--|---|--|---|
| None (default classifier)                              | None  | None                                       | System default configuration. No flows are lossless. To achieve lossless behavior for the default fcoe and no-loss forwarding classes, you must configure a CNP to enable PFC on their IEEE 802.1p code points (011 and 100 respectively).  |
| Classifier with no lossless forwarding classes         | None  | None                                       | No lossless traffic flows are configured; all traffic is best effort.   |
| Classifier with at least one lossless forwarding class | None  | None                                       | Because no CNP is attached to interfaces, PFC is not enabled on the code point of the lossless traffic and no headroom buffer is allocated to the lossless queue, so packets can drop during periods of congestion. This configuration does not achieve lossless behavior.        |
| None (default classifier)                              | PFC enabled on the fcoe and no-loss forwarding class code points (priorities) | None                                       | The default classifier classifies traffic into two lossless forwarding classes, fcoe and no-loss. The CNP enables PFC on the priorities mapped to both lossless forwarding classes, resulting in lossless behavior for traffic mapped to the fcoe and no-loss forwarding classes. |

Table 499: Results of Lossless Priority Configuration (*continued*)

| Classifier Configuration                               | Congestion Notification Profile Configuration                         | Ethernet PAUSE (IEEE 802.3X) Configuration   | Result   |
|--|---|--|--|
| None (default classifier)                              | None  | Flow control enabled   | The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length. The system does not calculate buffer headroom for individual output queues. Because Ethernet PAUSE is enabled on the link instead of PFC being enabled on the lossless priorities, the entire link is paused during periods of congestion. This configuration results in lossless behavior for all of the forwarding classes on the link, but because all traffic is paused, this can cause greater overall network congestion. |
| Classifier with at least one lossless forwarding class | PFC enabled on the lossless forwarding class code points (priorities) | None   | Headroom buffer allocated only to priorities that are mapped to the lossless forwarding classes and on which PFC is enabled. This configuration achieves lossless behavior for the lossless forwarding classes.  |
| Classifier with no lossless forwarding classes         | None  | Flow control enabled   | The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.   |
| Classifier with at least one lossless forwarding class | None  | Flow control enabled   | The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.   |
| Classifier with at least one lossless forwarding class | PFC enabled on the lossless forwarding class code points (priorities) | Flow control enabled on a <i>different</i> interface than the interface with the CNP | The system checks the available buffer space for both the PFC-enabled priorities and for the other link. If sufficient buffer space is available, the lossless forwarding classes configured with PFC on one interface and also all of the traffic on the link with Ethernet PAUSE enabled achieve lossless behavior.  |



**NOTE:** If you attempt to configure both PFC and Ethernet PAUSE on a link, the system returns a commit error. PFC and Ethernet PAUSE are mutually exclusive configurations on an interface.

### Backward Compatibility with Junos OS Releases Earlier Than Release 12.3

The addition of the no-loss packet drop attribute to forwarding class configuration means that when you upgrade from an earlier release to Junos OS Release 12.3, the new software might not preserve the lossless forwarding class configuration of the fcoe and no-loss forwarding classes.

If you used the default forwarding class configuration for the fcoe and no-loss forwarding classes, the CoS configuration is backward-compatible. You do not have to do anything to preserve the lossless behavior of traffic that uses those forwarding classes when you upgrade to Junos OS Release 12.3. (This is because the default configuration of these two forwarding classes includes the no-loss packet drop attribute.)

However, if you explicitly configured the fcoe or the no-loss forwarding class by including the **set forwarding-classes class forwarding-class-name queue-num queue-number** statement at the **[edit class-of-service]** hierarchy level, then those forwarding classes are no longer lossless, they are lossy. (They are lossy because explicit configuration in releases earlier than Junos OS Release 12.3 did not use the no-loss packet drop attribute.) In Junos OS Release 12.3 and later, you must include the no-loss packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

For example, before Junos OS Release 12.3, the following explicit configuration resulted in a lossless forwarding class:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3
```

However, in Junos OS Release 12.3, this configuration is lossy because it does not include the no-loss packet drop attribute. To preserve lossless behavior, after upgrading to Junos OS Release 12.3, you need to add the no-loss drop attribute:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3 no-loss
```

Alternatively, you can delete the explicit configuration before you upgrade to Junos OS Release 12.3 so that the system uses the default forwarding class, which is lossless:

```
user@switch# delete class-of-service forwarding-classes class fcoe queue-num 3
```



**NOTE:** The explicit configuration of other forwarding classes does not affect the lossless (or lossy) state of the fcoe and no-loss forwarding classes, because only the fcoe and no-loss forwarding classes were lossless forwarding classes before Junos OS Release 12.3. For example, if you explicitly configured the best-effort forwarding class but you used the default fcoe and no-loss forwarding classes in Junos OS Release 12.2, then when you upgrade to Junos OS Release 12.3, the fcoe and no-loss forwarding classes are still lossless (and the best-effort forwarding classes retains its explicit configuration).





**NOTE:** To achieve lossless behavior for the traffic belonging to any forwarding class, you must also use a CNP to enable PFC on the IEEE 802.1p priority mapped to the forwarding class and apply the CNP to the relevant interfaces, and ensure that DCBX exchanges the protocol TLVs for the application with the connected peer.

### Configuration Rules and Recommendations

Keep in mind the following configuration rules and recommendations when you configure lossless traffic flows:

- You can configure a maximum of six lossless forwarding classes (forwarding classes with the no-loss packet drop attribute).
- All forwarding classes that you map to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes must be lossless).
- You cannot configure flow control to pause a multidestination output queue. You can configure PFC flow control only to pause unicast output queues.
- Forwarding classes mapped to multidestination queues (queues 8 through 11) cannot have the no-loss packet drop attribute. (Multidestination forwarding classes cannot be configured as lossless forwarding classes.)
- Do not configure weighted random early detection (WRED) on lossless forwarding classes. (Do not associate a drop profile with a forwarding class that has the no-loss packet drop attribute.)

#### Related Documentation

- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS Buffer Configuration on page 5891](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 5950](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)

- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)
- [Example: Configuring PFC Across Layer 3 Interfaces on page 6138](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)

## Understanding Default CoS Scheduling and Classification

If you do not configure hierarchical scheduling on an interface, the switch uses the default classifiers for ingress traffic and the default schedulers for egress traffic. Default scheduling and classification handle all traffic types (best-effort, FCoE, no-loss, network-control, and multidestination traffic).

Hierarchical scheduling groups egress queues (priorities, configured as forwarding classes) into priority groups (forwarding class sets). If you use only the default traffic scheduling and classification, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default classifier settings. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.



**NOTE:** If you explicitly configure one or more priority groups on an interface, any forwarding class that is not assigned to a priority group on that interface receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class (priority) that you want to forward traffic on that interface must belong to a forwarding class set (priority group).

The following sections describe:

- [Default Classification on page 5856](#)
- [Default Scheduling on page 5859](#)
- [Default DCBX Advertisement on page 5861](#)
- [Default Scheduling and Classification Summary on page 5862](#)

### Default Classification

---

The default classifiers assign unicast and multicast best-effort and network-control ingress traffic to forwarding classes and loss priorities. The switch applies default unicast IEEE 802.1, unicast DSCP, and multidestination classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

Table 500 on page 5857 shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

**Table 500: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

Table 501 on page 5857 shows the default mapping of IEEE 802.1p code-point values to unicast forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

**Table 501: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

Table 502 on page 5858 shows the default mapping of IEEE 802.1 code-point values to multdestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

**Table 502: Default IEEE 802.1 Multidestination Classifiers**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | mcast            | low           |
| be1 (001)  | mcast            | low           |
| ef (010)   | mcast            | low           |
| ef1 (011)  | mcast            | low           |
| af11 (100) | mcast            | low           |
| af12 (101) | mcast            | low           |
| nc1 (110)  | mcast            | low           |
| nc2 (111)  | mcast            | low           |

Table 503 on page 5858 shows the default mapping of DSCP code-point values to unicast forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

**Table 503: Default DSCP IP and IPv6 Unicast Classifiers**

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| ef (101110)   | best-effort      | low           |
| af11 (001010) | best-effort      | low           |
| af12 (001100) | best-effort      | low           |
| af13 (001110) | best-effort      | low           |
| af21 (010010) | best-effort      | low           |
| af22 (010100) | best-effort      | low           |
| af23 (010110) | best-effort      | low           |
| af31 (011010) | best-effort      | low           |
| af32 (011100) | best-effort      | low           |
| af33 (011110) | best-effort      | low           |

Table 503: Default DSCP IP and IPv6 Unicast Classifiers (*continued*)

| Code Point    | Forwarding Class | Loss Priority |
|---------------|------------------|---------------|
| af41 (100010) | best-effort      | low           |
| af42 (100100) | best-effort      | low           |
| af43 (100110) | best-effort      | low           |
| be (000000)   | best-effort      | low           |
| cs1 (001000)  | best-effort      | low           |
| cs2 (010000)  | best-effort      | low           |
| cs3 (011000)  | best-effort      | low           |
| cs4 (100000)  | best-effort      | low           |
| cs5 (101000)  | best-effort      | low           |
| nc1 (110000)  | network-control  | low           |
| nc2 (111000)  | network-control  | low           |



**NOTE:** There are no default DSCP IP or IPv6 multdestination classifiers for multdestination traffic. DSCP IPv6 multdestination classifiers are not supported for multdestination traffic.

### Default Scheduling

The default schedulers allocate egress bandwidth resources to unicast and multicast egress traffic as shown in [Table 504 on page 5859](#):

Table 504: Default Scheduler Configuration

| Default Scheduler and Queue Number  | Transmit Rate (Minimum Guaranteed Bandwidth) | Shaping Rate (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|-------------------------------------|--|----------------------------------|--------------------------|----------|-------------|
| Best-effort scheduler (queue 0)     | 5%   | None                             | 5%                       | low      | 5%          |
| FCoE scheduler (queue 3)            | 35%  | None                             | 35%                      | low      | 35%         |
| No-loss scheduler (queue 4)         | 35%  | None                             | 35%                      | low      | 35%         |
| Network-control scheduler (queue 7) | 5%   | None                             | 5%                       | low      | 5%          |

Table 504: Default Scheduler Configuration (*continued*)

| Default Scheduler and Queue Number   | Transmit Rate (Minimum Guaranteed Bandwidth) | Shaping Rate (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|--------------------------------------|--|----------------------------------|--------------------------|----------|-------------|
| Multidestination scheduler (queue 8) | 20%  | None                             | 20%                      | low      | 20%         |



**NOTE:** The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth (transmit rate) of each queue.

By default, only the five default schedulers shown in [Table 504 on page 5859](#) have traffic mapped to them. Only the queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

Default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multidestination traffic. By default, unicast traffic consists of queue 0 (**best-effort** forwarding class), queue 3 (**fcoe** forwarding class), queue 4 (**no-loss** forwarding class), and queue 7 (**network-control** forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multidestination traffic (**mcast** queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multidestination traffic receives 2-Gbps of bandwidth.



**NOTE:** Multidestination queue 11 also receives a small amount of default bandwidth from the multidestination scheduler. CPU-generated multidestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multidestination traffic to a unicast forwarding class, that traffic uses queue 11.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

You should explicitly map traffic to non-default (unconfigured) queues and create schedulers to allocate bandwidth to those queues if you want to use them to forward traffic. By default, unicast queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)



**NOTE:** All four multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

If the other unicast queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

### Default DCBX Advertisement

When you configure hierarchical scheduling on an interface, DCBX advertises each priority group, the priorities in each priority group, and the bandwidth properties of each priority and priority group.

If you do not configure hierarchical scheduling on an interface, DCBX advertises the automatically created default priority group and its priorities. DCBX also advertises the default bandwidth allocation of the priority group, which is 100 percent of the port bandwidth.

## Default Scheduling and Classification Summary

---

If you do not configure hierarchical scheduling on an interface:

- Default classifiers classify ingress traffic.
- Default schedulers schedule egress traffic.
- DCBX advertises a single default priority group with 100 percent of the port bandwidth allocated to that priority group. All priorities (forwarding classes) are assigned to the default priority group and receive bandwidth based on their default schedulers. The default priority group is generated automatically and is not user-configurable.

### Related Documentation

- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding Default CoS Settings on page 5796](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Example: Configuring Queue Schedulers on page 6081](#)

## Understanding CoS Hierarchical Port Scheduling (ETS)

Scheduling defines the class-of-service (CoS) properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the queue priority, and the drop profiles associated with the queue.

Hierarchical port scheduling is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues. Hierarchical scheduling includes the Junos OS implementation of enhanced transmission selection (ETS, described in IEEE 802.1Qaz).



Video: [What is Enhanced Transmission Selection?](#)

---

This topic describes:

- [Hierarchical Scheduling Tiers on page 5863](#)
- [Hierarchical Scheduling and ETS on page 5863](#)
- [ETS Advertisement in DCBX on page 5864](#)
- [Hierarchical Scheduling Process on page 5865](#)



- [Strict-High Priority Queues and Hierarchical Scheduling on page 5866](#)
- [Default Hierarchical Scheduling on page 5866](#)

### Hierarchical Scheduling Tiers

The two tiers used in hierarchical scheduling are priorities and priority groups, as shown in [Table 505 on page 5863](#).

**Table 505: Hierarchical Scheduling Tiers**

| Junos OS Configuration Construct | Equivalent ETS Construct | Description  |
|----------------------------------|--------------------------|--|
| Forwarding class                 | Priority                 | <p>Think about priorities (forwarding classes) as output queues. You map forwarding classes to queues, so each forwarding class is in essence an output queue.</p> <p>When you use a classifier to map a forwarding class to an IEEE 802.1p code point, the code point identifies that traffic's priority for priority-based flow control (PFC). Thus the forwarding class, the queue mapped to the forwarding class, and the priority mapped to the forwarding class all identify the same traffic.</p> |
| Forwarding class set             | Priority group           | <p>Priority groups (forwarding class sets) are groups of priorities. Forwarding class membership in a forwarding class set defines the priority group to which each priority belongs.</p> <p>You can configure up to three unicast priority groups and one multicast forwarding class set.</p>   |

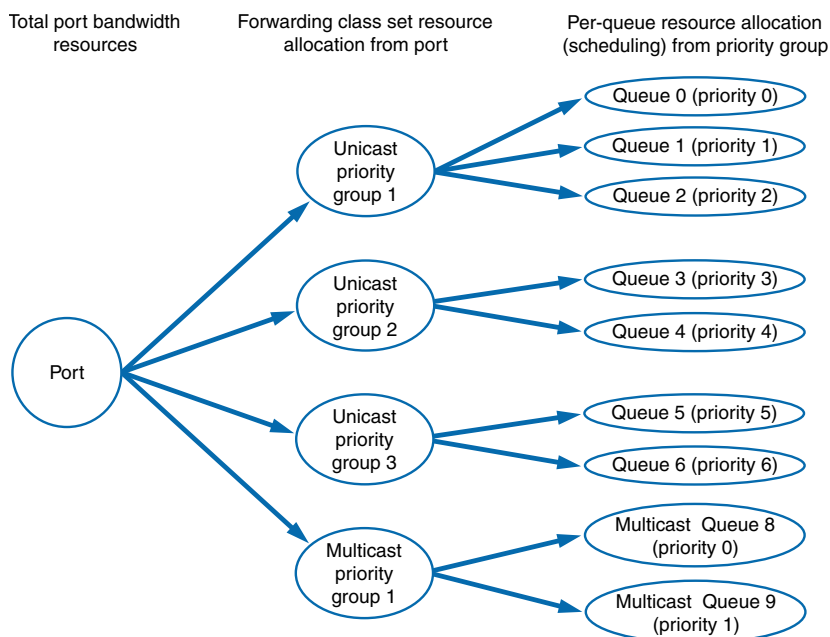


**NOTE:** If you explicitly configure one or more priority groups on an interface, any priority that is not assigned to a priority group on that interface is assigned to an automatically generated default priority group and receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class that you want to forward traffic on that interface must belong to a forwarding class set.

### Hierarchical Scheduling and ETS

Two-tier hierarchical scheduling enables you to manage bandwidth efficiently by enabling you to define the CoS properties for each priority group and for each priority. One tier of the hierarchical scheduler allocates port bandwidth to a priority group. The other tier of the hierarchical scheduler determines the portion of the priority group bandwidth that a queue can use.

The CoS properties you configure for a priority group define the port bandwidth resources available to the queues in that priority group. The CoS properties you configure for each queue specify the portion or percentage of the total bandwidth configured for the priority group that is available to the queue. [Figure 204 on page 5864](#) shows the relationship of port resource allocation to priority groups and priority group resource allocation to queues (priorities).

**Figure 204: Hierarchical Scheduling Tiers**

If a queue is not using its allocated bandwidth, ETS shares the unused bandwidth among the other queues in the priority group. If link bandwidth is available or if a priority group on a link is not using its allocated bandwidth, ETS shares the unused bandwidth with other priority groups on the link. In this way ETS improves link bandwidth utilization and provides each queue with the maximum bandwidth. Priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.



**NOTE:** The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows.

### ETS Advertisement in DCBX

When you configure hierarchical scheduling on a port, Data Center Bridging Capability Exchange Protocol (DCBX) advertises:

- Each priority group
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

When you configure hierarchical scheduling on a port, any priority that is not part of an explicitly configured priority group is assigned to the automatically generated default priority group and receives no bandwidth. The default priority group is transparent. It does not appear in the configuration.

---

## Hierarchical Scheduling Process

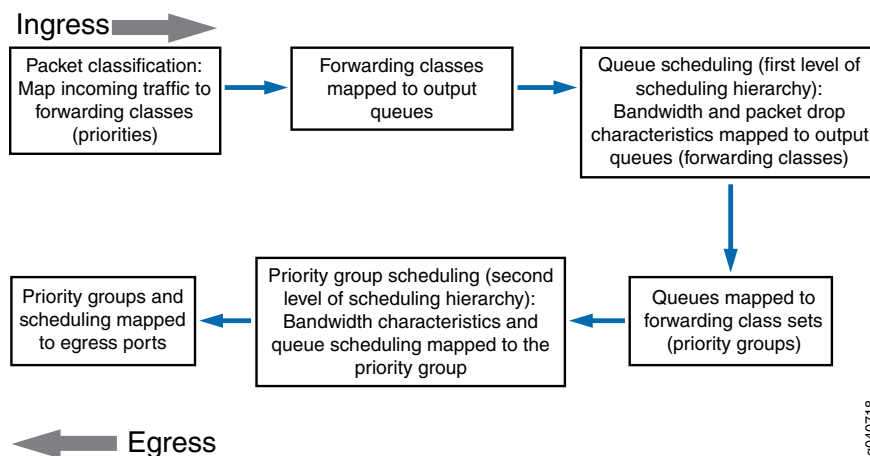
---

Hierarchical scheduling consists of multiple configuration steps that create the priorities and the priority groups, schedule their resources, and assign them to interfaces. The steps below correspond to the six blocks in the packet flow diagram shown in [Figure 205 on page 5866](#):

1. Packet classification:
  - Classify incoming traffic into priorities. This consists of either using the default classifiers or configuring classifiers to map IEEE 802.1p code points and loss priorities to the forwarding classes.
  - Apply the classifiers to ingress interfaces. This groups incoming traffic into forwarding classes (priorities) by mapping code points to forwarding classes and loss priorities on the specified interface.
2. Configure the output queues for the forwarding classes (priorities). This consists of either using the default forwarding classes and forwarding-class-to-queue mapping or creating your own forwarding classes and mapping them to queues.
3. Allocate resources to the forwarding classes:
  - Define resources for the priorities. This consists of configuring schedulers to set minimum guaranteed bandwidth, maximum bandwidth, drop profiles for Weighted Random Early Detection (WRED), and bandwidth priority to apply to a forwarding class. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.
  - Map resources to priorities. This consists of mapping forwarding classes to schedulers by using a scheduler map.
4. Configure priority groups. This consists of mapping forwarding classes (priorities) to forwarding class sets (priority groups) to define the priorities that belong to each priority group.
5. Define resources for the priority groups. This consists of configuring traffic control profiles to set minimum guaranteed bandwidth and maximum bandwidth for a priority group. Traffic control profiles also specify a scheduler map, which defines the resources (schedulers) for the priorities in the priority group. Extra port bandwidth is shared among priority groups in proportion to the minimum guaranteed bandwidth of each priority group.

The traffic control profile bandwidth settings determine the port resources available to the priority group, and the schedulers specified in the scheduler map determine the amount of the priority group resources that each priority receives.
6. Apply the hierarchical scheduling to a port. This consists of attaching one or more priority groups to a port interface. For each priority group, you also attach a traffic control profile. Different priority groups on the same port can use different traffic control profiles.

Figure 205: Hierarchical Scheduling Packet Flow



### Strict-High Priority Queues and Hierarchical Scheduling

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multdestination forwarding class set.
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.



**NOTE:** On a QFabric system, if a fabric (fte) interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best effort, lossless, strict-high priority, and multdestination traffic.

### Default Hierarchical Scheduling

If you do not explicitly configure hierarchical scheduling, the switch uses the default settings:

- The switch automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The switch assigns 100 percent of the port output bandwidth to the default forwarding class set. The default forwarding class set is

transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement.

- Ingress traffic is classified based on the default classifier settings.
- The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings.

**Related  
Documentation**

- [Understanding CoS Packet Flow on page 5793](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Priority Group Scheduling on page 5877](#)
- *Benefits of Configuring CoS Hierarchical Port Scheduling*
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Default CoS Scheduling and Classification on page 5856](#)
- *Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports*
- *Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)*
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)

## Understanding CoS Output Queue Schedulers

Output queue scheduling defines the class-of-service (CoS) properties of output queues (output queues are mapped to forwarding classes and IEEE 802.1p priorities). Queue scheduling works with priority group scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates port bandwidth to a group of queues called a priority group (forwarding class set), and queue scheduling determines the portion of the priority group's bandwidth that a particular queue can use.

Scheduler maps associate queue schedulers with forwarding classes, which are mapped to output queues. You can associate each scheduler map with a traffic control profile, and then associate each traffic control profile with a forwarding class set (priority group) and a port interface. In conjunction with the priority group scheduling configured in the traffic control profile, queue scheduling configures the output queues, packet schedulers, and weighted random early detection (WRED) packet drop processes that operate according to this mapping.



**NOTE:** When you configure bandwidth for a queue (forwarding class) or a priority group (forwarding class set), the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- [Output Queue Scheduling Components on page 5868](#)
- [Default Schedulers on page 5869](#)
- [Transmit Rate \(Minimum Guaranteed Bandwidth\) on page 5872](#)
- [Sharing Extra Bandwidth on page 5872](#)
- [Shaping Rate \(Maximum Bandwidth\) on page 5873](#)
- [Scheduling Priority on page 5873](#)
- [Scheduler Drop-Profile Maps on page 5874](#)
- [Buffer Size on page 5874](#)
- [Explicit Congestion Notification on page 5875](#)
- [Scheduler Maps on page 5876](#)

---

### Output Queue Scheduling Components

[Table 506 on page 5869](#) provides a quick reference to the scheduler components you can configure to determine the bandwidth properties of output queues, and [Table 507 on page 5869](#) provides a quick reference to some related scheduling configuration components.

Table 506: Output Queue Scheduler Components

| Output Queue Scheduler Component | Description  |
|----------------------------------|--|
| Buffer size                      | Sets the size of the queue buffer.   |
| Drop profile map                 | Maps a drop profile to a loss priority. Drop profile map components include: <ul style="list-style-type: none"> <li>Drop profile—Sets the probability of dropping packets as the queue fills up.</li> <li>Loss priority—Sets the traffic loss priority to which a drop profile applies.</li> </ul> |
| Explicit congestion notification | Enables explicit congestion notification (ECN) on the queue.   |
| Priority                         | Sets the scheduling priority applied to the queue.   |
| Shaping rate                     | Sets the maximum bandwidth the queue can consume.  |
| Transmit rate                    | Sets the minimum guaranteed bandwidth for the queue. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.   |

Table 507: Other Scheduling Components

| Other Scheduling Components | Description   |
|-----------------------------|---|
| Scheduler map               | Maps schedulers to forwarding classes (forwarding classes are mapped to queues, so a forwarding class represents a queue)   |
| Forwarding class            | Maps traffic to a queue. Classifiers map forwarding classes to IEEE 802.1p priorities. A forwarding class, an output queue, and an IEEE 802.1p priority are mapped to each other and identify the same traffic. (The IEEE 802.1p priority identifies incoming traffic, which is classified into the forwarding class, and the forwarding class is in turn mapped to an output queue for the traffic.) |
| Traffic control profile     | Configures scheduling for the forwarding class set (priority group), and associates a scheduler map with the forwarding class set to apply queue scheduling to the forwarding classes in the forwarding class set. Extra port bandwidth is shared among forwarding class sets in proportion to the minimum guaranteed bandwidth of each forwarding class set.   |
| Forwarding class set        | Name of a priority group. You map forwarding classes to forwarding class sets. A forwarding class set consists of one or more forwarding classes.   |

### Default Schedulers

Each forwarding class requires an associated scheduler. The default configuration uses only five forwarding classes, unicast best-effort (queue 0), fcoe (queue 3), no-loss (queue 4), network-control (queue 7), and multidestination (queue 8). You can use the default schedulers or you can define new schedulers for these five forwarding classes. For any other forwarding class, you must explicitly configure a scheduler.

[Table 508 on page 5870](#) shows the default schedulers.

Table 508: Default Schedulers

| Default Scheduler and Queue Number   | Guaranteed Rate (Minimum Bandwidth) | Shaping Rate (Maximum Bandwidth) | Excess Bandwidth Sharing | Priority | Buffer Size |
|--------------------------------------|-------------------------------------|----------------------------------|--------------------------|----------|-------------|
| Best-effort scheduler (queue 0)      | 5%                                  | None                             | 5%                       | Low      | 5%          |
| FCoE scheduler (queue 3)             | 35%                                 | None                             | 35%                      | Low      | 35%         |
| No-loss scheduler (queue 4)          | 35%                                 | None                             | 35%                      | Low      | 35%         |
| Network-control scheduler (queue 7)  | 5%                                  | None                             | 5%                       | Low      | 5%          |
| Multidestination scheduler (queue 8) | 20%                                 | None                             | 20%                      | Low      | 20%         |



**NOTE:** The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

By default, only the five default schedulers shown in [Table 508 on page 5870](#) have traffic mapped to them. Only the queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

Default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multidestination traffic. By default, unicast traffic consists of queue 0 (**best-effort** forwarding class), queue 3 (**fcoe** forwarding class), queue 4 (**no-loss** forwarding class), and queue 7 (**network-control** forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multidestination traffic (**mcast** queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multidestination traffic receives 2-Gbps of bandwidth.



**NOTE:** Multidestination queue 11 also receives a small amount of default bandwidth from the multidestination scheduler. CPU-generated multidestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multidestination traffic to a unicast forwarding class, that traffic uses queue 11.



Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

You should explicitly map traffic to non-default (unconfigured) queues if you want to use them to forward traffic. By default, unicast queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)



**NOTE:** All four multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

If the other unicast queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

### Transmit Rate (Minimum Guaranteed Bandwidth)

---

The transmit rate determines the minimum guaranteed bandwidth for each forwarding class. It also determines how much excess (extra) bandwidth each low-priority queue can share; each queue shares extra bandwidth in proportion to its transmit rate. You specify the rate in bits per second as a fixed value such as 1 Mbps or as a percentage of the total forwarding class set minimum guaranteed bandwidth (the guaranteed rate set in the traffic control profile). Either the default scheduler or a scheduler you configure allocates a portion of the outgoing interface bandwidth to each forwarding class.



**NOTE:** For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.

The allocated bandwidth can exceed the configured minimum rate if additional bandwidth is available from other queues in the forwarding class set. In case of congestion, the configured transmit rate is guaranteed for the queue. This property enables you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.



**NOTE:** Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

### Sharing Extra Bandwidth

---

Extra bandwidth is available to low-priority queues when the minimum guaranteed bandwidth of the queues does not use the full amount of forwarding class set bandwidth. This extra bandwidth is shared among the forwarding classes in the set based on the minimum guaranteed bandwidth of each queue.

For example, in a forwarding class set, Queue A has a transmit rate of 1 Gbps, Queue B has a transmit rate of 1 Gbps, and Queue C has a transmit rate of 2 Gbps. After servicing the minimum guaranteed bandwidth of these queues, the forwarding class set has an extra 2 Gbps of bandwidth available, and all three queues still have packets to forward. The queues receive the extra bandwidth in proportion to their transmit rates, so Queue A

receives an extra 500 Mbps, Queue B receives an extra 500 Mbps, and Queue C receives an extra 1 Gbps.

### Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth each forwarding class can consume. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total forwarding class set maximum bandwidth (the shaping rate set in the traffic control profile).

The maximum bandwidth for a queue depends on the total bandwidth available to the forwarding class set to which the queue belongs and how much bandwidth the other queues in the forwarding class set consume.



**NOTE:** On QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. A strict-high priority queue (or several queues with higher priorities than the starved queue) can consume all of the port bandwidth and prevent another queue from transmitting packets. To prevent a queue from being starved for bandwidth, you can configure a shaping rate on the queue or queues to prevent them from consuming all of the port bandwidth.



**NOTE:** We recommend that you always configure a shaping rate in the scheduler for strict-high priority queues to prevent them from starving other queues.

### Scheduling Priority

Scheduling priority determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic receive better access to the outgoing interface. The priority setting in the scheduler determines the priority for the queue.

Two levels of scheduling priority are supported:

- **Low**—Low-priority queues transmit traffic based on the weighted round robin (WRR) algorithm. The scheduler first determines if an individual queue is within its defined bandwidth profile. The scheduler then regularly reevaluates whether each individual queue is within its defined bandwidth profile and compares the amount of data the queue transmits to the amount of bandwidth the scheduler allocates to the queue. When the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when its transmitted amount is larger than its allocated amount. Out of profile queue data is transmitted only if bandwidth is available. Otherwise, it is buffered if buffer space is available. If no buffer space is available, the traffic may be dropped.
- **Strict-high**—You can configure only one queue as **strict-high** priority. The other 11 queues are **low** priority.

The **strict-high** priority queue receives preferential treatment over the low-priority queues. The **strict-high** priority queue receives all of its configured bandwidth before low-priority queues are serviced. Low-priority queues do not transmit traffic until the strict-high priority queue is empty. Carefully consider how much bandwidth you want to allocate to the **strict-high** priority queue to avoid starving the low-priority queues.

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multideestination forwarding class set.
- You cannot configure a minimum guaranteed bandwidth for a strict-high priority queue. (You cannot configure a transmit rate for a strict-high priority queue scheduler, and you cannot configure a guaranteed rate for a forwarding class set that has a strict-high priority queue.)
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Junos OS performs priority queueing using the following steps:

1. Services the strict-high priority queue before any other queues are served
2. Services the minimum bandwidth (transmit rate) of low-priority queues until the minimum is met or the queues are empty
3. Services all other low-priority queues and needs that exceed the minimum bandwidth

---

### Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. A drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type:

- PLP—Low, medium-high, high. You configure the PLP during classifier configuration. When you use a scheduler map to associate a forwarding class with a scheduler, you can use a drop-profile map to map different drop profiles to the forwarding class for different PLPs.
- Protocol type—Drop profiles match all protocol types.

---

### Buffer Size

Most of the total system buffer space is divided into two buffer pools, shared buffers and dedicated buffers. Shared buffers are a global pool that the ports share dynamically as needed. Dedicated buffers are a reserved portion of the buffer pool that is distributed

evenly to all of the ports. Each port receives an equal allocation of dedicated buffer space. The dedicated buffer allocation to ports is not configurable because it is reserved for the ports.

The queue buffers are allocated from the dedicated buffer pool assigned to the port. By default, ports divide their allocation of dedicated buffers among the egress queues in the same proportion as the default scheduler sets the minimum guaranteed transmission rates (**transmit-rate**) for traffic. Only the queues included in the default scheduler receive dedicated buffers.

If you do not use the default configuration, you can explicitly configure the queue buffer size in either of two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is mapped to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining port dedicated buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit scheduler means no dedicated buffer allocation for the queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.



**NOTE:** The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

For a complete discussion about queue buffer configuration in the context of ingress and egress port buffer configuration, see [“Understanding CoS Buffer Configuration” on page 5891](#).

### Explicit Congestion Notification

ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic because other traffic types use different methods of congestion notification—lossless traffic uses priority-based flow control (PFC) and strict-high priority traffic receives all of the port bandwidth it requires up to the point of a configured maximum rate.

## Scheduler Maps

---

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map, associate the scheduler map with a traffic control profile, and then associate the traffic control profile with an interface and a forwarding class set.

You can associate up to four user-defined scheduler maps with traffic control profiles.

### Related Documentation

- [Understanding Junos CoS Components on page 5789](#)
- [Understanding CoS Priority Group Scheduling on page 5877](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Buffer Configuration on page 5891](#)
- [Understanding CoS Explicit Congestion Notification on page 5926](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5886](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring Queue Scheduling Priority on page 6087](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Example: Configuring ECN on page 6090](#)

## Understanding CoS Priority Group Scheduling

Priority group scheduling defines the class-of-service (CoS) properties of a group of output queues (priorities). Priority group scheduling works with output queue scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates bandwidth to a group of queues (a priority group, called a forwarding class set in Junos OS configuration). Queue scheduling determines the portion of the priority group bandwidth that the particular queue can use.

You configure priority group scheduling in a traffic control profile and then associate the traffic control profile with a forwarding class set and an interface. You attach a scheduler map to the traffic control profile to specify the queue scheduling characteristics.



**NOTE:** When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- [Priority Group Scheduling Components on page 5877](#)
- [Default Traffic Control Profile on page 5878](#)
- [Guaranteed Rate \(Minimum Guaranteed Bandwidth\) on page 5878](#)
- [Sharing Extra Bandwidth on page 5878](#)
- [Shaping Rate \(Maximum Bandwidth\) on page 5879](#)
- [Scheduler Maps on page 5879](#)

### Priority Group Scheduling Components

[Table 509 on page 5877](#) provides a quick reference to the traffic control profile components you can configure to determine the bandwidth properties of priority groups, and [Table 510 on page 5878](#) provides a quick reference to some related scheduling configuration components.

**Table 509: Priority Group Scheduler Components**

| Traffic Control Profile Component | Description  |
|-----------------------------------|--|
| Guaranteed rate                   | Sets the minimum guaranteed port bandwidth for the priority group. Extra port bandwidth is shared among priority groups in proportion to the guaranteed rate of each priority group on the port. |
| Shaping rate                      | Sets the maximum port bandwidth the priority group can consume.  |
| Scheduler map                     | Maps schedulers to queues (forwarding classes, also called priorities). This determines the portion of the priority group bandwidth that a queue receives.                                       |

Table 510: Other Scheduling Components

| Other Scheduling Components | Description   |
|-----------------------------|---|
| Forwarding class            | Maps traffic to a queue (priority).   |
| Forwarding class set        | Name of a priority group. You map forwarding classes to priority groups. A forwarding class set consists of one or more forwarding classes. |
| Scheduler                   | Sets the bandwidth and scheduling priority of individual queues (forwarding classes).   |

### Default Traffic Control Profile

There is no default traffic control profile.

### Guaranteed Rate (Minimum Guaranteed Bandwidth)

The guaranteed rate determines the minimum guaranteed bandwidth for each priority group. It also determines how much excess (extra) port bandwidth the priority group can share; each priority group shares extra port bandwidth in proportion to its guaranteed rate. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total port bandwidth.

The minimum transmission bandwidth can exceed the configured rate if additional bandwidth is available from other priority groups on the port. In case of congestion, the configured guaranteed rate is guaranteed for the priority group. This property enables you to ensure that each priority group receives the amount of bandwidth appropriate to its level of service.



**NOTE:** Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

You cannot configure a guaranteed rate for forwarding class sets that include strict-high priority queues.

### Sharing Extra Bandwidth

Extra bandwidth is available to priority groups when the priority groups do not use the full amount of available port bandwidth. This extra port bandwidth is shared among the priority groups based on the minimum guaranteed bandwidth of each priority group.



For example, Port A has three priority groups: fc-set-1, fc-set-2, and fc-set-3. Fc-set-1 has a guaranteed rate of 2 Gbps, fc-set-2 has a guaranteed rate of 2 Gbps, and fc-set-3 has a guaranteed rate of 4 Gbps. After servicing the minimum guaranteed bandwidth of these priority groups, the port has an extra 2 Gbps of available bandwidth, and all three priority groups have still have packets to forward. The priority groups receive the extra bandwidth in proportion to their guaranteed rates, so fc-set-1 receives an extra 500 Mbps, fc-set-2 receives an extra 500 Mbps, and fc-set-3 receives an extra 1 Gbps.

### Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth the priority group can consume. You specify the rate in bits per second as a fixed value such as 5 Mbps or as a percentage of the total port bandwidth.

The maximum bandwidth for a priority group depends on the total bandwidth available on the port and how much bandwidth the other priority groups on the port consume.

### Scheduler Maps

A scheduler map maps schedulers to queues. When you associate a scheduler map with a traffic control profile, then associate the traffic control profile with an interface and a forwarding class set, the scheduling defined by the scheduler map determines the portion of the priority group resources that each individual queue can use.

You can associate up to four user-defined scheduler maps with traffic control profiles.

#### Related Documentation

- [Understanding Junos CoS Components on page 5789](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5886](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)

## Understanding CoS Traffic Control Profiles

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth that you assign to the forwarding class set in the traffic control profile.

This two-tier hierarchical scheduling architecture provides flexibility in allocating resources among queues and:

- Assigns a portion of port bandwidth to a priority group. You define the port resources for the priority group in a traffic control profile.
- Allocates priority group bandwidth among the queues that belong to the priority group. A scheduler map attached to the traffic control profile defines the amount of the priority group's resources that each queue can use.

Attaching a priority group and traffic control profile to a port defines the hierarchical scheduling properties of the group and the queues that belong to the group.

The ability to create priority groups supports enhanced transmission selection (ETS, described in IEEE 802.1Qaz). When a priority group does not use its allocated port bandwidth, ETS shares the excess port bandwidth among other priority groups on the port in proportion to their guaranteed minimum bandwidth (guaranteed rate). This utilizes the port bandwidth better than scheduling schemes that require setting strict priorities that reserve bandwidth for all groups whether it is needed or not. ETS allows traffic groups that need extra bandwidth to use it if the bandwidth is available, while preserving the ability to specify the minimum guaranteed bandwidth for traffic groups.

Traffic control profiles define the following CoS properties for priority groups:

- Minimum guaranteed bandwidth—Also known as the committed information rate (CIR). This is the minimum amount of port bandwidth the priority group receives. Priorities in the priority group receive their minimum guaranteed bandwidth as a portion of the priority group's minimum guaranteed bandwidth. The **guaranteed-rate** statement defines the minimum guaranteed bandwidth.



**NOTE:** You cannot apply a traffic control profile with a minimum guaranteed bandwidth to a priority group that includes strict-high priority queues.

- Shared excess (extra) bandwidth—When the priority groups on a port do not consume the full amount of bandwidth allocated to them or there is unallocated link bandwidth available, priority groups can contend for that extra bandwidth if they need it. Priorities in the priority group contend for extra bandwidth as a portion of the priority group's extra bandwidth. The amount of extra bandwidth for which a priority group can contend is proportional to the priority group's guaranteed minimum bandwidth (guaranteed rate).

- **Maximum bandwidth**—Also known as peak information rate (PIR). This is the maximum amount of port bandwidth the priority group receives. Priorities in the priority group receive their maximum bandwidth as a portion of the priority group's maximum bandwidth. The **shaping-rate** statement defines the maximum bandwidth.
- **Queue scheduling**—Each traffic control profile includes a scheduler map. The scheduler map maps priorities (forwarding classes) to schedulers to define the scheduling characteristics of the individual priorities in the priority group. The resources scheduled for each priority represent portions of the resources that the traffic control profile schedules for the entire priority group, not portions of the total link bandwidth. The **scheduler-maps** statement defines the mapping of forwarding classes to schedulers.

#### Related Documentation

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)

## Understanding CoS Priority Group and Queue Guaranteed Rates (Minimum Bandwidth)

You can set a guaranteed minimum bandwidth for individual forwarding classes (queues) and for groups of forwarding classes called forwarding class sets (priority groups). Setting a minimum guaranteed bandwidth ensures that priority groups and queues receive the bandwidth required to support the expected traffic.

This topic covers:

- [Guaranteeing Bandwidth Using Hierarchical Scheduling on page 5881](#)
- [Priority Group Guaranteed Rate \(Minimum Bandwidth\) on page 5883](#)
- [Queue Transmit Rate \(Minimum Bandwidth\) on page 5883](#)


### Guaranteeing Bandwidth Using Hierarchical Scheduling

The **guaranteed-rate** value for the priority group defines the minimum amount of bandwidth allocated to a forwarding class set on a port, whereas the **transmit-rate** value of the queue defines the minimum amount of bandwidth allocated to a particular queue in a priority group. The queue bandwidth is a portion of the priority group bandwidth.



**NOTE:** You cannot configure a minimum guaranteed bandwidth (transmit rate) for a forwarding class that is mapped to a strict-high priority queue, and you cannot configure a minimum guaranteed bandwidth (guaranteed rate) for a priority group that includes strict-high priority queues.

Figure 206 on page 5882 shows how the total port bandwidth is allocated to priority groups (forwarding class sets) based on the guaranteed rate of each priority group. It also shows how the guaranteed bandwidth of each priority group is allocated to the queues in the priority group based on the transmit rate of each queue.

**Figure 206: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling**

The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue transmit rates cannot exceed the total guaranteed rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.



**NOTE:** You must set both the priority group **guaranteed-rate** value and the queue **transmit-rate** value in order to configure the minimum bandwidth for individual queues. If you set the **transmit-rate** value but do not set the **guaranteed-rate** value, the configuration fails.

You can set the **guaranteed-rate** value for a priority group without setting the **transmit-rate** value for individual queues in the priority group. However, queues that do not have a configured **transmit-rate** value can become starved for bandwidth if other higher-priority queues need the priority group's bandwidth. To avoid starving a queue, it is a good practice to configure a **transmit-rate** value for most queues.

If you configure the guaranteed rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

### Priority Group Guaranteed Rate (Minimum Bandwidth)

Setting a priority group **guaranteed-rate** enables you to reserve a portion of the port bandwidth for the forwarding classes (queues) in that forwarding class set. The minimum bandwidth (**guaranteed-rate**) that you configure for a priority group sets the minimum bandwidth available to all of the forwarding classes in the forwarding class set.

The combined **guaranteed-rate** value of all of the forwarding class sets associated with an interface cannot exceed the amount of bandwidth available on that interface.

You configure the priority group **guaranteed-rate** in the traffic control profile. You cannot apply a traffic control profile that has a guaranteed rate to a priority group that includes strict-high priority queues.

### Queue Transmit Rate (Minimum Bandwidth)

Setting a queue **transmit-rate** enables you to reserve a portion of the priority group bandwidth for the individual queue. For example, a queue that handles Fibre Channel over Ethernet (FCoE) traffic might require a minimum rate of 4 Gbps to ensure the class of service that storage area network (SAN) traffic requires.

The priority group **guaranteed-rate** sets the aggregate minimum amount of bandwidth available to the queues that belong to the priority group. The cumulative total minimum bandwidth the queues consume cannot exceed the minimum bandwidth allocated to the priority group to which they belong. (The combined transmit rates of the queues in a priority group cannot exceed the priority group's guaranteed rate.)

You must configure the **guaranteed-rate** value of the priority group in order to set a **transmit-rate** value for individual queues that belong to the priority group. The reason is that if there is no guaranteed bandwidth for a priority group, there is no way to guarantee bandwidth for queues in that priority group.

You configure the queue **transmit-rate** in the scheduler configuration. You cannot configure a transmit rate for strict-high priority queues.

**Related  
Documentation**

- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)

## Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)

If the amount of traffic on an interface exceeds the maximum bandwidth of the interface, it leads to congestion. You can use priority group (forwarding class set) shaping and queue shaping to manage the excess traffic and avoid congestion.

The maximum bandwidth sets the most bandwidth a priority group or a queue can use after all of the priority group and queue minimum bandwidth requirements are met, even if more bandwidth is available.

This topic covers:

- [Priority Group Shaping on page 5884](#)
- [Queue Shaping on page 5884](#)
- [Shaping Maximum Bandwidth Using Hierarchical Scheduling on page 5885](#)

### Priority Group Shaping

---

Priority group shaping enables you to shape the aggregate traffic of a forwarding class set on a port to a maximum rate that is less than the line or port rate. The maximum bandwidth (**shaping-rate**) that you configure for a priority group sets the maximum bandwidth available to all of the forwarding classes (queues) in the forwarding class set.

If a port has more than one priority group and the combined **shaping-rate** value of the priority groups is greater than the amount of port bandwidth available, the bandwidth is shared proportionally among the priority groups.

You configure the priority group **shaping-rate** in the traffic control profile.

### Queue Shaping

---

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-high priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues.



**NOTE:** We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Similarly, for any queue, you can configure queue shaping (**shaping-rate**) to set the maximum bandwidth for a particular queue.

The **shaping-rate** value of the priority group sets the aggregate maximum amount of bandwidth available to the queues that belong to the priority group. The cumulative total bandwidth the queues consume cannot exceed the maximum bandwidth of the priority group to which they belong on a port.

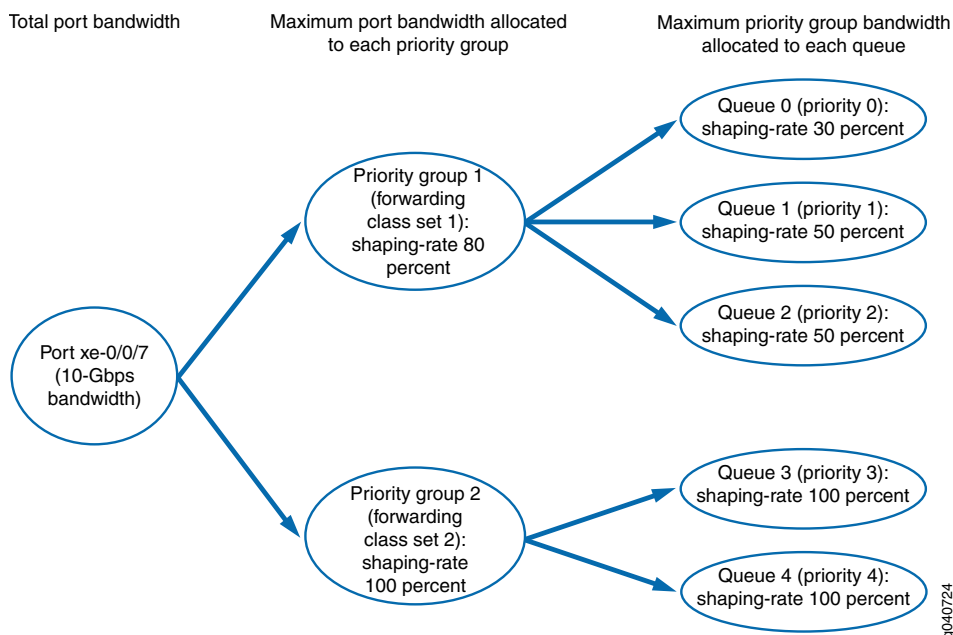
If a priority group has more than queue and the combined **shaping-rate** value of the queues is greater than the amount of bandwidth available to the priority group, the bandwidth is shared proportionally among the queues.

You configure the queue **shaping-rate** in the scheduler configuration.

### Shaping Maximum Bandwidth Using Hierarchical Scheduling

Priority group shaping defines the maximum bandwidth allocated to a forwarding class set on a port, whereas queue shaping defines a limit on maximum bandwidth usage per queue. The queue bandwidth is a portion of the priority group bandwidth.

[Figure 207 on page 5886](#) shows how the port bandwidth is allocated to priority groups (forwarding class sets) based on the shaping rate of each priority group, and how the bandwidth of each priority group is allocated to the queues in the priority group based on the shaping rate of each queue.

**Figure 207: Setting Maximum Bandwidth Using Hierarchical Scheduling****Related Documentation**

- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)

**Understanding CoS Scheduling Behavior and Configuration Considerations**

Many factors affect scheduling configuration and bandwidth requirements, including:

- When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.
- When you define a forwarding class that will be used on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:
  - Mapping a scheduler to the forwarding class in a scheduler map
  - Including the forwarding class in a forwarding class set



- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface
- On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.
- For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.
- Configuring the minimum guaranteed bandwidth (**transmit-rate**) for a queue (forwarding class) does not work unless you also configure the minimum guaranteed bandwidth (**guaranteed-rate**) for the priority group (forwarding class set) in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.) If you configure transmit rates whose sum exceeds the guaranteed rate of the forwarding class set, the commit check fails and the system rejects the configuration.

- The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.
- If you configure the **guaranteed-rate** of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.
- There are several factors to consider if you want to configure strict-high priority queues:
  - You cannot configure a minimum guaranteed bandwidth (**transmit-rate**) for a strict-high priority queue. You cannot configure a minimum guaranteed bandwidth (**guaranteed-rate**) for a forwarding class set that includes a strict-high priority queue.
  - You must create a separate forwarding class set for the strict-high priority queue.
  - Only one forwarding class set can contain strict-high priority queues.
  - Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
  - A strict-high priority queue cannot belong to a multidestination forwarding class set.
  - We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. Failure of a queue to transmit packets for 12 consecutive seconds may be due to:
  - A strict-high priority queue consuming all of the port bandwidth
  - Several queues consuming all of the port bandwidth
  - Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
  - Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

If the cause is a strict-high priority queue consuming all of the port bandwidth, use rate shaping to configure a maximum rate for the strict-high priority queue and prevent it from using all of the port bandwidth. To configure rate shaping, include the **shaping-rate (rate | percent percentage)** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level and apply the shaping rate to the strict-high priority scheduler. We recommend that you always apply a shaping rate to strict-high priority traffic to prevent the strict-high priority queue from starving other queues.

If several queues consume all of the port bandwidth, you can use a scheduler to rate shape those queues and prevent them from using all of the port bandwidth.

- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.
- When you set the maximum bandwidth for a queue or for a priority group (**shaping-rate**) at 100 Kbps or lower, the traffic shaping behavior is accurate only within +/– 20 percent of the configured **shaping-rate**.
- Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



**NOTE:** Do not configure drop profiles for the **fcoe** and **no-loss** forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

- On an ingress port, do not configure classifiers that map the same IEEE 802.1p code point to both a multdestination traffic flow and a lossless unicast traffic flow (such as the default lossless **fcoe** or **no-loss** forwarding classes). Any code point used for multdestination traffic on a port should not be used to classify unicast traffic into a lossless forwarding class on the same port.

If a multdestination traffic flow and a lossless unicast traffic flow use the same code point on a port, the multdestination traffic is treated the same way as the lossless traffic. For example, if priority-based flow control (PFC) is applied to the lossless traffic, the multdestination traffic of the same code point is also paused. During periods of congestion, treating multdestination traffic the same as lossless unicast traffic can create ingress port congestion for the multdestination traffic and affect the multdestination traffic on all of the egress ports the multdestination traffic uses.

For example, the following configuration can cause ingress port congestion for the multdestination flow:

1. For unicast traffic, IEEE 802.1p code point 011 is classified into the **fcoe** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class fcoe
loss-priority low code-points 011
```

2. For multdestination traffic, IEEE 802.1p code point 011 is classified into the **mcast** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast
loss-priority low code-points 011
```

3. The unicast classifier that maps traffic with code point 011 to the **fcoe** forwarding class is mapped to interface **xe-0/0/1**:

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl
```

4. The multdestination classifier that maps traffic with code point 011 to the **mcast** forwarding class is mapped to all interfaces (multdestination traffic maps to all interfaces and cannot be mapped to individual interfaces):

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl
```

Because the same code point (011) maps unicast traffic to a lossless traffic flow and also maps multdestination traffic to a multdestination traffic flow, the multdestination traffic flow might experience ingress port congestion during periods of congestion.

To avoid ingress port congestion, do not map the code point used by the multdestination traffic to lossless unicast traffic. For example:

1. Instead of classifying code point **011** into the **fcoe** forwarding class, classify code point **011** into the **best-effort** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class  
best-effort loss-priority low code-points 011
```

2. user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast  
loss-priority low code-points 011

3. user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl

4. user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl

Because the code point **011** does not map unicast traffic to a lossless traffic flow, the multidestination traffic flow does not experience ingress port congestion during periods of congestion.

The best practice is to classify unicast traffic with IEEE 802.1p code points that are also used for multidestination traffic into best-effort forwarding classes.

#### Related Documentation

- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Priority Group Scheduling on page 5877](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [\*Benefits of Configuring CoS Hierarchical Port Scheduling\*](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)

## Understanding CoS Buffer Configuration

Each QFX3500 and QFX3600 switch has 9 MB of Packet Forwarding Engine (PFE) wide common packet buffer memory that is used to store packets on interface queues. Each QFX5100 and EX4600 switch has 12 MB of PFE wide common packet buffer memory. The buffer memory has separate ingress and egress accounting to make accept, drop, or pause decisions. Because the switch has a single pool of memory with separate ingress and egress accounting, the full amount of buffer memory is available from both the ingress and the egress perspective. Packets are accounted for as they enter and leave the switch, but there is no concept of a packet arriving at an ingress buffer and then being moved to an egress buffer.

The buffers are divided into two pools from both an ingress and an egress perspective:

1. *Shared buffers* are a global memory pool that the switch allocates dynamically to ports as needed, so the buffers are shared among the switch ports.
2. *Dedicated buffers* are a memory pool divided equally among the switch ports. Each port receives a minimum guaranteed amount of buffer space, dedicated to each port, not shared among ports.



**NOTE:** Lossless traffic is traffic on which you enable priority-based flow control (PFC) to ensure lossless transport. Lossless traffic does not refer to best-effort traffic on a link enabled for Ethernet PAUSE (IEEE 802.3x).

The switch reserves nonconfigurable buffer space to ensure that ports and queues receive a minimum memory allocation. You can configure how the system uses the rest of the buffer space to optimize the allocation for your mix of network traffic. You can configure the percentage of available buffer space used as shared buffer space versus dedicated buffer space. You can also configure how shared buffer space is allocated to different types of traffic. You can optimize the buffer settings for the traffic on your network.

The default buffer configuration is designed for networks that have a balance of best-effort and lossless traffic.

The default class-of-service configuration provides two lossless forwarding classes (**fcoe** and **no-loss**), a best-effort unicast forwarding class, a network control traffic forwarding class, and one multidestination (multicast, broadcast, and destination lookup fail) forwarding class. Each default forwarding class maps to a different default output queue. The default configuration allocates the buffers in a manner that supports a moderate amount of lossless traffic while still providing the ability to absorb bursts in best-effort traffic transmission.

Changing the buffer settings changes the abilities of the buffers to absorb traffic bursts and handle lossless traffic. For example, networks with mostly best-effort traffic require allocating most of the shared buffer space to best-effort buffers. This provides deep, flexible buffers that can absorb traffic bursts with minimal packet loss, at the expense of buffer availability for lossless traffic.

Conversely, networks with mostly lossless traffic require allocating most of the shared buffer space to lossless headroom buffers. This prevents packet loss on lossless flows at the expense of absorbing bursty best-effort traffic efficiently.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

This topic describes the buffer architecture and settings:

- [Buffer Pools on page 5892](#)
- [Default Buffer Pool Values on page 5900](#)
- [Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios on page 5903](#)
- [Optimizing Buffer Configuration on page 5906](#)
- [General Buffer Configuration Rules and Considerations on page 5908](#)

---

## Buffer Pools

From both an ingress and an egress perspective, the PFE buffer is split into two main pools, a shared buffer pool and a dedicated buffer pool that ensures a minimum allocation to each port. You can configure the amount of buffer space allocated to each of the two pools. A portion of the buffer space is reserved so that there is always a minimum amount of shared and dedicated buffer space available to each port.

- **Shared buffer pool**—A global memory space that all of the ports on the switch share dynamically as they need buffers. The shared buffer pool is further partitioned into buffers for best-effort unicast, best-effort multdestination (broadcast, multicast, and destination lookup fail), and PFC (lossless) traffic types. You can allocate global shared memory space to buffer partitions to better support different mixes of network traffic. The larger the shared buffer pool, the better the switch can absorb traffic bursts because more shared memory is available for the traffic.
- **Dedicated buffer pool**—A reserved global memory space allocated equally to each port. The switch reserves a minimum dedicated buffer pool that is not user-configurable. You can divide the dedicated buffer allocation for a port among the port queues on a per-port, per-queue basis. (For example, this enables you to dedicate more buffer space to queues that transport lossless traffic.)

A larger dedicated buffer pool means a larger amount of dedicated buffer space for each port, so congestion on one port is less likely to affect traffic on another port because the traffic does not need to use as much shared buffer space. However, the larger the dedicated buffer pool, the less bursty traffic the switch can handle because there is less dynamic shared buffer memory.

You can configure the way the available unreserved portion of the buffer space is allocated to the global shared buffer pool and to the dedicated shared buffer pool by configuring the ingress and egress shared buffer percentages.

By default, 100 percent of the available unreserved buffer space is allocated to the shared buffer pool. If you change the percentage of space allocated to the shared buffer, the available buffer space that is not allocated to the shared buffer is allocated to the dedicated buffer. For example, if you configure the ingress shared buffer pool as 80 percent, the remaining 20 percent of the available buffer space is allocated to the dedicated buffer pool and divided equally across the ports.



**NOTE:** When 100 percent of the available (user-configurable) buffers are allocated to the shared buffer pool, the switch still reserves a minimum dedicated buffer pool.

You can separately configure ingress and egress shared buffer pool allocations. You can also partition the ingress and egress shared buffer pool to allocate percentages of the shared buffer pool to specific types of traffic. If you do not use the default configuration or one of the recommended configurations, pay particular attention to the ingress configuration of the lossless and lossless headroom buffers (these buffers handle PFC pause during periods of congestion) and to the egress configuration of the best-effort buffers to handle incast congestion (multiple synchronized sources sending data to the same receiver in parallel).

In addition to the shared buffer pool and the dedicated buffer pool, there is also a small ingress global headroom buffer pool that is reserved and is not configurable.

When contention for buffer space occurs, the switch uses an internal algorithm to ensure that the buffer pools are distributed fairly among competing flows. When traffic for a given flow exceeds the amount of dedicated port buffer reserved for that flow, the flow begins to consume memory from the dynamic shared buffer pool. Competing flows compete for shared buffer memory with other flows that also have exhausted their dedicated buffers. When there is no congestion, there are no competing flows.

- [Buffer Handling of Lossless Flows \(PFC\) Versus Ethernet PAUSE on page 5893](#)
- [Shared Buffer Pool and Partitions on page 5894](#)
- [Dedicated Port Buffer Pool and Buffer Allocation to Queues on page 5895](#)
- [Trade-off Between Shared Buffer Space and Dedicated Buffer Space on page 5899](#)
- [Order of Buffer Consumption on page 5899](#)

### ***Buffer Handling of Lossless Flows (PFC) Versus Ethernet PAUSE***

When we discuss lossless buffers in the following sections, we mean buffers that handle traffic on which you enable PFC to ensure lossless transport. The lossless buffers are not used for best-effort traffic on a link on which you enable Ethernet PAUSE (IEEE 802.3x). The lossless ingress and egress shared buffers, and the ingress lossless headroom shared buffer, are used only for traffic on which you enable PFC.



**NOTE:** To support lossless flows, you must configure the appropriate data center bridging capabilities (PFC, DCBX, or ETS) and scheduling properties.

### ***Shared Buffer Pool and Partitions***

The shared buffer pool is a global memory space that all of the ports on the switch share dynamically as they need buffers. The switch uses the shared buffer pool to absorb traffic bursts after the dedicated buffer pool for a port is exhausted.

You can divide both the ingress shared buffer pool and the egress shared buffer pool into three partitions to allocate percentages of each buffer pool to different types of traffic. When you partition the ingress or egress shared buffer pool:

- If you explicitly configure one ingress shared buffer partition, you must explicitly configure all three ingress shared buffer partitions. (You either explicitly configure all three ingress partitions or you use the default setting for all three ingress partitions.)

If you explicitly configure one egress shared buffer partition, you must explicitly configure all three egress shared buffer partitions. (You either explicitly configure all three egress partitions or you use the default setting for all three egress partitions.)

The switch returns a commit error if you do not explicitly configure all three partitions when configuring the ingress or egress shared buffer partitions.

- The combined percentages of the three ingress shared buffer partitions must total exactly 100 percent.

The combined percentages of the three egress shared buffer partitions must total exactly 100 percent.

When you explicitly configure ingress or egress shared buffer partitions, the switch returns a commit error if the total percentage of the three partitions does not equal 100 percent.

- If you explicitly partition one set of shared buffers, you do not have to explicitly partition the other set of shared buffers. For example, you can explicitly configure the ingress shared buffer partitions and use the default egress shared buffer partitions. However, if you change the buffer partitions for the ingress buffer pool to match the expected types of traffic flows, you would probably also want to change the buffer partitions for the egress buffer pool to match those traffic flows.

You can configure the percentage of available unreserved buffer space allocated to the shared buffer pool. Space that you do not allocate to the shared buffer pool is added to the dedicated buffer pool and divided equally among the ports. The default configuration allocates 100 percent of the unreserved ingress and egress buffer space to the shared buffers.

Configuring the ingress and egress shared buffer pool partitions enables you to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic.

### ***Ingress Shared Buffer Pool Partitions***

You can configure three ingress buffer pool partitions:

- Lossless buffers—Shared buffer pool for all lossless ingress traffic. The recommended minimum value for lossless buffers is 5 percent.



- **Lossless headroom buffers**—Shared buffer pool for packets received while a pause is asserted. If PFC is enabled on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers for which the recommended value can be less than 5 percent.)
- **Lossy buffers**—Shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The recommended minimum value for best-effort buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and best-effort buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. If you explicitly configure an ingress shared buffer partition, you must explicitly configure all three ingress buffer partitions, even if the lossless headroom buffer partition has a value of 0 (zero) percent.

### ***Egress Shared Buffer Pool Partitions***

You can configure three egress buffer pool partitions:

- **Lossless buffers**—Shared buffer pool for all lossless egress queues. The recommended minimum value for lossless buffers is 5 percent.
- **Lossy buffers**—Shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The recommended minimum value for best-effort buffers is 5 percent.
- **Multicast buffers**—Shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The recommended minimum value for multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and should have a value of at least 5 percent. If you explicitly configure an egress shared buffer partition, you must explicitly configure all three egress buffer partitions, and each partition should have a value of at least 5 percent.

### ***Dedicated Port Buffer Pool and Buffer Allocation to Queues***

The global dedicated buffer pool is memory that is allocated equally to each port, so each port receives a guaranteed minimum amount of buffer space. Dedicated buffers are not shared among ports. Each port receives an equal proportion of the dedicated buffer pool.

The amount of dedicated buffer space is not user-configurable and depends on the percentage of available nonreserved buffers allocated to the shared buffers. (The dedicated buffer space is equal to the minimum reserved port buffers plus the remainder of the available nonreserved buffers that are not allocated to the shared buffer pool.)

When traffic enters and exits the switch, the switch ports use their dedicated buffers to store packets. If the dedicated buffers are not sufficient to handle the traffic, the switch uses shared buffers. The only way to increase the dedicated buffer pool is to decrease the shared buffer pool from its default value of 100 percent of available unreserved buffers.



**NOTE:** If 100 percent of the available unreserved buffers are allocated to the shared buffer pool, the switch still reserves a minimum dedicated buffer pool.

The larger the shared buffer pool, the better the burst absorption across the ports. The larger the dedicated buffer pool, the larger the amount of dedicated buffer space for each port. The greater the dedicated buffer space, the less likely that congestion on one port can affect traffic on another port, because the traffic does not need to use as much shared buffer space.

### *Allocating Dedicated Port Buffers to Queues*

You can divide the dedicated buffer allocation for an egress port among the port queues by including the **buffer-size** statement in the scheduler configuration. This enables you to control the egress port dedicated buffer allocation on a per-port, per-queue basis. (For example, this enables you to dedicate more buffer space to queues that transport lossless traffic, or to stop the port from reserving buffers for queues that do not carry traffic.) Egress dedicated port buffer allocation is a hierarchical structure that allocates a global dedicated buffer pool evenly among ports, and then divides the allocation for each port among the port queues.

By default, ports divide their allocation of dedicated buffers among their egress queues in the same proportion as the default scheduler sets the minimum guaranteed transmission rates (the **transmit-rate** option) for traffic. Only the queues included in the default scheduler receive bandwidth and dedicated buffers, in the proportions shown in [Table 511 on page 5896](#):

**Table 511: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler)**

| Forwarding Class | Queue | Minimum Guaranteed Bandwidth (transmit-rate) | Proportion of Reserved Dedicated Port Buffers |
|------------------|-------|--|---|
| best-effort      | 0     | 5%   | 5%  |
| fcoe             | 3     | 35%  | 35%   |
| no-loss          | 4     | 35%  | 35%   |
| network-control  | 7     | 5%   | 5%  |
| mcast            | 8     | 20%  | 20%   |

In the default configuration, no egress queues other than the ones shown in [Table 511 on page 5896](#) receive an allocation of dedicated port buffers.



**NOTE:** The switch uses hierarchical scheduling to control port and queue bandwidth allocation, as described in “[Understanding CoS Hierarchical Port Scheduling \(ETS\)](#)” on page 5862 and shown in “[Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)](#)” on page 5966. For egress queue buffer size configuration, when you attach a traffic control profile (includes the queue scheduler information) to a port, the dedicated egress buffers on the port are divided among the queues as configured in the scheduler.

If you do not want to use the default allocation of dedicated port buffers to queues, use the **buffer-size** option in the scheduler that is attached to the port to configure the queue allocation. You can configure the dedicated buffer allocation to queues in two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is attached to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining dedicated port buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit scheduler for a queue means no dedicated buffer allocation for that queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.



**NOTE:** The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

### ***Configuring Dedicated Port Buffer Allocation to Queues***

In a port configuration that includes multiple forwarding class sets, with multiple forwarding classes mapped to multiple schedulers, the allocation of port dedicated buffers to queues depends on the mix of queues with buffer sizes configured as explicit percentages and queues configured with (or defaulted to) the **remainder** option.

The best way to demonstrate how using the percentage and remainder options affects dedicated port buffer allocation to queues is by showing an example of queue buffer allocation, and then showing how the queue buffer allocation changes when you add another forwarding class (queue) to the port.

[Table 512 on page 5898](#) shows an initial configuration that includes four forwarding class sets, the five default forwarding classes (mapped to the five default queues for those forwarding classes), the **buffer-size** option configuration, and the resulting buffer allocation for each queue. [Table 513 on page 5898](#) shows the same configuration after we add another forwarding class (best-effort-2, mapped to queue 1) to the best-effort forwarding class set. Comparing the buffer allocations in each table shows you how adding another queue

affects buffer allocation when you use remainders and explicit percentages to configure the buffer allocation for different queues.

**Table 512: Egress Queue Dedicated Buffer Allocation (Example 1)**

| Forwarding Class Set (Priority Group) | Forwarding Class | Queue | Scheduler Buffer Size Configuration | Buffer Allocation per Queue (Percentage) |
|---------------------------------------|------------------|-------|-------------------------------------|--|
| fc-set-be                             | best-effort      | 0     | 10%                                 | 10%                                      |
| fc-set-lossless                       | fcoe             | 3     | 20%                                 | 20%                                      |
|                                       | no-loss          | 4     | 40%                                 | 40%                                      |
| fc-set-strict-high                    | network-control  | 7     | remainder                           | 15%                                      |
| fc-set-mcast                          | mcast            | 8     | remainder                           | 15%                                      |

In this first example, 70 percent of the egress port dedicated buffer pool is explicitly allocated to the best-effort, fcoe, and no-loss queues. The remaining 30 percent of the port dedicated buffer pool is split between the two queues that use the **remainder** option (network-control and mcast), so each queue receives 15 percent of the dedicated buffer pool.

Now we add another forwarding class (queue) to the best-effort priority group (fc-set-be) and configure it with a buffer size of *remainder* instead of configuring a specific percentage. Because a third queue now shares the remaining dedicated buffers, the queues that share the remainder receive fewer dedicated buffers, as shown in [Table 513 on page 5898](#). The queues with explicitly configured percentages receive the configured percentage of dedicated buffers.

**Table 513: Egress Queue Dedicated Buffer Allocation with Another Remainder Queue (Example 2)**

| Priority Group (fc-set) | Forwarding Class | Queue | Scheduler Buffer Size Configuration | Buffer Allocation per Queue (Percentage) |
|-------------------------|------------------|-------|-------------------------------------|--|
| fc-set-be               | best-effort      | 0     | 10%                                 | 10%                                      |
|                         | best-effort-2    | 1     | remainder                           | 10%                                      |
| fc-set-lossless         | fcoe             | 3     | 20%                                 | 20%                                      |
|                         | no-loss          | 4     | 40%                                 | 40%                                      |
| fc-set-strict-high      | network-control  | 7     | remainder                           | 10%                                      |
| fc-set-mcast            | mcast            | 8     | remainder                           | 10%                                      |

The two tables show how the port divides the dedicated buffer space that remains after servicing the queues that have an explicitly configured percentage of dedicated buffer space.

### ***Trade-off Between Shared Buffer Space and Dedicated Buffer Space***

The trade-off between shared buffer space and dedicated buffer space is:

- Shared buffers provide better absorption of traffic bursts because there is a larger pool of dynamic buffers that ports can use as needed to handle the bursts. However, all flows that exhaust their dedicated buffer space compete for the shared buffer pool. A larger shared buffer pool means a smaller dedicated buffer pool, and therefore more competition for the shared buffer pool because more flows exhaust their dedicated buffer allocation. Too much shared buffer space results in no single flow receiving very much shared buffer space, to maintain fairness when many flows contend for that space.
- Dedicated buffers provide guaranteed buffer space to each port. The larger the dedicated buffer pool, the less likely that congestion on one port affects traffic on another port, because the traffic does not need to use as much shared buffer space. However, less shared buffer space means less ability to dynamically absorb traffic bursts.

For optimal burst absorption, the switch needs enough dedicated buffer space to avoid persistent competition for the shared buffer space. When fewer flows compete for the shared buffers, the flows that need shared buffer space to absorb bursts receive more of the shared buffer because fewer flows exhaust their dedicated buffer space.

The default configuration and all of the recommended configurations allocate 100 percent of the user-configurable memory space to the global shared buffer pool because the amount of space reserved for dedicated buffers provides enough space to avoid persistent competition for dynamic shared buffers. This results in fewer flows competing for the shared buffers, so the competing flows receive more of the buffer space.

### ***Order of Buffer Consumption***

The total buffer pool is divided into ingress and egress shared buffer pools and dedicated buffer pools. When traffic flows through the switch, the buffer space is used in a particular order that depends on the type of traffic.

On ingress, the order of buffer consumption is:

- Best-effort unicast traffic:
  1. Dedicated buffers
  2. Shared buffers
  3. Global headroom buffers (very small)
- Lossless unicast traffic:
  1. Dedicated buffers
  2. Shared buffers

3. Lossless headroom buffers
  4. Global headroom buffers (very small)
- Multidestination traffic:
    1. Dedicated buffers
    2. Shared buffers
    3. Global headroom buffers (very small)

On egress, the order of buffer consumption is the same for unicast best-effort, lossless unicast, and multidestination traffic:

- Dedicated buffers
- Shared buffers

In all cases on all ports, the switch uses the dedicated buffer pool first and the shared buffer pool only after the dedicated buffer pool for the port or queue is exhausted. This reserves the maximum amount of dynamic shared buffer space to absorb traffic bursts.

### Default Buffer Pool Values

---

You can view the default or configured ingress and egress buffer pool values in KB units using the **show class-of-service shared-buffer** operational command. You can view the configured shared buffer pool values in percent units using the **show configuration class-of-service shared-buffer** operational command.

This section provides the default total buffer, shared buffer, and dedicated buffer values.

- [Total Buffer Pool Size on page 5900](#)
- [Shared Buffer Pool Default Values on page 5900](#)
- [Dedicated Buffer Pool Default Values on page 5902](#)

#### **Total Buffer Pool Size**

The total buffer pool is common memory that has separate ingress and egress accounting, so the full buffer pool is available from both the ingress and egress perspective. The total buffer pool consists of the dedicated buffer space and the shared buffer space. The size of the total buffer pool is not user-configurable, but the allocation of buffer space to the dedicated and shared buffer pools is user-configurable.

On QFX3500 and QFX3600 switches, the combined total size of the ingress and egress buffer pools is approximately 9 MB (exactly 9360 KB).

On QFX5100 and EX4600 switches, the combined total size of the ingress and egress buffer pools is approximately 12 MB (exactly 12480 KB).

#### **Shared Buffer Pool Default Values**

The QFX5100 and EX4600 switches have a larger shared buffer pool (12 MB) than QFX3500 and QFX3600 switches (9 MB). However, the allocation of shared buffer space to the individual ingress and egress buffer pools is the same on a percentage basis, even

though the absolute values are different. For example, the default ingress lossless buffer is 9 percent of the total shared ingress buffer space on QFX5100, EX4600, QFX3500, and QFX3600 switches, even though the default absolute value of the ingress lossless buffer is 861.05KB on QFX5100 and EX4600 switches, and 648.18KB on QFX3500 and QFX3600 switches.

This section describes the default values in percent and in KB for the shared ingress and shared egress buffers.

- [Shared Ingress Buffer Default Values on page 5901](#)
- [Shared Egress Buffer Default Values on page 5901](#)

#### **Shared Ingress Buffer Default Values**

The QFX5100 and EX4600 switches have a larger shared ingress buffer than the QFX3500 and QFX3600 switches. [Table 514 on page 5901](#) shows the default ingress shared buffer allocation values in KB units for QFX5100 and EX4600 switches.

**Table 514: QFX5100 and EX4600 Switch Default Shared Ingress Buffer Values (KB)**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 9567.19 KB                  | 861.05 KB       | 4305.23 KB               | 4400.91 KB   |

[Table 515 on page 5901](#) shows the default ingress shared buffer allocation values in KB units for QFX3500 and QFX3600 switches.

**Table 515: QFX3500 and QFX3600 Switch Default Shared Ingress Buffer Values (KB)**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 7202 KB                     | 648.18 KB       | 3240.9 KB                | 3312.92 KB   |

[Table 516 on page 5901](#) shows the default ingress shared buffer allocation values as percentages for QFX5100, EX4600, QFX3500, and QFX3600 switches. (If you change the default shared buffer allocation, you configure the change as a percentage.)

**Table 516: Default Shared Ingress Buffer Values (Percentage)**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 100%                        | 9%              | 45%                      | 46%          |

#### **Shared Egress Buffer Default Values**

The QFX5100 and EX4600 switches have a larger shared egress buffer than the QFX3500 and QFX3600 switches. [Table 517 on page 5902](#) shows the default egress shared buffer allocation values in KB units for QFX5100 and EX4600 switches.

**Table 517: QFX5100 and EX4600 Switch Default Shared Egress Buffer Values (KB)**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 8736 KB                    | 4368 KB         | 2708.16 KB   | 1659.84 KB       |

Table 518 on page 5902 shows the default egress shared buffer allocation values in KB units.

**Table 518: QFX3500 and QFX3600 Switch Default Shared Egress Buffer Values (KB)**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 6656 KB                    | 3328 KB         | 2063.36 KB   | 1264.64 KB       |

Table 519 on page 5902 shows the default egress shared buffer allocation values as percentages.

**Table 519: Default Shared Egress Buffer Values (Percentage)**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 50%             | 31%          | 19%              |

#### ***Dedicated Buffer Pool Default Values***

The system reserves ingress and egress dedicated buffer pools that are divided equally among the switch ports. By default, the system allocates 100 percent of the available unreserved buffer space to the shared buffer pool. If you reduce the percentage of available unreserved buffer space allocated to the shared buffer pool, the remaining unreserved buffer space is added to the dedicated buffer pool allocation. You configure the amount of dedicated buffer pool space by reducing (or increasing) the percentage of buffer space allocated to the shared buffer pool. You do not directly configure the dedicated buffer pool allocation.

The default dedicated buffer pool values for QFX3500 and QFX3600 switches in KB units are:

- Ingress dedicated buffer—2158 KB
- Egress dedicated buffer—2704.0 KB

The default dedicated buffer pool values for QFX5100 switches in KB units are:

- Ingress dedicated buffer—2912.81 KB
- Egress dedicated buffer—3744 KB



## Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios

The way you configure the shared buffer pool depends on the mix of traffic on your network. This section provides shared buffer configuration recommendations for five basic network traffic scenarios:

- **Balanced traffic**—The network carries a balanced mix of unicast best-effort, lossless, and multicast traffic. (This is the default configuration.)
- **Best-effort unicast traffic**—The network carries mostly unicast best-effort traffic.
- **Best-effort traffic with Ethernet PAUSE (IEEE 802.3X) enabled**—The network carries mostly best-effort traffic with Ethernet PAUSE enabled on the links.
- **Best-effort multicast traffic**—The network carries mostly multicast best-effort traffic.
- **Lossless traffic**—The network carries mostly lossless traffic (traffic on which PFC is enabled).



**NOTE:** Lossless traffic is defined as traffic on which you enable PFC to ensure lossless transport. Lossless traffic does not refer to best-effort traffic on a link on which you enable Ethernet PAUSE. Start with the recommended profiles for each network traffic scenario, and adjust them if necessary for your network traffic conditions.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete. This includes changing the default configuration to one of the recommended configurations.

Because you configure buffer allocations in percentages, the recommended allocations for each network traffic scenario are valid for all QFX Series switches and the EX4600 switch. Use one of the following recommended shared buffer configurations for your network traffic conditions. Start with a recommended configuration, then make small adjustments to the buffer allocations to fine-tune the buffers if necessary as described in [“Optimizing Buffer Configuration” on page 5906](#).

- [Balanced Traffic \(Default Configuration\) on page 5903](#)
- [Best-Effort Unicast Traffic on page 5904](#)
- [Ethernet PAUSE Traffic on page 5905](#)
- [Best-Effort Multicast \(Multidestination\) Traffic on page 5905](#)
- [Lossless Traffic on page 5906](#)

### ***Balanced Traffic (Default Configuration)***

The default shared buffer configuration is optimized for networks that carry a balanced mix of best-effort unicast, lossless, and multidestination (multicast, broadcast, and

destination lookup fail) traffic. The default class-of-service (CoS) configuration is also optimized for networks that carry a balanced mix of traffic.

We recommend that you use the default shared buffer configuration for networks that carry a balanced mix of traffic, especially if you are using the default CoS settings.

[Table 520 on page 5904](#) shows the default ingress shared buffer allocations:

**Table 520: Default Ingress Shared Buffer Configuration**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 100%                        | 9%              | 45%                      | 46%          |

[Table 521 on page 5904](#) shows the default egress shared buffer allocations:

**Table 521: Default Egress Shared Buffer Configuration**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 50%             | 31%          | 19%              |

#### ***Best-Effort Unicast Traffic***

If your network carries mostly best-effort (lossy) unicast traffic, then the default shared buffer configuration allocates too much buffer space to support lossless transport. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 522 on page 5904](#)) and egress shared buffer settings (see [Table 523 on page 5904](#)):

**Table 522: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 100%                        | 5%              | 0%                       | 95%          |

**Table 523: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 5%              | 75%          | 20%              |

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic](#)” on page 6104 for an example that shows you how to configure the recommended buffer settings shown in [Table 522 on page 5904](#) and [Table 523 on page 5904](#).

**Ethernet PAUSE Traffic**

If your network carries mostly best-effort (lossy) traffic *and* enables Ethernet PAUSE on links, then the default shared buffer configuration allocates too much buffer space to the shared ingress buffer (Ethernet PAUSE traffic uses the dedicated buffers instead of shared buffers) and not enough space to the lossless-headroom buffers. We recommend that you use the following ingress shared buffer settings (see [Table 524 on page 5905](#)) and egress shared buffer settings (see [Table 525 on page 5905](#)):

**Table 524: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 70%                         | 5%              | 80%                      | 15%          |

**Table 525: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 5%              | 75%          | 20%              |

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled](#)” on page 6110 for an example that shows you how to configure the recommended buffer settings shown in [Table 522 on page 5904](#) and [Table 523 on page 5904](#).

**Best-Effort Multicast (Multidestination) Traffic**

If your network carries mostly best-effort (lossy) multicast traffic, then the default shared buffer configuration allocates too much buffer space to support lossless transport. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 526 on page 5905](#)) and egress shared buffer settings (see [Table 527 on page 5906](#)):

**Table 526: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 100%                        | 5%              | 0%                       | 95%          |

**Table 527: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 5%              | 20%          | 75%              |

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic](#)” on page 6116 for an example that shows you how to configure the recommended buffer settings shown in [Table 526 on page 5905](#) and [Table 527 on page 5906](#).

#### ***Lossless Traffic***

If your network carries mostly lossless traffic, then the default shared buffer configuration allocates too much buffer space to support best-effort traffic. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 528 on page 5906](#)) and egress shared buffer settings (see [Table 529 on page 5906](#)):

**Table 528: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Lossless Traffic**

| Total Shared Ingress Buffer | Lossless Buffer | Lossless-Headroom Buffer | Lossy Buffer |
|-----------------------------|-----------------|--------------------------|--------------|
| 100%                        | 15%             | 80%                      | 5%           |

**Table 529: Recommended Egress Shared Buffer Configuration for Networks with Mostly Lossless Traffic**

| Total Shared Egress Buffer | Lossless Buffer | Lossy Buffer | Multicast Buffer |
|----------------------------|-----------------|--------------|------------------|
| 100%                       | 90%             | 5%           | 5%               |

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic](#)” on page 6122 for an example that shows you how to configure the recommended buffer settings shown in [Table 528 on page 5906](#) and [Table 529 on page 5906](#).

#### **Optimizing Buffer Configuration**

Starting from the default configuration or from a recommended buffer configuration, you can further optimize the buffer allocation to best support the mix of traffic on your network. Adjust the settings gradually to fine-tune the shared buffer allocation. Use caution when adjusting the shared buffer configuration, not just when you fine-tune the ingress and egress buffer partitions, but also when you fine-tune the total ingress and egress shared buffer percentage. (Remember that if you allocate less than 100 percent of the available buffers to the shared buffers, the remaining buffers are added to the dedicated buffers). Tuning the buffers incorrectly can cause problems such as ingress port congestion.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

The relationship between the sizes of the ingress buffer pool and the egress buffer pool affects when and where packets are dropped. The buffer pool sizes include the shared buffers and the dedicated buffers. In general, if there are more ingress buffers than egress buffers, the switch can experience ingress port congestion because egress queues fill before ingress queues can empty.

Use the `show class-of-service shared-buffer` operational command to see the sizes in kilobytes (KB) of the dedicated and shared buffers and of the shared buffer partitions.

For best-effort traffic (unicast and multideestination), the combined ingress lossy shared buffer partition and ingress dedicated buffers must be *less than* the combined egress lossy and multicast shared buffer partitions plus the egress dedicated buffers. This prevents ingress port congestion by ensuring that egress best-effort buffers are deeper than ingress best-effort buffers, and ensures that if packets are dropped, they are dropped at the egress queues. (Packets dropping at the ingress prevents the egress schedulers from working properly.)

For lossless traffic (traffic on which you enable PFC), the combined ingress lossless shared buffer partition and a reasonable portion of the ingress headroom buffer partition, plus the dedicated buffers, must be *less than* the total egress lossless shared buffer partition and dedicated buffers. (A reasonable portion of the ingress headroom buffer is approximately 20 to 25 percent of the buffer space, but this varies depending on how much buffer headroom is required to support the lossless traffic.) When these conditions are met, if there is ingress port congestion, the ingress port congestion triggers PFC on the ingress port to prevent packet loss. If the total lossless ingress buffers exceed the total lossless egress buffers, packets could be dropped at the egress instead of PFC being applied at the ingress to prevent packet loss.



**NOTE:** If you commit a buffer configuration for which the switch does not have sufficient resources, the switch might log an error instead of returning a commit error. After you commit a buffer configuration, check the syslog messages to ensure that the new buffer configuration did not fail to commit.

If the buffer configuration commits but you receive a syslog message that indicates the configuration cannot be implemented, you can:

- Reconfigure the buffers or reconfigure other parameters (for example, the PFC configuration, which affects the need for lossless headroom buffers and lossless buffers—the more priorities you pause, the more lossless and lossless headroom buffer space you need), then attempt the commit operation again.
- Roll back the switch to the last successful configuration.

If you receive a syslog message that says the buffer configuration cannot be implemented, you must take corrective action. If you do not fix the configuration or roll back to a previous successful configuration, the system behavior is unpredictable.

---

### General Buffer Configuration Rules and Considerations

---

Keep the following rules and considerations in mind when you configure the buffers:

- Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.
- If you configure the ingress or egress shared buffer percentages as less than 100 percent, the remaining percentage of buffer space is added to the dedicated buffer pool.
- The sum of all of the ingress shared buffer partitions must equal 100 percent. Each partition must be configured with a value of at least 5 percent except the lossless headroom buffer, which can have a value of 0 percent.
- The sum of all of the egress shared buffer partitions must equal 100 percent. Each partition must be configured with a value of at least 5 percent.
- Lossless and lossless headroom shared buffers serve traffic on which you enable PFC, and do not serve traffic subject to Ethernet PAUSE.
- The switch uses the dedicated buffer pool first and the shared buffer pool only after the dedicated buffer pool for a port or queue is exhausted.
- Too little dedicated buffer space results in too much competition for shared buffer space.
- Too much dedicated buffer space results in poorer burst absorption because there is less available shared buffer space.
- Always check the syslog messages after you commit a new buffer configuration.
- The optimal buffer configuration for your network depends on the types of traffic on the network. If your network carries less traffic of a certain type (for example, lossless

traffic), then you can reduce the size of the buffers allocated to that type of traffic (for example, you can reduce the sizes of the lossless and lossless headroom buffers).

**Related Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 6110](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)

## Understanding CoS WRED Drop Profiles

When the number of packets queued is greater than the ability of the switch to empty an output queue, the queue requires a method for determining which packets to drop to relieve the congestion. Weighted random early detection (WRED) drop profiles define the drop probability of packets as the output queue fills. During periods of congestion, as the output queue fills, the switch drops incoming packets as determined by a drop profile until the output queue becomes less congested.

Depending on the drop probabilities, a drop profile can drop many packets long before the buffer becomes full, or it can drop only a few packets even if the buffer is almost full.

You configure drop profiles in the drop profile section of the class-of-service (CoS) configuration hierarchy. You apply drop profiles using a drop profile map in each scheduler configuration. For each scheduler, you can configure separate drop profiles for each combination of loss priority (low, medium-high, and high) and protocol.

Drop profiles define the meaning of each of the loss priorities by setting the values for when to drop packets and the probability that packets will drop.

You can configure a maximum of 32 drop profiles.



**NOTE:** You cannot apply drop profiles to multidestination (multicast) queues.

Do not apply drop profiles to lossless flows such as FCoE traffic, because the corresponding queues require lossless behavior. Use priority-based flow control (PFC) to prevent packet drop.

- [Drop Profile Parameters on page 5910](#)
- [Default Drop Profile on page 5911](#)
- [Packet Drop Method on page 5911](#)

- [Drop Profile Maps on page 5912](#)
- [Congestion Prevention on page 5912](#)
- [Configuring a WRED Drop Profile and Applying it to an Output Queue on page 5913](#)

### Drop Profile Parameters

Drop profiles specify two values:


- **Fill level**—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.
- **Drop probability**—The percentage value that corresponds to the likelihood that an individual packet is dropped.

You set two queue fill levels and two drop probabilities in each drop profile. The two fill levels and the two drop probabilities create two pairs of values. The first fill level and the first drop probability create one value pair and the second fill level and the second drop probability create the second value pair.

The first fill level value specifies the percentage of queue fullness at which packets begin to drop, known as the drop start point. Until the queue reaches this level of fullness, no packets are dropped. The second fill level value specifies the percentage of queue fullness at which all packets are dropped, known as the drop end point.

The first drop probability value is always 0 (zero). This pairs with the drop start point and specifies that until the queue fullness level reaches the first fill level, no packets drop. When the queue fullness exceeds the drop start point, packets begin to drop until the queue exceeds the second fill level, when all packets drop. The second drop probability value, known as the maximum drop rate, specifies the likelihood of dropping packets when the queue fullness reaches the drop end point. As the queue fills from the drop start point to the drop end point, packets drop in a smooth, linear pattern (called an interpolated graph) as shown in [Figure 208 on page 5910](#). After the drop end point, all packets drop.

**Figure 208: WRED-Drop Profile Packet Drop**





The thick line in [Figure 208 on page 5910](#) shows the packet drop characteristics for a sample WRED profile. At the drop start point, the queue reaches a fill level of 30 percent. At the drop end point, the queue fill level reaches 50 percent, and the maximum drop rate is 80 percent.

No packets drop until the queue fill level reaches the drop start point of 30 percent. When the queue reaches the 30 percent fill level, packets begin to drop. As the queue fills, the percentage of packets dropped increases in a linear fashion. When the queue fills to the drop end point of 50 percent, the rate of packet drop has increased to the maximum drop rate of 80 percent. When the queue fill level exceeds the drop end point of 50 percent, all of the packets drop until the queue fill level drops below 50 percent.

### Default Drop Profile

If you do not configure default profiles and apply them to queue schedulers, the switch uses the default drop profile for lossy traffic classes. In the default drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent. As soon as packets arrive on a queue, the default profile might begin to drop packets.

### Packet Drop Method

When a packet reaches the head of the queue, the switch generates a random number between 0 and 100. The switch plots the random number against the drop profile using the current fullness of the queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped.

To create the linear drop pattern from the drop start point to the drop end point, the drop probabilities are derived using a linear approximation with eight sections, or steps, from the minimum queue fill level to the maximum queue fill level. The fill levels are divided into the eight sections equally, starting at the minimum fill level and ending at the maximum fill level. As the queue fills, the percentage of dropped packets increases. The percentage of packets dropped is based on the maximum drop rate.

For example, the default drop profile (which specifies a maximum drop rate of 100 percent) has the following drop probabilities at each section, or step, in the eight-section linear drop pattern:

- First section—The minimum drop probability is 6.25 percent of the maximum drop rate. The maximum drop probability is 12.5 percent of the maximum drop rate.
- Second section—The minimum drop probability is 18.75 percent of the maximum drop rate. The maximum drop probability is 25 percent of the maximum drop rate.
- Third section—The minimum drop probability is 30.25 percent of the maximum drop rate. The maximum drop probability is 37.5 percent of the maximum drop rate.
- Fourth section—The minimum drop probability is 43.75 percent of the maximum drop rate. The maximum drop probability is 50 percent of the maximum drop rate.
- Fifth section—The minimum drop probability is 56.25 percent of the maximum drop rate. The maximum drop probability is 62 percent of the maximum drop rate.

- Sixth section—The minimum drop probability is 68.75 percent of the maximum drop rate. The maximum drop probability is 75.5 percent of the maximum drop rate.
- Seventh section—The minimum drop probability is 81.25 percent of the maximum drop rate. The maximum drop probability is 87.5 percent of the maximum drop rate.
- Eighth section—The minimum drop probability is 92.75 percent of the maximum drop rate. The maximum drop probability is 100 percent of the maximum drop rate.

Packets drop even when there is no congestion, because packet drops begin at the drop start point regardless of whether congestion exists on the port. The default drop profile example represents the worst-case scenario, because the drop start point fill level is 0 percent, so packet drop begins when the queue starts to receive packets.

You can specify when packets begin to drop by configuring a drop start point at a fill level greater than 0 percent. For example, if you configure a drop profile that has a drop start point of 30 percent, packets do not drop until the queue is 30 percent full. We recommend that you configure drop profiles that are appropriate to your network traffic conditions.

The smaller the gap between the minimum drop rate (which is always 0) and the maximum drop rate, the smaller the gap between the minimum drop probability and the maximum drop probability at each section (step) of the linear drop pattern. The default drop profile, which has the maximum gap between the minimum drop rate (0 percent) and the maximum drop rate (100 percent), has the highest gap between the minimum drop probability and the maximum drop probability at each step. Configuring a lower maximum drop rate for a drop profile reduces the gap between the minimum drop probability and the maximum drop probability.

---

### Drop Profile Maps

Drop profile maps are part of scheduler configuration. A drop profile map maps a drop profile to a loss priority and a protocol. Specifying the drop profile map in a scheduler associates the drop profile with the queues (forwarding classes) that you map to the scheduler in a scheduler map.

You configure loss priority for a queue in the classifier section of the CoS configuration hierarchy, and the loss priority is applied to the queue at the ingress interface.

Each scheduler can have multiple drop profile maps, one for each combination of loss priority and protocol.

---

### Congestion Prevention

Configuring drop profiles on output queues prevents them from impacting other queues on the egress ports. If you do not configure drop profiles and map them to output queues, output queues without drop profiles can impact output queues on other egress ports, even if those queues are not experiencing congestion.

For example, if an ingress port forwards traffic to more than one egress port, and at least one of the egress ports experiences congestion, that can cause ingress port congestion. Ingress port congestion (ingress buffer exceeds its resource allocation) can cause frames to drop at the ingress port instead of at the egress port. Ingress port frame drop affects

all of the egress ports to which the congested ingress port forwards traffic, not just the congested egress port.



**NOTE:** Do not configure drop profiles for the `fcoe` and `no-loss` forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

### Configuring a WRED Drop Profile and Applying it to an Output Queue

To configure a WRED packet drop profile and apply it to an output queue (using hierarchical scheduling):

1. Configure a drop profile using the statement **`set class-of-service drop-profiles profile-name interpolate fill-level drop-start-point fill-level drop-end-point drop-probability 0 drop-probability percentage`**.
2. Map the drop profile to a queue scheduler using the statement **`set class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high | high) protocol any drop-profile profile-name`**. The name of the drop-profile is the name of the WRED profile configured in step 1.
3. Map the scheduler, which step 2 associates with the drop profile, to the output queue using the statement **`set class-of-service scheduler-maps map-name forwarding-class forwarding-class-name scheduler scheduler-name`**. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in step 2.
4. Associate the scheduler map with a traffic control profile using the statement **`set class-of-service traffic-control-profiles tcp-name scheduler-map map-name`**. The scheduler map name is the name configured in step 3.
5. Associate the traffic control profile with an interface using the statement **`set class-of-service interface interface-name forwarding-class-set forwarding-class-set-name output-traffic-control-profile tcp-name`**. The output traffic control profile name is the name of the traffic control profile configured in step 4.

The interface uses the scheduler map in the traffic control profile to apply the drop profile (and other attributes) to the output queue (forwarding class) on that interface. Because you can use different traffic control profiles to map different schedulers to different interfaces, the same queue number on different interfaces can handle traffic in different ways.

#### Related Documentation

- [Understanding Junos CoS Components on page 5789](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)

- [Configuring CoS WRED Drop Profiles on page 6163](#)
- [Configuring CoS Drop Profile Maps on page 6164](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)

## Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. Rewrite rules set the value of the CoS bits within the header of the outgoing packet. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header, replacing the old CoS value. Rewrite rules must be assigned to an interface for rewrites to be activated.

You can apply (bind) one DSCP or DSCP IPv6 rewrite rule and one IEEE 802.1p rewrite rule to each interface as described in [“Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces” on page 5820](#). You can also bind EXP rewrite rules to **family mpls** logical interfaces to rewrite the CoS bits of MPLS traffic.

You cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Each interface supports only one DSCP rewrite rule. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6. You can apply an EXP rewrite rule on an interface that has DSCP or IEEE rewrite rules. Only MPLS traffic on **family mpls** interfaces uses the EXP rewrite rule.



**NOTE:** There are no default rewrite rules.

---

You can look at behavior aggregate (BA) classifiers and rewrite rules as two sides of the same coin. A BA classifier reads the CoS bits of incoming packets and classifies the packets into forwarding classes, then the system applies the CoS configured for the forwarding class to those packets. Rewrite rules change (rewrite) the CoS bits just before the packets leave the system so that the next switch or router can apply the appropriate level of CoS to the packets. When you apply a rewrite rule to an interface, the rewrite rule is the last CoS action performed on the packet before it is forwarded.

Rewrite rules alter CoS values in outgoing packets on the outbound interfaces of an edge switch to accommodate the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

---

The switch does not have default rewrite rules. If you want to apply a rewrite rule to outgoing packets, you must explicitly configure the rewrite rule.



**NOTE:** Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

MPLS EXP rewrite rules apply only to **family mpls** logical interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

#### Related Documentation

- [Understanding Junos CoS Components on page 5789](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)

## Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

Two methods of peer-to-peer flow control are supported:

- IEEE 802.3X Ethernet PAUSE
- IEEE 802.1Qbb priority-based flow control (PFC)

Ethernet PAUSE and PFC are link-level flow control mechanisms.



**NOTE:** For end-to-end congestion control, see [“Understanding CoS Explicit Congestion Notification” on page 5926](#).

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority is mapped to a 3-bit IEEE 802.1p CoS code point flag in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).



**Video:** [Why Use PFC in a Data Center Network?](#)

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on specified types of traffic (for example, Fibre Channel over Ethernet traffic).



**NOTE:** Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error. Ethernet PAUSE and PFC are mutually exclusive configurations on an interface.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

- [Ethernet PAUSE on page 5917](#)

- [PFC on page 5921](#)
- [Lossless Transport Support Summary on page 5924](#)

## Ethernet PAUSE

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

*Symmetric flow control* means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

*Asymmetric flow control* allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, the PFC configuration overrides Ethernet PAUSE flow control.) Both symmetric and asymmetric flow control are supported.

- [Symmetric Flow Control on page 5918](#)
- [Asymmetric Flow Control on page 5918](#)

### ***Symmetric Flow Control***

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

### ***Asymmetric Flow Control***

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 449 on page 5561](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

**Table 530: Asymmetric Ethernet PAUSE Flow Control Configuration**

| Receive (Rx) Buffer | Transmit (Tx) Buffer | Configured Flow Control State   |
|---------------------|----------------------|---|
| On                  | Off                  | Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).         |
| Off                 | On                   | Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.) |
| On                  | On                   | Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.   |
| Off                 | Off                  | Ethernet PAUSE flow control is disabled.  |

The configured flow control is the Ethernet PAUSE state configured on the interface.



On 1-Gigabit Ethernet interfaces, autonegotiation of Ethernet PAUSE with the connected peer is supported. (Autonegotiation on 10-Gigabit Ethernet interfaces is not supported.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected peer using a combination of the Ethernet PAUSE and ASM\_DIR bits, as described in [Table 450 on page 5562](#):

**Table 531: Flow Control State Advertised to the Connected Peer (Autonegotiation)**

| Rx Buffer State | Tx Buffer State | PAUSE Bit | ASM_DIR Bit | Description  |
|-----------------|-----------------|-----------|-------------|--|
| Off             | Off             | 0         | 0           | The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.   |
| On              | On              | 1         | 0           | The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).  |
| On              | Off             | 0         | 1           | The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).   |
| Off             | On              | 1         | 1           | The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.) |

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control behavior (resolution) between them, as shown in [Table 451 on page 5563](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series or EX4600 switch and on the connected peer (also known as the link partner). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations

on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



**NOTE:** In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

**Table 532: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces**

| Local Interface (QFX Series or EX4600 Switch) |             | Peer Interface |             | Local Resolution  | Peer Resolution   |
|---|-------------|----------------|-------------|---|---|
| PAUSE Bit                                     | ASM_DIR Bit | PAUSE Bit      | ASM_DIR Bit |   |   |
| 0   | 0           | Don't care     | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0   | 1           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0   | 1           | 1              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 0   | 1           | 1              | 1           | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive | Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive |
| 1   | 0           | 0              | Don't care  | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1   | 0           | 1              | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |
| 1   | 1           | 0              | 0           | Disable Ethernet PAUSE transmit and receive                       | Disable Ethernet PAUSE transmit and receive                       |
| 1   | 1           | 0              | 1           | Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit | Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive |
| 1   | 1           | Don't care     | Don't care  | Enable Ethernet PAUSE transmit and receive                        | Enable Ethernet PAUSE transmit and receive                        |



**NOTE:** For your convenience, [Table 451 on page 5563](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

## PFC

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- **Input**—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- **Output**—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic with the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as Fibre Channel over Ethernet (FCoE), LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of link-level flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)

- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 452 on page 5565](#):

**Table 533: Default PFC Priority to Queue and Forwarding Class Mapping**

| IEEE 802.1p Priority (Code Point) | Queue | Forwarding Class |
|-----------------------------------|-------|------------------|
| 0 (000)                           | 0     | best-effort      |
| 1 (001)                           | 1     | best-effort      |
| 2 (010)                           | 2     | best-effort      |
| 3 (011)                           | 3     | fcoe             |
| 4 (100)                           | 4     | no-loss          |
| 5 (101)                           | 5     | best-effort      |
| 6 (110)                           | 6     | network-control  |
| 7 (111)                           | 7     | network-control  |

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



**NOTE:** By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default forwarding class configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 802.1p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5837](#) and [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#).

You enable PFC on a priority by:

1. Specifying the IEEE 802.1p code point to pause in the input stanza of a CNP
2. Applying the CNP to the ingress interfaces on which you want to pause the traffic



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion of the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
  1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
  2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
  3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

Although unicast traffic and multideestination (multicast, broadcast, and destination lookup fail) traffic must use different classifiers, you can map a unicast queue (queue 0 through 7) and a multideestination queue (queue 8, 9, 10, or 11) to the same PFC priority so that both unicast and multicast traffic use that priority. Do not map multideestination traffic to lossless priorities. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.



**NOTE:** You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone switches, you can create two CNPs with an explicitly configured output stanza.

When a switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 5837 describes configuring output flow control.

---

### Lossless Transport Support Summary

---

The switch supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



**CAUTION:** Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



**NOTE:** Junos OS Release 12.2 introduced changes to the way the switch handles lossless forwarding classes (including the default fcoe and no-loss forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the fcoe and no-loss forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the fcoe or the no-loss forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the fcoe or the no-loss forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit fcoe and no-loss forwarding class configuration before you upgrade to Junos OS Release 12.2.

See *Overview of CoS Changes Introduced in Junos OS Release 12.2* for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the fcoe and no-loss forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new no-loss packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 5837](#) provides detailed information about the explicit configuration of lossless priorities and about the default configuration of lossless priorities, including the input and output stanzas of the CNP.



**NOTE:** PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

#### Related Documentation

- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)

- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding CoS Explicit Congestion Notification on page 5926](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)

## Understanding CoS Explicit Congestion Notification

Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.

ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic because other traffic types use different methods of congestion notification—lossless traffic uses priority-based flow control (PFC) and strict-high priority traffic receives all of the port bandwidth it requires up to the point of a configured maximum rate.

You enable ECN on individual output queues (as represented by forwarding classes) by enabling ECN in the queue scheduler configuration, mapping the scheduler to forwarding classes (queues), and then applying the scheduler to interfaces. For ECN to work on a queue, you must also apply a weighted random early detection (WRED) packet drop profile to the queue.

- [How ECN Works on page 5926](#)
- [WRED Drop Profile Control of ECN Thresholds on page 5931](#)
- [Support, Limitations, and Notes on page 5932](#)

---

### How ECN Works

Without ECN, switches respond to network congestion by dropping TCP/IP packets. Dropped packets signal the network that congestion is occurring. Devices on the IP network respond to TCP packet drops by reducing the packet transmission rate in order to allow the congestion to clear. However, the packet drop method of congestion notification and management has some disadvantages. For example, packets are dropped and must be retransmitted. And bursty traffic can cause the network to reduce the transmission rate too much, resulting in inefficient bandwidth utilization.



Instead of dropping packets to signal network congestion, ECN marks packets to signal network congestion, without dropping the packets. For ECN to work, all of the switches in the path between two ECN-enabled endpoints must have ECN enabled. ECN is negotiated during the establishment of the TCP connection between the endpoints.

ECN-enabled switches determine the queue congestion state based on the WRED packet drop profile configuration applied to the queue, so each ECN-enabled queue must also have a WRED drop profile. If a queue fills to the level that the WRED profile uses to begin dropping non-ECN-capable best-effort packets, the switch may mark a packet as experiencing congestion.

ECN communicates whether or not congestion is experienced by marking the two least-significant bits in the differentiated services (DiffServ) field in the IP header. The most significant six bits in the DiffServ field contain the differentiated services code point (DSCP bits). The state of the two ECN bits signals whether or not the packet is an ECN-capable packet and whether or not congestion has been experienced.

ECN-capable senders mark packets as ECN-capable. If a sender is not ECN-capable, it marks packets as not not ECN-capable. If an ECN-capable packet experiences congestion at the egress queue of a switch, the switch marks the packet as experiencing congestion. When the packet reaches the ECN-capable receiver (destination endpoint), the receiver echoes the congestion indicator to the sender (source endpoint) by sending a packet marked to indicate congestion.

After receiving the congestion indicator from the receiver, the source endpoint reduces the transmission rate to relieve the congestion. This is similar to the result of TCP congestion notification and management, but instead of dropping the packet to signal network congestion, ECN marks the packet and the receiver echoes the congestion notification to the sender. Because the packet is not dropped, the packet does not need to be retransmitted.

- [ECN Bits in the DiffServ Field on page 5927](#)
- [End-to-End ECN Behavior on page 5928](#)
- [ECN Compared to PFC and Ethernet PAUSE on page 5930](#)

#### ***ECN Bits in the DiffServ Field***

The two ECN bits in the DiffServ field provide four codes that determine if a packet is an ECN-capable transport (ECT) packet and if there is congestion experienced (CE), as shown in [Table 534 on page 5927](#):

**Table 534: ECN Bit Codes**

| ECN Bits (Code) | Meaning                             |
|-----------------|-------------------------------------|
| 00              | Non-ECT—Packet does not support ECN |
| 01              | ECT(1)—Packet supports ECN          |
| 10              | ECT(0)—Packet supports ECN          |
| 11              | CE—Congestion experienced           |

Codes 01 and 10 have the same meaning: ECN is supported and the packet is ECN-capable (ECT is set). There is no difference between these codes.

### End-to-End ECN Behavior

After the sending and receiving endpoints negotiate ECN, the sending endpoint marks packets as ECN-capable by setting the DiffServ ECN field to ECT(1) (01) or ECT(0) (10). Every intermediate switch between the endpoints must have ECN enabled or it does not work.

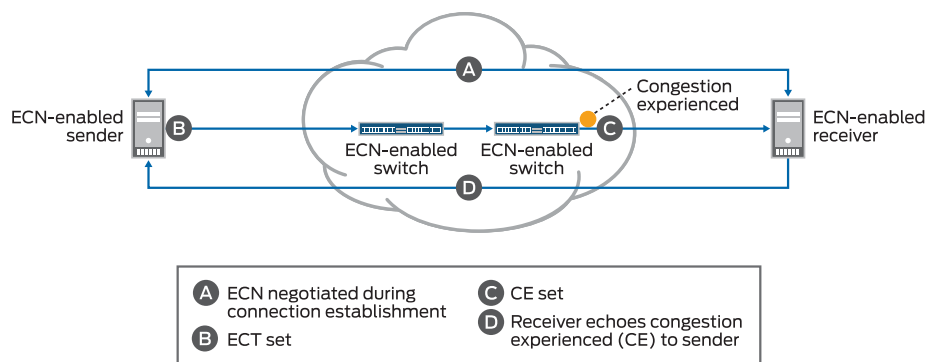
When a packet traverses a switch and experiences congestion at an output queue that uses the WRED packet drop mechanism, the switch marks the packet as experiencing congestion by setting the DiffServ ECN field to CE (11). Instead of dropping the packet (as with TCP congestion notification), the switch forwards the packet.



**NOTE:** At the egress queue, the WRED algorithm determines whether or not a packet is drop eligible based on the queue fill level (how full the queue is). If a packet is drop eligible and marked as ECN-capable, the packet can be marked CE and forwarded. If a packet is drop eligible and is not marked as ECN-capable, it might be dropped. See [“WRED Drop Profile Control of ECN Thresholds” on page 5931](#) for more information about the WRED algorithm.

When the packet reaches the receiver endpoint, the CE mark tells the receiver that there is network congestion. The receiver then sends (echoes) a message to the sender that indicates there is congestion on the network. The sender acknowledges the congestion notification message and reduces its transmission rate. [Figure 209 on page 5928](#) summarizes how ECN works to mitigate network congestion:

**Figure 209: Explicit Congestion Notification**



8042495

End-to-end ECN behavior includes:

1. The ECN-capable sender and receiver negotiate ECN capability during the establishment of their connection.
2. After successful negotiation of ECN capability, the ECN-capable sender sends IP packets with the ECT field set to the receiver.



**NOTE:** All of the intermediate devices in the path between the sender and the receiver must be ECN-enabled.

3. If the WRED algorithm on a switch egress queue determines that the queue is experiencing congestion and the packet is drop eligible, the switch can mark the packet as “congestion experienced” (CE) to indicate to the receiver that there is congestion on the network. If the packet has already been marked CE (congestion has already been experienced at the egress of another switch), the switch forwards the packet with CE marked.

If there is no congestion at the switch egress queue, the switch forwards the packet and does not change the ECT-enabled marking of the ECN bits, so the packet is still marked as ECN-capable but not as experiencing congestion.

Packets that are not marked as ECN-capable (ECT) are treated according to the WRED drop profile configuration and may be dropped during periods of congestion.

4. The receiver receives a packet marked CE to indicate that congestion was experienced along the congestion path.
5. The receiver echoes (sends) a packet back to the sender with the ECE bit (bit 9) marked in the flag field of the TCP header. The ECE bit is the ECN echo flag bit, which notifies the sender that there is congestion on the network.
6. The sender reduces the data transmission rate and sends a packet to the receiver with the CWR bit (bit 8) marked in the flag field of the TCP header. The CWR bit is the congestion window reduced flag bit, which acknowledges to the receiver that the congestion experienced notification was received.
7. When the receiver receives the CWR flag, the receiver stops setting the ECE bit in replies to the sender.

Table 535 on page 5929 summarizes the behavior of traffic on ECN-enabled queues.

**Table 535: Traffic Behavior on ECN-Enabled Queues**

| Incoming IP Packet Marking of ECN Bits | ECN Configuration on the Output Queue | Action if WRED Algorithm Determines Packet is Drop Eligible        | Outgoing Packet Marking of ECN Bits           |
|--|---------------------------------------|--|---|
| Non-ECT (00)                           | Does not matter                       | Drop   | Packet dropped—no ECN bits marked             |
| ECT (10 or 01)                         | ECN disabled                          | Drop   | Packet dropped—no ECN bits marked             |
| ECT (10 or 01)                         | ECN enabled                           | Do not drop. Mark packet as experiencing congestion (CE, bits 11). | Packet marked ECT (11) to indicate congestion |
| CE (11)                                | ECN disabled                          | Drop   | Packet dropped—no ECN bits marked             |

Table 535: Traffic Behavior on ECN-Enabled Queues (*continued*)

| Incoming IP Packet Marking of ECN Bits | ECN Configuration on the Output Queue | Action if WRED Algorithm Determines Packet is Drop Eligible  | Outgoing Packet Marking of ECN Bits           |
|--|---------------------------------------|--|---|
| CE (11)                                | ECN enabled                           | Do not drop. Packet is already marked as experiencing congestion, forward packet without changing the ECN marking. | Packet marked ECT (11) to indicate congestion |

When an output queue is not experiencing congestion as defined by the WRED drop profile mapped to the queue, all packets are forwarded, and no packets are dropped.

#### ***ECN Compared to PFC and Ethernet PAUSE***

ECN is an end-to-end network congestion notification mechanism for IP traffic. Priority-based flow control (PFC; IEEE 802.1Qbb) and Ethernet PAUSE (IEEE 802.3X) are different types of congestion management mechanisms. ECN requires that an output queue must also have an associated WRED packet drop profile. Output queues used for traffic on which PFC is enabled should not have an associated WRED drop profile. Interfaces on which Ethernet PAUSE is enabled should not have an associated WRED drop profile.

PFC is a peer-to-peer flow control mechanism to support lossless traffic. PFC enables connected peer devices to pause flow transmission during periods of congestion. PFC enables you to pause traffic on a specified type of flow on a link instead of on all traffic on a link. For example, you can (and should) enable PFC on lossless traffic classes such as the **fcoe** forwarding class. Ethernet PAUSE is also a peer-to-peer flow control mechanism, but instead of pausing only specified traffic flows, Ethernet PAUSE pauses all traffic on a physical link.

With PFC and Ethernet PAUSE, the sending and receiving endpoints of a flow do not communicate congestion information to each other across the intermediate switches. Instead, PFC controls flows between two PFC-enabled peer devices (for example, switches) that support data center bridging (DCB) standards. PFC works by sending a pause message to the connected peer when the flow output queue becomes congested. Ethernet PAUSE simply pauses all traffic on a link during periods of congestion and does not require DCB.

PFC works this way: if a switch output queue fills to a certain threshold, the switch sends a PFC pause message to the connected peer device that is transmitting data. The pause message tells the transmitting switch to pause transmission of the flow. When the congestion clears, the switch sends another PFC message to tell the connected peer to resume transmission. (If the output queue of the transmitting switch also reaches a certain threshold, that switch can in turn send a PFC pause message to the connected peer that is transmitting to it. In this way, PFC can propagate a transmission pause back through the network.)

See [“Understanding CoS Flow Control \(Ethernet PAUSE and PFC\)” on page 5559](#) for more information. For QFX5100 and EX4600 switches only, you can also refer to [“Understanding PFC Functionality Across Layer 3 Interfaces” on page 5950](#).

### WRED Drop Profile Control of ECN Thresholds

WRED drop profiles applied to output queues control how the switch marks ECN-capable packets. Drop profiles contain four configurable parameters that control when packets drop and the rate of packet drop:

- Drop start point—The queue fill level (threshold) at which the queue begins to drop packets.
- Drop end point—The queue fill level (threshold) at which the queue drops all packets that are not ECN-capable packets.



**NOTE:** Lossless and strict-high priority queues do not use drop profiles. Lossless queues use PFC to control the flow of traffic. Strict-high priority queues receive all of the port bandwidth they require up to the configured maximum bandwidth limit (scheduler shaping-rate).

- Drop probability at drop start point—The probability that a packet is dropped when the queue fill level reaches the drop start point. This is always zero (0) percent because until the queue fills to the drop start point, no packets are scheduled to drop.
- Drop probability at drop end point—The probability that a non-ECN packet is dropped when the queue fill level reaches the drop end point. All non-ECN packets are dropped after a queue reaches the drop end point.

As a queue fills from the drop start point to the drop end point, the probability of non-ECN packets being dropped increases. The more full a queue becomes, the greater the probability that non-ECN packets are dropped. (See [“Understanding CoS WRED Drop Profiles” on page 5909](#) for detailed information about how packet drop probabilities are calculated and applied to a queue.) However, ECN packets are not dropped unless the queue becomes completely full.

The drop profile configuration affects ECN packets as follows:

- Drop start point—ECN-capable packets might be marked as congestion experienced (CE).
- Drop end point—ECN-capable packets are always marked CE.

As a queue fills from the drop start point to the drop end point, the probability that an ECN packet is marked CE is the same as the probability that a non-ECN packet is dropped. As the queue fills, the probability of an ECN packet being marked CE increases, just as the probability of a non-ECN packet being dropped increases.

At the drop end point, all ECN packets are marked CE, but the ECN packets are not dropped. When the queue fill level exceeds the drop end point, all ECN packets are marked CE and all non-ECN packets are dropped. ECN packets are only dropped if the queue fills completely.

To configure a WRED packet drop profile and apply it to an output queue (using hierarchical scheduling):

1. Configure a drop profile using the statement **set class-of-service drop-profiles *profile-name* interpolate fill-level *drop-start-point* fill-level *drop-end-point* drop-probability 0 drop-probability *percentage***.
2. Map the drop profile to a queue scheduler using the statement **set class-of-service schedulers *scheduler-name* drop-profile-map loss-priority (low | medium-high | high) protocol any drop-profile *profile-name***. The name of the drop-profile is the name of the WRED profile configured in step 1.
3. Map the scheduler, which step 2 associates with the drop profile, to the output queue using the statement **set class-of-service scheduler-maps *map-name* forwarding-class *forwarding-class-name* scheduler *scheduler-name***. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in step 2.
4. Associate the scheduler map with a traffic control profile using the statement **set class-of-service traffic-control-profiles *tcp-name* scheduler-map *map-name***. The scheduler map name is the name configured in step 3.
5. Associate the traffic control profile with an interface using the statement **set class-of-service interface *interface-name* forwarding-class-set *forwarding-class-set-name* output-traffic-control-profile *tcp-name***. The output traffic control profile name is the name of the traffic control profile configured in step 4.

The interface uses the scheduler map in the traffic control profile to apply the drop profile (and other attributes, including the enable ECN attribute) to the output queue (forwarding class) on that interface. Because you can use different traffic control profiles to map different schedulers to different interfaces, the same queue number on different interfaces can handle traffic in different ways.

---

### Support, Limitations, and Notes

If the WRED algorithm that is mapped to a queue does not find a packet drop eligible, then the ECN configuration and ECN bits marking does not matter. The packet transport behavior is the same as when ECN is not enabled.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic, and you do not enable ECN on queues that handle lossless traffic or strict-high priority traffic.

ECN supports the following:

- IPv4 and IPv6 packets
- Untagged, single-tagged, and double-tagged packets
- The outer IP header of IP tunneled packets (but not the inner IP header)

ECN does not support the following:

- IP packets with MPLS encapsulation
- The inner IP header of IP tunneled packets (however, ECN works on the outer IP header)
- Multicast, broadcast, and destination lookup fail (DLF) traffic
- Non-IP traffic

**Related  
Documentation**

- [Understanding CoS WRED Drop Profiles on page 5909](#)
- [Example: Configuring ECN on page 6090](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Example: Configuring Queue Schedulers on page 6081](#)

## Understanding DCB Features and Requirements

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.



Video: [What is Data Center Bridging?](#)

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series and EX4600 switch support the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 5934](#)
- [ETS on page 5935](#)
- [DCBX on page 5936](#)

---

### Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.



This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 5935](#)
- [Buffer Management on page 5935](#)
- [Physical Interfaces on page 5935](#)

### ***PFC***

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

### ***Buffer Management***

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

### ***Physical Interfaces***

The switch supports 10-Gbps, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps (or faster) Ethernet interfaces.

### ***ETS***

---

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).
- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.

- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows. We recommend that you always configure a shaping rate to limit the amount of bandwidth a strict-high priority flow can consume by including the [shaping-rate](#) statement in the [\[edit class-of-service schedulers\]](#) hierarchy on the strict-high priority scheduler. This prevents a strict-high priority from starving other queues on the port.

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

---

## DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

### Related Documentation

- [Overview of Fibre Channel on page 5508](#)
- [Understanding FCoE on page 5518](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding DCBX on page 5580](#)
- [Understanding Fibre Channel Terminology on page 5569](#)

- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)

## Understanding DCBX

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.



Video: [What is DCBX Protocol?](#)

This topic describes:

- [DCBX Basics on page 5937](#)
- [DCBX Modes and Support on page 5938](#)
- [DCBX Attribute Types on page 5941](#)
- [DCBX Application Protocol TLV Exchange on page 5942](#)
- [DCBX and PFC on page 5943](#)
- [DCBX and ETS on page 5943](#)

### DCBX Basics

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that

feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



**NOTE:** When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

---

## DCBX Modes and Support

This section describes DCBX support:

- [DCBX Modes \(Versions\) on page 5938](#)
- [Autonegotiation on page 5940](#)
- [CNA Support for DCBX Modes on page 5941](#)
- [Interface Support for DCBX on page 5941](#)

### ***DCBX Modes (Versions)***

The two most common DCBX modes are supported:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



**NOTE:** The switch does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The switch drops LLDP frames that contain pre-CEE DCBX TLVs.

Table 454 on page 5582 summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:

**Table 536: Summary of Differences Between IEEE DCBX and DCBX Version 1.01**

| Characteristic   | IEEE DCBX  | DCBX Version 1.01  |
|--|--|--|
| OUI  | 0x0080c2   | 0x001b21   |
| Frame Format   | Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs   | Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer.   |
| Symmetric/asymmetric configuration with peer                                     | Asymmetric or symmetric  | Symmetric only   |
| Differences in the <b>show dcbx interface interface-name</b> operational command | <ul style="list-style-type: none"> <li>• Synchronization information is not shown because symmetric configuration is not required.</li> <li>• Operational state information is not shown because the operational states do not have to be symmetric.</li> <li>• TLV type is shown because unique TLVs are sent for each DCBX attribute.</li> <li>• ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs.</li> </ul> | <ul style="list-style-type: none"> <li>• Synchronization information is shown because symmetric configuration is required.</li> <li>• Operational state information is shown because the operational states do have to be symmetric.</li> <li>• TLV type is not shown because one TLV is used for all attribute information.</li> <li>• Recommendation TLV is not sent (DCBX Version 1.01 uses the “willing” bit to determine whether or not an interface uses the peer interface configuration).</li> </ul> |

For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—<http://www.ieee802.org/1/files/public/docs2008/az-wedekar-dcbx-capability-exchange-discovery-protocol-1108-v1.01.pdf>
- For IEEE DCBX—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>



**NOTE:** As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.
- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



**NOTE:** On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

---

### ***Autonegotiation***

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on standalone switches compared to QFabric systems:

- Standalone switches—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.
- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



**NOTE:** If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

### **CNA Support for DCBX Modes**

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on switch interfaces depends on the DCBX features that the CNAs in your network support.

### **Interface Support for DCBX**

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

### **DCBX Attribute Types**

DCBX has three attribute types:

- **Informational**—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- **Asymmetric**—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- **Symmetric**—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 5941](#)
- [Symmetric Attributes on page 5942](#)

### **Asymmetric Attributes**

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing,” the configuration on the interface cannot be overridden by the peer interface configuration.
- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface

sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

### ***Symmetric Attributes***

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

### **DCBX Application Protocol TLV Exchange**

---

DCBX advertises the switch’s capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 5942](#)
- [FCoE Application Protocol TLV Exchange on page 5942](#)
- [Disabling Application Protocol TLV Exchange on page 5943](#)

### ***Application Protocol TLV Exchange***

For all applications, DCBX advertises the application’s state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application’s TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

### ***FCoE Application Protocol TLV Exchange***

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map



**NOTE:** If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

---



If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.



**NOTE:** If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE, DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

### ***Disabling Application Protocol TLV Exchange***

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

## **DCBX and PFC**

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.

You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

## **DCBX and ETS**

This section describes:

- [Default DCBX ETS Advertisement on page 5944](#)
- [ETS Advertisement and Peer Configuration on page 5944](#)
- [ETS Recommendation TLV on page 5944](#)

### ***Default DCBX ETS Advertisement***

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

### ***ETS Advertisement and Peer Configuration***

DCBX does not control the switch's ETS (hierarchical scheduling) operational state. If the connected peer is configured as "willing," DCBX pushes the switch's ETS configuration to the switch's peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority (essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

### ***ETS Recommendation TLV***

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is "willing," it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.



**NOTE:** You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

#### Related Documentation

- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding FCoE on page 5518](#)
- [Configuring the DCBX Mode on page 5668](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)

## Understanding DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

---

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 5590](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case (see [“Application Maps” on page 5590](#)) and only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic.

This topic describes:

- [Applications on page 5946](#)
- [Application Maps on page 5947](#)
- [Classifying and Prioritizing Application Traffic on page 5948](#)
- [Enabling Interfaces to Exchange Application Protocol Information on page 5949](#)
- [Disabling DCBX Application Protocol Exchange on page 5949](#)

---

### Applications

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise, except FCoE if FCoE is the only application that you want the interface to advertise.



**NOTE:** If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

### Application Maps

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.

To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

---

### Classifying and Prioritizing Application Traffic

---

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.

### Enabling Interfaces to Exchange Application Protocol Information

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



**NOTE:** You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

### Disabling DCBX Application Protocol Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

#### Related Documentation

- [Understanding DCBX on page 5580](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)

- [Example: Configuring Unicast Classifiers on page 6066](#)

## QFX5100 Switches Only

---

- [Understanding PFC Functionality Across Layer 3 Interfaces on page 5950](#)

### Understanding PFC Functionality Across Layer 3 Interfaces

Priority-based flow control (PFC) allows you to select traffic flows within a link and pause them, so that the output queues associated with the flows do not overflow and drop packets. (PFC is more granular than Ethernet PAUSE, which pauses all traffic on a physical link.) PFC helps you configure lossless transport for traffic flows across a data center bridging network.

However, you might want to create a traffic flow that losslessly traverses the Layer 2 data center bridging network *and* also losslessly traverses a Layer 3 network that connects Ethernet hosts in different Layer 2 networks. On a QFX5100 or EX4600 switch running the Enhanced Layer 2 Software (ELS) CLI, in addition to configuring PFC on Layer 2 (bridging) interfaces, you can configure PFC on traffic that traverses Layer 3 interfaces. This enables you to preserve the lossless characteristics that PFC provides on traffic, even when the traffic crosses Layer 3 interfaces that connect two Layer 2 networks.



Video: [Preserving Lossless Behavior on an SDN or Overlay Network](#)

PFC works the same way across Layer 3 interfaces as it works across Layer 2 interfaces. When an output queue buffer reaches a certain fill level threshold, the switch sends a PFC pause message to the connected peer to pause transmission of the traffic on which PFC is enabled. Pausing the incoming traffic prevents the queue buffer from overflowing and dropping packets, just as on Layer 2 interfaces. When the queue buffer fill level decreases below a certain threshold, the interface sends a message to the connected peer to restart traffic transmission.

Although PFC is a data center bridging technology, PFC also works on Layer 3 interfaces because PFC operates at the queue level. When you use an IEEE 802.1p classifier to classify incoming traffic (map incoming traffic to a forwarding class and a loss priority based on the IEEE 802.1p code point in the Ethernet frame header) and you enable PFC on the appropriate priority (IEEE 802.1p code point), PFC works on Layer 2 and Layer 3 interfaces.



**NOTE:** Lossless traffic on Layer 3 interfaces *must* use an IEEE 802.1p classifier to classify incoming traffic, because PFC does not use DSCP or DSCP IPv6 code points to identify traffic for flow control. PFC cannot pause traffic flows unless the incoming traffic is classified by an IEEE 802.1p classifier. Do not apply a DSCP (or a DSCP IPv6) classifier to Layer 3 traffic on which you want to enable PFC.



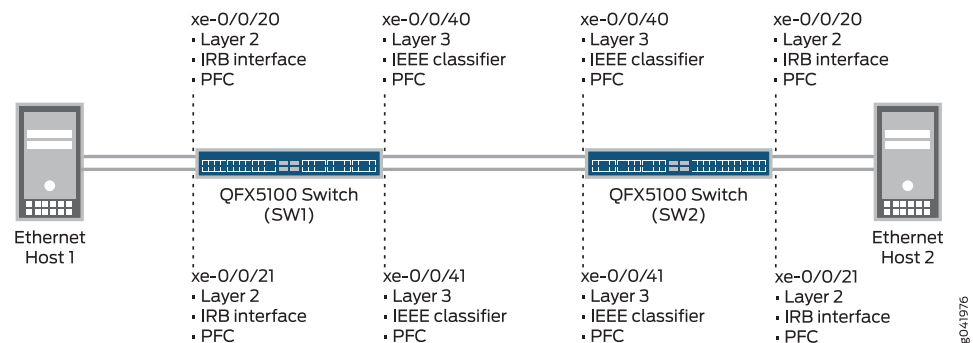
Because PFC functionality relies on the mapping (classifying) of traffic to IEEE 802.1p code points and on enabling PFC on the correct code point(s) at each interface, you must ensure that incoming traffic has the correct 3-bit IEEE 802.1p code point (priority) in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits).



**NOTE:** Layer 3 interfaces do not support FCoE traffic. FCoE traffic must use Layer 2 interfaces and cannot use Layer 3 interfaces. Therefore, you cannot enable PFC on FCoE traffic across Layer 3 interfaces.

Figure [Figure 210 on page 5951](#) shows a topology in which two Ethernet hosts in Layer 2 networks communicate across a Layer 3 network, with PFC enabled on all of the Layer 2 and Layer 3 switch interfaces.

**Figure 210: Enabling PFC Across Layer 3 Interface Hops**



The Ethernet host-facing interfaces (xe-0/0/20 and xe-0/0/21 on both switches) and the Layer 3 network-facing interfaces (interfaces xe-0/0/40 and xe-0/0/41 on both switches) require different interface configurations to enable PFC on the Layer 3 interfaces. In addition, the class of service (CoS) for each interface must be configured correctly, including enabling PFC on the traffic that you want to treat as lossless traffic:

Ethernet-host facing interfaces (xe-0/0/20 and xe-0/0/21) require the following configuration:

- Set interfaces as family ethernet-switching
- Set the interface mode as trunk mode
- Create VLANs to carry the traffic
- Create IRB interfaces to place the Layer 2 VLAN traffic on Layer 3 for transport between IP networks
- Create an IEEE 802.1p classifier to classify incoming traffic into the correct forwarding class, based on the IEEE 802.1p code point
- Create a congestion notification profile (CNP) to configure PFC on the IEEE 802.1p code point of the traffic that you want treat as lossless traffic

- Apply the classifier and the CNP to the Layer 2 interfaces
- Configure CoS: lossless forwarding classes, hierarchical port scheduling (also known as enhanced transmission selection) and apply it to the Layer 2 interfaces

Layer 3 IP network-facing interfaces (xe-0/0/40 and xe-0/0/41) require the following configuration:

- Set interfaces as family inet
- Set VLAN tagging on the interfaces
- Create VLANs to carry the traffic
- Create an IEEE 802.1p classifier to classify incoming traffic into the correct forwarding class, based on the IEEE 802.1p code point (do not use a DSCP or DSCP IPv6 classifier)
- Create a congestion notification profile (CNP) to configure PFC on the IEEE 802.1p code point of the traffic that you want treat as lossless traffic on the Layer 3 interfaces
- Apply the IEEE 802.1p classifier and the CNP to the Layer 3 interfaces
- Configure CoS: lossless forwarding classes, hierarchical port scheduling (enhanced transmission selection) and apply it to the Layer 3 interfaces



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

When you configure the Layer 2 and Layer 3 interfaces correctly, the switch enables PFC on the traffic between Ethernet Host 1 and Ethernet Host 2 across the entire path between the two hosts. If any output queue in the path on which PFC is enabled experiences congestion, PFC pauses the traffic and prevents packet loss for the flow.

#### Related Documentation

- [Example: Configuring PFC Across Layer 3 Interfaces on page 6138](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding Integrated Routing and Bridging on page 1539](#) (Topic also applies to IRB interfaces.)

---

## QFX3500 and QFX3600 Virtual Chassis Only

- [CoS on Virtual Chassis Switch Ports on page 5953](#)

## CoS on Virtual Chassis Switch Ports

QFX Series and EX4600 Virtual Chassis devices have access ports to connect to external peer devices. Virtual Chassis devices also have Virtual Chassis ports (VCPs) to interconnect members of the Virtual Chassis, in a similar way that QFabric system Node devices have fabric (fte) ports to connect to the QFabric system Interconnect device. VCPs are not used for external access.

Class of service (CoS) on Virtual Chassis access ports is the same as CoS on these devices when they are in standalone mode or used as QFabric system Node devices. However, CoS on VCPs differs in several ways from CoS on QFabric system Node device fabric ports.

This topic describes CoS support on Virtual Chassis access interfaces and on VCPs:

- [Access Interface CoS Support on page 5953](#)
- [VCP Interface CoS Support on page 5955](#)
- [CPU-Generated Host Outbound Traffic on page 5956](#)

### Access Interface CoS Support

CoS on Virtual Chassis access interfaces is the same as CoS on standalone device and Node device access interfaces, except for shared buffer settings. The documentation for QFX Series and EX4600 switch CoS on access interfaces applies to Virtual Chassis access interfaces, except some of the shared buffer documentation.

- [Similarities in CoS Support on Virtual Chassis Access Interfaces Compared to Standalone Device \(or QFabric system Node device\) Access Interfaces on page 5954](#)
- [Differences in CoS Support on Virtual Chassis Access Interfaces Compared to Standalone Device \(or QFabric system Node device\) Access Interfaces on page 5955](#)

***Similarities in CoS Support on Virtual Chassis Access Interfaces Compared to Standalone Device (or QFabric system Node device) Access Interfaces***

Virtual Chassis access interfaces support the following CoS features in the same way as access interfaces on standalone devices and QFabric system Node devices:

- Forwarding classes—The default forwarding classes, queue mapping, and packet drop attributes ([Table 537 on page 5954](#)) are the same:

**Table 537: Default Forwarding Class Configuration**

| Default Forwarding Class | Default Queue Mapping | Default Packet Drop Attribute |
|--------------------------|-----------------------|-------------------------------|
| best-effort (be)         | 0                     | drop                          |
| fcoe                     | 3                     | no-loss                       |
| no-loss                  | 4                     | no-loss                       |
| network-control (nc)     | 7                     | drop                          |
| mcast                    | 8                     | drop                          |

- Packet classification—Classifier default settings and configuration are the same. Support for behavior aggregate, multifield, multidestination, and fixed classifiers is the same.
- Enhanced transmission selection (ETS)—This data center bridging (DCB) feature that supports hierarchical scheduling has the same defaults and user configuration, including forwarding class set (priority group) and traffic control profile configuration.
- Priority-based flow control (PFC)—This DCB feature that supports lossless transport has the same defaults and user configuration, including support for six lossless priorities (forwarding classes).
- Ethernet PAUSE—This feature has the same defaults and configuration.
- Queue scheduling—This feature has the same defaults, configuration, and scheduler-to-forwarding-class mapping. Queue scheduling is a subset of hierarchical scheduling.
- Priority group (forwarding class set) scheduling—This feature has the same defaults and configuration. Priority group scheduling is a subset of hierarchical scheduling.
- WRED profiles—This feature has the same defaults and configuration.
- Code-point aliases—This feature has the same defaults and configuration.
- Rewrite rules—This feature has the same defaults and configuration (no default rewrite rules applied to egress traffic).
- Host outbound traffic—This feature has the same defaults and configuration.

### ***Differences in CoS Support on Virtual Chassis Access Interfaces Compared to Standalone Device (or QFabric system Node device) Access Interfaces***

The default shared buffer settings and the way in which you configure shared buffers are the same on Virtual Chassis access interfaces as on standalone and QFabric system Node devices. The difference is that on Virtual Chassis access interfaces, the shared buffer configuration is global and applies to all access ports on all members of the Virtual Chassis, while on standalone or QFabric system Node devices, you can configure different buffer settings on different access interfaces.

You cannot configure different shared buffer settings for different Virtual Chassis members. All members of a Virtual Chassis use the same shared buffer configuration.

### **VCP Interface CoS Support**

CoS on the VCP interfaces that connect the Virtual Chassis members is similar to CoS on the fabric interfaces of QFabric system Node devices, but there are several important differences:

- [Similarities in CoS Support on VCP Interfaces and QFabric System Node Device Fabric Interfaces on page 5955](#)
- [Differences in CoS Support on VCP Interfaces and QFabric System Node Device Fabric Interfaces on page 5956](#)

### ***Similarities in CoS Support on VCP Interfaces and QFabric System Node Device Fabric Interfaces***

VCP interfaces support full hierarchical scheduling (ETS). ETS includes the following CoS features. VCP interfaces support no other CoS features.

- Creating forwarding class sets (priority groups) and mapping forwarding classes to forwarding class sets.
- Scheduling individual output queues. The scheduler defaults and configuration are the same as the scheduler on access interfaces.
- Scheduling priority groups (forwarding class sets) using a traffic control profile. The defaults and configuration are the same as on access interfaces.



**NOTE:** You cannot attach classifiers, congestion notification profiles, or rewrite rules to VCP interfaces. Also, you cannot configure buffer settings on VCP interfaces. Similar to fabric interfaces on QFabric system Node devices, you can only attach forwarding class sets and traffic control profiles to VCP interfaces.

The behavior of lossless traffic across 40-Gigabit VCP interfaces is the same as the behavior of lossless traffic across QFabric system Node device fabric ports. The system automatically enables flow control for lossless forwarding classes (priorities). The system dynamically calculates buffer headroom that is allocated from the global lossless-headroom buffer for the lossless forwarding classes on each 40-Gigabit VCP interface. If there is not enough global lossless-headroom buffer space to support the

number of lossless flows on a 40-Gigabit VCP interface, the system generates a syslog message.



**NOTE:** After you configure lossless transport on a Virtual Chassis, check the syslog messages to ensure that there is sufficient buffer space to support the configuration.



**NOTE:** If you break out a 40-Gigabit VCP interface into 10-Gigabit VCP interfaces, lossless transport is not supported on the 10-Gigabit VCP interfaces. Lossless transport is supported only on 40-Gigabit VCP interfaces. (10-Gigabit access interfaces support lossless transport.)

---

### ***Differences in CoS Support on VCP Interfaces and QFabric System Node Device Fabric Interfaces***

Although most of the CoS behavior on VCP interfaces is similar to CoS behavior on the fabric ports of QFabric system Node devices, there are some important differences:

- Hierarchical scheduling (queue and priority group scheduling)—On QFabric system Node device fabric interfaces, you can apply a different hierarchical scheduler (traffic control profile) to different priority groups (forwarding class sets) on different interfaces. However, on VCP interfaces, the schedulers that you apply to priority groups are global to all VCP interfaces. One hierarchical scheduler controls scheduling for a priority group on all VCP interfaces.

You attach a scheduler to VCP interfaces using the global identifier (*vcp-\**) for VCP interfaces. For example, if you want to apply a traffic control profile (traffic control profiles contain both queue and priority group scheduling configuration) named *vcp-hpc-tcp* to a forwarding class set named *vcp-hpc-fcset*, you include the following statement in the configuration:

```
[edit]
user@switch# set class-of-service interfaces vcp-* forwarding-class-set vcp-hpc-fcset
output-traffic-control-profile vcp-hpc-tcp
```

The system applies the hierarchical scheduler *vcp-hpc-tcp* to the traffic mapped to the priority group *vcp-hpc-fcset* on all VCP interfaces.

- You cannot attach classifiers, congestion notification profiles, or rewrite rules to VCP interfaces. Also, you cannot configure buffer settings on VCP interfaces. Similar to QFabric system Node device fabric interfaces, you can only attach forwarding class sets and traffic control profiles to VCP interfaces.
- Lossless transport is supported only on 40-Gigabit VCP interfaces. If you break out a 40-Gigabit VCP interface into 10-Gigabit VCP interfaces, lossless transport is not supported on the 10-Gigabit VCP interfaces.

---

### **CPU-Generated Host Outbound Traffic**

CPU-generated host outbound traffic is forwarded on the network-control forwarding class, which is mapped to queue 7. If you use the default scheduler, the network-control

queue receives a guaranteed minimum bandwidth (transmit rate) of 5 percent of port bandwidth. The guaranteed minimum bandwidth is more than sufficient to ensure lossless transport of host outbound traffic.

However, if you configure and apply a scheduler instead of using the default scheduler, you must ensure that the network-control forwarding class (or whatever forwarding class you configure for host outbound traffic) receives sufficient guaranteed bandwidth to prevent packet loss.



**TIP:** If you configure a scheduler instead of using the default scheduler, we recommend that you configure the network-control queue (or the queue you configure for host outbound traffic if it is not the network-control queue) as a strict-high priority queue. Strict-high priority queues receive the bandwidth required to transmit their entire queues before other queues are served. To limit the amount of bandwidth a strict-high priority queue can consume (and to prevent the strict-high priority queue from starving other queues), apply a shaping rate to the strict-high priority traffic in the scheduler configuration.

As with all strict-high priority traffic, if you configure the network-control queue (or any other queue) as a strict-high priority queue, you must also create a separate forwarding class set (priority group) that contains only strict-high priority traffic, and apply the strict-high priority forwarding class set and its traffic control profile (hierarchical scheduler) to the VCP interfaces.

#### Related Documentation

- [Understanding Default CoS Settings on page 5796](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Buffer Configuration on page 5891](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806](#)

### Virtual Chassis Fabric Only

- [CoS on Virtual Chassis Fabric \(VCF\) EX4300 Leaf Devices \(Mixed Mode\) on page 5958](#)

## CoS on Virtual Chassis Fabric (VCF) EX4300 Leaf Devices (Mixed Mode)

A Virtual Chassis Fabric (VCF) uses QFX5100 switches as spine devices and can use QFX5100, QFX3500, QFX3600, and EX4300 switches as leaf devices. When a VCF includes more than one type of leaf device (mixed mode), the CoS feature support on the VCF depends on the capability of the lowest-featured device. In mixed mode, the supported CoS features are the “lowest common denominator” of the features supported by the leaf devices. If one leaf device does not support a particular feature, that feature is not supported on the VCF even if every other leaf device supports the feature.



**NOTE:** EX4300 leaf devices do not support several CoS features that are supported on QFX5100, QFX3600, and QFX3500 devices. However, even when a VCF includes an EX4300 leaf device, other leaf devices might support those CoS features.

- [VCF CoS in Mixed Mode with an EX4300 Leaf Device on page 5958](#)
- [Scheduling on an EX4300 VCF Leaf Device on page 5960](#)

### VCF CoS in Mixed Mode with an EX4300 Leaf Device

In mixed mode, if all of the leaf devices are QFX5100, QFX3500, and QFX3600 switches, the full QFX Series CoS feature set is available, including data center bridging (DCB) features such as enhanced transmission selection (ETS, IEEE 802.1Qaz), priority-based flow control (PFC, IEEE 802.1Qbb), and Data Center Bridging Exchange Protocol (DCBX, an extension of LLDP, IEEE 802.1AB).

However, the EX4300 leaf device does not support DCB standards (ETS, PFC, DCBX). The lack of support for DCB standards means that the EX4300 leaf device does not support lossless transport. So a VCF that includes an EX4300 as a leaf device does not support lossless storage traffic such as Fibre Channel over Ethernet (FCoE).

In addition, a VCF with an EX4300 leaf device either does not support or has limited support for some other CoS features that the QFX Series switches support, including some buffer configuration features, some packet rewrite features, and Ethernet PAUSE (IEEE 802.3X).

[Table 538 on page 5958](#) summarizes the CoS support on a VCF in mixed mode with one or more EX4300 leaf devices.

**Table 538: Support of QFX CoS Features on a VCF in Mixed Mode with an EX4300 Leaf Device**

| QFX Series CoS Feature | Support in Mixed Mode with an EX4300 Leaf Device   |
|------------------------|--|
| Forwarding Classes     | The EX4300 leaf device uses the QFX Series default forwarding classes, the default QFX Series forwarding class to queue mapping, and the QFX Series maximum number of supported forwarding classes (12). |



**Table 538: Support of QFX CoS Features on a VCF in Mixed Mode with an EX4300 Leaf Device (*continued*)**

| QFX Series CoS Feature                                   | Support in Mixed Mode with an EX4300 Leaf Device  |
|--|---|
| Lossless Forwarding Classes                              | <p>Not supported.</p> <p>For example, the QFX Series default lossless forwarding classes <b>fcoe</b> and <b>no-loss</b> are not treated as lossless forwarding classes. Traffic mapped to lossless forwarding classes (default lossless forwarding classes or user-defined lossless forwarding classes) is treated as best-effort traffic.</p>  |
| Shared buffer configuration                              | <p>Ingress shared buffer configuration is not supported. Egress shared buffer configuration does not support partitioning into three buffer pools.</p> <p>If there is a shared buffer configuration, only the total egress shared buffer configuration is used. Ingress shared buffer configuration and egress buffer partitioning configuration is ignored.</p>  |
| Classifier on a Layer 2 interface                        | One classifier per protocol is supported on a port. On a physical port, for a particular protocol, the same Layer 2 classifier is used on all of the logical interfaces.  |
| Classifier on a Layer 3 interface                        | Supported.  |
| Multi-destination classifier                             | <p>Supported.</p> <p>The EX4300 leaf device uses the same default classifier as the QFX5100 spine device. As on QFX Series switches, a multi-destination classifier is global and is applied to all VCF interfaces. Multi-destination classifiers are valid only for multicast forwarding classes. You can configure two multi-destination classifiers, one for IEEE 802.1p traffic and one for DSCP traffic (the DSCP multi-destination classifier applies to both IPv4 and IPv6 traffic).</p>                                       |
| Congestion notification profile                          | <p>Not supported.</p> <p>If a congestion notification profile is configured on the QFX5100 spine device, it is ignored because the EX4300 leaf device does not support lossless transport, so end-to-end lossless behavior is not possible.</p>   |
| Ethernet PAUSE (IEEE 802.3X)                             | <p>Not supported.</p> <p>If Ethernet PAUSE is configured, it is ignored.</p>  |
| Hierarchical scheduling (ETS)                            | <p>Translated into port-based scheduling.</p> <p>The EX4300 device does not support ETS scheduling. A VCF translates ETS scheduling configured on a QFX5100 spine device into port scheduling on an EX4300 leaf device. The hierarchical structure of mapping forwarding classes into forwarding class sets (fc-sets) is ignored.</p> <p><a href="#">“Scheduling on an EX4300 VCF Leaf Device” on page 5960</a> provides details on how a VCF translates QFX Series ETS scheduling into port scheduling on an EX4300 leaf device.</p> |
| Hierarchical scheduling (ETS) on a spine device VCP port | On QFX5100 VCP ports, the hierarchical mapping of forwarding classes to forwarding class sets is supported. However, scheduling on an EX4300 leaf device is translated into port scheduling.  |

**Table 538: Support of QFX CoS Features on a VCF in Mixed Mode with an EX4300 Leaf Device** (*continued*)

| QFX Series CoS Feature               | Support in Mixed Mode with an EX4300 Leaf Device  |
|--------------------------------------|---|
| Drop profile (WRED)                  | <p>QFX Series drop profiles are supported. The EX4300 device as a standalone switch supports four packet loss priorities. However, as part of a mixed mode VCF, the EX4300 leaf device supports only the three packet loss priorities that the QFX Series switches support:</p> <ul style="list-style-type: none"> <li>• low</li> <li>• medium-high</li> <li>• high</li> </ul> <p>Supporting only three packet loss priorities means that the behavior of the EX4300 switch as a leaf device is different from the behavior as a standalone switch.</p> |
| Rewrite rules on a Layer 2 interface | Supported, but with a limit of one rewrite rule per physical interface. All traffic uses the same rewrite rule.   |
| Rewrite rules on a Layer 3 interface | Supported, but with a limit of one rewrite rule per physical interface. The same rewrite rule is used on all traffic on the interface.  |
| Rewrite value for FCoE traffic       | <p>Not supported.</p> <p>If a rewrite value for FCoE traffic, is configured, it is ignored. (A mixed mode VCF does not support lossless traffic.)</p>   |

In addition to the CoS limitations shown in [Table 538 on page 5958](#), using wild cards in a LAG configuration is not supported in mixed mode with one or more EX4300 leaf devices.

### Scheduling on an EX4300 VCF Leaf Device

Because the EX4300 leaf device does not support ETS, the VCF translates the ETS scheduling configuration into the port scheduling configuration that the EX4300 device supports. The QFX5100 spine device uses two-tier ETS scheduling, as described in detail in [“Understanding CoS Hierarchical Port Scheduling \(ETS\)” on page 5862](#).

Briefly, ETS allocates port bandwidth into forwarding class sets (priority groups) and forwarding classes (priorities) in a hierarchical manner. Each forwarding class set consists of individual forwarding classes, with each forwarding class mapped to an output queue.

Port bandwidth (minimum guaranteed bandwidth and maximum bandwidth) is allocated to each forwarding class set. Forwarding class set bandwidth is in turn allocated to the forwarding classes in the forwarding class set. If a forwarding class does not use its bandwidth allocation, other forwarding classes within the same forwarding class set can share the unused bandwidth. If the forwarding classes in a forwarding class set do not use the bandwidth allocated to that forwarding class set, other forwarding class sets on the port can share the unused bandwidth. (This is how ETS increases port bandwidth utilization, by sharing unused bandwidth among forwarding classes and forwarding class sets.)

However, the EX4300 leaf device supports port scheduling, not ETS. Port scheduling is a “flat” scheduling method that allocates bandwidth directly to forwarding classes in a non-hierarchical manner.

The VCF translates the two tiers of the ETS scheduling configuration (forwarding class sets and forwarding classes) into a single port scheduling configuration as follows:

- The bandwidth allocated to a forwarding class set is divided equally among the forwarding classes in the forwarding class set. (Traffic control profiles schedule bandwidth allocation to forwarding class sets.) The minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth limit (**shaping-rate**) of the forwarding class set determine the guaranteed minimum bandwidth and the maximum bandwidth the forwarding classes receive, *unless* those values are different in the forwarding class scheduler configuration.
- If there is an explicit forwarding class bandwidth scheduler configuration, it overrides the forwarding class set configuration. Bandwidth scheduling values that are not explicitly configured in a forwarding class scheduler use the values from the forwarding class set (the traffic control profile configuration). Forwarding class schedulers control the minimum guaranteed bandwidth (**transmit-rate**), the maximum bandwidth (**shaping-rate**), and the priority (**priority**) for each forwarding class (output queue). Because the priority value is not configured at the forwarding class set level, the priority configured in the forwarding class scheduler is always used.

The following two scenarios illustrate how a VCF translates an ETS configuration into a port scheduling configuration:

#### Scenario 1

A forwarding class set named **fc-set-1** has a configured guaranteed minimum bandwidth (**guaranteed-rate**) of 4G, and a configured maximum bandwidth (**shaping-rate**) of 5G.

Forwarding class set **fc-set-1** consists of two forwarding classes, named **fc-1** and **fc-2**:

- Forwarding class **fc-1** has a guaranteed minimum bandwidth (**transmit-rate**) of 2.5G. There is no configured maximum bandwidth (**shaping-rate**).
- Forwarding class **fc-2** has a guaranteed minimum bandwidth (**transmit-rate**) of 1.5G. There is no configured maximum bandwidth (**shaping-rate**).

On the EX4300 leaf device, the ETS configuration above is translated approximately to the following port scheduling configuration:

- Guaranteed minimum bandwidth—Because guaranteed minimum bandwidth has been explicitly configured in the forwarding class scheduler, forwarding class **fc-1** receives a transmit rate of 2.5G and forwarding class **fc-2** receives a transmit rate of 1.5G.



**NOTE:** If there had been no forwarding class scheduler **transmit-rate** configuration, then the forwarding class set minimum guaranteed bandwidth of 4G would have been split evenly between the forwarding classes, with each forwarding class receiving a minimum guaranteed bandwidth rate of 2G.

- Maximum bandwidth—Because there is no explicit maximum bandwidth (**shaping-rate**) configuration for the forwarding classes, the forwarding classes that belong to the

forwarding class set receive an equal share of the maximum bandwidth configured at the forwarding class set level in the traffic control profile. Because the forwarding class set maximum bandwidth is 5G, forwarding classes **fc-1** and **fc-2** each receive a maximum bandwidth of 2.5G.

In this scenario, the minimum guaranteed bandwidth and the maximum bandwidth configured at the forwarding class set hierarchy level are achieved on the forwarding classes that belong to the forwarding class set. (This does not always happen, as Scenario 2 shows.) However, unused bandwidth is not shared the same way. For example, if forwarding class **fc-1** experienced a burst of traffic at 3.5G, it would be limited to a maximum of 2.5G and traffic would be dropped. Using ETS, if forwarding class **fc-2** was not using its allocated maximum bandwidth, then **fc-1** could use (share) that unused bandwidth. But flat port scheduling does not share the unused bandwidth.

## Scenario 2

A forwarding class set named **fc-set-2** has a configured guaranteed minimum bandwidth (**guaranteed-rate**) of 6G, and a configured maximum bandwidth (**shaping-rate**) of 9G.

Forwarding class set **fc-set-2** consists of three forwarding classes, named **fc-3**, **fc-4**, and **fc-5**:

- Forwarding class **fc-3** has a guaranteed minimum bandwidth (**transmit-rate**) of 1G. There is no configured maximum bandwidth (**shaping-rate**).
- Forwarding class **fc-4** has a maximum bandwidth (**shaping-rate**) of 2G. There is no configured guaranteed minimum bandwidth (**transmit-rate**).
- Forwarding class **fc-5** has a guaranteed minimum bandwidth (**transmit-rate**) of 3G. There is no configured maximum bandwidth (**shaping-rate**).

On the EX4300 leaf device, the ETS configuration above is translated approximately to the following port scheduling configuration:

- Guaranteed minimum bandwidth—Two forwarding classes (**fc-3** and **fc-5**) have an explicitly configured transmit rate, and one forwarding class (**fc-4**) does not. Forwarding classes **fc-3** and **fc-5** receive the minimum guaranteed bandwidth configured in their schedulers, so forwarding class **fc-3** receives 1G guaranteed minimum bandwidth and forwarding class **fc-5** receives 3G guaranteed minimum bandwidth.

Forwarding class **fc-4** does not have an explicitly configured transmit rate, so the port derives the minimum guaranteed bandwidth from the forwarding class set guaranteed rate. Forwarding class set **fc-set-2** has a minimum guaranteed bandwidth (**guaranteed-rate**) of 6G, and there are three forwarding classes in the forwarding class set. Forwarding class **fc-4** receives an equal share (one third) of the forwarding class set minimum guaranteed bandwidth. So forwarding class **fc-4** is allocated a guaranteed minimum bandwidth (**transmit-rate**) of 2G (6G divided by 3 forwarding classes = 2G).

- Maximum bandwidth—Forwarding class **fc-4** has an explicitly configured shaping rate, and forwarding classes **fc-3** and **fc-5** do not. Forwarding class **fc-4** receives the maximum bandwidth configured in its scheduler, so forwarding class **fc-4** receives a maximum bandwidth of 2G.

Forwarding classes **fc-3** and **fc-5** do not have explicitly configured shaping rates, so the port derives the maximum bandwidth from the forwarding class set shaping rate. Forwarding class set **fc-set-2** has a maximum bandwidth (**shaping-rate**) of 9G, and there are three forwarding classes in the forwarding class set. Forwarding classes **fc-3** and **fc-5** each receive an equal share (one third) of the forwarding class set shaping rate. So forwarding classes **fc-3** and **fc-5** are allocated a maximum bandwidth of 3G each (9G divided by 3 forwarding classes = 3G).

Forwarding class **fc-4** receives less maximum bandwidth than forwarding classes **fc-3** and **fc-5** because the explicitly configured shaping rate for forwarding class **fc-4** is only 2G, and the explicit forwarding class configuration overrides the forwarding class set configuration.



**NOTE:** Scenario 2 shows that in some cases, the guaranteed minimum bandwidth (**guaranteed-rate**) and the maximum bandwidth (**shaping-rate**) configured for a forwarding class set might not be achieved at the forwarding class (queue) level. In Scenario 2, forwarding class set **fc-set-2** has a shaping rate of 9G, but the sum of the implemented forwarding class shaping rates is only 8G [(3G for **fc-3**) + (2G for **fc-4**) + (3G for **fc-5**)].

#### Related Documentation

- [Virtual Chassis Fabric Software Features Overview](#)
- [Understanding Default CoS Settings on page 5796](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)

## Learn About Technology

- [Data Center Technology Overview Videos on page 5963](#)

### Data Center Technology Overview Videos

Juniper Information Experience (iX) videos provide brief, high-level overviews of data center technologies and concepts. Each video runs approximately one-and-a-half to two minutes in length. This document contains SDN-related videos and links to conceptual documents that contain other data center technology videos:

- [Learn About Video: Why Do We Need an IP Fabric? on page 5963](#)
- [Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric? on page 5964](#)
- [Learn About Video: Why Use an Overlay Network in a Data Center? on page 5964](#)
- [Conceptual Documents That Contain Technology Overview Videos on page 5964](#)

#### Learn About Video: Why Do We Need an IP Fabric?

The video *Why Do We Need an IP Fabric?* presents a brief overview of IP Fabric use cases.



Video: [Why Do We Need an IP Fabric?](#)

---

### **Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?**

---

The video *What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?* presents a brief overview of the arguments for using Border Gateway Protocol (BGP) as the data center IP fabric control plane protocol.



Video: [What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?](#)

---

### **Learn About Video: Why Use an Overlay Network in a Data Center?**

---

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of data center overlay networks.



Video: [Why Use an Overlay Network in a Data Center?](#)

---

### **Conceptual Documents That Contain Technology Overview Videos**

---

The following conceptual documents include brief video overviews of the technology:

- [Understanding DCB Features and Requirements on page 5515](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding DCBX on page 5580](#)
- [Understanding PFC Functionality Across Layer 3 Interfaces on page 5950](#)
- [Virtual Chassis Fabric Overview on page 7033](#)
- [“Understanding In-Service Software Upgrade \(ISSU\)” on page 25 and “In-Service Software Upgrade \(ISSU\) System Requirements” on page 13 \(same video\)](#)

## CHAPTER 73

# Configuration

- [Configuration Examples on page 5965](#)
- [Configuration Examples \(QFX5100 Switches Only\) on page 6138](#)
- [Configuration Tasks on page 6156](#)
- [Configuration Statements on page 6192](#)

### Configuration Examples

---

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5987](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5995](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Example: Configuring Forwarding Classes on page 6075](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Queue Scheduling Priority on page 6087](#)
- [Example: Configuring ECN on page 6090](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 6110](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 6128](#)

### Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Hierarchical port scheduling defines the class-of-service (CoS) properties of output queues, which are mapped to forwarding classes (forwarding classes are mapped to IEEE 802.1p priorities, so mapping queues to forwarding classes also maps queues to priorities). Hierarchical port scheduling enables you to group priorities that require similar CoS resources into priority groups. You define the port bandwidth resources for a priority group, and you define the amount of the priority group's resources that each priority in the group can use.

Hierarchical port scheduling is the Junos OS implementation of enhanced transmission selection (ETS, described in IEEE 802.1Qaz). One major benefit of hierarchical port scheduling is greater port bandwidth utilization. If a priority group on a port does not use all of its allocated bandwidth, other priority groups on that port can use that bandwidth. Also, if a priority within a priority group does not use its allocated bandwidth, other priorities within that priority group can use that bandwidth.

Configuring hierarchical scheduling is a multistep procedure that includes:

- Mapping forwarding classes to queues
- Defining forwarding class sets (priority groups)
- Defining behavior aggregate classifiers
- Configuring priority-based flow control (PFC) for lossless priorities (queues)
- Applying classifiers and PFC configuration to ingress interfaces
- Defining drop profiles
- Defining schedulers
- Mapping forwarding classes to schedulers
- Defining traffic control profiles
- Assigning priority groups and traffic control profiles to egress ports

This example describes how to configure hierarchical scheduling:

- [Requirements on page 5967](#)
- [Overview on page 5967](#)



- [Configuration on page 5970](#)
- [Verification on page 5978](#)

---

## Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

---

## Overview

Keep the following considerations in mind when you plan the port bandwidth allocation for priority groups and for individual priorities:

- How much traffic and what types of traffic you expect to traverse the system.
- How you want to divide different types of traffic into priorities (forwarding classes, also called queues) to apply different CoS treatment to the traffic. Dividing traffic into priorities includes:
  - Mapping the code points of ingress traffic to forwarding classes using behavior aggregate (BA) classifiers. This classifies incoming traffic into the appropriate forwarding class.
  - Mapping forwarding classes to output queues. This defines the output queue for each type of traffic.
  - Attaching the BA classifier to the desired ingress interfaces so that incoming traffic maps to the desired forwarding classes and queues.
- How you want to organize priorities into priority groups (forwarding class sets).

Traffic that requires similar treatment usually belongs in the same priority group. To do this, place forwarding classes that require similar bandwidth, loss, and other characteristics in the same forwarding class set. For example, you can map all types of best-effort traffic forwarding classes into one forwarding class set.

- How much of the port bandwidth you want to allocate to each priority group and to each of the priorities in each priority group. The following considerations apply to bandwidth allocation:
  - Estimate how much traffic you expect in each forwarding class (output queue) and how much traffic you expect in each forwarding class set (the aggregate amount of traffic in the forwarding classes that belong to the forwarding class set).
  - The combined minimum guaranteed bandwidth of the priorities (forwarding classes) in a priority group should not exceed the minimum guaranteed bandwidth of the priority group. The transmit rate scheduler parameter defines the minimum guaranteed bandwidth for forwarding classes. Scheduler maps associate schedulers with forwarding classes.
  - The combined minimum guaranteed bandwidth of the priority groups (forwarding class sets) on a port should not exceed the port's total bandwidth. Traffic control profiles define the minimum bandwidth for a forwarding class set. Associating a

scheduler map with a traffic control profile sets the scheduling for the individual forwarding classes in the forwarding class set.

This example creates hierarchical port scheduling by defining priority groups for best effort, guaranteed delivery, and high-performance computing (HPC) traffic. Each priority group includes priorities that need to receive similar CoS treatment. Each priority group and each priority within each priority group receive the CoS resources needed to service their flows. Lossless priorities use PFC to prevent packet loss when the network experiences congestion.

### Topology

Table 539 on page 5968 shows the configuration components for this example.

**Table 539: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology**

| Property   | Settings  |
|--|---|
| Hardware   | QFX3500 switch  |
| Mapping of forwarding classes (priorities) to queues   | <p><b>best-effort</b> to queue 0</p> <p><b>be</b> to queue 1</p> <p><b>fcoe</b> (Fibre Channel over Ethernet) to queue 3</p> <p><b>no-loss</b> to queue 4</p> <p><b>hpc</b> (high-performance computing) to queue 5</p> <p><b>network-control</b> to queue 7</p> <p><b>NOTE:</b> If you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless <b>fcoe</b> and <b>no-loss</b> forwarding classes. If you explicitly configure the default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (<b>best-effort</b>) traffic and does <i>not</i> receive lossless treatment.</p> <p>In Junos OS Release 12.3 and later, you can include the <i>no-loss</i> packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.</p> |
| Forwarding class sets (priority groups)  | <p><b>best-effort-pg:</b> contains forwarding classes <b>best-effort</b>, <b>be</b>, and <b>network control</b></p> <p><b>guar-delivery-pg:</b> contains forwarding classes <b>fcoe</b> and <b>no-loss</b></p> <p><b>hpc-pg:</b> contains forwarding class <b>hpc</b></p>   |
| Behavior aggregate classifier (maps forwarding classes and loss priorities to incoming packets by IEEE 802.1 code point) | <p>Name—<b>hsclassifier1</b></p> <p>Code point mapping:</p> <ul style="list-style-type: none"> <li>• <b>000</b> to forwarding class <b>best-effort</b> and loss priority <b>low</b></li> <li>• <b>001</b> to forwarding class <b>be</b> and loss priority <b>high</b></li> <li>• <b>011</b> to forwarding class <b>fcoe</b> and loss priority <b>low</b></li> <li>• <b>100</b> to forwarding class <b>no-loss</b> and loss priority <b>low</b></li> <li>• <b>101</b> to forwarding class <b>hpc</b> and loss priority <b>low</b></li> <li>• <b>110</b> to forwarding class <b>network-control</b> and loss priority <b>low</b></li> </ul>   |

**Table 539: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology** (*continued*)

| Property                              | Settings  |
|---------------------------------------|---|
| PFC                                   | <p>Congestion notification profile name—<b>gd-cnp</b></p> <p>PFC enabled on code points: <b>011</b> (<b>fcoe</b> priority), <b>010</b> (<b>no-loss</b> priority)</p>  |
| Drop profiles                         | <p><b>dp-be-low</b>: drop start point <b>25</b>, drop end point <b>50</b>, maximum drop rate <b>80</b></p> <p><b>dp-be-high</b>: drop start point <b>10</b>, drop end point <b>40</b>, maximum drop rate <b>100</b></p> <p><b>dp-hpc</b>: drop start point <b>75</b>, drop end point <b>90</b>, maximum drop rate <b>75</b></p> <p><b>dp-nc</b>: drop start point <b>80</b>, drop end point <b>100</b>, maximum drop rate <b>100</b></p>  |
| Queue schedulers                      | <p><b>be-sched</b>: minimum bandwidth <b>3g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b>, drop profiles <b>dp-be-low</b> and <b>dp-be-high</b></p> <p><b>fcoe-sched</b>: minimum bandwidth <b>2.5g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b></p> <p><b>hpc-sched</b>: minimum bandwidth <b>2g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b>, drop profile <b>dp-hpc</b></p> <p><b>nc-sched</b>: minimum bandwidth <b>500m</b>, maximum bandwidth <b>100%</b>, priority <b>low</b>, drop profile <b>dp-nc</b></p> <p><b>nl-sched</b>: minimum bandwidth <b>2g</b>, maximum bandwidth <b>100%</b>, priority <b>low</b></p> |
| Forwarding class-to-scheduler mapping | <p>Scheduler map <b>be-map</b>:</p> <p>Forwarding class <b>best-effort</b>, scheduler <b>be-sched</b></p> <p>Forwarding class <b>be</b>, scheduler <b>be-sched</b></p> <p>Forwarding class <b>network-control</b>, scheduler <b>nc-sched</b></p> <p>Scheduler map <b>gd-map</b>:</p> <p>Forwarding class <b>fcoe</b>, scheduler <b>fcoe-sched</b></p> <p>Forwarding class <b>no-loss</b>, scheduler <b>nl-sched</b></p> <p>Scheduler map <b>hpc-map</b>:</p> <p>Forwarding class <b>hpc</b>, scheduler <b>hpc-sched</b></p>   |
| Traffic control profiles              | <p><b>be-tcp</b>: scheduler map <b>be-map</b>, minimum bandwidth <b>3.5g</b>, maximum bandwidth <b>100%</b></p> <p><b>gd-tcp</b>: scheduler map <b>gd-map</b>, minimum bandwidth <b>4.5g</b>, maximum bandwidth <b>100%</b></p> <p><b>hpc-tcp</b>: scheduler map <b>hpc-map</b>, minimum bandwidth <b>2g</b>, maximum bandwidth <b>100%</b></p>   |
| Interfaces                            | <p>This example configures hierarchical port scheduling on interfaces <b>xe-0/0/20</b> and <b>xe-0/0/21</b>. Because traffic is bidirectional, you apply the ingress and egress configuration components to both interfaces:</p> <ul style="list-style-type: none"> <li>• Classifier Name—<b>hsclassifier1</b></li> <li>• Forwarding class sets—<b>best-effort-pg</b>, <b>guar-deliver-pg</b>, <b>hpc-pg</b></li> <li>• Congestion notification profile—<b>gd-cnp</b></li> </ul>  |

Figure 211 on page 5970 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example. You can perform the configuration steps in a different sequence if you want.

**Figure 211: Hierarchical Port Scheduling Components Block Diagram**


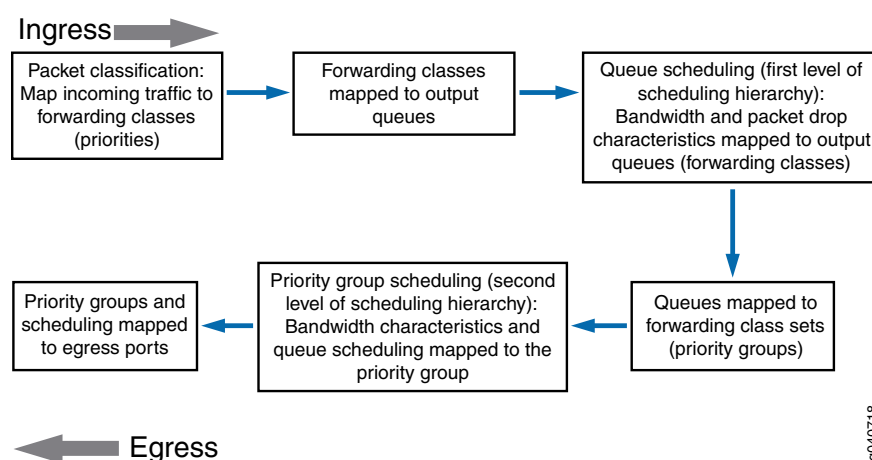


Figure 212 on page 5970 shows a block diagram of the hierarchical scheduling packet flow from ingress to egress.

**Figure 212: Hierarchical Port Scheduling Packet Flow Block Diagram**



### Configuration

#### CLI Quick Configuration

To quickly configure hierarchical port scheduling, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network

configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets guar-delivery-pg class fcoe
set forwarding-class-sets guar-delivery-pg class no-loss
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points 000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority low code-points 011
set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority low code-points 100
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low code-points 110
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011 pfc
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100 pfc
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability 80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability 100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability 100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability 75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers fcoe-sched priority low transmit-rate 2500m
set schedulers fcoe-sched shaping-rate percent 100
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set schedulers nl-sched priority low transmit-rate 2g
set schedulers nl-sched shaping-rate percent 100
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles be-tcp shaping-rate percent 100
set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate 4500m
set traffic-control-profiles gd-tcp shaping-rate percent 100
```

```

set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp

```

### Step-by-Step Procedure

To perform a step-by-step configuration of the forwarding classes (priorities), forwarding class sets (priority groups), classifiers, queue schedulers, PFC, traffic control profiles, and interfaces to set up hierarchical port scheduling (ETS):

1. Configure the forwarding classes (priorities) and map them to unicast output queues (do not explicitly map the **fcoe** and **no-loss** forwarding classes to output queues; use the default configuration):

```

[edit class-of-service]
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class be2 queue-num 1
user@switch# set forwarding-classes class hpc queue-num 5
user@switch# set forwarding-classes class network-control queue-num 7

```

2. Configure forwarding class sets (priority groups) to group forwarding classes (priorities) that require similar CoS treatment:

```

[edit class-of-service]
user@switch# set forwarding-class-sets best-effort-pg class best-effort
user@switch# set forwarding-class-sets best-effort-pg class be2
user@switch# set forwarding-class-sets best-effort-pg class network-control
user@switch# set forwarding-class-sets guar-delivery-pg class fcoe
user@switch# set forwarding-class-sets guar-delivery-pg class no-loss
user@switch# set forwarding-class-sets hpc-pg class hpc

```

3. Configure a classifier to set the loss priority and IEEE 802.1 code points assigned to each forwarding class at the ingress:

```

[edit class-of-service]
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort
loss-priority low code-points 000
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority
high code-points 001
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority
low code-points 011
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority
low code-points 100
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low
code-points 101
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control
loss-priority low code-points 110

```

4. Configure a congestion notification profile to enable PFC on the FCoE and no-loss queue IEEE 802.1 code points:

```

[edit class-of-service]
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011
pfc
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100
pfc

```

5. Assign the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
```

6. Apply the PFC configuration to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
```

7. Configure the drop profile for the best-effort low loss-priority queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50
drop-probability 0 drop-probability 80
```

8. Configure the drop profile for the best-effort high loss-priority queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40
drop-probability 0 drop-probability 100
```

9. Configure the drop profile for the network-control queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability
0 drop-probability 100
```

10. Configure the drop profile for the high-performance computing queue:

```
[edit class-of-service]
user@switch# set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability
0 drop-probability 75
```

11. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profiles for the best-effort queue:

```
[edit class-of-service]
user@switch# set schedulers be-sched priority low transmit-rate 3g
user@switch# set schedulers be-sched shaping-rate percent 100
user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any
drop-profile dp-be-low
user@switch# set schedulers be-sched drop-profile-map loss-priority high protocol any
drop-profile dp-be-high
```

12. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 2500m
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

13. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the high-performance computing queue:

```
[edit class-of-service]
user@switch# set schedulers hpc-sched priority low transmit-rate 2g
user@switch# set schedulers hpc-sched shaping-rate percent 100
user@switch# set schedulers hpc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-hpc
```

14. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the network-control queue:  

```
[edit class-of-service]
user@switch# set schedulers nc-sched priority low transmit-rate 500m
user@switch# set schedulers nc-sched shaping-rate percent 100
user@switch# set schedulers nc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-nc
```
15. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the no-loss queue:  

```
[edit class-of-service]
user@switch# set schedulers nl-sched priority low transmit-rate 2g
user@switch# set schedulers nl-sched shaping-rate percent 100
```
16. Map the schedulers to the appropriate forwarding classes (queues):  

```
[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class be2 scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class network-control scheduler
nc-sched
user@switch# set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
user@switch# set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
```
17. Define the traffic control profile for the best-effort priority group (queue scheduler to mapping, minimum guaranteed bandwidth, and maximum bandwidth):  

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate
3500m
user@switch# set traffic-control-profiles be-tcp shaping-rate percent 100
```
18. Define the traffic control profile for the guaranteed delivery priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):  

```
[edit class-of-service]
user@switch# set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate
4500m
user@switch# set traffic-control-profiles gd-tcp shaping-rate percent 100
```
19. Define the traffic control profile for the high-performance computing priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):  

```
[edit class-of-service]
user@switch# set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate
2g
user@switch# set traffic-control-profiles hpc-tcp shaping-rate percent 100
```
20. Apply the three priority groups (forwarding class sets) and the appropriate traffic control profiles to the egress ports:  

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```



```

user@switch# set interfaces xe-0/0/21 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp

```

### Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** and **no-loss** lossless forwarding classes):

```

user@switch> show configuration class-of-service
classifiers {
  ieee-802.1 hsclassifier1 {
    forwarding-class best-effort {
      loss-priority low code-points 000;
    }
    forwarding-class be2 {
      loss-priority high code-points 001;
    }
    forwarding-class fcoe {
      loss-priority low code-points 011;
    }
    forwarding-class no-loss {
      loss-priority low code-points 100;
    }
    forwarding-class hpc {
      loss-priority low code-points 101;
    }
    forwarding-class network-control {
      loss-priority low code-points 110;
    }
  }
}
drop-profiles {
  dp-be-low {
    interpolate {
      fill-level [ 25 50 ];
      drop-probability [ 0 80 ];
    }
  }
  dp-be-high {
    interpolate {
      fill-level [ 10 40 ];
      drop-probability [ 0 100 ];
    }
  }
  dp-hpc {
    interpolate {
      fill-level [ 75 90 ];
      drop-probability [ 0 75 ];
    }
  }
  dp-nc {
    interpolate {

```

```
        fill-level [ 80 100 ];
        drop-probability [ 0 100 ];
    }
}
forwarding-classes {
    class best-effort queue-num 0;
    class be2 queue-num 1;
    class hpc queue-num 5;
    class network-control queue-num 7;
}
traffic-control-profiles {
    be-tcp {
        scheduler-map be-map;
        shaping-rate percent 100;
        guaranteed-rate 3500000000;
    }
    gd-tcp {
        scheduler-map gd-map;
        shaping-rate percent 100;
        guaranteed-rate 4500000000;
    }
    hpc-tcp {
        scheduler-map hpc-map;
        shaping-rate percent 100;
        guaranteed-rate 2g;
    }
}
forwarding-class-sets {
    guar-delivery-pg {
        class fcoe;
        class no-loss;
    }
    best-effort-pg {
        class best-effort;
        class be2;
        class network-control;
    }
    hpc-pg {
        class hpc;
    }
}
congestion-notification-profile {
    gd-cnp {
        input {
            ieee-802.1 {
                code-point 011 {
                    pfc;
                }
                code-point 100 {
                    pfc;
                }
            }
        }
    }
}
```

```

interfaces {
  xe-0/0/20 {
    forwarding-class-set {
      best-effort-pg {
        output-traffic-control-profile be-tcp;
      }
      guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
      }
      hpc-pg {
        output-traffic-control-profile hpc-tcp;
      }
    }
    congestion-notification-profile gd-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 hsclassifier1;
      }
    }
  }
  xe-0/0/21 {
    forwarding-class-set {
      best-effort-pg {
        output-traffic-control-profile be-tcp;
      }
      guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
      }
      hpc-pg {
        output-traffic-control-profile hpc-tcp;
      }
    }
    congestion-notification-profile gd-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 hsclassifier1;
      }
    }
  }
}
scheduler-maps {
  be-map {
    forwarding-class best-effort scheduler be-sched;
    forwarding-class network-control scheduler nc-sched;
    forwarding-class be2 scheduler be-sched;
  }
  gd-map {
    forwarding-class fcoe scheduler fcoe-sched;
    forwarding-class no-loss scheduler nl-sched;
  }
  hpc-map {
    forwarding-class hpc scheduler hpc-sched;
  }
}
schedulers {
  be-sched {

```

```
    transmit-rate 3g;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-be-low;
    drop-profile-map loss-priority high protocol any drop-profile dp-be-high;
  }
  fcoe-sched {
    transmit-rate 2500000000;
    shaping-rate percent 100;
    priority low;
  }
  hpc-sched {
    transmit-rate 2g;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-hpc;
  }
  nc-sched {
    transmit-rate 500m;
    shaping-rate percent 100;
    priority low;
    drop-profile-map loss-priority low protocol any drop-profile dp-nc;
  }
  nl-sched {
    transmit-rate 2g;
    shaping-rate percent 100;
    priority low;
  }
}
```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command, and then copy the hierarchy and paste it into the switch terminal window.

---

## Verification

---

To verify that the hierarchical port scheduling components have been created and are operating properly, perform these tasks:

- [Verifying That the Forwarding Classes \(Priorities\) Have Been Created on page 5979](#)
- [Verifying That the Forwarding Class Sets \(Priority Groups\) Have Been Created on page 5979](#)
- [Verifying That the Classifier Has Been Created on page 5980](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5980](#)
- [Verifying That the Output Queue Schedulers Have Been Created on page 5981](#)
- [Verifying That the Drop Profiles Have Been Created on page 5984](#)

- [Verifying That the Priority Group Output Schedulers \(Traffic Control Profiles\) Have Been Created on page 5985](#)
- [Verifying the Interface Configuration on page 5986](#)

### *Verifying That the Forwarding Classes (Priorities) Have Been Created*

**Purpose** Verify that the forwarding classes have been created and mapped to the correct queues. (The system shows only the explicitly configured forwarding classes. It does not show default forwarding classes such as **fcoe** and **no-loss**.)

**Action** List the forwarding classes using the operational mode command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| be2              | 1  | 3     | normal            | Disabled |
| hpc              | 2  | 4     | normal            | Disabled |
| network-control  | 3  | 7     | normal            | Disabled |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command lists all of the configured forwarding classes, the internal identification number of each forwarding class, the queues that are mapped to the forwarding classes, the policing priority, and whether the forwarding class is lossless (no-loss packet drop attribute enabled) or lossy forwarding class (no-loss packet drop attribute disabled). The command output shows that:

- Forwarding class **best-effort** maps to queue **0** and is lossy
- Forwarding class **be2** maps to queue **1** and is lossy
- Forwarding class **hpc** maps to queue **5** and is lossy
- Forwarding class **network-control** maps to queue **7** and is lossy

In addition, the command lists the default multicast (multidestination) forwarding class and the default queue to which it is mapped.

### *Verifying That the Forwarding Class Sets (Priority Groups) Have Been Created*

**Purpose** Verify that the priority groups have been created and that the correct priorities (forwarding classes) belong to the appropriate priority group.

**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set**:

```
user@switch> show class-of-service forwarding-class-set
```

```
Forwarding class set: best-effort-pg, Type: normal-type, Forwarding class set
index: 19907
```

| Forwarding class | Index |
|------------------|-------|
| best-effort      | 0     |

|                 |   |
|-----------------|---|
| be2             | 1 |
| network-control | 5 |

Forwarding class set: guar-delivery-pg, Type: normal-type, Forwarding class set index: 43700

|                  |       |
|------------------|-------|
| Forwarding class | Index |
| fcoe             | 2     |
| no-loss          | 3     |

Forwarding class set: hpc-pg, Type: normal-type, Forwarding class set index: 60758

|                  |       |
|------------------|-------|
| Forwarding class | Index |
| hpc              | 4     |

**Meaning** The **show class-of-service forwarding-class-set** command lists all of the configured forwarding class sets (priority groups), the forwarding classes (priorities) that belong to each priority group, and the internal index number of each priority group. The command output shows that:

- The forwarding class set **best-effort-pg** includes the forwarding classes **best-effort**, **be2**, and **network-control**.
- The forwarding class set **guar-delivery-pg** includes the forwarding classes **fcoe** and **no-loss**.
- The forwarding class set **hpc-pg** includes the forwarding class **hpc**.

#### *Verifying That the Classifier Has Been Created*

**Purpose** Verify that the classifier maps forwarding classes to the correct IEEE 802.1p code points and packet loss priorities.

**Action** List the classifier configured for hierarchical port scheduling using the operational mode command **show class-of-service classifier name hsclassifier1**:

```
user@switch> show class-of-service classifier name hsclassifier1
Classifier: hsclassifier1, Code point type: ieee-802.1, Index: 43607
Code point      Forwarding class      Loss priority
000             best-effort           low
001             be2                   high
011             fcoe                  low
100             no-loss               low
101             hpc                   low
110             network-control      low
```

**Meaning** The **show class-of-service classifier name hsclassifier1** command lists all of the IEEE 802.1p code points and the loss priorities mapped to all of the forwarding classes in the classifier. The command output shows that the forwarding classes **best-effort**, **be2**, **no-loss**, **fcoe**, **hpc**, and **network-control** have been created and mapped to IEEE 802.1p code points and loss priorities.

#### *Verifying That Priority-Based Flow Control Has Been Enabled*

**Purpose** Verify that PFC is enabled on the correct priorities for lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: gd-cnp, Index: 51687
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Enabled  | 2500 |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      | 0                   |
| 010      | 1                   |
| 011      | 2                   |
| 100      | 3                   |
| 101      | 4                   |
| 110      | 5                   |
| 111      | 6                   |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** command lists all of the congestion notification profiles and the IEEE 802.1p code points with PFC enabled. The command output shows that PFC is enabled for code points **011** (**fcoe** priority and queue) and **100** (**no-loss** priority and queue) for the **gd-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

#### *Verifying That the Output Queue Schedulers Have Been Created*

**Purpose** Verify that the output queue schedulers have been created with the correct bandwidth parameters and priorities, mapped to the correct queues, and mapped to the correct drop profiles.

**Action** List the scheduler maps using the operational mode command **show class-of-service scheduler-map**:

```
user@switch> show class-of-service scheduler-map
```

```
Scheduler map: be-map, Index: 64023
```

```
Scheduler: be-sched, Forwarding class: best-effort, Index: 13005
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
```

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 55387 | dp-be-low              |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 4369  | dp-be-high             |

Scheduler: be-sched, Forwarding class: be2, Index: 13005

Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 55387 | dp-be-low              |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 4369  | dp-be-high             |

Scheduler: nc-sched, Forwarding class: network-control, Index: 45740

Transmit rate: 5000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | dp-nc                  |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler map: gd-map, Index: 61447

Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289

Transmit rate: 25000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler: nl-sched, Forwarding class: no-loss, Index: 29359

Transmit rate: 20000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler map: hpc-map, Index: 56941



```
Scheduler: hpc-sched, Forwarding class: hpc, Index: 55900
Transmit rate: 2000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       57716  dp-hpc
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>
```

**Meaning** The **show class-of-service scheduler-map** command lists all of the configured scheduler maps. For each scheduler map, the command output includes:

- The name of the scheduler map (**scheduler-map** field)
- The name of the scheduler (**scheduler** field)
- The forwarding classes mapped to the scheduler (**forwarding-class** field)
- The minimum guaranteed queue bandwidth (**transmit-rate** field)
- The scheduling priority (**priority** field)
- The maximum bandwidth in the priority group the queue can consume (**shaping-rate** field)
- The drop profile loss priority (**loss priority** field) for each drop profile name (**name** field)

The command output shows that:

- The scheduler map **be-map** has been created and has these properties:
  - There are two schedulers, **be-sched** and **nc-sched**.
  - The scheduler **be-sched** has two forwarding classes, **best-effort** and **be2**.
  - Scheduler **be-sched** forwarding classes **best-effort** and **be2** share a minimum guaranteed bandwidth of **3000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and use the drop profile **dp-be-low** for low loss-priority traffic, the default drop profile for medium-high loss-priority traffic, and the drop profile **dp-be-high** for high loss-priority traffic.
  - The scheduler **nc-sched** has one forwarding class, **network-control**.
  - The **network-control** forwarding class has a minimum guaranteed bandwidth of **5000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-nc** for low loss-priority traffic and the default drop profile for medium-high and high loss priority traffic.
- The scheduler map **gd-map** has been created and has these properties:
  - There are two schedulers, **fcoe-sched** and **nl-sched**.
  - The scheduler **fcoe-sched** has one forwarding class, **fcoe**.

- The **fcoe** forwarding class has a minimum guaranteed bandwidth of **2500000000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler **nl-sched** has one forwarding class, **no-loss**.
- The **no-loss** forwarding class has a minimum guaranteed bandwidth of **2000000000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler map **hpc-map** has been created and has these properties:
  - There is one scheduler, **hpc-sched**.
  - The scheduler **hpc-sched** has one forwarding class, **hpc**.
  - The **hpc** forwarding class has a minimum guaranteed bandwidth of **2000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-hpc** for low loss-priority traffic and the default drop profile for medium-high and high loss-priority traffic.

#### *Verifying That the Drop Profiles Have Been Created*

**Purpose** Verify that the drop profiles **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc** have been created with the correct fill levels and drop probabilities.

**Action** List the drop profiles using the operational mode command **show configuration class-of-service drop-profiles**:

```
user@switch> show configuration class-of-service drop-profiles
dp-be-low {
    interpolate {
        fill-level [ 25 50 ];
        drop-probability [ 0 80 ];
    }
}
dp-be-high {
    interpolate {
        fill-level [ 10 40 ];
        drop-probability [ 0 100 ];
    }
}
dp-hpc {
    interpolate {
        fill-level [ 75 90 ];
        drop-probability [ 0 75 ];
    }
}
dp-nc {
    interpolate {
        fill-level [ 80 100 ];
        drop-probability [ 0 100 ];
    }
}
```

**Meaning** The **show configuration class-of-service drop-profiles** command lists the drop profiles and their properties. The command output shows that there are four drop profiles configured, **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc**. The output also shows that:

- For **dp-be-low**, the drop start point (the first fill level) is when the queue is 25 percent filled, the drop end point (the second fill level) occurs when the queue is 50 percent filled, and the drop probability at the drop end point is 80 percent.
- For **dp-be-high**, the drop start point (the first fill level) is when the queue is 10 percent filled, the drop end point (the second fill level) occurs when the queue is 40 percent filled, and the drop probability at the drop end point is 100 percent.
- For **dp-hpc**, the drop start point (the first fill level) is when the queue is 75 percent filled, the drop end point (the second fill level) occurs when the queue is 90 percent filled, and the drop probability at the drop end point is 75 percent.
- For **dp-nc**, the drop start point (the first fill level) is when the queue is 80 percent filled, the drop end point (the second fill level) occurs when the queue is 100 percent filled, and the drop probability at the drop end point is 100 percent.

***Verifying That the Priority Group Output Schedulers (Traffic Control Profiles) Have Been Created***

**Purpose** Verify that the traffic control profiles **be-tcp**, **gd-tcp**, and **hpc-tcp** have been created with the correct bandwidth parameters and scheduler mapping.

**Action** List the traffic control profiles using the operational mode command **show class-of-service traffic-control-profile**:

```
user@switch> show class-of-service traffic-control-profile
Traffic control profile: be-tcp, Index: 40535
  Shaping rate: 100 percent
  Scheduler map: be-map
  Guaranteed rate: 3500000000

Traffic control profile: gd-tcp, Index: 37959
  Shaping rate: 100 percent
  Scheduler map: gd-map
  Guaranteed rate: 4500000000

Traffic control profile: hpc-tcp, Index: 47661
  Shaping rate: 100 percent
  Scheduler map: hpc-map
  Guaranteed rate: 2000000000
```

**Meaning** The **show class-of-service traffic-control-profile** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**traffic-control-profile**)
- The maximum port bandwidth the priority group can consume (**shaping-rate**)
- The scheduler map associated with the traffic control profile (**scheduler-map**)
- The minimum guaranteed priority group port bandwidth (**guaranteed-rate**)

The command output shows that:

- The traffic control profile **be-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **be-map**, and has a minimum guaranteed bandwidth of **3500000000 bps**.
- The traffic control profile **gd-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **gd-map**, and has a minimum guaranteed bandwidth of **4500000000 bps**.
- The traffic control profile **hpc-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **hpc-map**, and has a minimum guaranteed bandwidth of **2000000000 bps**.

### *Verifying the Interface Configuration*

**Purpose** Verify that the classifier, the congestion notification profile, and the forwarding class sets are configured on interfaces **xe-0/0/20** and **xe-0/0/21**.

**Action** List the interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
forwarding-class-set {
    best-effort-gp {
        output-traffic-control-profile be-tcp;
    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd_cnp;
unit 0 {
    classifiers {
        ieee-802.1 hscclassifier1;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
forwarding-class-set {
    best-effort-gp {
        output-traffic-control-profile be-tcp;
    }
    guar-delivery-pg {
        output-traffic-control-profile gd-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
congestion-notification-profile gd_cnp;
unit 0 {
    classifiers {
        ieee-802.1 hscclassifier1;
    }
}
```

**Meaning** The `show configuration class-of-service interfaces interface-name` command shows that each interface includes the forwarding class sets **best-effort-pg**, **guar-delivery-pg**, and **hpc-pg**, congestion notification profile **gd-cnp**, and the IEEE 802.1p classifier **hsclassifier1**.

- Related Documentation**
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
  - [Benefits of Configuring CoS Hierarchical Port Scheduling](#)
  - [Assigning CoS Components to Interfaces on page 6185](#)
  - [Example: Configuring WRED Drop Profiles on page 6071](#)
  - [Example: Configuring Drop Profile Maps on page 6073](#)
  - [Example: Configuring Forwarding Classes on page 6075](#)
  - [Example: Configuring Forwarding Class Sets on page 6078](#)
  - [Example: Configuring Queue Schedulers on page 6081](#)
  - [Example: Configuring Queue Scheduling Priority on page 6087](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
  - [Example: Configuring Maximum Output Bandwidth on page 6101](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
  - [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
  - [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
  - [Understanding CoS Scheduling Behavior and Configuration Considerations on page 5886](#)
  - [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#)
  - [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)

## Example: Configuring CoS PFC for FCoE Traffic

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS flag in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

To configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic.
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.
- Configure the bandwidth scheduling for the FCoE forwarding class output queue.
- Create a forwarding class set (priority group) that includes the FCoE forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).
- Configure the bandwidth scheduling for the FCoE priority group.
- Apply the scheduling to the egress interfaces.



**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding classes for the lossless fcoe forwarding class. If you explicitly configure default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 5988](#)
- [Overview on page 5988](#)
- [Configuration on page 5990](#)
- [Verification on page 5993](#)

---

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

---

### Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to:

- Assign FCoE traffic to the FCoE priority at the ingress.
- Create and apply CoS for the FCoE traffic using ETS (hierarchical port scheduling).

- Apply PFC to the FCoE traffic.
- Apply the configuration to ingress and egress interfaces.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Each interface in this example is configured as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and port scheduling are applied to all of the interfaces.

### Topology


Table 458 on page 5607 shows the configuration components for this example.

**Table 540: Components of the PFC for FCoE Traffic Configuration Topology**

| Component   | Settings  |
|---|---|
| Hardware  | QFX3500 switch  |
| Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point) | Code point <b>011</b> to forwarding class <b>fcoe</b> and loss priority <b>low</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b> |
| PFC congestion notification profile   | <b>fcoe-cnp:</b><br>Code point <b>011</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>  |
| FCoE queue scheduler  | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>   |
| Forwarding class-to-scheduler mapping   | Scheduler map <b>fcoe-map:</b><br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b>   |
| Forwarding class set (FCoE priority group)  | <b>fcoe-pg:</b><br>Forwarding class <b>fcoe</b><br>Egress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>                                     |
| Traffic control profile   | <b>fcoe-tcp:</b><br>Scheduler map <b>fcoe-map</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b>   |

Figure 195 on page 5608 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 213: PFC for FCoE Traffic Configuration Components Block Diagram



### Configuration

#### CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100
set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

#### Step-by-Step Procedure

To configure the FCoE forwarding class (priority), ingress classifier, output queue scheduling, forwarding class set (priority group) and its output port scheduling, PFC application, and interfaces to set up PFC for FCoE traffic:

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:
 

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
```
2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:



```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```

3. Apply the PFC configuration to the ingress interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
```

4. Assign the classifier to the ingress interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
```

5. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

6. Map the FCoE forwarding class to the FCoE scheduler:

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

7. Configure the forwarding class set for the FCoE traffic:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```

8. Define the traffic control profile for the FCoE forwarding class set:

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

9. Apply the FCoE forwarding class set and traffic control profile to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

## Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class):

```
user@switch> show configuration class-of-service
classifiers {
```

```
ieee-802.1 fcoe-classifier {
  forwarding-class fcoe {
    loss-priority low code-points 011;
  }
}
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 3000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
}
interfaces {
  xe-0/0/31 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-classifier;
      }
    }
  }
  xe-0/0/32 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-classifier;
      }
    }
  }
  xe-0/0/33 {
```

```

congestion-notification-profile fcoe-cnp;
forwarding-class-set {
    fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
    }
}
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}
}
xe-0/0/34 {
    congestion-notification-profile fcoe-cnp;
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    unit 0 {
        classifiers {
            ieee-802.1 fcoe-classifier;
        }
    }
}
}
scheduler-maps {
    fcoe-map {
        forwarding-class fcoe scheduler fcoe-sched;
    }
}
schedulers {
    fcoe-sched {
        transmit-rate 3000000000;
        shaping-rate percent 100;
        priority low;
    }
}
}

```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 5994](#)
- [Verifying the Ingress Interface PFC Configuration on page 5994](#)

### *Verifying That Priority-Based Flow Control Has Been Enabled*

**Purpose** Verify that PFC is enabled on the FCoE queue to enable lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: fcoe-cnp, Index: 51697
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      |                     |
|          | 0                   |
| 001      |                     |
|          | 1                   |
| 010      |                     |
|          | 2                   |
| 011      |                     |
|          | 3                   |
| 100      |                     |
|          | 4                   |
| 101      |                     |
|          | 5                   |
| 110      |                     |
|          | 6                   |
| 111      |                     |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point **011** for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### *Verifying the Ingress Interface PFC Configuration*

**Purpose** Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.

**Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
```

```

congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe-classifier;
    }
}

```

**Meaning** The `show configuration class-of-service interfaces` commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
  - [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

### Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).



**NOTE:** This example uses Junos OS without support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see [“Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG” on page 5614](#).

You can use an MC-LAG to provide a redundant aggregation layer for Fibre Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.



**NOTE:** This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the two switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see [“Example: Configuring Multichassis Link Aggregation” on page 2471](#). However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.

QFX Series switches and EX4600 switches support MC-LAGs. QFabric system Node devices do not support MC-LAGs.

This topic describes:

- [Requirements on page 5996](#)
- [Overview on page 5997](#)
- [Configuration on page 6002](#)
- [Verification on page 6010](#)

---

## Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX3500 Switches that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX3500 Switches that provide FCoE server access in transit switch mode and that connect to the MC-LAG switches. These switches can be standalone QFX3500 switches or they can be Node devices in a QFabric system.
- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 12.2 or later for the QFX Series.

## Overview

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX3500 switches that form the MC-LAG, including priority-based flow control (PFC) and enhanced transmission selection (ETS; hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group).



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Disable IGMP snooping on the FCoE VLAN.



**NOTE:** This is only necessary if IGMP snooping is enabled on the VLAN. Before Junos OS Release 13.2, IGMP snooping was enabled by default on VLANs. Beginning with Junos OS Release 13.2, IGMP snooping is enabled by default only on the default VLAN.

- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

## Topology

Switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 196 on page 5616](#).

Figure 214: Supported Topology for an MC-LAG on an FCoE Transit Switch

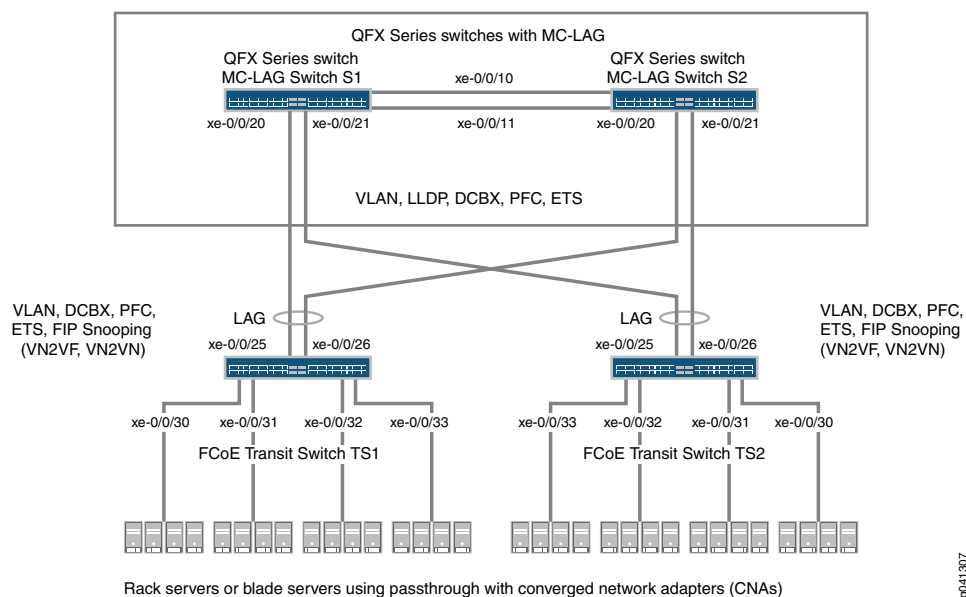


Table 459 on page 5616 shows the configuration components for this example.

Table 541: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

| Component  | Settings   |
|--|--|
| Hardware   | Four QFX3500 switches (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access). |
| Forwarding class (all switches)  | Default <b>fcoe</b> forwarding class.  |
| Classifier (forwarding class mapping of incoming traffic to IEEE priority) | Default IEEE 802.1p trusted classifier on all FCoE interfaces.   |



**Table 541: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (continued)**

| Component  | Settings   |
|--|--|
| LAGs and MC-LAG  | <p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S1 to Switch S2. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S2 to Switch S1. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p><b>NOTE:</b> Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> |
| FCoE queue scheduler (all switches)                      | <p><b>fcoe-sched:</b><br/>           Minimum bandwidth <b>3g</b><br/>           Maximum bandwidth <b>100%</b><br/>           Priority <b>low</b></p>   |
| Forwarding class-to-scheduler mapping (all switches)     | <p>Scheduler map <b>fcoe-map</b>:<br/>           Forwarding class <b>fcoe</b><br/>           Scheduler <b>fcoe-sched</b></p>   |
| Forwarding class set (FCoE priority group, all switches) | <p><b>fcoe-pg:</b><br/>           Forwarding class <b>fcoe</b></p> <p>Egress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>  |
| Traffic control profile (all switches)                   | <p><b>fcoe-tcp:</b><br/>           Scheduler map <b>fcoe-map</b><br/>           Minimum bandwidth <b>3g</b><br/>           Maximum bandwidth <b>100%</b></p>   |

Table 541: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (*continued*)

| Component  | Settings  |
|--|---|
| PFC congestion notification profile (all switches) | <b>fcoe-cnp:</b><br>Code point 011<br><br>Ingress interfaces: <ul style="list-style-type: none"> <li>• S1—LAG ae0 and MC-LAG ae1</li> <li>• S2—LAG ae0 and MC-LAG ae1</li> <li>• TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33</li> <li>• TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33</li> </ul>  |
| FCoE VLAN name and tag ID                          | Name— <b>fcoe_vlan</b><br>ID—100<br><br>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.<br><br>Disable IGMP snooping on the interfaces that belong to the FCoE VLAN on all four switches.   |
| FIP snooping                                       | Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.<br><br>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration. |



**NOTE:** This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode and tagged access mode ports if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is queue 3.



**NOTE:** In Junos OS Release 12.2, traffic mapped to explicitly configured forwarding classes, even lossless forwarding classes such as **fcoe**, is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment. To receive lossless treatment in Release 12.2, traffic must use one of the default lossless forwarding classes (**fcoe** or **no-loss**).

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.

- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk and tagged-access port modes. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (011) to the FCoE forwarding class. If you choose to configure the BA classifier instead of using the default classifier, you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point 011 in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure enhanced transmission selection (ETS, also known as hierarchical scheduling) on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic.
- Apply the ETS scheduling to the interfaces.
- Configure the port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers and how to disable IGMP snooping on the FCoE VLAN. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required

to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. ([“Example: Configuring Multichassis Link Aggregation” on page 2471](#) describes how to configure MC-LAGs.)
- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. ([“Configuring Link Aggregation” on page 2593](#) describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

### Configuration

---

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [Configuring MC-LAG Switches S1 and S2 on page 6004](#)
- [Configuring FCoE Transit Switches TS1 and TS2 on page 6005](#)
- [Results on page 6007](#)

#### CLI Quick Configuration

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the **[edit]** hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

#### Switch S1 and Switch S2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
```

```
set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the **[edit]** hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

## Switch TS1 and Switch TS2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
set ethernet-switching-options secure-access-port vlan fcoe_vlan examine-fip examine-vn2v2
beacon-period 90000
```

### *Configuring MC-LAG Switches S1 and S2*

**Step-by-Step Procedure** To configure CoS resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:  

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):  

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:  

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
5. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
```
6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
7. Apply the PFC configuration to the LAG and MC-LAG interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```
8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):  

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
9. Disable IGMP snooping on the FCoE VLAN:  

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```

10. Add the member interfaces to the LAG between the two MC-LAG switches:

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```

11. Add the member interfaces to the MC-LAG:

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```

12. Configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and MC-LAG interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```

14. Set the LAG and MC-LAG interfaces as FCoE trusted ports. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

### *Configuring FCoE Transit Switches TS1 and TS2*

#### **Step-by-Step Procedure**

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:

- ```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
  5. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```
  6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point 011:
 

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
  7. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:
 

```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile
fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile
fcoe-cnp
```
  8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):
 

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
  9. Disable IGMP snooping on the FCoE VLAN:
 

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```
  10. Add the member interfaces to the LAG:
 

```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```
  11. On the LAG (**ae1**), configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**):



```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

12. On the FCoE access interfaces (xe-0/0/30, xe-0/0/31, xe-0/0/32, xe-0/0/33), configure the port mode as **tagged-access** and membership in the FCoE VLAN (**fcoe\_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and FCoE access interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

14. Set the LAG interface as an FCoE trusted port. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```



**NOTE:** Access ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are not configured as FCoE trusted ports. The access ports remain in the default state as untrusted ports because they connect directly to FCoE devices and must perform FIP snooping to ensure network security.

15. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses VN2VN\_Port FIP snooping; the example is equally valid if you use VN2VF\_Port FIP snooping):

```
[edit]
user@switch# set ethernet-switching-options secure-access-port vlan fcoe_vlan
examine-fip examine-vn2vn beacon-period 90000
```

### Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are the same):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
```

```
fcoe-tcp {
  scheduler-map fcoe-map;
  shaping-rate percent 100;
  guaranteed-rate 3000000000;
}
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
}
interfaces {
  ae0 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  ae1 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
```



**NOTE:** The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see [“Example: Configuring Multichassis Link Aggregation” on page 2471](#).

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are the same):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 30000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  xe-0/0/30 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/31 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/32 {
```

```
forwarding-class-set {
  fcoe-pg {
    output-traffic-control-profile fcoe-tcp;
  }
}
congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
```



**NOTE:** The forwarding class and classifier configurations are not shown because the `show` command does not display default portions of the configuration.

---

## Verification

---

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default **fcoe** forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 6011](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created on page 6012](#)

- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created on page 6012](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 6012](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 6013](#)
- [Verifying That the Interfaces Are Correctly Configured on page 6015](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 6017](#)
- [Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 on page 6018](#)
- [Verifying That IGMP Snooping Is Disabled on the FCoE VLAN on page 6019](#)

### ***Verifying That the Output Queue Schedulers Have Been Created***

**Purpose** Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

**Action** List the scheduler map using the operational mode command **show class-of-service scheduler-map fcoe-map**:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
  Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
  Buffer Limit: none, Priority: low
  Excess Priority: unspecified
  Shaping rate: 100 percent,
  drop-profile-map-set-type: mark
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low            any        1      <default-drop-profile>
    Medium high    any        1      <default-drop-profile>
    High           any        1      <default-drop-profile>
```

**Meaning** The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

***Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created***

**Purpose** Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.

**Action** List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:

```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
  Shaping rate: 100 percent
  Scheduler map: fcoe-map
  Guaranteed rate: 3000000000
```

**Meaning** The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:

- The name of the traffic control profile (**fcoe-tcp**)
- The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
- The scheduler map associated with the traffic control profile (**fcoe-map**)
- The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

***Verifying That the Forwarding Class Set (Priority Group) Has Been Created***

**Purpose** Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.

**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:

```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
  Forwarding class          Index
  fcoe                      1
```

**Meaning** The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

***Verifying That Priority-Based Flow Control Has Been Enabled***

**Purpose** Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

**Action** List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
  Priority    PFC      MRU
  000        Disabled
  001        Disabled
  010        Disabled
  011        Enabled    2500
  100        Disabled
  101        Disabled
  110        Disabled
  111        Disabled
Type: Output
  Priority    Flow-Control-Queues
  000
  001        0
  010        1
  011        2
  100        3
  101        4
  110        5
  111        6
  111        7
```

**Meaning** The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011** (**fcoe** queue) for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

#### *Verifying That the Interface Class of Service Configuration Has Been Created*

**Purpose** Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

**Action** List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
```

```
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}  
  
ae1 {  
  forwarding-class-set {  
    fcoe-pg {  
      output-traffic-control-profile fcoe-tcp;  
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces  
xe-0/0/30 {  
  forwarding-class-set {  
    fcoe-pg {  
      output-traffic-control-profile fcoe-tcp;  
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}  
xe-0/0/31 {  
  forwarding-class-set {  
    fcoe-pg {  
      output-traffic-control-profile fcoe-tcp;  
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}  
xe-0/0/32 {  
  forwarding-class-set {  
    fcoe-pg {  
      output-traffic-control-profile fcoe-tcp;  
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}  
xe-0/0/33 {  
  forwarding-class-set {  
    fcoe-pg {  
      output-traffic-control-profile fcoe-tcp;  
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}  
ae1 {  
  forwarding-class-set {  
    fcoe-pg {  
      output-traffic-control-profile fcoe-tcp;  
    }  
  }  
  congestion-notification-profile fcoe-cnp;  
}
```



**Meaning** The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



**NOTE:** Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33**, which are not members of a LAG.

#### *Verifying That the Interfaces Are Correctly Configured*

**Purpose** Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

**Action** List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
xe-0/0/10 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/11 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/20 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/21 {
    ether-options {
        802.3ad ae1;
    }
}
```

```
ae0 {
  mtu 2180;
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members fcoe_vlan;
      }
    }
  }
}
ae1 {
  mtu 2180;
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
      vlan {
        members fcoe_vlan;
      }
    }
  }
}
```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```
user@switch> show configuration interfaces
xe-0/0/25 {
  ether-options {
    802.3ad ae1;
  }
}
xe-0/0/26 {
  ether-options {
    802.3ad ae1;
  }
}
xe-0/0/30 {
  mtu 2180;
  unit 0 {
    family ethernet-switching {
      port-mode tagged-access;
      vlan {
        members fcoe_vlan;
      }
    }
  }
}
xe-0/0/31 {
  mtu 2180;
  unit 0 {
    family ethernet-switching {
      port-mode tagged-access;
      vlan {
        members fcoe_vlan;
      }
    }
  }
}
xe-0/0/32 {
```

```

mtu 2180;
unit 0 {
    family ethernet-switching {
        port-mode tagged-access;
        vlan {
            members fcoe_vlan;
        }
    }
}
xe-0/0/33 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}

```

**Meaning** The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The port mode (**trunk** mode for interfaces that connect two switches, **tagged-access** mode for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe\_vlan**)

***Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces***

**Purpose** Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches

TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the port security configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration ethernet-switching-options secure-access-port**:

```
user@switch> show configuration ethernet-switching-options secure-access-port
interface ae1.0 {
    fcoe-trusted;
}
vlan fcoe_vlan {
    examine-fip {
        examine-vn2vn {
            beacon-period 90000;
        }
    }
}
```

**Meaning** The **show configuration ethernet-switching-options secure-access-port** command lists port security information, including whether a port is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- FIP snooping is enabled (**examine-fip**) on the FCoE VLAN (**fcoe\_vlan**), the type of FIP snooping is VN2VN\_Port FIP snooping (**examine-vn2vn**) and the beacon period is set to 90000 milliseconds. On Transit Switches TS1 and TS2, all interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Transit Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

#### ***Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2***

**Purpose** Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

**Action** List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan,      Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
...
```



**NOTE:** The output has been truncated to show only the relevant information.

**Meaning** The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe\_vlan**
- The FIP snooping mode is VN2VN\_Port FIP snooping (**VN2VN Snooping**)

#### *Verifying That IGMP Snooping Is Disabled on the FCoE VLAN*

**Purpose** Verify that IGMP snooping is disabled on the FCoE VLAN on all four switches.

**Action** List the IGMP snooping protocol information on each of the four switches using the **show configuration protocols igmp-snooping** command:

```
user@switch> show configuration protocols igmp-snooping
vlan fcoe_vlan {
    disable;
}
```

**Meaning** The **show configuration protocols igmp-snooping** command lists the IGMP snooping configuration for the VLANs configured on the switch. The command output shows that IGMP snooping is disabled on the FCoE VLAN (**fcoe\_vlan**).

- Related Documentation**
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
  - [Configuring Link Aggregation on page 2593](#)
  - [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Understanding Multichassis Link Aggregation on page 2411](#)
  - [Understanding MC-LAGs on an FCoE Transit Switch on page 5555](#)

### **Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)**

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your network uses a different priority than 3 for FCoE traffic, you need to configure lossless FCoE transport on that priority. This example shows you how to

configure lossless FCoE transport on a converged Ethernet network that uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic instead of using priority 3.

- [Requirements on page 6020](#)
- [Overview on page 6020](#)
- [Configuration on page 6022](#)
- [Verification on page 6024](#)

---

## Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

---

## Overview

Although FCoE traffic typically uses IEEE 802.1p priority 3 on converged Ethernet networks, some networks use a different priority for FCoE traffic. Regardless of the priority used, FCoE traffic must receive lossless treatment. Supporting lossless behavior for FCoE traffic when your network does not use priority 3 requires configuring:

- A lossless forwarding class for FCoE traffic.
- A behavior aggregate (BA) classifier to map the FCoE forwarding class to the appropriate IEEE 802.1p priority.
- A congestion notification profile (CNP) to enable PFC on the FCoE code point at the interface ingress and to configure flow control on the interface egress. Flow control on the interface egress enables the interface to respond to PFC messages received from the connected peer and pause the correct IEEE 802.1p priority on the correct output queue.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

---

- A DCBX application and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priority. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that

the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

### Topology

This example shows how to configure one lossless FCoE traffic class, map it to a priority other than priority 3, and configure flow control to ensure lossless behavior on the interfaces. This example uses two Ethernet interfaces, xe-0/0/25 and xe-0/0/26. The interfaces connect to a converged Ethernet network that uses IEEE 802.1p priority 5 (code point 101) for FCoE traffic.

The configuration on the two interfaces is the same. Both interfaces use the same explicitly configured lossless FCoE forwarding class and the same ingress classifier. Both interfaces enable PFC on priority 5 and enable flow control on the same output queue (which is mapped to the lossless FCoE forwarding class).

Table 542 on page 6021 shows the configuration components for this example.

**Table 542: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3**

Component	Settings
Hardware	QFX3500 switch
Forwarding class	Name— <b>fcoe1</b>  Queue mapping—queue 5  Packet drop attribute— <b>no-loss</b>  <b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.
BA classifier	Name— <b>fcoe_p5</b>  FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b> .

**Table 542: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3 (continued)**

Component	Settings
PFC configuration (CNPs)	CNP name— <b>fcoe_p5_cnp</b>  Input CNP code point— <b>101</b>  MRU— <b>2240</b> bytes  Cable length— <b>100</b> meters  Output CNP code point— <b>101</b>  Output CNP flow control queue— <b>5</b>  <b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.
DCBX application mapping	Application name— <b>fcoe_p5_app</b>  Application EtherType— <b>0x8906</b>  Application map name— <b>fcoe_p5_app_map</b>  Application map code points— <b>101</b>  <b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

### Configuration

#### CLI Quick Configuration

To quickly configure a lossless FCoE forwarding class that uses a different priority than IEEE 802.1p priority 3 for FCoE traffic on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 100
```



```

set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map

```

### *Configuring A Lossless FCoE Forwarding Class On IEEE 802.1p Priority 5*

#### **Step-by-Step Procedure**

To configure a lossless forwarding class for FCoE traffic on IEEE 802.1p priority 5 (code point 101), classify FCoE traffic into the lossless forwarding class, configure a congestion notification profile to enable PFC on the FCoE priority and output queue, and configure DCBX application protocol TLV exchange for traffic on the FCoE priority:

1. Configure the lossless forwarding class (named **fcoe1** and mapped to output queue **5**) for FCoE traffic on IEEE 802.1p priority 5:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```

2. Configure the ingress classifier (**fcoe\_p5**). The classifier maps the FCoE priority (code point 101) to the lossless FCoE forwarding class **fcoe1**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101

```

3. Apply the classifier to interfaces **xe-0/0/25** and **xe-0/0/26**:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5

```

4. Configure the CNP. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5

```

5. Apply the CNP to the interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp

```

6. Configure the DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priority 5 instead of on the default priority 3:

```

[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906

```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priority:

```
[edit]
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101
```

8. Apply the application map to the Ethernet interfaces so that DCBX exchanges FCoE application TLVs on the correct code point:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map
```

### Verification

To verify the configuration and proper operation of the lossless forwarding class and IEEE 802.1p priority, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6024](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6025](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6025](#)
- [Verifying the Interface Configuration on page 6026](#)
- [Verifying the DCBX Application Configuration on page 6026](#)
- [Verifying the DCBX Application Map Configuration on page 6026](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6027](#)

#### Verifying the Forwarding Class Configuration

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

#### *Verifying the Behavior Aggregate Classifier Configuration*

**Purpose** Verify that the classifier maps the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifier configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_p5, Code point type: ieee-802.1, Index: 63065
  Code point      Forwarding class      Loss priority
  101             fcoe1                     low
```

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe\_p5** maps code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**, and all other priorities to the **best-effort** forwarding class with a packet loss priority of **high**.

#### *Verifying the PFC Flow Control Configuration (CNP)*

**Purpose** Verify that PFC is enabled on the correct input priority and that flow control is configured on the correct output queue in the CNP.

**Action** Display the congestion notification profile using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_p5_cnp, Index: 12137
Type: Input
Cable Length: 100 m
  Priority      PFC      MRU
  000          Disabled
  001          Disabled
  010          Disabled
  011          Disabled
  100          Disabled
  101          Enabled    2240
  110          Disabled
  111          Disabled
Type: Output
  Priority      Flow-Control-Queues
  101
  5
```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the configured CNPs.

The **fcoe\_p5\_cnp** CNP input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

### *Verifying the Interface Configuration*

**Purpose** Verify that the correct classifier and congestion notification profile are configured on the interfaces.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/25** and **show configuration class-of-service interfaces xe-0/0/26**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/25
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p5;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/26
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p5;
    }
}
```

**Meaning** Both the **show configuration class-of-service interfaces xe-0/0/25** command and the **show configuration class-of-service interfaces xe-0/0/26** command show that the congestion notification profile **fcoe\_p5\_cnp** is configured on each interface, and that the IEEE 802.1p classifier associated with each interface is **fcoe\_p5**.

### *Verifying the DCBX Application Configuration*

**Purpose** Verify that the DCBX application for FCoE is configured.

**Action** List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application fcoe_p5_app {
    ether-type 0x8906;
```

**Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe\_p5\_app** is configured with an EtherType of **0x8906**.

### *Verifying the DCBX Application Map Configuration*

**Purpose** Verify that the application map is configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
fcoe_p5_app_map {
    application fcoe_p5_app code-points 101;
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map **fcoe\_p5\_app\_map** consists of the application named **fcoe\_p5\_app**, which is mapped to IEEE 802.1p code point 101.

### *Verifying the DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application map is applied to the correct interfaces.

**Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/25.0 {
    application-map fcoe_p5_app_map;
}
interface xe-0/0/26.0 {
    application-map fcoe_p5_app_map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/25.0** and **xe-0/0/26.0** use application map **fcoe\_p5\_app\_map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)
  - [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
  - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
  - [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
  - [Example: Configuring Unicast Classifiers on page 6066](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
  - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your converged Ethernet network uses more than one priority for FCoE traffic, you need to configure lossless transport for each FCoE priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses both priority 3 (IEEE 802.1p code point 011) and priority 5 (IEEE 802.1p code point 101) for FCoE traffic.

- [Requirements on page 6028](#)
- [Overview on page 6028](#)
- [Configuration on page 6031](#)
- [Verification on page 6032](#)

---

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

---

### Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, a converged Ethernet network might include two separate FCoE networks that use different priorities to identify traffic. Interfaces that carry traffic for both FCoE networks need to support lossless FCoE transport on both priorities.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class).
- A behavior aggregate (BA) classifier to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities).
- A congestion notification profile (CNP) to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifier, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

### Topology

This example shows how to configure two lossless FCoE traffic classes on an interface, map them to two different priorities, and configure flow control to ensure lossless behavior. This example uses two Ethernet interfaces, xe-0/0/20 and xe-0/0/21, that are connected to the converged Ethernet network. Both interfaces transport FCoE traffic on priorities 3 (011) and 5 (101), and must support lossless transport of that traffic.

Table 543 on page 6029 shows the configuration components for this example.

**Table 543: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology**

Component	Settings
Hardware	QFX3500 switch
Forwarding classes	<p>Name—<b>fcoe1</b>  Queue mapping—queue 5  Packet drop attribute—<b>no-loss</b></p> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—<b>fcoe</b>  This is the default lossless FCoE forwarding class, so no configuration required. The <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b>.</p>

**Table 543: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (*continued*)**

Component	Settings
BA classifier	<p>Name—<b>fcoe_classifier</b></p> <p>FCoE priority mapping for forwarding class <b>fcoe</b>—mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>.</p> <p>FCoE priority mapping for forwarding class <b>fcoe1</b>—mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>.</p>
PFC configuration (CNP)	<p>CNP name—<b>fcoe_cnp</b></p> <p>Input CNP code points—<b>011</b> and <b>101</b></p> <p>MRU—2240 bytes</p> <p>Cable length—100 meters</p> <p>Output CNP code points—<b>011</b> and <b>101</b></p> <p>Output CNP flow control queues—<b>3</b> and <b>5</b></p> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for PFC pause unless they are included in the explicitly configured output CNP. In this example, because the explicit output CNP overwrites the default output CNP, we must explicitly configure flow control on queue 3.</p>
DCBX application mapping	<p>Application name—<b>fcoe_app</b></p> <p>Application EtherType—<b>0x8906</b></p> <p>Application map name—<b>fcoe_app_map</b></p> <p>Application map code points—<b>011</b> and <b>101</b></p> <p><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>
Interfaces	<p>Interfaces <b>xe-0/0/20</b> and <b>xe-0/0/21</b> use the same configuration:</p> <ul style="list-style-type: none"> <li>Classifier—<b>fcoe_classifier</b></li> <li>CNP—<b>fcoe_cnp</b></li> <li>DCBX application map—<b>fcoe_app_map</b></li> </ul>



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.



## Configuration

### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch interface, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1p fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1p fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1p fcoe_classifier
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1p fcoe_classifier
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1p code-point 011 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1p code-point 101 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1p code-point 011
pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1p code-point 101
pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps fcoe_app_map application fcoe_app code-points [011 101]
set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

### Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic on the same interface, classify FCoE traffic into the forwarding classes, configure CNPs to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```



**NOTE:** This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier. The classifier maps the FCoE priorities (IEEE 802.1p code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1p fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1p fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101
```

3. Apply the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier
```

4. Configure the CNP. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points 011 and 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 101 pfc flow-control-queue 5
```

5. Apply the CNP to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
```

6. Configure a DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on both of the IEEE 802.1p priorities used for FCoE transport:

```
[edit]
user@switch# set applications application fcoe_app ether-type 0x8906
```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priorities:

```
[edit]
user@switch# set policy-options application-maps fcoe_app_map application fcoe_app code-points [011 101]
```

8. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

---

## Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6033](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6033](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6034](#)
- [Verifying the Interface Configuration on page 6034](#)
- [Verifying the DCBX Application Configuration on page 6035](#)

- [Verifying the DCBX Application Map Configuration on page 6035](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6035](#)

### *Verifying the Forwarding Class Configuration*

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

### *Verifying the Behavior Aggregate Classifier Configuration*

**Purpose** Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifiers using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe\_classifier, Code point type: ieee-802.1p, Index: 10964

Code point	Forwarding class	Loss priority
011	fcoe	low
101	fcoe1	low

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe\_classifier** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

### *Verifying the PFC Flow Control Configuration (CNP)*

**Purpose** Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities.

**Action** List the CNPs using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_cnp, Index: 46504
Type: Input
Cable Length: 100 m
  Priority   PFC      MRU
  000       Disabled
  001       Disabled
  010       Disabled
  011       Enabled   2240
  100       Disabled
  101       Enabled   2240
  110       Disabled
  111       Disabled
Type: Output
  Priority   Flow-Control-Queues
  011
  101
  3
  5
```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the CNP.

The CNP **fcoe\_cnp** input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the interface cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

### *Verifying the Interface Configuration*

**Purpose** Verify that the classifier and congestion notification profile are configured on the interfaces. Both interfaces should show the same configuration.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_classifier;
  }
}

user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_cnp;
unit 0 {
```

```

        classifiers {
            ieee-802.1 fcoe_classifier;
        }
    }

```

**Meaning** The **show configuration class-of-service interfaces xe-0/0/20** command shows that the congestion notification profile **fcoe\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_classifier**.

The **show configuration class-of-service interfaces xe-0/0/21** command shows that the congestion notification profile **fcoe\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_classifier**.

#### *Verifying the DCBX Application Configuration*

**Purpose** Verify that the DCBX application for FCoE is configured.

**Action** List the DCBX applications by using the configuration mode command **show applications**:

```

user@switch# show applications
application fcoe_app {
    ether-type 0x8906;
}

```

**Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe\_app** is configured with an EtherType of **0x8906**.

#### *Verifying the DCBX Application Map Configuration*

**Purpose** Verify that the application map is configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```

user@switch# show policy-options application-maps
fcoe_app_map {
    application fcoe_app code-points [011 101];
}

```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map **fcoe\_app\_map** consists of the application named **fcoe\_app**, which is mapped to IEEE 802.1p code points **011** and **101** (priorities 3 and 5, respectively).

#### *Verifying the DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application map is applied to the interfaces.

**Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```

user@switch# show protocols dcbx

```

```
interface xe-0/0/20.0 {  
    application-map fcoe_app_map;  
}  
interface xe-0/0/21.0 {  
    application-map fcoe_app_map;  
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/20.0** and **xe-0/0/21.0** use application map **fcoe\_app\_map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)
  - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
  - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
  - [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
  - [Example: Configuring Unicast Classifiers on page 6066](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
  - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic if it uses different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and map them to two different priorities on an FCoE transit switch.

- [Requirements on page 6036](#)
- [Overview on page 6037](#)
- [Configuration on page 6041](#)
- [Verification on page 6044](#)

---

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

## Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, when the switch acts as a transit switch, it could be connected to two QFX3500 switches in FCoE-FC gateway mode. Each of the gateway switches could connect a set of FCoE clients to a different SAN, and each set of FCoE clients could use a different priority for FCoE traffic to avoid fate sharing and maintain separation of the two FCoE networks. In this case, you need to configure two forwarding classes for FCoE traffic, each mapped to a different output queue and a different priority.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- Behavior aggregate (BA) classifiers to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

## Topology

This example shows how to configure two lossless FCoE traffic classes, map them to two different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses three Ethernet interfaces, xe-0/0/20, xe-0/0/21, and xe-0/0/22:

- Interface xe-0/0/20 connects to an FCoE-FC gateway that connects to Fibre Channel (FC) SAN 1. FCoE traffic to and from FC SAN 1 uses the default **fcoe** forwarding class and the default mapping to priority 3 (IEEE 802.1p code point 011) and output queue 3.
- Interface xe-0/0/21 connects to another FCoE-FC gateway that connects to Fibre Channel (FC) SAN 2. FCoE traffic to and from FC SAN-2 uses an explicitly configured FCoE forwarding class that is mapped to priority 5 (code point 101) and output queue 5.
- Interface xe-0/0/22 connects to FCoE devices on the converged Ethernet network and handles traffic destined for FC SAN 1 and FC SAN 2. Interface xe-0/0/22 must properly handle lossless FCoE traffic of both priorities (both FCoE forwarding classes), including pausing the traffic on ingress or egress as required.

Figure 215 on page 6038 shows the topology for this example, and Table 544 on page 6038 shows the configuration components for this example.

Figure 215: Topology of the Two Lossless FCoE Priorities Example

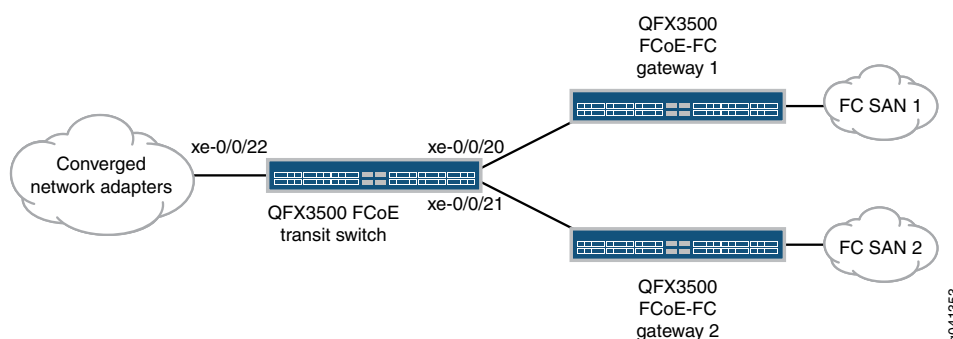


Table 544: Components of the Two Lossless FCoE Priorities Configuration Topology

Component	Settings
Hardware	QFX3500 switch
Forwarding classes	<p>Name—<b>fcoe1</b>  Queue mapping—queue 5  Packet drop attribute—<b>no-loss</b></p> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—<b>fcoe</b>  This is the default lossless FCoE forwarding class, so no configuration required. The <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b></p>



**Table 544: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)**

Component	Settings
BA classifiers	<p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/20 classifier: Name—<b>fcoe_p3</b> FCoE priority mapping—Forwarding class <b>fcoe</b> mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/21 classifier: Name—<b>fcoe_p5</b> FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/22 classifier: Name—<b>fcoe_p3_p5</b> FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> and a packet loss priority of <b>low</b>, and forwarding class <b>fcoe</b> mapped to code point <b>011</b> and a packet loss priority of <b>low</b>.</li> </ul>

Table 544: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)

Component	Settings
PFC configuration (CNPs)	<p>Each interface requires a different CNP because each interface handles a different subset of FCoE traffic and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/20 CNP:  CNP name—<b>fcoe_p3_cnp</b>  Input CNP code point—<b>011</b>  MRU—2240 bytes  Cable length—100 meters    <b>NOTE:</b> Because interface xe-0/0/20 uses the default FCoE configuration, output queue 3 is paused by default and you do not need to configure the output stanza of the CNP.</li> <li>Interface xe-0/0/21 CNP:  CNP name—<b>fcoe_p5_cnp</b>  Input CNP code point—<b>101</b>  MRU—2240 bytes  Cable length—150 meters  Output CNP code point—<b>101</b>  Output CNP flow control queue—<b>5</b></li> <li>Interface xe-0/0/22 CNP:  CNP name—<b>fcoe_p3_p5_cnp</b>  Input CNP code points—<b>011</b> and <b>101</b>  MRU—2240 bytes (both priorities)  Cable length—100 meters  Output CNP code points—<b>011</b> (for queue 3) and <b>101</b> (for queue 5)  Output CNP flow control queues—<b>3</b> for priority 3 (code point 011) and <b>5</b> for priority 5 (code point 101)</li> </ul> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>

**Table 544: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)**

Component	Settings
DCBX application mapping	<p>Interface xe-0/0/20 does not need an application map because DCBX exchanges application protocol TLVs only on the default FCoE priority (priority 3).</p> <p>Interface xe-0/0/21 requires an application map that enables DCBX application protocol TLV exchange on priority 5 (code point 101) for FCoE traffic. Interface xe-0/0/22 requires an application map that enables DCBX application protocol TLV exchange both on priority 3 (code point 011) and on priority 5 (code point 101) for FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/21 DCBX application mapping: Application name—<b>fcoe_p5_app</b> Application ether-type—<b>0x8906</b> Application map name—<b>fcoe_p5_app_map</b> Application map code points—<b>101</b></li> <li>Interface xe-0/0/22 DCBX application mapping: Application name—<b>fcoe_all_app</b> Application ether-type—<b>0x8906</b> Application map name—<b>fcoe_all_app_map</b> Application map code points—<b>011</b> and <b>101</b></li> </ul> <p><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

### Configuration

#### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
```

```

set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set applications application fcoe_all_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set policy-options application-maps fcoe_all_app_map application fcoe_all_app code-points [011
101]
set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

### Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic on different interfaces, classify FCoE traffic into the forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```



**NOTE:** This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier (**fcoe\_p3**) for interface **xe-0/0/20**. The classifier maps the FCoE priority (IEEE 802.1p code point **011**) to lossless FCoE forwarding class **fcoe**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low code-points
011

```

3. Configure the ingress classifier (**fcoe\_p5**) for interface **xe-0/0/21**. The classifier maps the FCoE priority (IEEE 802.1p code point **101**) to lossless FCoE forwarding class **fcoe1**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101

```

4. Configure the ingress classifier (**fcoe\_p3\_p5**) for interface **xe-0/0/22**. The classifier maps the two FCoE priorities (IEEE 802.1p code points **011** and **101**) to the two lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

5. Apply each classifier to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
```

6. Configure the CNP input stanza for interface **xe-0/0/20** to enable PFC on the FCoE priority (IEEE 802.1p code point **011**), set the MRU value (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queue 3 is paused by default on priority 3, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100
```

7. Configure the CNP for interface **xe-0/0/21**. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point **101**), sets the MRU value (2240 bytes), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

8. Configure the CNP for interface **xe-0/0/22**. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points **011** and **101**), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5
```

9. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
```

```

user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp

```

10. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/21. Interface xe-0/0/21 uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on priority 5 on interface xe-0/0/21. Configure an application named **fcoe\_p5\_app** for FCoE traffic (EtherType **0x8906**) and configure an application map named **fcoe\_p5\_app\_map** to map the application to code point 101:

```

[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101

```



**NOTE:** Interface xe-0/0/20 uses the default FCoE configuration (priority 3). DCBX exchanges protocol TLVs for the FCoE application by default, so you do not need to configure DCBX explicitly on interface xe-0/0/20.

11. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/22. Interface xe-0/0/22 uses both priority 3 (IEEE 802.1p code point 011) and priority 5 for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on both priority 3 and priority 5. Configure an application named **fcoe\_all\_app** for FCoE traffic (EtherType **0x8906**) and configure an application map named **fcoe\_all\_app\_map** to map the application to code points 011 and 101:

```

[edit]
user@switch# set applications application fcoe_all_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_all_app_map application
fcoe_all_app code-points [011 101]

```

12. Apply the application maps to the interfaces xe-0/0/21 and xe-0/0/22 so that DCBX exchanges FCoE application TLVs on the correct code points on each interface:

```

[edit]
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

## Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6045](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6045](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6046](#)
- [Verifying the Interface Configuration on page 6048](#)
- [Verifying the DCBX Application Configuration on page 6048](#)
- [Verifying the DCBX Application Map Configuration on page 6049](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6049](#)

**Verifying the Forwarding Class Configuration**

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

**Verifying the Behavior Aggregate Classifier Configuration**

**Purpose** Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe\_p3, Code point type: ieee-802.1, Index: 13913

Code point	Forwarding class	Loss priority
011	fcoe	low

Classifier: fcoe\_p5, Code point type: ieee-802.1, Index: 63065

Code point	Forwarding class	Loss priority
101	fcoe1	low

Classifier: fcoe\_p3\_p5, Code point type: ieee-802.1, Index: 10964

Code point	Forwarding class	Loss priority
011	fcoe	low
101	fcoe1	low

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are three classifiers, **fcoe\_p3**, **fcoe\_p5**, and **fcoe\_p3\_p5**.

Classifier **fcoe\_p3** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**.

Classifier **fcoe\_p5** maps code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe\_p3\_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

### ***Verifying the PFC Flow Control Configuration (CNP)***

**Purpose** Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	0
001	1
010	2
011	3
100	4
101	5
110	6
111	7

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240



```

100      Disabled
101      Enabled      2240
110      Disabled
111      Disabled
Type: Output
Priority  Flow-Control-Queues
011
3
101
5

Name: fcoe_p5_cnp, Index: 12133
Type: Input
Cable Length: 150 m
Priority  PFC      MRU
000      Disabled
001      Disabled
010      Disabled
011      Disabled
100      Disabled
101      Enabled   2240
110      Disabled
111      Disabled
Type: Output
Priority  Flow-Control-Queues
101
5

```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the three CNPs. For CNP **fcoe\_p3\_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows the default mapping of priorities to output queues.



**NOTE:** By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queue 3 responds to pause messages from the connected peer on interfaces that use CNP **fcoe\_p3\_cnp**, because the input stanza enables PFC priority 3 only.

For CNP **fcoe\_p3\_p5\_cnp**, the input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

For CNP **fcoe\_p5\_cnp**, the input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **150** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

### *Verifying the Interface Configuration*

**Purpose** Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20**, **show configuration class-of-service interfaces xe-0/0/21**, and **show configuration class-of-service interfaces xe-0/0/22**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
  congestion-notification-profile fcoe_p3_cnp;
  unit 0 {
    classifiers {
      ieee-802.1p fcoe_p3;
    }
  }
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
  congestion-notification-profile fcoe_p5_cnp;
  unit 0 {
    classifiers {
      ieee-802.1p fcoe_p5;
    }
  }
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/22
  congestion-notification-profile fcoe_p3_p5_cnp;
  unit 0 {
    classifiers {
      ieee-802.1p fcoe_p3_p5;
    }
  }
```

**Meaning** The **show configuration class-of-service interfaces xe-0/0/20** command shows that the congestion notification profile **fcoe\_p3\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3**.

The **show configuration class-of-service interfaces xe-0/0/21** command shows that the congestion notification profile **fcoe\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p5**.

The **show configuration class-of-service interfaces xe-0/0/22** command shows that the congestion notification profile **fcoe\_p3\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3\_p5**.

### *Verifying the DCBX Application Configuration*

**Purpose** Verify that the two DCBX applications for FCoE are configured.

**Action** List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
  application fcoe_all_app {
    ether-type 0x8906;
```

```
application fcoe_p5_app {
    ether-type 0x8906;
```

**Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe\_all\_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic) and that the application **fcoe\_p5\_app** is also configured with an EtherType of **0x8906**.

### *Verifying the DCBX Application Map Configuration*

**Purpose** Verify that the application maps are configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
fcoe_all_app_map {
    application fcoe_all_app code-points [011 101];
}
fcoe_p5_app_map {
    application fcoe_p5_app code-points 101;
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there are two application maps.

Application map **fcoe\_all\_app\_map** consists of the application named **fcoe\_all\_app** mapped to IEEE 802.1p code points **011** (priority 3) and **101** (priority 5).

Application map **fcoe\_p5\_app\_map** consists of the application named **fcoe\_p5\_app** mapped to IEEE 802.1p code point **101** (priority 5).

### *Verifying the DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application maps are applied to the correct interfaces.

**Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/21.0 {
    application-map fcoe_p5_app_map;
}
interface xe-0/0/22.0 {
    application-map fcoe_all_app_map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interface **xe-0/0/21.0** uses application map **fcoe\_p5\_app\_map** and interface **xe-0/0/22.0** uses application map **fcoe\_all\_app\_map**.



**NOTE:** Because interface xe-0/0/20 uses the default lossless FCoE configuration, you do not configure application mapping to interface xe-0/0/20. The default configuration automatically exchanges application protocol TLVs for the default FCoE configuration on priority 3 (IEEE 802.1p code point 011).

**Related Documentation**

- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic on different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, and map the forwarding classes to three different priorities. (The converged Ethernet network includes two FCoE networks, each of which uses a different priority to identify FCoE traffic, and an iSCSI network.)

- [Requirements on page 6050](#)
- [Overview on page 6051](#)
- [Configuration on page 6055](#)
- [Verification on page 6059](#)

### Requirements

---

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode

- Junos OS Release 12.3 or later for the QFX Series

## Overview

Some converged Ethernet networks support FCoE on more than one IEEE 802.1p priority and also require supporting other lossless traffic classes. Interfaces that carry multiple lossless forwarding classes need to support lossless behavior for the priorities mapped to those forwarding classes. To support the two FCoE forwarding classes and the iSCSI forwarding class used in this example, you need to configure:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- A lossless forwarding class for iSCSI traffic
- Behavior aggregate (BA) classifiers to map the lossless forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE and iSCSI code points at the interface ingress, and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the FCoE and iSCSI traffic on the configured lossless priorities. By default, DCBX is enabled on all Ethernet interfaces for FCoE, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.


## Topology

This example shows how to configure two lossless FCoE traffic classes and one lossless iSCSI traffic class, map them to three different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses four Ethernet interfaces, xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34:

- Interface xe-0/0/31 handles FCoE traffic on priority 3 (IEEE 802.1p code point 011) and iSCSI traffic on priority 4 (code point 100).
- Interface xe-0/0/32 handles FCoE traffic on priority 5 (code point 101) and iSCSI traffic on priority 4.
- Interface xe-0/0/33 handles FCoE traffic on priority 3 and priority 5.
- Interface xe-0/0/34 handles iSCSI traffic on priority 4.

Figure 216 on page 6052 shows the topology for this example, and Table 545 on page 6052 shows the configuration components for this example.

**Figure 216: Topology of the Lossless FCoE and iSCSI Priorities Example**



**Table 545: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology**

Component	Settings
Hardware	QFX3500 switch

**Table 545: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

Component	Settings
Forwarding classes	<p>This example uses one explicitly configured lossless FCoE forwarding class, the default lossless FCoE forwarding class, and one explicitly configured iSCSI forwarding class.</p> <ul style="list-style-type: none"> <li>iSCSI forwarding class: Name—<b>iscsi</b> Queue mapping—queue 4 Packet drop attribute—<b>no-loss</b></li> <li>FCoE forwarding class (explicitly configured): Name—<b>fcoe1</b> Queue mapping—queue 5 Packet drop attribute—<b>no-loss</b></li> </ul> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <ul style="list-style-type: none"> <li>FCoE forwarding class (default) Name—<b>fcoe</b> The default <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b>.</li> </ul>
BA classifiers	<p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/31 classifier: Name—<b>fcoe_p3_iscsi</b> FCoE priority mapping—Forwarding class <b>fcoe</b> mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>. iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/32 classifier: Name—<b>fcoe_p5_iscsi</b> FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>. iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/33 classifier: Name—<b>fcoe_p3_p5</b> FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (priority 5) and a packet loss priority of <b>low</b>, and forwarding class <b>fcoe</b> mapped to code point <b>011</b> and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/34 classifier: Name—<b>iscsi_classifier</b> iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> </ul>

**Table 545: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (continued)**

Component	Settings
PFC configuration (CNPs)	<p>Each interface requires a different CNP because each interface handles a different subset of FCoE and iSCSI traffic, and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/31 CNP:  CNP name—<b>fcoe_p3_cnp</b>  Input CNP code points—<b>011</b> and <b>100</b>  MRU—2240 bytes for code point <b>011</b>, default value (2500 bytes) for code point <b>100</b>  Cable length—100 meters    <b>NOTE:</b> On interface xe-0/0/31, the FCoE forwarding class is mapped to queue 3 and priority 3 (code point 011), and the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Therefore, interface xe-0/0/31 does not require an output CNP configuration because queue 3 and queue 4 are enabled for PFC flow control by default on code points 011 and 100, respectively.</li> <li>Interface xe-0/0/32 CNP:  CNP name—<b>fcoe_p5_cnp</b>  Input CNP code points—<b>100</b> and <b>101</b>  MRU—Default value (2500 bytes) for code point <b>100</b>, <b>2240</b> bytes for code point <b>101</b>  Cable length—150 meters  Output CNP code points—<b>100</b> and <b>101</b>  Output CNP flow control queues—<b>4</b> and <b>5</b></li> <li>Interface xe-0/0/33 CNP:  CNP name—<b>fcoe_p3_p5_cnp</b>  Input CNP code points—<b>011</b> and <b>101</b>  MRU—<b>2240</b> bytes (both priorities)  Cable length—100 meters  Output CNP code points—<b>011</b> and <b>101</b>  Output CNP flow control queues—<b>3</b> and <b>5</b></li> <li>Interface xe-0/0/34 CNP:  CNP name—<b>iscsi_cnp</b>  Input CNP code point—<b>100</b>  MRU—<b>2500</b> bytes (default value)  Cable length—100 meters    <b>NOTE:</b> On interface xe-0/0/34, the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Interface xe-0/0/34 does not require an output CNP configuration because queue 4 is enabled for PFC flow control by default on code point 100.</li> </ul> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>



**Table 545: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

Component	Settings
DCBX application mapping	<p>This example requires configuring applications for FCoE and iSCSI, including them in the same application map, and applying the application map to all four interfaces.</p> <p>Application map name—<b>dcbx_iscsi_fcoe_app_map</b></p> <ul style="list-style-type: none"> <li>FCoE application name—<b>fcoe_app</b> Application ether-type—<b>0x8906</b> Application map code points—<b>011</b> and <b>101</b></li> <li>iSCSI application name—<b>iscsi_app</b> Application protocol type—<b>tcp</b> Application destination port—<b>3260</b> Application map code point—<b>100</b></li> </ul> <p><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

### Configuration

#### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes and one lossless iSCSI forwarding class and map them to different priorities, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class iscsi queue-num 4 no-loss
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100
set class-of-service interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
set class-of-service interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
set class-of-service interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5set
class-of-service interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier
```

```

set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile iscsi_cnp input ieee-802.1 code-point 100 pfc
set class-of-service congestion-notification-profile iscsi_cnp input cable-length 100
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
set class-of-service interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
set applications application iscsi_app protocol tcp destination-port 3260
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps dcbx_iscsi_fcoe_app_map application iscsi_app code-points
100
set policy-options application-maps dcbx_iscsi_fcoe_app_map application fcoe_app code-points
[011 101]
set protocols dcbx interface xe-0/0/31 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/32 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/33 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/34 application-map dcbx_iscsi_fcoe_app_map

```

### Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, classify the traffic into the three forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding classes **iscsi** for iSCSI traffic and **fcoe1** for FCoE traffic (this example uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class) and map them to output queues:

```

[edit class-of-service]
user@switch# set forwarding-classes class iscsi queue-num 4 no-loss
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```

2. Configure the ingress classifier (**fcoe\_p3\_iscsi**) for interface **xe-0/0/31**. The classifier maps the FCoE priority (code point **011**) to lossless FCoE forwarding class **fcoe** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]

```

```

user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100

```

3. Configure the ingress classifier (**fcoe\_p5\_iscsi**) for interface **xe-0/0/32**. The classifier maps the FCoE priority (code point **101**) to lossless FCoE forwarding class **fcoe1** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101

```

4. Configure the ingress classifier (**fcoe\_p3\_p5**) for interface **xe-0/0/33**. The classifier maps the two FCoE priorities (code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101

```

5. Configure the ingress classifier (**iscsi\_classifier**) for interface **xe-0/0/34**. The classifier maps the iSCSI priority (code point **101**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100

```

6. Apply each classifier to the appropriate interface:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier

```

7. Configure the CNP input stanza for interface **xe-0/0/31** to enable PFC on the FCoE and iSCSI priorities that the interface handles (code points **011** and **100**), set the MRU value for the FCoE traffic (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queues 3 and 4 are paused by default on priorities 3 and 4, respectively, and we are not explicitly configuring output queue flow control for any other queues.

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100

```

8. Configure the CNP for interface **xe-0/0/32**. The input stanza enables PFC on the FCoE priority (code point **101**), sets the MRU value for FCoE traffic (2240 bytes), enables PFC on the iSCSI priority (code point **100**), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority and on output queue 4 on the iSCSI priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

9. Configure the CNP for interface xe-0/0/33. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points 011 and 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5
```

10. Configure the CNP input stanza for interface xe-0/0/34 to enable PFC on the iSCSI priority (code point 100) and set the cable length value (100 meters). No output stanza is needed because queue 4 is paused by default on priority 4, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile iscsi_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile iscsi_cnp input cable-length 100
```

11. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
```

12. Configure the DCBX applications for FCoE and iSCSI to map to the interfaces so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priorities used for FCoE and iSCSI traffic:

```
[edit]
user@switch# set applications application fcoe_app ether-type 0x8906
user@switch# set applications application iscsi_app protocol tcp destination-port 3260
```

13. Configure a DCBX application map to map the FCoE and iSCSI applications to the correct priorities:

```
[edit]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
fcoe_app code-points [011 101]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
iscsi_app code-points 100
```

14. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/31 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/32 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/33 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/34 application-map
dcbx_iscsi_fcoe_app_map
```

### Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 6059](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 6060](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 6060](#)
- [Verifying the Interface Configuration on page 6063](#)
- [Verifying the DCBX Application Configuration on page 6064](#)
- [Verifying the DCBX Application Map Configuration on page 6064](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 6065](#)

#### *Verifying the Forwarding Class Configuration*

**Purpose** Verify that the lossless forwarding classes **iscsi** and **fcoe1** have been created and that the default lossless forwarding class **fcoe** is still enabled for lossless transport.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
iscsi	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **iscsi** and **fcoe1** forwarding classes are configured on output queues 4 and 5, respectively, with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default **fcoe** forwarding class, it remains in its default state (lossless configuration).

### *Verifying the Behavior Aggregate Classifier Configuration*

**Purpose** Verify that the four classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_p3_iscsi, Code point type: ieee-802.1, Index: 13915
  Code point  Forwarding class  Loss priority
  011         fcoe             low
  100         iscsi            low
```

```
Classifier: fcoe_p5_iscsi, Code point type: ieee-802.1, Index: 62035
  Code point  Forwarding class  Loss priority
  100         iscsi            low
  101         fcoe1            low
```

```
Classifier: fcoe_p3_p5, Code point type: ieee-802.1, Index: 17774
  Code point  Forwarding class  Loss priority
  011         fcoe             low
  101         fcoe1            low
```

```
Classifier: iscsi_classifier, Code point type: ieee-802.1, Index: 31635
  Code point  Forwarding class  Loss priority
  100         iscsi            low
```

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are four classifiers, **fcoe\_p3\_iscsi**, **fcoe\_p5\_iscsi**, **fcoe\_p3\_p5**, and **iscsi\_classifier**.

Classifier **fcoe\_p3\_iscsi** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and code point **100** (priority 4) to explicitly configured lossless forwarding class **iscsi**.

Classifier **fcoe\_p5\_iscsi** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**, and code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe\_p3\_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **iscsi\_classifier** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**.

### *Verifying the PFC Flow Control Configuration (CNP)*

**Purpose** Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Enabled	9216
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	
	0
001	
	1
010	
	2
011	
	3
100	
	4
101	
	5
110	
	6
111	
	7

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
011	
	3
101	
	5

```
Name: fcoe_p5_cnp, Index: 12133
```

```
Type: Input
```

```
Cable Length: 150 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	

```
011      Disabled
100      Enabled      9216
101      Enabled      2240
110      Disabled
111      Disabled
Type: Output
100
      4
101
      5

Name: iscsi_cnp, Index: 19342
Type: Input
Cable Length: 100 m
Priority  PFC      MRU
000      Disabled
001      Disabled
010      Disabled
011      Disabled
100      Enabled      9216
101      Disabled
110      Disabled
111      Disabled
Type: Output
Priority  Flow-Control-Queues
000
      0
001
      1
010
      2
011
      3
100
      4
101
      5
110
      6
111
      7
```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the four CNPs.

For CNP **fcoe\_p3\_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3) with an MRU of **2240** bytes, and cable length of **100** meters. The input stanza also shows that PFC is enabled on code point **100** (priority 4) with the default MRU value of **9216** bytes. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.





**NOTE:** By default, only queues 3 and 4 are enabled respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queues 3 and 4 respond to pause messages from the connected peer on interfaces that use CNP `fcoe_p3_cnp` because the input stanza enables PFC only on priorities 3 and 4.

For CNP `fcoe_p3_p5_cnp`, the input stanza shows that PFC is enabled on code points 011 and 101 (priority 5), the MRU is 2240 bytes on both priorities, and the cable length is 100 meters. The CNP output stanza shows that output flow control is configured on queues 3 and 5 for code points 011 and 101, respectively.

For CNP `fcoe_p5_cnp`, the input stanza shows that PFC is enabled on code points 100 and 101. The MRU for code point 101 (FCoE traffic) is 2240 bytes and the MRU for code point 100 is 9216. The interface cable length is 150 meters. The CNP output stanza shows that output flow control is configured on queue 4 for code point 100 and on queue 5 for code point 101.

For CNP `iscsi_cnp`, the input stanza shows that PFC is enabled on code point 100, the MRU value is 9216 bytes, and the interface cable length is 100 meters. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.

### *Verifying the Interface Configuration*

**Purpose** Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

**Action** List the ingress interfaces using the operational mode commands `show configuration class-of-service interfaces xe-0/0/31`, `show configuration class-of-service interfaces xe-0/0/32`, `show configuration class-of-service interfaces xe-0/0/33`, and `show configuration class-of-service interfaces xe-0/0/34`:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3_iscsi;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p5_iscsi;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_p3_p5;
  }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile iscsi_cnp;
unit 0 {
  classifiers {
    ieee-802.1 iscsi_classifier;
  }
}
```

**Meaning** The **show configuration class-of-service interfaces xe-0/0/31** command shows that the congestion notification profile **fcoe\_p3\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3\_iscsi**.

The **show configuration class-of-service interfaces xe-0/0/32** command shows that the congestion notification profile **fcoe\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p5\_iscsi**.

The **show configuration class-of-service interfaces xe-0/0/33** command shows that the congestion notification profile **fcoe\_p3\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3\_p5**.

The **show configuration class-of-service interfaces xe-0/0/34** command shows that the congestion notification profile **iscsi\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **iscsi\_classifier**.

#### *Verifying the DCBX Application Configuration*

**Purpose** Verify that the DCBX applications for FCoE and iSCSI are configured.

**Action** List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application iscsi_app {
  protocol tcp;
  destination-port 3260;
}
application fcoe_app {
  ether-type 0x8906;
```

**Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **iscsi\_app** is configured with a protocol value of **tcp** and a destination port value of **3260**, and that the application **fcoe\_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic).

#### *Verifying the DCBX Application Map Configuration*

**Purpose** Verify that the application map is configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
    application iscsi_app code-points 100;
    application fcoe_app code-points [011 101];
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there is one application map named **dcbx-iscsi-fcoe\_app\_map**. It consists of the application **iscsi\_app** mapped to code point **100** and the application **fcoe\_app** mapped to code points **011** and **101**.

### *Verifying the DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application maps are applied to the correct interfaces.

**Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/31.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/32.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/33.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/34.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that all four interfaces use the application map **dcbx-iscsi-fcoe-app-map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
  - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
  - [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)
  - [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
  - [Example: Configuring Unicast Classifiers on page 6066](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
  - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)

- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Example: Configuring Unicast Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. You apply classifiers to ingress interfaces.

- [Requirements on page 6066](#)
- [Overview on page 6066](#)
- [Configuring Unicast Classifiers on page 6067](#)
- [Verification on page 6067](#)

### Requirements

---

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

### Overview

---

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP or DSCP IPv6) value, IEEE 802.1p value, or MPLS EXP value. (EXP classifiers can be applied only to **family mpls** interfaces.)
- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.



**NOTE:** You must assign unicast traffic and multidestination (multicast, broadcast, and destination lookup fail) traffic to different classifiers. One classifier cannot include both unicast and multidestination forwarding classes. A unicast classifier can include only forwarding classes for unicast traffic.

---

This example describes how to configure a BA classifier called **ba-ucast-classifier** as the default IEEE 802.1 map and apply it to ingress interface **xe-0/0/10**. The BA classifier assigns loss priorities, as shown in [Table 546 on page 6067](#), to incoming packets in the four forwarding classes.

You can use the same procedure to set multifield classifiers (except that you use firewall filter rules).

**Table 546: ba-ucast-classifier Loss Priority Assignments**

Unicast Forwarding Class	For CoS Traffic Type	ba-ucast-classifier Loss Priority to IEEE 802.1p Code Point Mapping	Packet Drop Attribute
<b>be</b>	Best-effort traffic	Low loss priority code point: <b>000</b>	drop
<b>fcoe</b>	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic	Low loss priority code point: <b>011</b>	no-loss
<b>no-loss</b>	Guaranteed delivery for TCP traffic	Low loss priority code point: <b>100</b>	no-loss
<b>nc</b>	Network-control traffic	Low loss priority code point: <b>110</b>	drop

### Configuring Unicast Classifiers

To configure a unicast IEEE 802.1 BA classifier named **ba-ucast-classifier** as the default IEEE 802.1 map:

- Associate code point **000** with forwarding class **be** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier import default forwarding-class be
loss-priority low code-points 000
```
- Associate code point **011** with forwarding class **fcoe** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class fcoe loss-priority low
code-points 011
```
- Associate code point **100** with forwarding class **no-loss** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class no-loss loss-priority low
code-points 100
```
- Associate code point **110** with forwarding class **nc** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class nc loss-priority low
code-points 110
```
- Apply the unicast classifier to ingress interface **xe-0/0/10**:  

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/10 unit 0 classifiers ieee-802.1 ba-ucast-classifier
```

### Verification

To verify the unicast classifier configuration, perform these tasks:

- [Verifying the Unicast Classifier Configuration on page 6068](#)
- [Verifying the Ingress Interface Configuration on page 6068](#)

### *Verifying the Unicast Classifier Configuration*

**Purpose** Verify that you configured the unicast classifier with the correct forwarding classes, loss priorities, and code points.

**Action** List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-ucast-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-ucast-classifier
    forwarding-class be {
        loss-priority low code-points 000;
    }
    forwarding-class fcoe {
        loss-priority low code-points 011;
    }
    forwarding-class no-loss {
        loss-priority low code-points 100;
    }
    forwarding-class nc
        loss-priority low code-points 110;
    }
```

### *Verifying the Ingress Interface Configuration*

**Purpose** Verify that the unicast classifier **ba-ucast-classifier** is attached to ingress interface **xe-0/0/10**.

**Action** List the ingress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/10**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/10
congestion-notification-profile fcoe-cnp;
unit 0 {
    classifiers {
        ieee-802.1 ba-ucast-classifier;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
  - [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
  - [Configuring a Global MPLS EXP Classifier on page 4479](#)
  - [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
  - [Monitoring CoS Classifiers on page 6289](#)
  - [Understanding CoS Classifiers on page 5810](#)
  - [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)

## Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class.

- [Requirements on page 6069](#)
- [Overview on page 6069](#)
- [Configuring Multidestination Classifiers on page 6070](#)
- [Verification on page 6070](#)

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

### Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value or IEEE 802.1p value.



**NOTE:** DSCP IPv6 multidestination classifiers are not supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces.



**NOTE:** You must assign unicast traffic and multicast traffic to different classifiers. One classifier cannot include both unicast and multicast forwarding classes. A multidestination classifier can include only forwarding classes for multicast traffic.

The following example describes how to configure a BA classifier called **ba-mcast-classifier**, which is applied to all of the switch interfaces. The BA classifier assigns loss priorities, as shown in [Table 547 on page 6070](#), to incoming packets in the multdestination forwarding class.

You can use the same procedure to set multifield classifiers (except that you use firewall filter rules).

**Table 547: BA-mcast-classifier Loss Priority Assignments**

Multicast Forwarding Class	For CoS Traffic Type	ba-mcast-classifier Assignment
<b>mcast</b>	Best-effort multicast traffic	Low loss priority code point: 000

### Configuring Multidestination Classifiers

To configure a multicast IEEE 802.1 BA classifier named **ba-mcast-classifier**:

- Associate code point **000** with forwarding class **mcast** and loss priority **low**:  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-mcast-classifier forwarding-class mcast loss-priority low
code-points 000
```
- Configure the classifier as a multdestination classifier:  

```
[edit class-of-service]
user@switch# set multi-destination classifiers ieee-802.1 ba-mcast-classifier
```

### Verification

To verify the multdestination classifier configuration, perform these tasks:

- [Verifying the IEEE 802.1 Multidestination Classifier on page 6070](#)
- [Verifying the Multidestination Classifier Configuration on page 6070](#)

#### Verifying the IEEE 802.1 Multidestination Classifier

**Purpose** Verify that the classifier **ba-mcast-classifier** is configured as the IEEE 802.1 multdestination classifier:

**Action** Verify the results of the classifier configuration using the operational mode command **show configuration class-of-service multi-destination classifiers ieee-802.1**:

```
user@switch> show configuration class-of-service multi-destination classifiers ieee-802.1
ba-mcast-classifier;
```

#### Verifying the Multidestination Classifier Configuration

**Purpose** Verify that you configured the multdestination classifier with the correct forwarding classes, loss priorities, and code points.

**Action** List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier
```



```
forwarding-class mcast {
    loss-priority low code-points 000;
}
```

#### Related Documentation

- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 6162](#)
- [Monitoring CoS Classifiers on page 6289](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)

## Example: Configuring WRED Drop Profiles

You can configure an interpolated weighted random early detection (WRED) profile to control packet drop characteristics for different traffic loss priorities.



**NOTE:** You cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Also, do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

- [Requirements on page 6071](#)
- [Overview on page 6071](#)
- [Configuring a Drop Profile on page 6073](#)
- [Verification on page 6073](#)

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

### Overview

You associate a WRED profile with a loss priority in a scheduler. When you attach the scheduler to a forwarding class (queue), you apply the interpolated drop profile to traffic of the specified loss priority in that queue. *Interpolated* means that the switch creates a smooth drop curve from a drop start point to a drop end point, with a maximum drop rate that is reached at the drop end point:

- Drop start point—Percentage of average queue fill level when the WRED algorithm starts to drop packets. Before the drop start point, no packets are scheduled to drop.
- Drop end point—Average queue fill level at which all subsequently arriving packets are dropped. When the queue fill levels falls below the drop end point, packets begin to

be forwarded again. (At the drop end point, the packet drop probability becomes 100 percent.)

- Maximum drop rate—Drop probability when the average queue fill level reaches the drop end point.


You set the drop start point and the drop end point by specifying two queue fill level percentage values. The first value is the drop start point and the second value is the drop end point.

You set the maximum drop rate by specifying two drop probability percentage values. The first value is always zero (0), which is the minimum drop rate, the probability of dropping a packet at the drop start point. The second value is the maximum drop rate at the drop end point.

The drop rate is zero until the queue fill level reaches the drop start point. As the queue continues to fill, packets drop in smooth linear curve until the queue reaches the drop end point, when packets drop at the maximum drop rate. If the queue fills beyond the drop end point, all packets that match the drop profile are dropped.

Figure 217 on page 6072 shows the graph for a drop profile with a drop start point of 30 percent, a drop end point of 50 percent, and a maximum drop rate of 80 percent.

**Figure 217: WRED Drop Profile Packet Drop Example**



The graph shows that when the queue fill level is less than 30 percent, the packet drop rate is zero. When the queue fill level reaches 30 percent, packets begin to drop. As the queue fills, a higher percentage of packets drop. When the queue fill level reaches 50 percent, the packet drop rate has climbed to 80 percent. When the queue fill level exceeds 50 percent, all packets drop.

This example describes how to configure the drop profile shown in Figure 217 on page 6072. The drop profile will have:

- The name **be-dp1**
- 30 percent for the drop start point (first **fill-level** setting)

- 50 percent for the drop end point (second **fill-level** setting)
- 0 percent for the minimum drop rate (first **drop-probability** setting)
- 80 percent for the maximum drop rate (second **drop-probability** setting)

You apply a drop profile by configuring a drop profile map that maps the drop profile to a packet loss priority and associates the drop profile and packet loss priority with a scheduler. When you associate the scheduler with a forwarding class (queue), the switch applies the drop profile to the packets in the forwarding class that have a matching packet loss priority.

### Configuring a Drop Profile

1. Set the drop start point at **30** percent, the drop end point at **50** percent, the minimum drop rate at **0** percent, and the maximum drop rate at **80** percent for the drop profile **be-dp1**:

```
[edit class-of-service]
user@switch# set drop-profile be-dp1 interpolate fill-level 30 fill-level 50 drop-probability
0 drop-probability 80
```

### Verification

#### Verifying the Drop Profile Configuration

**Purpose** Verify that you configured the drop profile **be-dp1** with the correct drop start and end points and with the correct drop rates.

**Action** Verify the results of the drop profile configuration using the operational mode command **show configuration class-of-service drop-profiles be-dp1**:

```
user@switch> show configuration class-of-service drop-profiles be-dp1
interpolate {
    fill-level [ 30 50 ];
    drop-probability [ 0 80 ];
}
```

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Configuring CoS WRED Drop Profiles on page 6163](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)

### Example: Configuring Drop Profile Maps

A drop-profile map associates a WRED profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class,

the drop profile map associated with the scheduler applies the specified WRED profile to traffic in the forwarding class that matches the specified loss priority.

- [Requirements on page 6074](#)
- [Overview on page 6074](#)
- [Configuring a Drop Profile Map on page 6074](#)
- [Verification on page 6075](#)

---

## Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

---

## Overview

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

The following example describes how to configure a drop profile map for a scheduler named **mylan** that includes:

- A drop profile called **lp-profile** for low-priority traffic
- A drop profile called **mh-profile** for medium-high priority traffic
- A drop profile called **h-profile** for high-priority traffic

You apply the drop profiles in the drop profile map to a forwarding class by associating the scheduler **mylan** with a forwarding class in a scheduler map.

---

## Configuring a Drop Profile Map

1. Configure the drop profile for low-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority low protocol any
drop-profile lp-profile
```

2. Configure the drop profile for medium-high priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority medium-high protocol
any drop-profile mh-profile
```

3. Configure the drop profile for high-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority high protocol any
drop-profile h-profile
```

## Verification

### Verifying the Drop Profile Map Configuration

**Purpose** Verify that you configured the drop profile map for the scheduler **mylan** with the correct loss priorities and drop profiles.

**Action** Verify the results of the drop profile map configuration using the operational mode command **show configuration class-of-service schedulers mylan**:

```
user@switch> show configuration class-of-service schedulers mylan
transmit-rate 3g;
shaping-rate percent 100;
priority low;
drop-profile-map loss-priority low protocol any drop-profile lp-profile;
drop-profile-map loss-priority medium-high protocol any drop-profile mh-profile;
drop-profile-map loss-priority high protocol any drop-profile h-profile;
```



**NOTE:** This example does not include configuring scheduler bandwidth and priority. This information (transmit rate, shaping rate, and priority) is shown for completeness.

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Configuring CoS Drop Profile Maps on page 6164](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)

## Example: Configuring Forwarding Classes

Forwarding classes allow you to group packets for transmission. You assign packets to unicast or multideestination (multicast, broadcast, and destination lookup fail) output queues based on forwarding classes.

- [Requirements on page 6075](#)
- [Overview on page 6076](#)
- [Configuring Forwarding Classes on page 6077](#)
- [Verification on page 6078](#)

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

---

The switch supports a total of 12 forwarding classes. In order to forward traffic, you must map (assign) the forwarding classes to unicast or multideestination output queues. The switch has 12 queues. Queues 0 through 7 are for unicast traffic, and queues 8 through 11 are for multideestination traffic. The switch supports up to two lossless forwarding classes.

By default, four categories of unicast forwarding classes and one multideestination forwarding class are defined. You can define the remaining seven forwarding classes and configure them as unicast or multideestination by mapping them to unicast or multideestination queues. The type of queue, unicast or multideestination, determines the type of forwarding class.

The four default unicast forwarding classes are:

- **be**—Best-effort traffic
- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **nc**—Network control traffic

The default multideestination forwarding class is:

- **mcast**—Multideestination traffic

Map forwarding classes to queues using the **class** statement, which enables you to configure up to 12 forwarding classes. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue. The statement format is:

```
[edit class-of-service forwarding-classes]
user@switch# class class-name queue-num queue-number;
```



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

---



**NOTE:** Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

```
[edit class-of-service forwarding-classes]
user@switch# queue queue-number class-name
```

The `queue` statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the `queue` statement to map forwarding classes to queues, edit the configuration to replace the `queue` statement with the `class` statement.



**NOTE:** Hierarchical scheduling controls output queue forwarding. When you define a forwarding class that will carry traffic on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

Table 548 on page 6077 shows the configuration forwarding-class-to-queue mapping for this example:

**Table 548: Forwarding-Class-to-Queue Example Configuration**

Forwarding Class	Queue
best-effort	0
nc	7
mcast	8

### Configuring Forwarding Classes

To configure CoS forwarding classes, map the forwarding classes to queues:

1. Map the **best-effort** forwarding class to queue 0:

```
[edit class-of-service forwarding-classes]
user@switch# set class best-effort queue-num 0
```

2. Map the **nc** forwarding class to queue 7:

```
[edit class-of-service forwarding-classes]
```

```
user@switch# set class nc queue-num 7
```

3. Map the **mcast-be** forwarding class to queue 8:

```
[edit class-of-service forwarding-classes]  
user@switch# set class mcast-be queue-num 8
```

---

### Verification

#### *Verifying the Forwarding-Class-to-Queue Mapping*

**Purpose** Verify the forwarding-class-to-queue mapping. (The system shows only the explicitly configured forwarding classes; it does not show default forwarding classes such as **fcoe** and **no-loss**.)

**Action** Verify the results of the forwarding class configuration using the operational mode command **show configuration class-of-service forwarding-classes**:

```
user@switch> show configuration class-of-service forwarding-classes  
class best-effort queue-num 0;  
class network-control queue-num 7;  
class mcast queue-num 8;
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Defining CoS Forwarding Classes on page 6164](#)
  - [Monitoring CoS Forwarding Classes on page 6290](#)
  - [Overview of CoS Changes Introduced in Junos OS Release 11.3](#)
  - [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
  - [Understanding CoS Forwarding Classes on page 5830](#)

### Example: Configuring Forwarding Class Sets

A forwarding class set (fc-set) is a priority group for enhanced transmission selection (ETS) traffic control. Each fc-set consists of one or more forwarding classes (output queues).

ETS enables you to configure link resources (bandwidth and bandwidth sharing characteristics) for a priority group, and then allocate the priority group's resources among the forwarding classes that belong to the priority group. This is called two-tier, or hierarchical, scheduling. Traffic control profiles control the scheduling for the priority group, and schedulers control the scheduling for individual forwarding classes.

- [Requirements on page 6079](#)
- [Overview on page 6079](#)
- [Configuring Forwarding Class Sets on page 6080](#)
- [Verification on page 6080](#)



## Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

You can configure up to three unicast fc-sets and one multicast fc-set. A common way to configure unicast priority groups is to configure separate fc-sets for local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and then assign the appropriate forwarding classes to each fc-set.



**NOTE:** If you configure strict-high priority queues, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queues. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues. The multidestination fc-set cannot contain strict-high priority queues.

To apply ETS, you map one or more fc-sets to a physical egress port. You can map up to three forwarding class sets to each port. When you map an fc-set to a port, the port uses hierarchical scheduling to allocate port resources to the priority group (fc-set) and to allocate the priority group resources to the queues (forwarding classes) that belong to the priority group.

This example describes how to:

- Configure three fc-sets called **lan-pg**, **san-pg**, and **hpc-pg**.
- Assign forwarding classes to each of the fc-sets.
- Apply the fc-sets and their output traffic control profiles to an egress interface.

This example does not describe how to configure the forwarding classes assigned to the fc-sets or how to configure traffic control profiles. [Table 549 on page 6079](#) shows the configuration components for this example:

**Table 549: Components of the Forwarding Class Sets Configuration Example**

Component	Settings
Hardware	QFX3500 switch
LAN traffic priority group	Forwarding class set: <b>lan-pg</b> Forwarding classes: <b>best-effort-1</b> , <b>best-effort-2</b>
SAN traffic priority group	Forwarding class set: <b>san-pg</b> Forwarding classes: <b>fcoe</b> , <b>fcoe-2</b>

Table 549: Components of the Forwarding Class Sets Configuration Example (*continued*)

Component	Settings
HPC traffic priority group	Forwarding class set: <b>hpc-pg</b> Forwarding classes: <b>nc, high-perf</b>
Egress interface	<b>xe-0/0/7</b>

### Configuring Forwarding Class Sets

1. Define the **lan-pg** priority group (fc-set) and assign to it the forwarding classes **best-effort-1** and **best-effort-2**:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets lan-pg class best-effort-1
user@switch# set forwarding-class-sets lan-pg class best-effort-2
```
2. Define the **san-pg** priority group and assign to it the forwarding classes **fcoe** and **fcoe-2**:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets san-pg class fcoe
user@switch# set forwarding-class-sets san-pg class fcoe-2
```
3. Define the **hpc-pg** priority group and assign to it the forwarding classes **nc** and **high-perf**:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets hpc-pg class nc
user@switch# set forwarding-class-sets hpc-pg class high-perf
```
4. Map the three forwarding class sets to an interface (the output traffic control profiles associated with the forwarding class sets determine the class of service scheduling for the priority groups):  

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile lan-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set san-pg
output-traffic-control-profile san-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```

### Verification

To verify the priority group configuration, perform these tasks:

- [Verifying Forwarding Class Set Membership on page 6080](#)
- [Verifying the Egress Interface Configuration on page 6081](#)

#### Verifying Forwarding Class Set Membership

**Purpose** Verify that you configured the **lan-pg**, **san-pg**, and **hpc-pg** priority groups with the correct forwarding classes.

**Action** List the forwarding class set member configuration using the operational mode command **show configuration class-of-service forwarding-class-sets**:

```
user@switch> show configuration class-of-service forwarding-class-sets
lan-pg {
    class best-effort-1;
    class best-effort-2;
}
san-pg {
    class fcoe;
    class fcoe-2;
}
hpc-pg {
    class high-perf;
    class nc;
}
```

#### *Verifying the Egress Interface Configuration*

**Purpose** Verify that egress interface **xe-0/0/7** is associated with the **lan-pg**, **san-pg**, and **hpc-pg** priority groups and with the correct output traffic control profiles.

**Action** Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    lan-pg {
        output-traffic-control-profile lan-tcp;
    }
    san-pg {
        output-traffic-control-profile san-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Example: Configuring Queue Schedulers on page 6081](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
  - [Defining CoS Forwarding Class Sets on page 6166](#)
  - [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5835](#)

## Example: Configuring Queue Schedulers

Schedulers define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit congestion notification (ECN) is enabled on the queue, and the WRED packet drop profiles associated with the queue.

- [Requirements on page 6082](#)
- [Overview on page 6082](#)

- [Configuring a CoS Scheduler on page 6085](#)
- [Verification on page 6085](#)

---

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

---

## Overview

Scheduler parameters define the following characteristics for the queues mapped to the scheduler:

- **transmit-rate**—Minimum bandwidth, also known as the committed information rate (CIR). Each queue mapped to the scheduler receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue. You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.



**NOTE:** The **transmit-rate** setting works only if you also configure the **guaranteed-rate** in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the **guaranteed-rate**, the **transmit-rate** does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate. If you configure transmit rates whose sum exceeds the forwarding class set guaranteed rate, the commit check fails, and the system rejects the configuration.



**NOTE:** Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR). Each queue receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



**NOTE:** Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:

- **low**—The scheduler has low priority.
- **strict-high**—The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **drop-profile-map**—Mapping of a drop profile to a loss priority and protocol to apply WRED to the scheduler.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—Enables ECN on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.



**NOTE:** Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



**NOTE:** Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Scheduler maps associate schedulers with forwarding classes (queues). After defining schedulers and mapping them to queues in a scheduler map, to configure hardware queue scheduling (port scheduling) you:

1. Associate a scheduler map with a traffic control profile (a traffic control profile schedules resources for a group of forwarding classes, called a *forwarding class set* or *priority group*).
2. Attach a forwarding class and a traffic control profile to an interface.

You can associate up to four user-defined scheduler maps with forwarding class sets.

This process configures the hardware queues, packet schedulers, and WRED characteristics that operate according to the scheduler mapping. The traffic control profile uses the scheduler CoS properties to determine the resources that should be allocated to the individual output queues from the total resources available to the priority group.

Table 550 on page 6084 shows the configuration components for this example.

**Table 550: Components of the Queue Scheduler Configuration Example**

Component	Settings
Hardware	QFX3500 switch
Scheduler	Name: <b>be-sched</b> Transmit rate: <b>20%</b> Shaping rate: <b>40%</b> Buffer size: <b>20%</b> Priority: <b>low</b> Drop profile: <b>be-dp</b> ECN: <b>disable</b> (default)
Scheduler map	Name: <b>be-map</b> Forwarding class to associate with the <b>be-sched</b> scheduler: <b>best-effort</b>
Traffic control profile	Name: <b>be-tcp</b>  <b>NOTE:</b> This topic does not describe how to define a traffic control profile.
Forwarding class set	Name: <b>lan-pg</b>

## Configuring a CoS Scheduler

To configure a CoS scheduler using the CLI:

1. Create a scheduler (**be-sched**) with a minimum guaranteed bandwidth of 2 Gbps, a maximum bandwidth of 4 Gbps, low priority, and map it to the drop profile **be-dp**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate percent 20
user@switch# set be-sched shaping-rate percent 40
user@switch# set be-sched buffer-size percent 20
user@switch# set be-sched priority low
user@switch# set be-sched drop-profile-map loss-priority low protocol any drop-profile
be-dp
```



**NOTE:** Because ECN is disabled by default, no ECN configuration is shown.

2. Configure a scheduler map (**be-map**) that associates the scheduler (**be-sched**) with the forwarding class (**best-effort**):

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

3. Associate the scheduler map **be-map** with a traffic control profile (**be-tcp**):

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

4. Associate the traffic control profile **be-tcp** with a forwarding class set (**lan-pg**) and a 10-Gigabit Ethernet interface (**xe-0/0/7**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile be-tcp
```

5. Alternatively, you can assign the scheduler map (**be-map**) to all the 10-Gigabit Ethernet interfaces using wildcards (**xe-\***):

```
[edit class-of-service interfaces]
user@switch# set xe-* forwarding-class-set lan-pg output-traffic-control-profile be-tcp
```

## Verification

To verify that the queue scheduler has been created and is mapped to the correct interfaces, perform these tasks:

- [Verifying the Scheduler Configuration on page 6085](#)
- [Verifying the Scheduler Map Configuration on page 6086](#)
- [Verifying That the Scheduler Is Associated with the Interface on page 6086](#)

### Verifying the Scheduler Configuration

**Purpose** Verify that the queue scheduler **be-sched** has been created with a minimum guaranteed bandwidth of 2 Gbps, a maximum bandwidth of 4 Gbps, the priority set to **low**, and the drop profile **be-dp**.

**Action** Display the scheduler using the operational mode command **show configuration class-of-service schedulers be-sched**:

```
user@switch> show configuration class-of-service schedulers be-sched
transmit-rate percent 20;
shaping-rate percent 40;
buffer-size percent 20
priority low;
drop-profile-map loss-priority low protocol any drop-profile be-dp;
```

#### *Verifying the Scheduler Map Configuration*

**Purpose** Verify that the scheduler map **be-map** has been created and associates the forwarding class **best-effort** with the scheduler **be-sched**, and also that the scheduler map is attached to the traffic control profile **be-tcp**.

**Action** Display the scheduler map using the operational mode command **show configuration class-of-service scheduler-maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Display the traffic control profile to verify that the scheduler map **be-map** is attached using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```



**NOTE:** This topic does not describe how to configure a traffic control profile or its allocation of port bandwidth. Using a traffic control profile to configure the port resource allocation to the priority group is necessary to implement hierarchical scheduling.

---

#### *Verifying That the Scheduler Is Associated with the Interface*

**Purpose** Verify that the forwarding class set (**lan-pg**) and the traffic control profile (**be-tcp**) that are associated with the queue scheduler are attached to the interface **xe-0/0/7**.

**Action** List the interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    lan-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)



- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring ECN on page 6090](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Monitoring CoS Scheduler Maps on page 6293](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\)](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

## Example: Configuring Queue Scheduling Priority

You can configure the bandwidth scheduling priority of individual queues by specifying the priority in a scheduler, and then using a scheduler map to associate the scheduler with a queue.

- [Requirements on page 6087](#)
- [Overview on page 6087](#)
- [Configuring Queue Scheduling Priority on page 6088](#)
- [Verification on page 6089](#)

### Requirements

---

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

### Overview

---

Queues can have one of two bandwidth priorities:

- **strict-high**—You can configure only one queue as a strict-high or high-priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.



**NOTE:** If you configure strict-high priority queues, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queues. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues. The multidestination fc-set cannot contain strict-high priority queues.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **low**—Low priority. Traffic with this priority is serviced after any queue that has a **strict-high** priority.

Table 551 on page 6088 shows the configuration components for this example.

This example describes how to set the queue priority for two forwarding classes (queues) named **fcoe** and **no-loss**. Both queues have a priority of **low**. The scheduler for the **fcoe** queue is named **fcoe-sched** and the scheduler for the **no-loss** queue is named **nl-sched**. One scheduler map, **schedmap1**, associates the schedulers to the queues.

**Table 551: Components of the Queue Scheduler Priority Configuration Example**

Component	Settings
Hardware	QFX3500 switch
Schedulers	<b>fcoe-sched</b> for FCoE traffic <b>nl-sched</b> for no-loss traffic
Priority	<b>low</b> for FCoE traffic <b>low</b> for no-loss traffic
Scheduler map	<b>schedmap1</b> : FCoE mapping: scheduler <b>fcoe-sched</b> to forwarding class <b>fcoe</b> No-loss mapping: scheduler <b>nl-sched</b> to forwarding class <b>no-loss</b>

### Configuring Queue Scheduling Priority

To configure queue priority using the CLI:

1. Create the FCoE scheduler with **low** priority:  

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low
```
2. Create the no-loss scheduler with **low** priority:  

```
[edit class-of-service]
```

```
user@switch# set schedulers nl-sched priority low
```

3. Associate the schedulers with the desired queues in the scheduler map:

```
[edit class-of-service]
```

```
user@switch# set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
```

```
user@switch# set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

## Verification

To verify that you configured the queue scheduling priority for bandwidth and mapped the schedulers to the correct forwarding classes, perform these tasks:

- [Verifying the Queue Scheduling Priority on page 6089](#)
- [Verifying the Scheduler-to-Forwarding-Class Mapping on page 6089](#)

### Verifying the Queue Scheduling Priority

**Purpose** Verify that you configured the queue schedulers **fcoe-sched** and **nl-sched** with **low** queue scheduling priority.

**Action** Display the **fcoe-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers fcoe-sched priority**:

```
user@switch> show configuration class-of-service schedulers fcoe-sched priority
priority low;
```

Display the **nl-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers nl-sched priority**:

```
user@switch> show configuration class-of-service schedulers nl-sched priority
priority low;
```

### Verifying the Scheduler-to-Forwarding-Class Mapping

**Purpose** Verify that you configured the scheduler map **schedmap1** to map scheduler **fcoe-sched** to forwarding class **fcoe** and schedule **nl-sched** to forwarding class **no-loss**.

**Action** Display the scheduler map **schedmap1** using the operational mode command **show configuration class-of-service scheduler-maps schedmap1**:

```
user@switch> show configuration class-of-service scheduler-maps schedmap1
forwarding-class fcoe scheduler fcoe-sched;
forwarding-class no-loss scheduler nl-sched;
```

## Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Defining CoS Queue Scheduling Priority on page 6171](#)
- [Monitoring CoS Scheduler Maps on page 6293](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)

## Example: Configuring ECN

This example shows how to enable explicit congestion notification (ECN) on an output queue.

- [Requirements on page 6090](#)
- [Overview on page 6090](#)
- [Configuration on page 6091](#)
- [Verification on page 6093](#)

---

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX Series switch
- Junos OS Release 13.2X51-D25 or later for the QFX Series

---

### Overview

ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.

A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.

ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.

ECN is disabled by default. You can enable ECN on best-effort traffic. ECN should not be enabled on lossless traffic queues, which uses priority-based flow control (PFC) for congestion notification, and ECN should not be enabled on strict-high priority traffic queues.

To enable ECN on an output queue, you not only need to enable ECN in the queue scheduler, you also need to:

- Configure a WRED packet drop profile.
- Configure a queue scheduler that includes the WRED drop profile and enables ECN. (This example shows only ECN and drop profile configuration; you can also configure bandwidth, priority, and buffer settings in a scheduler.)
- Map the queue scheduler to a forwarding class (output queue).

- Add the forwarding class to a forwarding class set (priority group, for hierarchical scheduling).
- Associate the queue scheduler with a traffic control profile (priority group scheduler for hierarchical scheduling).
- Apply the traffic control profile and the forwarding class set to an interface. On that interface, the specified output queue (the queue mapped to the forwarding class) uses the scheduler that is mapped to the traffic control profile, which enables ECN on the queue and applies the WRED drop profile to the queue.

Table 552 on page 6091 shows the configuration components for this example.

**Table 552: Components of the ECN Configuration Example**

Component	Settings
Hardware	QFX Series switch
Drop profile	Name: <b>be-dp</b> Drop start fill level: <b>30</b> percent Drop end fill level: <b>75</b> percent Drop probability at drop start (minimum drop rate): <b>0</b> percent Drop probability at drop end (maximum drop rate): <b>80</b> percent
Scheduler	Name: <b>be-sched</b> ECN: enabled Drop profile: <b>be-dp</b> Transmit rate: <b>25%</b> Shaping rate: <b>50%</b> Buffer size: <b>25%</b> Priority: <b>low</b>
Scheduler map	Name: <b>be-map</b> Forwarding class: <b>best-effort</b> Scheduler: <b>be-sched</b>  <b>NOTE:</b> By default, the <b>best-effort</b> forwarding class is mapped to output queue <b>0</b> .
Forwarding class set	Name: <b>be-pg</b> Forwarding class: <b>best-effort</b> (queue 0)
Traffic control profile	Name: <b>be-tcp</b> Scheduler map: <b>be-map</b>
Interface	Name: <b>xe-0/0/20</b> Forwarding class set: <b>be-pg</b> (Output) traffic control profile: <b>be-tcp</b>

### Configuration

**CLI Quick Configuration** To quickly configure the drop profile, scheduler with ECN enabled, and to map the scheduler to an output queue on an interface, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network

configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service]
set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability 0 drop-probability 80
set schedulers be-sched explicit-congestion-notification
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
set schedulers be-sched transmit-rate percent 25
set schedulers be-sched shaping-rate percent 50
set schedulers be-sched buffer-size percent 25
set schedulers be-sched priority low
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set forwarding-class-sets be-pg class best-effort
set traffic-control-profiles be-tcp scheduler-map be-map
set interfaces xe-0/0/20 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

### Configuring ECN

#### Step-by-Step Procedure

To configure ECN using the CLI:

1. Configure the WRED packet drop profile **be-dp**. This example uses a drop start point of **30** percent, a drop end point of **75** percent, a minimum drop rate of **0** percent, and a maximum drop rate of **80** percent:

```
[edit class-of-service]
user@switch# set drop-profile be-dp interpolate fill-level 30 fill-level 75 drop-probability 0 drop-probability 80
```

2. Create the scheduler **be-sched** with ECN enabled and associate the drop profile **be-dp** with the scheduler:

```
[edit class-of-service]
user@switch# set schedulers be-sched explicit-congestion-notification
user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile be-dp
user@switch# set be-sched transmit-rate percent 25
user@switch# set be-sched shaping-rate percent 50
user@switch# set be-sched buffer-size percent 25
user@switch# set be-sched priority low
```

3. Map the scheduler **be-sched** to the **best-effort** forwarding class (output queue 0) using scheduler map **be-map**:

```
[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
```

4. Add the forwarding class **best-effort** to the forwarding class set **be-pg**:

```
[edit class-of-service]
user@switch# set forwarding-class-sets be-pg class best-effort
```

5. Associate the scheduler map **be-map** with the traffic control profile **be-tcp**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map
```

6. Associate the traffic control profile **be-tcp** and the forwarding class set **be-pg** with the interface on which you want to enable ECN on the best-effort queue:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

## Verification

To verify that ECN is enabled, perform the following task:

- [Verifying That ECN Is Enabled on page 6093](#)

### *Verifying That ECN Is Enabled*

**Purpose** Verify that ECN is enabled in the scheduler **be-sched** by showing the configuration for the scheduler map **be-map**.

**Action** Display the scheduler map configuration using the operational mode command **show class-of-service scheduler-map be-map**:

```
user@switch> show class-of-service scheduler-map be-map
Scheduler map: be-map, Index: 12240
```

```
Scheduler:be-sched, Forwarding class: best-effort, Index: 115
Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent,
Buffer Limit: none, Priority: low
Excess Priority: unspecified, Explicit Congestion Notification: enable
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       3312   be-dp
  Medium-high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>
```

**Meaning** The **show class-of-service scheduler-map** operational command shows the configuration of the scheduler associated with the scheduler map and the forwarding class mapped to that scheduler. The output shows that:

- The scheduler associated with the scheduler map is **be-sched**.
- The scheduler map applies to the forwarding class **best-effort** (output queue 0).
- The scheduler **be-sched** has a transmit rate of **25** percent, a queue buffer size of **25** percent, and a drop priority of **low**.
- Explicit congestion notification state is **enable**.
- The WRED drop profile used for low drop priority traffic is **be-dp**.

- Related Documentation**
- [Example: Configuring Queue Schedulers on page 6081](#)
  - [Example: Configuring WRED Drop Profiles on page 6071](#)
  - [Example: Configuring Forwarding Class Sets on page 6078](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Understanding CoS Explicit Congestion Notification on page 5926](#)

## Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

- [Requirements on page 6094](#)
- [Overview on page 6094](#)
- [Configuring a Traffic Control Profile on page 6095](#)
- [Verification on page 6096](#)

---

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

---

### Overview

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the committed information rate (CIR). Each priority group receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



**NOTE:** In order for the **transmit-rate** option (minimum bandwidth for a queue that you set using scheduler configuration) to work properly, you must configure the **guaranteed-rate** for the priority group. If a priority group does not have a guaranteed minimum bandwidth, the queues (forwarding classes) that belong to the priority group cannot have a guaranteed minimum bandwidth.



**NOTE:** Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR). Each priority group receives a maximum of the configured amount of absolute



bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



**NOTE:** Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the priority group bandwidth, not the amounts or percentages of total link bandwidth.)



**NOTE:** Because a port can have more than one priority group, when you assign resources to a priority group, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

For example, if you map three priority groups to a 10-Gigabit Ethernet port, the queues associated with all three of the priority groups share the 10-Gbps bandwidth as defined by the traffic control profiles. Therefore, the total combined guaranteed-rate value of the three priority groups should not exceed 10 Gbps. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue (forwarding class) transmit rates cannot exceed the total guaranteed-rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.

If you configure the guaranteed-rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

### Configuring a Traffic Control Profile

This example describes how to configure a traffic control profile named **san-tcp** with a scheduler map named **san-map1** and allocate to it a minimum bandwidth of 4 Gbps and a maximum bandwidth of 8 Gbps:

1. Create the traffic control profile and set the **guaranteed-rate** (minimum guaranteed bandwidth) to **4g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp guaranteed-rate 4g
```

2. Set the **shaping-rate** (maximum guaranteed bandwidth) to **8g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp shaping-rate 8g
```

3. Associate the scheduler map **san-map1** with the traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp scheduler-map san-map1
```

---

## Verification

### *Verifying the Traffic Control Profile Configuration*

**Purpose** Verify that the traffic control profile **san-tcp** has been created with a minimum guaranteed bandwidth of 4 Gbps, a maximum bandwidth of 8 Gbps, and the scheduler map **san-map1**.

**Action** List the traffic control profile using the operational mode command **show configuration class-of-service traffic-control-profiles san-tcp**:

```
user@switch> show configuration class-of-service traffic-control-profiles san-tcp
scheduler-map san-map1;
shaping-rate percent 8g;
guaranteed-rate 4g;
```

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)

## Example: Configuring Minimum Guaranteed Output Bandwidth

Scheduling the minimum guaranteed output bandwidth for a queue (forwarding class) requires configuring both tiers of the two-tier hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs.

- [Requirements on page 6096](#)
- [Overview on page 6097](#)
- [Configuring Guaranteed Minimum Bandwidth on page 6098](#)
- [Verification on page 6099](#)

---

## Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

The priority group minimum guaranteed bandwidth defines the minimum total amount of bandwidth available for all of the queues in the priority group to meet their minimum bandwidth requirements.

The **transmit-rate** setting in the scheduler configuration determines the minimum guaranteed bandwidth for an individual queue. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.

The **guaranteed-rate** setting in the traffic control profile configuration determines the minimum guaranteed bandwidth for a priority group. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



**NOTE:** You must configure both the **transmit-rate** value for the queue and the **guaranteed-rate** value for the priority group in order to set a valid minimum bandwidth guarantee for a queue. (If the priority group does not have a guaranteed minimum bandwidth, there is no guaranteed bandwidth pool from which the queue can take its guaranteed minimum bandwidth.)

The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)



**NOTE:** When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



**NOTE:** You cannot configure minimum guaranteed bandwidth on strict-high priority queues or on a priority group that contains strict-high priority queues.

This example describes how to:

- Configure a transmit rate (minimum guaranteed queue bandwidth) of 2 Gbps for queues in a scheduler named **be-sched**.
- Configure a guaranteed rate (minimum guaranteed priority group bandwidth) of 4 Gbps for a priority group in a traffic control profile named **be-tcp**.

- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the minimum guaranteed bandwidth scheduling to the egress interface **xe-0/0/7**.

Table 553 on page 6098 shows the configuration components for this example:

**Table 553: Components of the Minimum Guaranteed Output Bandwidth Configuration Example**

Component	Settings
Hardware	QFX3500 switch
Minimum guaranteed queue bandwidth	Transmit rate: <b>2g</b>
Minimum guaranteed priority group bandwidth	Guaranteed rate: <b>4g</b>
Scheduler	<b>be-sched</b>
Scheduler map	<b>be-map</b>
Traffic control profile	<b>be-tcp</b>
Forwarding class set (priority group)	<b>be-pg</b>
Queue (forwarding class)	<b>best-effort</b>
Egress interface	<b>xe-0/0/7</b>

### Configuring Guaranteed Minimum Bandwidth

To configure the minimum guaranteed bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the minimum guaranteed queue bandwidth of 2 Gbps for scheduler **be-sched**:  

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate 2g
```
2. Configure the minimum guaranteed priority group bandwidth of 4 Gbps for traffic control profile **be-tcp**:  

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp guaranteed-rate 4g
```
3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:  

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

- Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

- Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

- Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

## Verification

To verify the minimum guaranteed output bandwidth configuration, perform these tasks:

- Verifying the Minimum Guaranteed Queue Bandwidth on page 6099
- Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association on page 6099
- Verifying the Scheduler Map Configuration on page 6100
- Verifying Queue (Forwarding Class) Membership in the Priority Group on page 6100
- Verifying the Egress Interface Configuration on page 6100

### *Verifying the Minimum Guaranteed Queue Bandwidth*

**Purpose** Verify that you configured the minimum guaranteed queue bandwidth as **2g** in the scheduler **be-sched**.

**Action** Display the minimum guaranteed bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched transmit-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched transmit-rate
2g;
```

### *Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association*

**Purpose** Verify that the minimum guaranteed priority group bandwidth is **4g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

**Action** Display the minimum guaranteed bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate
4g;
```

Display the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

#### *Verifying the Scheduler Map Configuration*

**Purpose** Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

**Action** Display the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service schedulers maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

#### *Verifying Queue (Forwarding Class) Membership in the Priority Group*

**Purpose** Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

**Action** Display the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

#### *Verifying the Egress Interface Configuration*

**Purpose** Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

**Action** Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Example: Configuring Queue Schedulers on page 6081](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
  - [Example: Configuring Queue Scheduling Priority on page 6087](#)
  - [Example: Configuring Forwarding Class Sets on page 6078](#)
  - [Understanding CoS Traffic Control Profiles on page 5880](#)
  - [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)

## Example: Configuring Maximum Output Bandwidth

Scheduling the maximum output bandwidth for a queue (forwarding class) requires configuring both tiers of the hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs.

- [Requirements on page 6101](#)
- [Overview on page 6101](#)
- [Configuring Maximum Bandwidth on page 6102](#)
- [Verification on page 6103](#)

### Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

### Overview

The priority group maximum bandwidth defines the maximum total amount of bandwidth available for all of the queues in the priority group.

The **shaping-rate** setting in the scheduler configuration determines the maximum bandwidth for an individual queue.

The **shaping-rate** setting in the traffic control profile configuration determines the maximum bandwidth for a priority group.



**NOTE:** When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



**NOTE:** When you set the maximum bandwidth (**shaping-rate**) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within  $\pm 20$  percent of the configured **shaping-rate** value.

This example describes how to:

- Configure a maximum rate of 4 Gbps for queues in a scheduler named **be-sched**.
- Configure a maximum rate of 6 Gbps for a priority group in a traffic control profile named **be-tcp**.

- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the bandwidth scheduling to the interface **xe-0/0/7**.

Table 554 on page 6102 shows the configuration components for this example:

**Table 554: Components of the Maximum Output Bandwidth Configuration Example**

Component	Settings
Hardware	QFX3500 switch
Maximum queue bandwidth	Shaping rate: <b>4g</b>
Maximum priority group bandwidth	Shaping rate: <b>6g</b>
Scheduler	<b>be-sched</b>
Scheduler map	<b>be-map</b>
Traffic control profile	<b>be-tcp</b>
Forwarding class set (priority group)	<b>be-pg</b>
Queue (forwarding class)	<b>best-effort</b>
Egress interface	<b>xe-0/0/7</b>

### Configuring Maximum Bandwidth

To configure the maximum bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the maximum queue bandwidth of 4 Gbps for scheduler **be-sched**:  

```
[edit class-of-service schedulers]
user@switch# set be-sched shaping-rate 4g
```
2. Configure the maximum priority group bandwidth of 6 Gbps for traffic control profile **be-tcp**:  

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp shaping-rate 6g
```
3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:  

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```
4. Associate the scheduler map with the traffic control profile:



```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

- Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

- Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

## Verification

To verify the maximum output bandwidth configuration, perform these tasks:

- [Verifying the Maximum Queue Bandwidth on page 6103](#)
- [Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association on page 6103](#)
- [Verifying the Scheduler Map Configuration on page 6104](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 6104](#)
- [Verifying the Egress Interface Configuration on page 6104](#)

### *Verifying the Maximum Queue Bandwidth*

**Purpose** Verify that you configured the maximum queue bandwidth as **4g** in the scheduler **be-sched**.

**Action** List the maximum bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched shaping-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched shaping-rate
4g;
```

### *Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association*

**Purpose** Verify that the maximum priority group bandwidth is **6g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

**Action** List the maximum bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp shaping-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp shaping-rate
4g;
```

List the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

### *Verifying the Scheduler Map Configuration*

**Purpose** Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

**Action** List the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service schedulers maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

### *Verifying Queue (Forwarding Class) Membership in the Priority Group*

**Purpose** Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

**Action** List the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

### *Verifying the Egress Interface Configuration*

**Purpose** Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

**Action** List the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

### **Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 5862](#)

## **Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic**

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly best-effort (lossy) unicast traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 6105](#)
- [Overview on page 6105](#)
- [Configuration on page 6107](#)
- [Verification on page 6108](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause

message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)

- **Lossy buffers**—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- **Lossless buffers**—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- **Lossy buffers**—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- **Multicast buffers**—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly best-effort unicast traffic, more buffer space needs to be allocated to lossy buffers, and less buffer space should be allocated to lossless buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly unicast traffic.

### **Topology**

[Table 555 on page 6106](#) shows the configuration components for this example.

**Table 555: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies**

Component	Settings
Hardware	QFX3500 switch

**Table 555: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies (*continued*)**

Component	Settings
Ingress shared buffer	<p>Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%</p> <p>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%</p> <p>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 0%</p> <p>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 95%</p>
Egress shared buffer	<p>Percentage of available egress buffer space allocated to the egress shared buffer: 100%</p> <p>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%</p> <p>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 75%</p> <p>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 20%</p>

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly best-effort unicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 0
set ingress buffer-partition lossy percent 95
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 75
set egress buffer-partition multicast percent 20
```

#### *Configuring the Global Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic*

#### Step-by-Step Procedure

To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly best-effort unicast traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:
 

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:
 

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 0
user@switch# set ingress buffer-partition lossy percent 95
```

3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 75
user@switch# set egress buffer-partition multicast percent 20
```

### Results

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
    percent 100;
    buffer-partition lossless {
        percent 5;
    }
    buffer-partition lossy {
        percent 95;
    }
    buffer-partition lossless-headroom {
        percent 0;
    }
}
egress {
    percent 100;
    buffer-partition lossless {
        percent 5;
    }
    buffer-partition lossy {
        percent 75;
    }
    buffer-partition multicast {
        percent 20;
    }
}
```

---

### Verification

Verify that the shared buffer configuration has been created properly.

#### *Verifying the Shared Buffer Configuration*

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2158.00 KB
```

```

Shared Buffer      : 7202.00 KB
Lossless          : 360.10 KB
Lossless Headroom : 0.00 KB
Lossy             : 6841.90 KB

```

Lossless Headroom Utilization:

Node Device	Total	Used	Free
0	0.00 KB	0.00 KB	0.00 KB

Egress:

```

Total Buffer      : 9360.00 KB
Dedicated Buffer  : 2704.00 KB
Shared Buffer     : 6656.00 KB
Lossless         : 332.80 KB
Multicast        : 1331.20 KB
Lossy            : 4992.00 KB

```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 360.10 KB to lossless traffic
  - No space to lossless headroom traffic
  - 6841.90 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Because the lossless headroom buffer partition is set to 0 (zero) percent, the total amount of lossless headroom buffer space is 0 KB; therefore the amount of used and free lossless headroom buffer space is also 0 KB.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.

- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 1331.20 KB to multicast traffic
  - 4992 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 and EX4600 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 and EX4600 switches.

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 6110](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

### Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly best-effort (lossy) traffic on links with Ethernet PAUSE (IEEE 802.3X) enabled. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.



Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 6111](#)
- [Overview on page 6111](#)
- [Configuration on page 6113](#)
- [Verification on page 6114](#)

---

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

---

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly best-effort traffic on links enabled for Ethernet PAUSE, more buffer space needs to be allocated to ingress dedicated port buffers, and less buffer space should be allocated to ingress shared buffers. Also, more buffer space needs to be allocated to lossless-headroom buffers, and less space to ingress lossy buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly best-effort traffic on links enabled for Ethernet PAUSE.

### Topology

Table 556 on page 6112 shows the configuration components for this example.

**Table 556: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE**

Component	Settings
Hardware	QFX3500 switch
Ingress shared buffer	Percentage of available ingress buffer space allocated to the ingress shared buffer: 70%  Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%  Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 80%  Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 15%

**Table 556: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE (continued)**

Component	Settings
Egress shared buffer	Percentage of available egress buffer space allocated to the egress shared buffer: 100%
	Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%
	Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 75%
	Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 20%

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly best-effort unicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 70
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 80
set ingress buffer-partition lossy percent 15
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 75
set egress buffer-partition multicast percent 20
```

#### *Configuring the Global Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links Enabled for Ethernet PAUSE*

#### Step-by-Step Procedure

To configure the global ingress and egress shared buffer allocations and partitions:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 70
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 80
user@switch# set ingress buffer-partition lossy percent 15
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 75
```

```
user@switch# set egress buffer-partition multicast percent 20
```

### Results

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
    percent 70;
    buffer-partition lossless {
        percent 5;
    }
    buffer-partition lossy {
        percent 15;
    }
    buffer-partition lossless-headroom {
        percent 80;
    }
}
egress {
    percent 100;
    buffer-partition lossless {
        percent 5;
    }
    buffer-partition lossy {
        percent 75;
    }
    buffer-partition multicast {
        percent 20;
    }
}
```

---

### Verification

Verify that the shared buffer configuration has been created properly.

#### *Verifying the Shared Buffer Configuration*

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 4318.60 KB
  Shared Buffer     : 5041.40 KB
    Lossless       : 252.07 KB
    Lossless Headroom : 4033.12 KB
    Lossy          : 756.21 KB

Egress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2704.00 KB
  Shared Buffer     : 6656.00 KB
    Lossless       : 332.80 KB
```

```

Multicast      : 1331.20 KB
Lossy          : 4992.00 KB

```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 4318.6 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 70 percent of the available (user-configurable) buffer space.
- With the ingress shared buffer pool configured as 70 percent of the available buffers, the total size of the ingress shared buffer pool is 5041.4 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 252.07 KB to lossless traffic
  - 4033.12 KB to lossless headroom traffic
  - 756.21 KB to lossy unicast traffic

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 1331.20 KB to multicast traffic
  - 4992 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 and EX4600 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 and EX4600 switches.

**Related Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

## Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly multicast traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 6116](#)
- [Overview on page 6117](#)
- [Configuration on page 6118](#)
- [Verification on page 6120](#)

### Requirements

---

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

---

## Overview

---

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly multicast traffic, more buffer space needs to be allocated to lossy buffers, less buffer space should be allocated to lossless buffers, and more space needs to be allocated to egress multicast buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly multicast traffic.

### Topology

[Table 557 on page 6118](#) shows the configuration components for this example.

**Table 557: Components of the Recommended Shared Buffer Configuration for Multicast Network Topologies**

Component	Settings
Hardware	QFX3500 switch
Ingress shared buffer	Percentage of available ingress buffer space allocated to the ingress shared buffer: 100% Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5% Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 0% Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 95%
Egress shared buffer	Percentage of available egress buffer space allocated to the egress shared buffer: 100% Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5% Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 20% Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 75%

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly multicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 0
set ingress buffer-partition lossy percent 95
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 20
set egress buffer-partition multicast percent 75
```



*Configuring the Global Shared Buffer Pool for Networks with Mostly Multicast Traffic*

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly multicast traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 0
user@switch# set ingress buffer-partition lossy percent 95
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 20
user@switch# set egress buffer-partition multicast percent 75
```

**Results**

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
    percent 100;
    buffer-partition lossless {
        percent 5;
    }
    buffer-partition lossy {
        percent 95;
    }
    buffer-partition lossless-headroom {
        percent 0;
    }
}
egress {
    percent 100;
    buffer-partition lossless {
        percent 5;
    }
    buffer-partition lossy {
        percent 20;
    }
    buffer-partition multicast {
        percent 75;
    }
}
```

## Verification

---

Verify that the shared buffer configuration has been created properly.

### *Verifying the Shared Buffer Configuration*

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2158.00 KB
  Shared Buffer     : 7202.00 KB
    Lossless       : 360.10 KB
    Lossless Headroom : 0.00 KB
    Lossy          : 6841.90 KB

  Lossless Headroom Utilization:
  Node Device      Total      Used      Free
  0                0.00 KB    0.00 KB    0.00 KB

Egress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2704.00 KB
  Shared Buffer     : 6656.00 KB
    Lossless       : 332.80 KB
    Multicast      : 4992.00 KB
    Lossy          : 1331.20 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 360.10 KB to lossless traffic
  - No space to lossless headroom traffic

- 6841.90 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Because the lossless headroom buffer partition is set to 0 (zero) percent, the total amount of lossless headroom buffer space is 0 KB; therefore the amount of used and free lossless headroom buffer space is also 0 KB.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 4992 KB to multicast traffic
  - 1331.20 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 and EX4600 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 and EX4600 switches.

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 6110](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

## Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly lossless traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.



**NOTE:** When we discuss lossless buffers, we mean buffers that handle traffic on which you enable priority-based flow control (PFC) to ensure lossless transport. The lossless buffers are not used for best-effort traffic on a link on which you enable Ethernet PAUSE (IEEE 802.3x).

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 6122](#)
- [Overview on page 6123](#)
- [Configuration on page 6124](#)
- [Verification on page 6126](#)

---

### Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly lossless traffic, more buffer space needs to be allocated to lossless buffers, and less buffer space should be allocated to lossy buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly lossless traffic.

### Topology

[Table 558 on page 6124](#) shows the configuration components for this example.

**Table 558: Components of the Recommended Shared Buffer Configuration for Lossless Network Topologies**

Component	Settings
Hardware	QFX3500 switch
Ingress shared buffer	Percentage of available ingress buffer space allocated to the ingress shared buffer: 100% Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 15% Percentage of ingress buffer space allocated to lossless headroom traffic (lossless headroom buffer partition): 80% Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 5%
Egress shared buffer	Percentage of available egress buffer space allocated to the egress shared buffer: 100% Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 90% Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 5% Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 5%

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly lossless traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 15
set ingress buffer-partition lossless-headroom percent 80
set ingress buffer-partition lossy percent 5
set egress percent 100
set egress buffer-partition lossless percent 90
set egress buffer-partition lossy percent 5
set egress buffer-partition multicast percent 5
```

*Configuring the Global Shared Buffer Pool for Networks with Mostly Lossless Traffic*

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly lossless traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 15
user@switch# set ingress buffer-partition lossless-headroom percent 80
user@switch# set ingress buffer-partition lossy percent 5
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 90
user@switch# set egress buffer-partition lossy percent 5
user@switch# set egress buffer-partition multicast percent 5
```

**Results**

Display the results of the configuration:

```
rroot@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
  percent 100;
  buffer-partition lossless {
    percent 15;
  }
  buffer-partition lossy {
    percent 5;
  }
  buffer-partition lossless-headroom {
    percent 80;
  }
}
egress {
  percent 100;
  buffer-partition lossless {
    percent 90;
  }
  buffer-partition lossy {
    percent 5;
  }
  buffer-partition multicast {
    percent 5;
  }
}
```

## Verification

---

Verify that the shared buffer configuration has been created properly.

### *Verifying the Shared Buffer Configuration*

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2158.00 KB
  Shared Buffer     : 7202.00 KB
    Lossless       : 1080.30 KB
    Lossless Headroom : 5761.60 KB
    Lossy          : 360.10 KB

  Lossless Headroom Utilization:
  Node Device      Total          Used          Free
  0                5761.60 KB    0.00 KB      5761.60 KB

Egress:
  Total Buffer      : 9360.00 KB
  Dedicated Buffer  : 2704.00 KB
  Shared Buffer     : 6656.00 KB
    Lossless       : 5990.40 KB
    Multicast      : 332.80 KB
    Lossy          : 332.80 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 1080 KB to lossless traffic
  - 5761.60 KB to lossless headroom traffic



- 360.10 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Of the total available lossless headroom buffer space of 5761.60 KB, currently no buffer space is being used, so all 5761.60 KB of buffer space is free.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 5990.40 KB to lossless traffic
  - 332.80 KB to multicast traffic
  - 332.80 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 and EX4600 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 and EX4600 switches.

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 6110](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

## Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The switch handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the switch exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE priority (the FCoE IEEE 802.1p code point). The default priority mapping for the FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).
- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 6128](#)
- [Overview on page 6129](#)
- [Configuration on page 6132](#)
- [Verification on page 6134](#)

---

### Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX Series device

- Junos OS Release 12.1 or later for the QFX Series

## Overview

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



**NOTE:** DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 5669](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.
- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 455 on page 5597](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 456 on page 5597](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

**Table 559: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)**

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low

**Table 559: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier) (*continued*)**

Code Point	Forwarding Class	Loss Priority
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

**Table 560: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)**

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	best-effort	low
110	best-effort	low
111	best-effort	low

**Topology**

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



**NOTE:** You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 457 on page 5598 shows the configuration components for this example.

**Table 561: Components of DCBX Application Protocol Exchange Configuration Topology**

Component	Settings
Hardware	QFX Series device
LLDP	Enabled by default on Ethernet interfaces
DCBX	Enabled by default on Ethernet interfaces
iSCSI application (Layer 4)	Application name— <b>iscsi</b> protocol— <b>TCP</b> destination-port— <b>3260</b> code-points— <b>111</b>
PTP application (Layer 2)	Application name— <b>ptp</b> ether-type— <b>0x88F7</b> code-points— <b>001, 101</b>
FCoE application (Layer 2)	Application name— <b>fcoe</b> ether-type— <b>0x8906</b> code-points— <b>011</b>  <b>NOTE:</b> You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map.
Application maps	<b>dcbx-iscsi-fcoe-app-map</b> —Maps the iSCSI and FCoE applications to IEEE 802.1p code points  <b>dcbx-iscsi-ptp-app-map</b> —Maps iSCSI and PTP applications to IEEE 802.1p code points
Interfaces	<b>xe-0/0/10</b> —Configured to exchange FCoE and iSCSI application TLVs (uses application map <b>dcbx-iscsi-fcoe-app-map</b> , carries FCoE traffic, and has PFC enabled on the FCoE priority)  <b>xe-0/0/11</b> —Configured to exchange iSCSI and PTP application TLVs (uses application map <b>dcbx-iscsi-ptp-app-map</b> )
PFC congestion notification profile for FCoE application exchange	<b>fcoe-cnp:</b> <ul style="list-style-type: none"> <li>Code point—<b>011</b></li> <li>Interface—<b>xe-0/0/10</b></li> </ul>

**Table 561: Components of DCBX Application Protocol Exchange Configuration Topology (continued)**

Component	Settings
Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point)	<p><b>fcoe-iscsi-cl1:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>fcoe</b> forwarding class to the IEEE 802.1p code point used for the FCoE application (011) and a loss priority of <b>high</b></li> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of <b>high</b></li> <li>Applied to interface <b>xe-0/0/10</b></li> </ul> <p><b>iscsi-ntp-cl2:</b></p> <ul style="list-style-type: none"> <li>Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (111) and a loss priority of <b>low</b></li> <li>Maps the <b>best-effort</b> forwarding class to the IEEE 802.1p code points used for the PTP application (001 and 101) and a loss priority of <b>low</b></li> <li>Applied to interface <b>xe-0/0/11</b></li> </ul>



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

### Configuration

#### CLI Quick Configuration

To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ntp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ntp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ntp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1 iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1

```

```
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

### *Configuring DCBX Application Protocol TLV Exchange*

**Step-by-Step Procedure** To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.  

```
[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88F7
```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.  

```
[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ptp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ptp-app-map application PTP code-points [001 101]
```
4. Apply the iSCSI and FCoE application map to interface xe-0/0/10, and apply the iSCSI and PTP application map to interface xe-0/0/11.  

```
[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ptp-app-map
```
5. Create the congestion notification profile to enable PFC on the FCoE code point (011), and apply the congestion notification profile to interface xe-0/0/10.  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
```
6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.  

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority high code-points 011
user@switch# set ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class network-control loss-priority high code-points 111
```
7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.  

```
[edit class-of-service classifiers]
```

```
user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
```

8. Apply the classifiers to the appropriate interfaces.

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

---

### Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 6134](#)
- [Verifying the Application Map Configuration on page 6134](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 6135](#)
- [Verifying the PFC Configuration on page 6135](#)
- [Verifying the Classifier Configuration on page 6136](#)

#### *Verifying the Application Configuration*

**Purpose** Verify that DCBX applications have been configured.

**Action** List the applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application iSCSI {
    protocol tcp;
    destination-port 3260;
}

application fcoe {
    ether-type 0x8906;
}

application ptp {
    ether-type 0x88F7;
}
```

**Meaning** The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

#### *Verifying the Application Map Configuration*

**Purpose** Verify that the application maps have been configured.



**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
    application iSCSI code-points 111;
    application FCoE code-points 011;
}

dcbx-iscsi-ptp-app-map {
    application iSCSI code-points 111;
    application PTP code-points [001 101];
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ptp-app-map**.

The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the FCoE application, which is mapped to IEEE 802.1p code point 011.

The application map **dcbx-iscsi-ptp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the PTP application, which is mapped to IEEE 802.1p code points 001 and 101.

#### *Verifying DCBX Application Protocol Exchange Interface Configuration*

**Purpose** Verify that the application maps have been applied to the correct interfaces.

**Action** List the application maps by using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
    application-map dcbx-iscsi-ptp-app-map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ptp-app-map**.

#### *Verifying the PFC Configuration*

**Purpose** Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

**Action** Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
  input {
    ieee-802.1 {
      code-point 011 {
        pfc;
      }
    }
  }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
  congestion-notification-profile fcoe-cnp;
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

**Meaning** The **show class-of-service congestion-notification-profile** configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile **fcoe-cnp** has been configured and has enabled PFC on the IEEE 802.1p code point **011** (the default FCoE code point).

The **show class-of-service interfaces** configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile **fcoe-cnp**, which enables PFC on the FCoE code point, is applied to interface **xe-0/0/10**.

### *Verifying the Classifier Configuration*

**Purpose** Verify that the classifiers have been configured and applied to the correct interfaces.

**Action** Display the classifier configuration by using the configuration mode command **show class-of-service**:

```
user@switch# show class-of-service
classifiers {
  ieee-802.1 fcoe-iscsi-cl1 {
    import default;
    forwarding-class network-control {
      loss-priority high code-points 111;
    }
  }
  forwarding-class fcoe {
    loss-priority high code-points 011;
  }
}
```

```

    }
  }
  ieee-802.1 iscsi-ptp-cl2 {
    import default;
    forwarding-class network-control {
      loss-priority low code-points 111;
    }
    forwarding-class best-effort {
      loss-priority low code-points [ 001 101 ];
    }
  }
}
interfaces {
  xe-0/0/10 {
    congestion-notification-profile fcoe-cnp;
    unit 0 {
      classifiers {
        ieee-802.1 fcoe-iscsi-cl1;
      }
    }
  }
  xe-0/0/11 {
    unit 0 {
      classifiers {
        ieee-802.1 iscsi-ptp-cl2;
      }
    }
  }
}
}

```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

**Meaning** The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ptp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ptp-cl2** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).

- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ptp-cl2** is mapped to interface **xe-0/0/11.0**.

#### Related Documentation

- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [show dcbx on page 5723](#)
- [show dcbx neighbors on page 5724](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Using DCBX Protocol to Lower Costs](#)

## Configuration Examples (QFX5100 Switches Only)

---

- [Example: Configuring PFC Across Layer 3 Interfaces on page 6138](#)

### Example: Configuring PFC Across Layer 3 Interfaces

Priority-based flow control (PFC) helps ensure lossless transport across data center bridging interfaces by pausing incoming traffic when output queue buffers fill to a certain threshold. On a QFX5100 switch or an EX4600 switch running the Enhanced Layer 2 Software (ELS) CLI, in addition to configuring PFC on Layer 2 (bridging) interfaces, you can configure PFC on traffic that traverses Layer 3 interfaces. This enables you to preserve the lossless characteristics that PFC provides on traffic, even when the traffic crosses Layer 3 interfaces that connect two Layer 2 networks.

- [Requirements on page 6138](#)
- [Overview on page 6139](#)
- [Configuration on page 6142](#)
- [Verification on page 6149](#)

#### Requirements

---

This example uses the following hardware and software components:

- Two Juniper Networks QFX5100 Switches
- Junos OS Release 13.2 or later for the QFX Series

- Two Ethernet hosts

### Overview

On a network that uses two QFX5100 or EX4600 switches to connect hosts on two different Ethernet networks across a Layer 3 network, to configure PFC across the Layer 2 and Layer 3 interfaces, you must:

- Configure the Layer 2 and Layer 3 interfaces on the switches
- Configure VLANs to carry the traffic across the Layer 2 and Layer 3 networks
- Configure integrated routing and bridging (IRB) interfaces on the Layer 2 interfaces to move the Layer 2 VLAN traffic to Layer 3
- Configure and apply the appropriate classifiers to the interfaces
- Configure and apply congestion notification profiles (CNPs) on the interfaces to enable PFC on the traffic that you want to be lossless



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure lossless forwarding classes and hierarchical port scheduling (also known as enhanced transmission selection) on the interfaces



**NOTE:** PFC operates at the queue level, based on the IEEE 802.1p code point in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits). For this reason, traffic on Layer 3 interfaces on which you want to enable PFC must use an IEEE 802.1p classifier to map incoming traffic to forwarding classes (which are in turn mapped to output queues) and loss priorities. You cannot use a DSCP or DSCP IPv6 classifier to classify Layer 3 traffic if you want to enable PFC on traffic flows.

### Topology

Figure 218 on page 6140 shows the topology for this example.

Figure 218: Enabling PFC Across Layer 3 Interface Hops




Table 562 on page 6140 shows the configuration components for this example. On the two QFX5100 switches, the Ethernet host-facing interfaces use the same interface names and configuration, and the Layer 3 network-facing interfaces use the same interface names and configuration.

Table 562: Components of the PFC Across Layer 3 Interfaces Topology

Component	Settings
Hardware	Two QFX5100 switches Two Ethernet hosts
Layer 3 interfaces (xe-0/0/40 and xe-0/0/41) and VLANs	Interface xe-0/0/40: <ul style="list-style-type: none"> <li>Interface family—inet</li> <li>Interface IP address—100.103.1.2/24</li> <li>VLAN tagging—enabled</li> <li>Interface VLAN ID—103</li> </ul> Interface xe-0/0/41: <ul style="list-style-type: none"> <li>Interface family—inet</li> <li>Interface IP address—100.104.1.2/24</li> <li>VLAN tagging—enabled</li> <li>Interface VLAN ID—104</li> </ul>
Layer 2 interfaces (xe-0/0/20 and xe-0/0/21) and VLAN membership	Family: Ethernet switching Interface mode—trunk Interface xe-0/0/20 VLAN membership—vlan105 Interface xe-0/0/21 VLAN membership—vlan106
VLANs for the IRB interfaces	VLAN unit 105—family inet, IP address 100.105.1.1/24 VLAN unit 106—family inet, IP address 100.106.1.1/24

**Table 562: Components of the PFC Across Layer 3 Interfaces Topology (*continued*)**

Component	Settings
Layer 2 IRB interfaces	<p>Interface xe-0/0/20:</p> <ul style="list-style-type: none"> <li>IRB interface unit—105</li> <li>IRB interface family—inert</li> <li>IRB interface IP address—100.105.1.1/24</li> <li>IRB interface VLAN ID—105</li> <li>Layer 3 interface name—irb.105</li> </ul> <p>Interface xe-0/0/21:</p> <ul style="list-style-type: none"> <li>IRB interface unit—106</li> <li>IRB interface family—inert</li> <li>IRB interface IP address—100.106.1.1/24</li> <li>IRB interface VLAN ID—106</li> <li>Layer 3 interface name—irb.106</li> </ul>
Forwarding classes (both switches)	<p>Name—lossless-3 Queue mapping—queue 3 Packet drop attribute—no-loss</p> <p>Name—lossless-4 Queue mapping—queue 4 Packet drop attribute—no-loss</p> <p><b>NOTE:</b> Matching the forwarding class names (lossless-3 and lossless-4) to the queue number and to the classified IEEE 802.1p code point (priority) creates a configuration that is logical and easy to map because the forwarding class, queue, and priority all use the same number.</p> <p>Name—all-others Queue mapping—queue 0 Packet drop attribute—none</p> <p><b>NOTE:</b> The forwarding class <i>all-others</i> is for best-effort traffic that traverses the interfaces.</p>
Layer 2 interface behavior aggregate (BA) classifier	<p>Name—lossless-3-4-ieee Forwarding class lossless-3—mapped to code point 011 (IEEE 802.1p priority 3) and a packet loss priority of low Forwarding class lossless-4—mapped to code point 100 (IEEE 802.1p priority 4) and a packet loss priority of low</p> <p>Apply the Layer 2 IEEE 802.1p classifier to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41).</p>
Congestion notification profile (PFC, both switches)	<p>Name—lossless-cnp PFC enabled on IEEE 802.1p code points—011 (lossless-3 forwarding class and priority), 100 (lossless-4 forwarding class and priority)</p> <p>Apply the CNP to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41) to enable PFC on IEEE 802.1p code points 011 and 100.</p>

**Table 562: Components of the PFC Across Layer 3 Interfaces Topology (continued)**

Component	Settings
Hierarchical port scheduling (ETS)	<p>Hierarchical port scheduling (ETS) includes configuring:</p> <ul style="list-style-type: none"> <li>• Schedulers to assign bandwidth to traffic</li> <li>• Scheduler mapping to forwarding classes</li> <li>• Grouping of the forwarding classes (priorities) in forwarding class sets (priority groups)</li> <li>• A traffic control profile to assign bandwidth to the forwarding class set and to associate the forwarding class set with the scheduler mapping</li> </ul> <p>Hierarchical port scheduling also includes applying the hierarchical scheduler (defined in the traffic control profile) to the interfaces.</p> <p>This example focuses on configuring PFC across the Layer 2 and Layer 3 interfaces. To maintain this focus, this example includes the CLI statements needed to configure hierarchical port scheduling, but does not include descriptive explanations of the configuration. The <i>Related Documentation</i> section provides links to example documents that show how to configure hierarchical port scheduling.</p> <p>Apply the scheduling configuration to both the Layer 2 and the Layer 3 interfaces (xe-0/0/20, xe-0/0/21, xe-0/0/40, and xe-0/0/41).</p>

### Configuration

- [Step-by-Step Procedure on page 6144](#)
- [Results on page 6146](#)

#### CLI Quick Configuration

To configure PFC across Layer 3 interfaces, copy the following commands, paste them in a text file, remove the line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level. The same configuration applies to both QFX5100 Switch SW1 and QFX5100 Switch SW2:

```

set interfaces xe-0/0/40 vlan-tagging
set interfaces xe-0/0/40 unit 0 vlan-id 103
set interfaces xe-0/0/40 unit 0 family inet address 100.103.1.2/24
set interfaces xe-0/0/41 vlan-tagging
set interfaces xe-0/0/41 unit 0 vlan-id 104
set interfaces xe-0/0/41 unit 0 family inet address 100.104.1.2/24
set interfaces xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/20 unit 0 family ethernet-switching vlan members vlan105
set interfaces xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
set interfaces xe-0/0/21 unit 0 family ethernet-switching vlan members vlan106
set interfaces vlan unit 105 family inet address 100.105.1.1/24
set interfaces vlan unit 106 family inet address 100.106.1.1/24
set interfaces irb unit 105 family inet address 100.105.1.1/24
set interfaces irb unit 106 family inet address 100.106.1.1/24
set vlans vlan105 vlan-id 105
set vlans vlan106 vlan-id 106
set vlans vlan105 l3-interface irb.105
set vlans vlan106 l3-interface irb.106

```



```

set class-of-service forwarding-classes class lossless-3 queue-num 3 no-loss
set class-of-service forwarding-classes class lossless-4 queue-num 4 no-loss
set class-of-service forwarding-classes class all-others queue-num 0
set class-of-service classifiers ieee-802.1 lossless-3-4-ieee forwarding-class lossless-3 loss-priority
low code-points 011
set class-of-service classifiers ieee-802.1 lossless-3-4-ieee forwarding-class lossless-4 loss-priority
low code-points 100
set class-of-service congestion-notification-profile lossless-cnp input ieee-802.1 code-point 011
pfc
set class-of-service congestion-notification-profile lossless-cnp input ieee-802.1 code-point 100
pfc
set class-of-service schedulers lossless_sch transmit-rate 6g
set class-of-service schedulers lossless_sch shaping-rate percent 100
set class-of-service schedulers all-others_sch transmit-rate 4g
set class-of-service scheduler-maps lossless_map forwarding-class lossless-3 scheduler
lossless_sch
set class-of-service scheduler-maps lossless_map forwarding-class lossless-4 scheduler
lossless_sch
set class-of-service scheduler-maps all-others_map forwarding-class all-others scheduler
all-others_sch
set class-of-service forwarding-class-sets lossless_fc_set class lossless-3
set class-of-service forwarding-class-sets lossless_fc_set class lossless-4
set class-of-service forwarding-class-sets all-others_fc_set class all-others
set class-of-service traffic-control-profiles lossless_tcp scheduler-map lossless_map
set class-of-service traffic-control-profiles lossless_tcp guaranteed-rate percent 60
set class-of-service traffic-control-profiles lossless_tcp shaping-rate percent 100
set class-of-service traffic-control-profiles all-others_tcp scheduler-map all-others_map
set class-of-service traffic-control-profiles all-others_tcp guaranteed-rate percent 40
set class-of-service interfaces xe-0/0/20 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/20 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/20 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
set class-of-service interfaces xe-0/0/21 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/21 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
set class-of-service interfaces xe-0/0/40 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/40 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/40 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/40 classifiers ieee-802.1 lossless-3-4-ieee
set class-of-service interfaces xe-0/0/41 forwarding-class-set lossless_fc_set
output-traffic-control-profile lossless_tcp
set class-of-service interfaces xe-0/0/41 forwarding-class-set all-others_fc_set
output-traffic-control-profile all-others_tcp
set class-of-service interfaces xe-0/0/41 congestion-notification-profile lossless-cnp
set class-of-service interfaces xe-0/0/41 classifiers ieee-802.1 lossless-3-4-ieee

```

### Step-by-Step Procedure

**Step-by-Step Procedure** The following step-by-step procedure shows you how to configure the interfaces, VLANs, lossless forwarding classes, classifiers, and PFC settings to enable PFC across Layer 3 interfaces. For completeness, the class-of-service scheduling configuration (hierarchical port scheduling) is included in the procedure, but without explanatory text. See the *Related Documentation* links for detailed examples of the scheduling elements of the configuration.

1. Configure the Layer 3 interface VLANs and IP addresses:

```
[edit interfaces]
user@switch# set xe-0/0/40 vlan-tagging
user@switch# set xe-0/0/40 unit 0 vlan-id 103
user@switch# set xe-0/0/40 unit 0 family inet address 100.103.1.2/24
user@switch# set xe-0/0/41 vlan-tagging
user@switch# set xe-0/0/41 unit 0 vlan-id 104
user@switch# set xe-0/0/41 unit 0 family inet address 100.104.1.2/24
```

2. Configure the Layer 2 interface VLAN membership and interface mode:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family ethernet-switching interface-mode trunk
user@switch# set xe-0/0/20 unit 0 family ethernet-switching vlan members vlan105
user@switch# set xe-0/0/21 unit 0 family ethernet-switching interface-mode trunk
user@switch# set xe-0/0/21 unit 0 family ethernet-switching vlan members vlan106
```

3. Configure the VLANs for the IRB interfaces:

```
[edit interfaces]
user@switch# set vlan unit 105 family inet address 100.105.1.1/24
user@switch# set vlan unit 106 family inet address 100.106.1.1/24
```

4. Configure the IRB interfaces and VLANs to transport incoming Layer 2 traffic assigned to VLANs vlan105 (of which interface xe-0/0/20 is a member) and vlan106 (of which interface xe-0/0/21 is a member) across Layer 3:

```
[edit]
user@switch# set interfaces irb unit 105 family inet address 100.105.1.1/24
user@switch# set interfaces irb unit 106 family inet address 100.106.1.1/24
user@switch# set vlans vlan105 vlan-id 105
user@switch# set vlans vlan106 vlan-id 106
user@switch# set vlans vlan105 l3-interface irb.105
user@switch# set vlans vlan106 l3-interface irb.106
```

5. Configure the lossless forwarding classes and a best-effort forwarding class for any other traffic that might use the interfaces:

```
[edit class-of-service]
user@switch# set forwarding-classes class lossless-3 queue-num 3 no-loss
user@switch# set forwarding-classes class lossless-4 queue-num 4 no-loss
user@switch# set forwarding-classes class all-others queue-num 0
```

6. Configure the IEEE classifier for the Layer 2 and Layer 3 interfaces to classify incoming traffic into the lossless forwarding classes based on the IEEE 802.1p code point of the traffic:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 lossless-3-4-ieee forwarding-class lossless-3 loss-priority low code-points 011
```

```
user@switch# set ieee-802.1 lossless-3-4-ieee forwarding-class lossless-4 loss-priority
low code-points 100
```

7. Configure the CNP to enable PFC on the lossless priorities (the lossless forwarding classes mapped to IEEE 802.1p code points 3 and 4):

```
[edit class-of-service congestion-notification-profile]
user@switch# set lossless-cnp input ieee-802.1 code-point 011 pfc
user@switch# set lossless-cnp input ieee-802.1 code-point 100 pfc
```

8. Configure hierarchical scheduling to support the lossless configuration (included here for completeness; see the *Related Documentation* links for detailed examples of scheduling configuration) and apply it to the Layer 2 and Layer 3 interfaces:

```
[edit class-of-service]
set schedulers lossless_sch transmit-rate 6g
set schedulers lossless_sch shaping-rate percent 100
set schedulers all-others_sch transmit-rate 4g
set scheduler-maps lossless_map forwarding-class lossless-3 scheduler lossless_sch
set scheduler-maps lossless_map forwarding-class lossless-4 scheduler lossless_sch
set scheduler-maps all-others_map forwarding-class all-others scheduler all-others_sch
set forwarding-class-sets lossless_fc_set class lossless-3
set forwarding-class-sets lossless_fc_set class lossless-4
set forwarding-class-sets all-others_fc_set class all-others
set traffic-control-profiles lossless_tcp scheduler-map lossless_map
set traffic-control-profiles lossless_tcp guaranteed-rate percent 60
set traffic-control-profiles lossless_tcp shaping-rate percent 100
set traffic-control-profiles all-others_tcp scheduler-map all-others_map
set traffic-control-profiles all-others_tcp guaranteed-rate percent 40
set interfaces xe-0/0/20 forwarding-class-set lossless_fc_set output-traffic-control-profile
lossless_tcp
set interfaces xe-0/0/20 forwarding-class-set all-others_fc_set output-traffic-control-profile
all-others_tcp
set interfaces xe-0/0/21 forwarding-class-set lossless_fc_set output-traffic-control-profile
lossless_tcp
set interfaces xe-0/0/21 forwarding-class-set all-others_fc_set output-traffic-control-profile
all-others_tcp
set interfaces xe-0/0/40 forwarding-class-set lossless_fc_set output-traffic-control-profile
lossless_tcp
set interfaces xe-0/0/40 forwarding-class-set all-others_fc_set output-traffic-control-profile
all-others_tcp
set interfaces xe-0/0/41 forwarding-class-set lossless_fc_set output-traffic-control-profile
lossless_tcp
set interfaces xe-0/0/41 forwarding-class-set all-others_fc_set output-traffic-control-profile
all-others_tcp
```

9. Apply the Layer 2 IEEE 802.1p classifier and the CNP to the Layer 3 interfaces:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/40 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# set xe-0/0/40 congestion-notification-profile lossless-cnp
user@switch# set xe-0/0/41 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# set xe-0/0/41 congestion-notification-profile lossless-cnp
```

10. Apply the Layer 2 IEEE 802.1p classifier and the CNP to the Layer 2 interfaces:

```
[edit class-of-service interfaces]
user@switch# xe-0/0/20 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# xe-0/0/20 congestion-notification-profile lossless-cnp
user@switch# xe-0/0/21 unit 0 classifiers ieee-802.1 lossless-3-4-ieee
user@switch# xe-0/0/21 congestion-notification-profile lossless-cnp
```

### Results

Display the results of the interface, VLAN, and class-of-service configurations (the system shows only the explicitly configured parameters; it does not show default parameters). The results are valid for both QFX5100 Switch SW1 and QFX5100 Switch SW2 because the same configuration is used on both switches.

Display the results of the interface configuration:

```
user@switch# show configuration interfaces
```

```
xe-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan105;
      }
    }
  }
}
xe-0/0/21 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan106;
      }
    }
  }
}
xe-0/0/40 {
  vlan-tagging;
  unit 0 {
    vlan-id 103;
    family inet {
      address 100.103.1.2/24;
    }
  }
}
xe-0/0/41 {
  vlan-tagging;
  unit 0 {
    vlan-id 104;
    family inet {
      address 100.104.1.2/24;
    }
  }
}
irb {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
  unit 106 {
    family inet {
      address 100.106.1.1/24;
    }
  }
}
```

```

vlan {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
  unit 106 {
    family inet {
      address 100.106.1.1/24;
    }
  }
}

```

Display the results of the vlan configuration:

```

user@switch# show configuration vlans
vlan105 {
  vlan-id 105;
  l3-interface irb.105;
}
vlan106 {
  vlan-id 106;
  l3-interface irb.106;
}

```

Display the results of the class-of-service configuration:

```

user@switch# show configuration class-of-service
classifiers {
  ieee-802.1 lossless-3-4-ieee {
    forwarding-class lossless-3 {
      loss-priority low code-points 011;
    }
    forwarding-class lossless-4 {
      loss-priority low code-points 100;
    }
  }
}
forwarding-classes {
  class lossless-3 queue-num 3 no-loss;
  class lossless-4 queue-num 4 no-loss;
  class all-others queue-num 0;
}
traffic-control-profiles {
  lossless_tcp {
    scheduler-map lossless_map;
    shaping-rate percent 100;
    guaranteed-rate percent 60;
  }
  all-others_tcp {
    scheduler-map all-others_map;
    guaranteed-rate percent 40;
  }
}
forwarding-class-sets {
  lossless_fc_set {
    class lossless-3;
    class lossless-4;
  }
  all-others_fc_set {

```

```
        class all-others;
    }
}
congestion-notification-profile {
    lossless-cnp {
        input {
            ieee-802.1 {
                code-point 011 {
                    pfc;
                }
                code-point 100 {
                    pfc;
                }
            }
        }
    }
}
}
interfaces {
    xe-0/0/20 {
        forwarding-class-set {
            lossless_fc_set {
                output-traffic-control-profile lossless_tcp;
            }
            all-others_fc_set {
                output-traffic-control-profile all-others_tcp;
            }
        }
        congestion-notification-profile lossless-cnp;
        unit 0 {
            classifiers {
                ieee-802.1 lossless-3-4-ieee;
            }
        }
    }
    xe-0/0/21 {
        forwarding-class-set {
            all-others_fc_set {
                output-traffic-control-profile all-others_tcp;
            }
            lossless_fc_set {
                output-traffic-control-profile lossless_tcp;
            }
        }
        congestion-notification-profile lossless-cnp;
        unit 0 {
            classifiers {
                ieee-802.1 lossless-3-4-ieee;
            }
        }
    }
    xe-0/0/40 {
        forwarding-class-set {
            lossless_fc_set {
                output-traffic-control-profile lossless_tcp;
            }
            all-others_fc_set {
                output-traffic-control-profile all-others_tcp;
            }
        }
        congestion-notification-profile lossless-cnp;
        classifiers {
```

```

        ieee-802.1 lossless-3-4-ieee;
    }
}
xe-0/0/41 {
    forwarding-class-set {
        lossless_fc_set {
            output-traffic-control-profile lossless_tcp;
        }
        all-others_fc_set {
            output-traffic-control-profile all-others_tcp;
        }
    }
    congestion-notification-profile lossless-cnp;
    classifiers {
        ieee-802.1 lossless-3-4-ieee;
    }
}
}
scheduler-maps {
    lossless_map {
        forwarding-class lossless-3 scheduler lossless_sch;
        forwarding-class lossless-4 scheduler lossless_sch;
    }
    all-others_map {
        forwarding-class all-others scheduler all-others_sch;
    }
}
schedulers {
    lossless_sch {
        transmit-rate 6g;
        shaping-rate percent 100;
    }
    all-others_sch {
        transmit-rate 4g;
    }
}
}

```



**TIP:** To quickly configure the switch, issue the `load merge terminal` command, and then copy the hierarchies and paste them into the switch terminal window.

## Verification

To verify that the PFC across Layer 3 interfaces configuration has been created and is operating properly, perform these tasks:

- [Verifying the Interface Configuration on page 6150](#)
- [Verifying the VLAN Configuration on page 6152](#)
- [Verifying the PFC Configuration \(Congestion Notification Profile\) on page 6152](#)
- [Verify the Forwarding Class Configuration on page 6153](#)

- [Verifying the Classifier Configuration on page 6153](#)
- [Verifying the Interface CoS Configuration \(Hierarchical Scheduling, PFC, and Classifier Mapping to Interfaces\) on page 6154](#)

### ***Verifying the Interface Configuration***

**Purpose** Verify that the Layer 2 Ethernet interfaces, Layer 3 IP interfaces, IRB interfaces, and VLAN interfaces have been created on the switch and are correctly configured.



**Action** Display the switch interface configuration using the **show configuration interfaces** command:

```

user@switch> show configuration interfaces
xe-0/0/20 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan105;
      }
    }
  }
}
xe-0/0/21 {
  unit 0 {
    family ethernet-switching {
      interface-mode trunk;
      vlan {
        members vlan106;
      }
    }
  }
}
xe-0/0/40 {
  vlan-tagging;
  unit 0 {
    vlan-id 103;
    family inet {
      address 100.103.1.2/24;
    }
  }
}
xe-0/0/41 {
  vlan-tagging;
  unit 0 {
    vlan-id 104;
    family inet {
      address 100.104.1.2/24;
    }
  }
}
irb {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
  unit 106 {
    family inet {
      address 100.106.1.1/24;
    }
  }
}
vlan {
  unit 105 {
    family inet {
      address 100.105.1.1/24;
    }
  }
}

```

```
    unit 106 {  
        family inet {  
            address 100.106.1.1/24;  
        }  
    }  
}
```

**Meaning** The **show configuration interfaces** command displays all of the interfaces configured on the switch. The command output shows that:

- Interfaces xe-0/0/20 and xe-0/0/21 are Ethernet interfaces (family ethernet-switching) in trunk interface mode. Interface xe-0/0/20 is a member of VLAN vlan105, and interface xe-0/0/21 is a member of VLAN vlan106.
- Interfaces xe-0/0/40 and xe-0/0/41 are IP interfaces (family inet) with VLAN tagging enabled. Interface xe-0/0/40 has an IP address of 100.103.1.2/24 and a VLAN ID of 103. Interface xe-0/0/41 has an IP address of 100.104.1.2/24 and a VLAN ID of 104.
- Two IRB interfaces are configured, IRB unit 105 with an IP address of 100.105.1.1/24 and IRB unit 106 with an IP address of 100.106.1.1/24.
- Two VLAN interfaces are configured, VLAN unit 105 with an IP address of 100.105.1.1/24 (for IRB interface unit 105) and VLAN unit 106 with an IP address of 100.106.1.1/24 (for IRB interface unit 106).

#### *Verifying the VLAN Configuration*

**Purpose** Verify that VLANs have been created on the switch and are correctly configured.

**Action** Display the VLAN configuration using the **show configuration vlans** command:

```
user@switch> show configuration vlans  
vlan105 {  
    vlan-id 105;  
    l3-interface irb.105;  
}  
vlan106 {  
    vlan-id 106;  
    l3-interface irb.106;  
}
```

**Meaning** The **show configuration vlans** command displays all of the VLANs configured on the switch. The command output shows that:

- VLAN vlan105 has been configured with VLAN ID 105 on IRB interface irb.105.
- VLAN vlan106 has been configured with VLAN ID 106 on IRB interface irb.106.

#### *Verifying the PFC Configuration (Congestion Notification Profile)*

**Purpose** Verify that PFC has been enabled on the correct IEEE 802.1p code points (priorities) in the CNP.

**Action** Display the PFC configuration using the **show configuration class-of-service congestion-notification-profile** command:

```
user@switch> show configuration class-of-service congestion-notification-profile
lossless-cnp {
    input {
        ieee-802.1 {
            code-point 011 {
                pfc;
            }
            code-point 100 {
                pfc;
            }
        }
    }
}
```

**Meaning** The **show configuration class-of-service congestion-notification-profile** command displays all of the CNPs configured on the switch. The command output shows that:

- The CNP named **lossless-cnp** is configured on the switch.
- The CNP **lossless-cnp** enables PFC on IEEE 802.1p code points 100 and 100.

#### *Verify the Forwarding Class Configuration*

**Purpose** Verify that the two lossless forwarding classes and the best-effort forwarding class have been configured on the switch.

**Action** Display the forwarding class configuration using the **show configuration class-of-service forwarding-classes** command:

```
user@switch> show configuration class-of-service forwarding-classes
class lossless-3 queue-num 3 no-loss;
class lossless-4 queue-num 4 no-loss;
class all-others queue-num 0;
```

**Meaning** The **show configuration class-of-service forwarding-classes** command displays all of the forwarding classes configured on the switch (default forwarding classes are not displayed). The command output shows that:

- Forwarding class **lossless-3** is mapped to queue 3 and is configured as a lossless forwarding class (the **no-loss** attribute is applied)
- Forwarding class **lossless-4** is mapped to queue 4 and is configured as a lossless forwarding class (the **no-loss** attribute is applied)
- Forwarding class **all-others** is mapped to queue 0. It is not a lossless forwarding class (the **no-loss** attribute is not applied).

#### *Verifying the Classifier Configuration*

**Purpose** Verify that the IEEE 802.1p classifier has been configured on the switch.

**Action** Display the classifier configuration using the **show configuration class-of-service classifiers** command:

```
user@switch> show configuration class-of-service classifiers
ieee-802.1 lossless-3-4-ieee {
  forwarding-class lossless-3 {
    loss-priority low code-points 011;
  }
  forwarding-class lossless-4 {
    loss-priority low code-points 100;
  }
}
```

**Meaning** The **show configuration class-of-service classifiers** command displays all of the classifiers configured on the switch. The command output shows that the Layer 2 IEEE 802.1p classifier **lossless-3-4-ieee** classifies traffic with the code point 011 into the **lossless-3** forwarding class with a loss priority of **low**, and classifies traffic with the code point 100 into the **lossless-4** forwarding class with a loss priority of **low**.

***Verifying the Interface CoS Configuration (Hierarchical Scheduling, PFC, and Classifier Mapping to Interfaces)***

**Purpose** Verify that the interfaces have the correct hierarchical scheduling, PFC, and classifier configurations.

**Action** Display the interface CoS configuration using the **show configuration class-of-service interfaces** command:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/20 {
  forwarding-class-set {
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
  unit 0 {
    classifiers {
      ieee-802.1 lossless-3-4-ieee;
    }
  }
}
xe-0/0/21 {
  forwarding-class-set {
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
  unit 0 {
    classifiers {
      ieee-802.1 lossless-3-4-ieee;
    }
  }
}
xe-0/0/40 {
  forwarding-class-set {
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
  classifiers {
    ieee-802.1 lossless-3-4-ieee;
  }
}
xe-0/0/41 {
  forwarding-class-set {
    lossless_fc_set {
      output-traffic-control-profile lossless_tcp;
    }
    all-others_fc_set {
      output-traffic-control-profile all-others_tcp;
    }
  }
  congestion-notification-profile lossless-cnp;
}
```

```
    classifiers {  
        ieee-802.1 lossless-3-4-ieee;  
    }  
}
```

**Meaning** The **show configuration class-of-service interfaces** command displays all of the CoS components configured on the switch interfaces. The command output shows that:

- The configuration on Layer 2 Ethernet interfaces xe-0/0/20 and xe-0/0/21 includes:
  - Hierarchical scheduling—The forwarding class set **lossless\_fc\_set** with the traffic control profile **lossless\_tcp** for the lossless traffic, and the forwarding class set **all-others\_fc\_set** with the traffic control profile **all-others\_tcp** for the best-effort traffic are applied to both interfaces.
  - PFC—The **lossless-cnp** congestion notification profile is applied to both interfaces.
  - Classifiers—The Layer 2 IEEE 802.1p classifier **lossless-3-4-ieee** is applied to both interfaces.
- The configuration on Layer 3 IP interfaces xe-0/0/40 and xe-0/0/41 includes:
  - Hierarchical scheduling—The forwarding class set **lossless\_fc\_set** with the traffic control profile **lossless\_tcp** for the lossless traffic, and the forwarding class set **all-others\_fc\_set** with the traffic control profile **all-others\_tcp** for the best-effort traffic are applied to both interfaces.
  - PFC—The **lossless-cnp** congestion notification profile is applied to both interfaces.
  - Classifiers—The Layer 2 IEEE 802.1p classifier **lossless-3-4-ieee** is applied to both interfaces. Traffic that would use a DSCP or a DSCP IPv6 classifier if it were configured uses the IEEE 802.1p classifier instead. Using the IEEE 802.1p classifier allows the interface to use PFC to pause traffic during periods of congestion to prevent packet loss.

**Related Documentation**

- [Understanding PFC Functionality Across Layer 3 Interfaces on page 5950](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Forwarding Classes on page 6075](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)

---

## Configuration Tasks

- [Configuring CoS on page 6157](#)
- [Defining CoS Code-Point Aliases on page 6159](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)

- [Configuring a Global MPLS EXP Classifier on page 6161](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 6162](#)
- [Configuring CoS WRED Drop Profiles on page 6163](#)
- [Configuring CoS Drop Profile Maps on page 6164](#)
- [Defining CoS Forwarding Classes on page 6164](#)
- [Defining CoS Forwarding Class Sets on page 6166](#)
- [Disabling the ETS Recommendation TLV on page 6167](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Defining CoS Queue Scheduling Priority on page 6171](#)
- [Changing the Host Outbound Traffic Default Queue Mapping on page 6172](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 6184](#)
- [Assigning CoS Components to Interfaces on page 6185](#)
- [Configuring the DCBX Mode on page 6186](#)
- [Configuring DCBX Autonegotiation on page 6187](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 6190](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 6191](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 6192](#)

## Configuring CoS

The class-of-service topics describe how to configure the Junos CoS components. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue. After defining the CoS components, you assign classifiers to the required physical and logical interfaces.

You can configure various CoS components individually or in combination to define CoS services.



**NOTE:** When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

The following topics describe how to configure CoS components :

- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Example: Configuring Forwarding Classes on page 6075](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Queue Scheduling Priority on page 6087](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 5995](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
- [Defining CoS Code-Point Aliases on page 6159](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\)](#)
- [Assigning CoS Components to Interfaces on page 6185](#)
- [Changing the Host Outbound Traffic Default Queue Mapping on page 6172](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)



- [Configuring the DCBX Mode on page 5668](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)

## Defining CoS Code-Point Aliases

You can use code-point aliases to streamline the process of configuring CoS features on your switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- DSCP or DSCP IPv6—Handles incoming IPv4 or IPv6 packets.
- IEEE 802.1p—Handles Layer 2 CoS.

To configure a code-point alias:

1. Specify a CoS marker type (IEEE 802.1 or DSCP).
2. Assign an alias.
3. Specify the code point that corresponds to the alias.

```
[edit class-of-service code-point-aliases]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) alias-name code-point-bits
```

For example, to configure a code-point alias for an IEEE 802.1 CoS marker type that has the alias name `fcoe1` and maps to the code-point bits 011:

```
[edit class-of-service code-point-aliases]
user@switch# set ieee-802.1 fcoe1 011
```

### Related Documentation

- [Monitoring CoS Value Aliases on page 6294](#)
- [Understanding CoS Code-Point Aliases on page 5808](#)

## Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP or DSCP IPv6) value, the IEEE 802.1p CoS value, or the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. (See [“Configuring a Global MPLS EXP Classifier” on page 4479](#) for how to define EXP classifiers for MPLS traffic.) BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.

Unicast traffic must use different classifiers than multidestination (multicast, broadcast, and destination lookup fail) traffic.

To configure a unicast DSCP, DSCP IPv6, or IEEE 802.1p BA classifier using the CLI:

1. Create a unicast BA classifier:

- To create a unicast DSCP, DSCP IPv6, or IEEE 802.1p BA classifier based on the default classifier, import the default DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name import default forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

- To create a unicast BA classifier that is not based on the default classifier, create a DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the unicast classifier to a specific 10-Gigabit Ethernet interface or to all 10-Gigabit Ethernet interfaces or to all Fibre Channel interfaces on the switch.

- To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set interface-name unit unit classifiers (dscp | ieee-802.1) classifier-name
```

- To apply the classifier to all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and the logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * classifiers (dscp | ieee-802.1) classifier-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 6162](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Monitoring CoS Classifiers on page 6289](#)

- [Understanding CoS Classifiers on page 5810](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)

## Configuring a Global MPLS EXP Classifier

EXP packet classification associates incoming packets with a particular MPLS CoS servicing level. EXP behavior aggregate (BA) classifiers examine the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. EXP BA classifiers allow you to set the forwarding class and loss priority of an MPLS packet based on the incoming CoS value.

You can configure as many EXP classifiers as you want, however, the switch uses only one MPLS EXP classifier as a global classifier, which is applied only on interfaces configured as **family mpls**. All **family mpls** switch interfaces use the global EXP classifier to classify MPLS traffic.

If an EXP classifier is configured, MPLS traffic on **family mpls** interfaces uses the EXP classifier. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.



**NOTE:** There is no default MPLS EXP classifier. If you want to use an MPLS EXP classifier, you must configure it. The MPLS EXP classifier is global and applies only to all **family mpls** interfaces on the switch. You can configure as many MPLS EXP classifiers as you want, but you can only use one MPLS EXP classifier on switch interfaces at any time.

To configure a unicast MPLS EXP classifier using the CLI:

1. Create an EXP classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1 | exp) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the EXP classifier to the switch interfaces:

```
[edit class-of-service]
user@switch# set system-defaults classifiers exp classifier-name
```

### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)

## Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP) value or IEEE 802.1p CoS value in the packet header to determine the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.



**NOTE:** DSCP IPv6 multidestination classifiers are not supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces.

Unicast and multidestination traffic must use different classifiers.

To configure a multidestination BA classifier using the CLI:

1. Create a DSCP or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name forwarding-class forwarding-class-name
loss-priority level code-points [aliases] [bit-patterns]
```

2. Configure the classifier as a multidestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers (dscp | ieee-802.1) classifier-name
```

### Related Documentation

- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 6069](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
- [Monitoring CoS Classifiers on page 6289](#)
- [Understanding CoS Classifiers on page 5810](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)

## Configuring CoS WRED Drop Profiles

You can configure an interpolated weighted random early detection (WRED) profile to control packet drop characteristics for different traffic loss priorities.



**NOTE:** You cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Also, do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

*Interpolated* means that the switch creates a smooth drop curve from a drop start point to a drop end point, with a maximum drop rate that is reached at the drop end point.

The drop start point is the average queue fill level when the WRED algorithm starts to drop packets. Before the drop start point, no packets are scheduled to drop. Specify the drop start point using the first of two **fill-level** statements.

The drop end point is the average queue fill level at which all subsequently arriving packets are dropped. When the queue fill levels falls below the drop end point, packets begin to be forwarded again. (At the drop end point, the packet drop probability becomes 100 percent.) Specify the drop end point using the second of two **fill-level** statements.

The minimum drop rate is always 0. Specify the minimum drop rate using the first of two **drop-probability** statements. The maximum drop rate is the drop probability when the average queue fill level reaches the drop end point. Specify the maximum drop rate using the second of two **drop-probability** statements.

The drop rate is zero until the queue fill level reaches the drop start point. As the queue continues to fill, packets drop in smooth linear curve until the queue reaches the drop end point, when packets drop at the maximum drop rate. If the queue fills beyond the drop end point, all packets that match the drop profile are dropped.

To configure a WRED profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring WRED Drop Profiles on page 6071](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Configuring CoS Drop Profile Maps on page 6164](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)

## Configuring CoS Drop Profile Maps

A drop-profile map associates a WRED profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class, the drop profile map associated with the scheduler applies the specified WRED profile to traffic in the forwarding class that matches the specified loss priority.

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

To configure a drop-profile map using the CLI:

- For the desired scheduler, configure the traffic loss priority and specify the drop profile you want to use to control the drop characteristics for traffic of that loss priority:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name drop-profile-map loss-priority level protocol
any drop-profile drop-profile-name
```

### Related Documentation

- [Example: Configuring Drop Profile Maps on page 6073](#)
- [Configuring CoS WRED Drop Profiles on page 6163](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Understanding CoS WRED Drop Profiles on page 5909](#)

## Defining CoS Forwarding Classes

Forwarding classes allow you to group packets for transmission. The switch supports a total of 12 forwarding classes. In order to forward traffic, you map (assign) the forwarding classes to unicast or multdestination (multicast, broadcast, and destination lookup fail) output queues.

The switch has 12 output queues. Queues 0 through 7 are for unicast traffic and queues 8 through 11 are for multicast traffic. Forwarding classes mapped to unicast queues must carry unicast traffic, and forwarding classes mapped to multdestination queues must carry multdestination traffic. There are four default unicast forwarding classes and one default multdestination forwarding class.

The default unicast forwarding classes are:

- **best-effort**—Best-effort traffic
- **fcoe**—Guaranteed delivery for FCoE traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **network-control**—Network control traffic

The default multidestination forwarding class is:

- **mcast**—Multidestination traffic

Map forwarding classes to queues using the **class** statement, which enables you to configure up to 12 forwarding classes. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. In addition, all forwarding classes mapped to a particular queue must be either lossless or lossy. You cannot mix lossless and lossy forwarding classes (traffic) on the same queue. Also, you cannot mix unicast and multicast forwarding classes on the same queue.

[edit class-of-service forwarding-classes]

```
user@switch# class class-name queue-num queue-number <no-loss>
```



**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment unless you include the optional no-loss packet drop attribute introduced in Junos OS Release 12.3 in the forwarding class configuration..



**NOTE:** Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

[edit class-of-service forwarding-classes]

```
user@switch# queue queue-number class-name
```

The **queue** statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the **queue** statement to map forwarding classes to queues, edit the configuration to replace the **queue** statement with the **class** statement.

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Forwarding Classes on page 6075](#)
- [Monitoring CoS Forwarding Classes on page 6290](#)
- [Understanding CoS Forwarding Classes on page 5830](#)

## Defining CoS Forwarding Class Sets

A forwarding class set is a priority group for enhanced transmission selection (ETS) traffic control. Each forwarding class set consists of one or more forwarding classes (priorities, which can also be considered as output queues).

You can configure up to three unicast forwarding class sets and one multicast forwarding class set.

To configure a forwarding class set using the CLI:

1. Assign one or more forwarding classes to the forwarding class set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets forwarding-class-set-name class
forwarding-class-name
```

2. Map the forwarding class set to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set forwarding-class-set-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 5835](#)



## Disabling the ETS Recommendation TLV

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



**NOTE:** Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# **set enhanced-transmission-selection no-recommendation-tlv**

### Related Documentation

- [Configuring the DCBX Mode on page 5668](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Understanding DCBX on page 5580](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

## Defining CoS Queue Schedulers

Schedulers define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit congestion notification (ECN) is enabled on the queue, the WRED packet drop profiles associated with the queue, and the queue buffer size.

The parameters you configure in a scheduler define the following characteristics for the queues mapped to the scheduler:

- **transmit-rate**—Minimum bandwidth, also known as the committed information rate (CIR), set as a percentage rate or as an absolute value in bits per second. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.



**NOTE:** You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR), set as a percentage rate or as an absolute value in bits per second.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:

- **low**—The scheduler has low priority.
- **strict-high**—The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **drop-profile-map**—Drop profile mapping to a loss priority and protocol to apply WRED to the scheduler.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—Enables ECN on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.



**NOTE:** Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



**NOTE:** Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

To apply scheduling properties to traffic, map schedulers to forwarding classes using a scheduler map, and then associate the scheduler map with the interfaces. This applies the configured scheduling to the traffic in the specified forwarding class on the associated interface. Using different scheduler maps, you can map different schedulers to the same traffic (the same forwarding class) to apply different scheduling to that traffic on different interfaces.

To configure a scheduler using the CLI:

1. Name the scheduler and define the minimum guaranteed bandwidth for the queue:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name transmit-rate (rate | percent percentage)
```

2. Define the maximum bandwidth for the queue:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set shaping-rate (rate | percent percentage)
```

3. Define the queue priority:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set priority level
```

4. Define the drop profile using a drop profile map:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set drop-profile-map loss-priority (low | medium-high | high) protocol protocol
drop-profile drop-profile-name
```

5. Configure the size of the port dedicated buffer space for the queue:

```
[edit class-of-service schedulers scheduler-name]
```

```
user@switch# set buffer-size percent 20
```

6. Enable ECN, if desired (queue should handle best-effort traffic):

```
[edit class-of-service schedulers scheduler-name]  
user@switch# set explicit-congestion-notification
```

7. Configure a scheduler map to map the scheduler to a forwarding class, which applies the scheduler's properties to the traffic in that forwarding class:

```
[edit class-of-service]  
user@switch# set scheduler-maps scheduler-map-name forwarding-class  
forwarding-class-name scheduler scheduler-name
```

8. Assign the scheduler map and its associated schedulers to one or more interfaces using hierarchical scheduling. See [“Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)” on page 5966](#) for a detailed example of hierarchical scheduling.

```
[edit class-of-service]  
user@switch# set traffic-control-profiles tcp-name scheduler-map scheduler-map-name  
user@switch# set interfaces interface-name forwarding-class-set fc-set-name  
output-traffic-control-profile tcp-name
```

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring ECN on page 6090](#)
- [Defining CoS Queue Scheduling Priority on page 6171](#)
- [Configuring CoS WRED Drop Profiles on page 6163](#)
- [Monitoring CoS Scheduler Maps on page 6293](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Priority Group Scheduling on page 5877](#)
- [Understanding CoS Buffer Configuration on page 5891](#)
- [Understanding CoS Explicit Congestion Notification on page 5926](#)

## Defining CoS Queue Scheduling Priority

You can configure the scheduling priority of individual queues by specifying the priority in a scheduler, and then associating the scheduler with a queue by using a scheduler map. Queues can have one of two bandwidth priorities:

- **strict-high** —The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **low**—Low priority. Traffic with this priority is serviced after any queue that has a **strict-high** priority.
- To configure queue priority using the CLI:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name priority level
```

### Related Documentation

- [Example: Configuring Queue Scheduling Priority on page 6087](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Monitoring CoS Scheduler Maps on page 6293](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)

## Changing the Host Outbound Traffic Default Queue Mapping

If you do not want to use the default mapping of host Routing Engine and CPU outbound traffic to queues, you can change the default output queue. You can also change the default DSCP bits used in the type of service (ToS) field of packets generated by the Routing Engine.

Configuring a queue for host outbound traffic maps all traffic that the host generates to one forwarding class (queue). The configuration is global and applies to all host-generated traffic on the switch. Configuring a forwarding class for host outbound traffic does not affect transit or incoming traffic.



**NOTE:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) packets generated by the CPU are always transmitted on the `fcoe` queue (queue 3), even if you configure a queue for host outbound traffic. This helps to ensure lossless behavior for FCoE traffic. QFabric systems classify FIP control packets into the same traffic class (`fcoe`) across the Interconnect device (fabric) and the egress Node device.

To change the host outbound traffic egress queue by including the **host-outbound-traffic** statement at the **[edit class-of-service]** hierarchy level:

```
[edit class-of-service]
host-outbound-traffic {
  forwarding-class class-name;
  dscp-code-point code-point;
}
```

For example, to map host outbound traffic to queue 7 (the network control forwarding class) and set the DSCP code point value to 101010:

```
[edit class-of-service]
host-outbound-traffic {
  forwarding-class network-control;
  dscp-code-point 101010
}
```

### Related Documentation

- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806](#)

## Defining CoS Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) contained in a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the committed information rate (CIR). The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



**NOTE:** You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR).
- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the priority group bandwidth, not the amounts or percentages of total link bandwidth.)



**NOTE:** Because a port can have more than one priority group, when you assign resources to a priority group, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

To configure a traffic control profile using the CLI:

1. Name the traffic control profile and define the minimum guaranteed bandwidth for the priority group:  

```
[edit class-of-service ]
user@switch# set traffic-control-profiles traffic-control-profile-name guaranteed-rate (rate | percent percentage)
```
2. Define the maximum bandwidth for the priority group:  

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set shaping-rate (rate | percent percentage)
```
3. Attach a scheduler map to the traffic control profile:  

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set scheduler-map scheduler-map-name
```

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)

## Configuring CoS PFC (Congestion Notification Profiles)

A congestion notification profile (CNP) enables priority-based flow control (PFC) on specified IEEE 802.1p priorities (code points). A CNP has two components:

- Input CNP:
  - Enable PFC on a specified priority.
  - Configure the maximum receive unit (MRU) on an interface for traffic that matches the PFC priority (optional).
  - Specify the length of the attached cable on the ingress interface (optional)
- Output CNP (optional): Configure flow control to enable PFC pause on specific output queues for specified priorities.



**NOTE:** By default, output queues 3 and 4 (which are mapped to default lossless forwarding classes `fcoe` and `no-loss`, respectively) are configured to respond to PFC pause messages received from the connected peer on priorities 3 and 4 (code points 011 and 100, respectively). If you explicitly configure flow control on any output queue, you must configure flow control on every output queue that you want to respond to pause messages. (The explicit configuration overrides the default configuration.)

To achieve lossless behavior, the output queue priorities on which you enable PFC flow control must match the PFC priorities on which you enable PFC on the input interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101) in the output component of the CNP, then you must also enable pause on priorities 3 and 5 on the input component of the CNP. (In addition, the forwarding classes mapped to the paused output queues must be lossless forwarding classes.)

Associating a CNP with an interface enables PFC on the ingress traffic that matches the priority specified in the input CNP, and programs the queues listed in the output CNP to pause when the interface receives a PFC pause message from the connected peer. Configure PFC on a priority end to end along the entire data path to create a lossless lane of traffic on the network.





**NOTE:** You must enable PFC on the priority used by FCoE traffic on ingress interfaces (input CNP). Enable PFC on the FCoE priority on every interface that carries FCoE traffic. By convention, FCoE traffic uses priority 3 (code point 011), which maps to queue 3. If your network uses priority 3 for FCoE traffic, the default forwarding class and classifier configuration support lossless transport, but you must still configure a CNP and apply it to the correct ingress interfaces to enable PFC and achieve lossless transport.

If your network does not use priority 3 for FCoE traffic, you need to configure a classifier that classifies FCoE traffic into a lossless forwarding class, based on the priority your network uses for FCoE traffic. If you are not using the default lossless forwarding class configuration, then you also need to ensure that the output queue mapped to the lossless FCoE forwarding class is programmed to pause.

You can attach only one CNP to an interface. There is no limit to the total number of CNPs you can create.

Configuring a CNP consists of:

- Naming the CNP.
- Specifying the IEEE 802.1 code point (priority) on which you want to enable PFC on ingress interfaces (input CNP).
- Optionally, specifying the MRU and the length of the attached cable on ingress interfaces (input CNP).
- Optionally, configuring flow control (PFC pause) on specified output queues if you want queues other than queues 3 and 4 to respond to pause messages received from the connected peer (output CNP).
- Mapping the CNP to an interface.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

1. Enable PFC on the desired priority in the input CNP and optionally configure the interface MRU for traffic on that priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input ieee-802.1 code-point
code-point bits pfc mru mru-value
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC on IEEE 802.1 code point **011** and configures an MRU value of **2240**:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011
pfc mru 2240
```

2. Configure the length of the cable attached to the ingress interface (optional):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input cable-length
cable-length-value
```

For example, to configure a CNP named **fcoe-cnp** that sets the length of the ingress interface cable to **100** meters:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input cable-length 100
```

3. (Optional) Configure flow control on output queues:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
code-point-bits flow-control-queue [queue | list-of-queues]
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC pause flow control on output queues 3 and 5 for FCoE traffic that uses priority 3 (code point **011**) and on output queue 4 for traffic that uses priority 4 (code point **100**):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
011 flow-control-queue [3 5]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
100 flow-control-queue 4
```

4. Map the CNP to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface congestion-notification-profile cnp-name
```

For example, to map the CNP **fcoe-cnp** to the interface **xe-0/0/7**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 congestion-notification-profile fcoe-cnp
```

#### Related Documentation

- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)
- [Assigning CoS Components to Interfaces on page 6185](#)
- [Monitoring Interfaces That Have CoS Components on page 6291](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)

## Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control

Ethernet PAUSE flow control is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link, including Ethernet links that belong to Ethernet link aggregated (LAG) interfaces. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to PAUSE messages it receives from the connected peer to stop sending traffic.

Symmetric flow control means that an interface has the same PAUSE configuration in both directions. The PAUSE generation and PAUSE response functions are both configured as enabled, or they are both disabled.

Asymmetric flow control allows you to configure the PAUSE functionality in each direction independently on an interface. The configuration for generating PAUSE messages and for responding to PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. If you do not want to PAUSE all of the traffic on a link, you can use priority-based flow control (PFC) to selectively pause traffic based on its IEEE 802.1p code point.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. If you attempt to configure both features, the switch returns a commit error. Ethernet PAUSE and PFC are also mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

By default, all flow control features are disabled. You enable symmetric flow control on the interfaces on which you want to PAUSE all of the traffic on a link.

- To enable symmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
user@switch# set flow-control
```

- To disable symmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
user@switch# set no-flow-control
```

### Related Documentation

- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Configuring CoS Asymmetric Ethernet PAUSE Flow Control

Ethernet PAUSE flow control is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link, including Ethernet links that belong to link aggregated (LAG) interfaces. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to PAUSE messages it receives from the connected peer to stop sending traffic.

Asymmetric flow control allows you to configure the PAUSE functionality in each direction independently on an interface. The configuration for generating PAUSE messages and for responding to PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction.

Symmetric flow control means that the interface has the same configuration in both directions. The PAUSE generation and PAUSE response functions are both configured as enabled or they are both disabled. If you do not want to PAUSE all of the traffic on a link, you can use priority-based flow control (PFC) to selectively pause traffic based on its IEEE 802.1p code point.

Asymmetric flow control provides the ability to configure the receive buffer and transmit buffer Ethernet PAUSE actions independently on an interface. The buffers perform the following actions:

- The receive buffers generate and send PAUSE messages to the connected peer to ask the peer to stop sending traffic for a time period specified in the PAUSE frame. The peer interface's buffers may store outgoing frames until the PAUSE period elapses and the interface can resume sending traffic.
- The transmit buffers respond to PAUSE messages received from the connected peer to stop sending traffic to the peer. The transmit buffer may store outgoing frames until the PAUSE period elapses and the interface can resume sending traffic.

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to PAUSE messages:

- Receive buffers on—Enable PAUSE transmission (generate and send PAUSE frames)
- Transmit buffers on—Enable PAUSE reception (respond to received PAUSE frames)

You must explicitly set both the receive buffer and the transmit buffer to configure asymmetric flow control.

- To configure asymmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
```

```
user@switch# set configured-flow-control rx-buffers (on | off) tx-buffers (on | off)
```

For example, to configure interface **xe-0/0/24** to generate and send PAUSE messages but not to respond to received PAUSE messages:

```
set interfaces xe-0/0/24 ether-options configured-flow-control rx-buffers on tx-buffers off
```

For example, to configure interface **xe-0/0/30** to respond to received PAUSE messages but not to generate and send PAUSE messages:

```
set interfaces xe-0/0/30 ether-options configured-flow-control rx-buffers off tx-buffers on
```



**NOTE:** If you configure both buffers to be on, that is equivalent to symmetric flow control. If you configure both buffers to be off, there is no flow control (flow control is disabled).

#### Related Documentation

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Configuring Global Ingress and Egress Shared Buffers

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of ingress and egress buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best-effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from one of the recommended configurations, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffer allocation and partitioning using the CLI:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent percent
```

2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent percent
user@switch# set ingress buffer-partition lossless-headroom percent percent
user@switch# set ingress buffer-partition lossy percent percent
```

3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set egress percent percent
```

4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent percent
user@switch# set egress buffer-partition lossy percent percent
user@switch# set egress buffer-partition multicast percent percent
```

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 6110](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

## Defining CoS Rewrite Rules

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on an interface (EXP rewrite rules can only be enabled on **family mpls** logical interfaces, not on physical interfaces). You can also apply an existing rewrite rule on an interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

To create rewrite rules and enable them on interfaces:

- To create an 802.1p rewrite rule named **customup-rw** in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point 000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high code-point 001
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point 010
user@switch# set ieee-802.1 customup-rw forwarding-class fcoe loss-priority low code-point 011
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority low code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef-no-loss loss-priority high code-point 101
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority low code-point 110
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority high code-point 111
```



- To enable an 802.1p rewrite rule named **customup-rw** on a Layer 2 interface:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/7 unit 0 rewrite-rules ieee-802.1
customup-rw
```



**NOTE:** All forwarding classes assigned to port xe-0/0/7 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

- To enable an 802.1p rewrite rule named **customup-rw** on all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical interface (unit) number:

```
[edit]
user@switch# set class-of-service interfaces xe-* unit * rewrite-rules customup-rw
```



**NOTE:** In this case, *all* forwarding classes assigned to *all* 10-Gigabit Ethernet ports must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

#### Related Documentation

- [Monitoring CoS Rewrite Rules on page 6292](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)

## Configuring Rewrite Rules for MPLS EXP Classifiers

You configure EXP rewrite rules to alter CoS values in outgoing MPLS packets on the outbound **family mpls** interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure an EXP CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on a logical **family mpls** interface. EXP rewrite rules can only be enabled on logical **family mpls** interfaces, not on physical interfaces or on interfaces of other family types. You can also apply an existing EXP rewrite rule on a logical interface.



**NOTE:** There are no default rewrite rules.

You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch. On a given **family mpls** logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

To create an EXP rewrite rule for MPLS traffic and enable it on a logical interface:

1. Create an EXP rewrite rule:

```
user@switch# set class-of-service rewrite-rules exp rewrite-rule-name forwarding-class forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

For example, to configure an EXP rewrite rule named **exp-rr-1** for a forwarding class named **mpls-1** with a loss priority of **low** that rewrites the EXP code point value to **001**:

```
user@switch# set class-of-service rewrite-rules exp exp-rr-1 forwarding-class mpls-1 loss-priority low code-points 001
```

2. Apply the rewrite rule to a logical interface:

```
user@switch # set class-of-service interfaces interface-name unit logical-unit rewrite-rules exp rewrite-rule-name
```

For example, to apply a rewrite rule named **exp-rr-1** to logical interface **xe-0/0/10.0**:

```
user@switch# set class-of-service interfaces xe-0/0/10 unit 0 rewrite-rules exp exp-rr-1
```



**NOTE:** In this example, all forwarding classes assigned to port xe-0/0/10 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same interface.

#### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)
- [Monitoring CoS Rewrite Rules on page 6292](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)

## Assigning CoS Components to Interfaces

After you define the following CoS components, you assign them to physical or logical interfaces. Components that you assign to physical interfaces are valid for all of the logical interfaces configured on the physical interface. Components that you assign to a logical interface are valid only for that logical interface.

- Classifiers—Assign only to logical interfaces.
- Congestion notification profiles—Assign only to physical interfaces.
- Forwarding classes—Assign to interfaces by mapping to forwarding class sets.
- Forwarding class sets—Assign only to physical interfaces.
- Output traffic control profiles—Assign only to physical interfaces (with a forwarding class set).
- Rewrite rules—Assign only to logical interfaces.

You can assign a CoS component to a single interface or to multiple interfaces using wildcards. You can also assign a congestion notification profile or a forwarding class set globally to all interfaces.

To assign CoS components to interfaces:

Assign CoS components to a single interface by associating a CoS component (for example a forwarding class set named **san-priority-group**) with an interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set san-priority-group
```

Assign a CoS component to multiple interfaces by associating a CoS component (for example, a rewrite rule named **customup-rw**) to all 10-Gigabit Ethernet interfaces on the switch, use wildcard characters for the interface name and logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * rewrite-rules ieee-802.1 customup-rw
```

Assign a congestion notification profile or a forwarding class set globally to all interfaces using the **set class-of-service interfaces all** statement. For example, to assign a forwarding class set named **be\_fcset** to all interfaces:

```
[edit class-of-service interfaces]
user@switch# set all forwarding-class-set be_fcset
```



**NOTE:** If there is an existing CoS configuration of any type on an interface, the global configuration is not applied to that particular interface. The global configuration is applied to all interfaces that do not have an existing CoS configuration.

For example, if you configure a rewrite rule, assign it to interfaces **xe-0/0/20.0** and **xe-0/0/22.0**, and then configure a congestion notification profile and apply it to all interfaces, the congestion notification profile is applied to every interface except **xe-0/0/20** and **xe-0/0/22**.

**Related  
Documentation**

- [Monitoring Interfaces That Have CoS Components on page 6291](#)
- [Understanding Junos CoS Components on page 5789](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820](#)

## Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. Three DCBX modes are supported:

- Autonegotiation—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- IEEE DCBX—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- DCBX Version 1.01—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric system Node devices other than QFX3500 switches come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



**NOTE:** Pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00 are not supported. If an interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface **xe-0/0/21**:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

#### Related Documentation

- [Configuring DCBX Autonegotiation on page 5669](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Understanding DCBX on page 5580](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [show dcbx neighbors on page 5724](#)

## Configuring DCBX Autonegotiation

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default forwarding class. If ETS is not enabled on an interface (no priority groups are configured), all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



**NOTE:** If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE-FC gateway, it does not send or receive FCoE Initialization Protocol (FIP) packets.
- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.

To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# set enhanced-transmission-selection no-recommendation-tlv

#### Related Documentation

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Disabling the ETS Recommendation TLV on page 5672](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)

## Defining an Application for DCBX Application Protocol TLV Exchange

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

### Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Configuring DCBX Autonegotiation on page 5669](#)



- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [show dcbx neighbors on page 5724](#)

## Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [ aliases ] [ bit-patterns ]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[ 001 101 ]
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675](#)
- [Configuring DCBX Autonegotiation on page 5669](#)

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5724](#)

## Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

```
[edit protocols]
```

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674](#)
- [Configuring DCBX Autonegotiation on page 5669](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 5724](#)

---

## Configuration Statements

- [application \(Application Maps\) on page 6195](#)
- [application \(Applications\) on page 6196](#)

- [application-map](#) on page 6197
- [application-maps](#) on page 6198
- [applications \(Applications\)](#) on page 6199
- [applications \(DCBX\)](#) on page 6200
- [buffer-partition \(Egress\)](#) on page 6201
- [buffer-partition \(Ingress\)](#) on page 6203
- [buffer-size](#) on page 6205
- [cable-length \(Congestion Notification\)](#) on page 6207
- [class-of-service](#) on page 6208
- [class \(Forwarding Classes\)](#) on page 6212
- [class \(Forwarding Class Sets\)](#) on page 6213
- [classifiers](#) on page 6214
- [code-point \(Input Congestion Notification\)](#) on page 6215
- [code-point \(Output Congestion Notification\)](#) on page 6216
- [code-point \(Rewrite Rules\)](#) on page 6217
- [code-point-aliases](#) on page 6217
- [code-points \(Application Maps\)](#) on page 6218
- [code-points \(CoS\)](#) on page 6218
- [configured-flow-control](#) on page 6219
- [congestion-notification-profile](#) on page 6220
- [dcbx](#) on page 6222
- [dcbx-version](#) on page 6223
- [destination-port \(Applications\)](#) on page 6224
- [disable \(DCBX\)](#) on page 6225
- [drop-probability](#) on page 6226
- [drop-profile](#) on page 6227
- [drop-profile-map](#) on page 6227
- [drop-profiles](#) on page 6228
- [dscp](#) on page 6229
- [dscp-ipv6](#) on page 6231
- [dscp-code-point](#) on page 6232
- [egress \(Buffer Configuration\)](#) on page 6233
- [enhanced-transmission-selection](#) on page 6234
- [ether-type](#) on page 6235
- [exp](#) on page 6236
- [explicit-congestion-notification](#) on page 6237
- [fill-level](#) on page 6238

- [flow-control](#) on page 6239
- [flow-control-queue \(Output Congestion Notification\)](#) on page 6240
- [forwarding-class](#) on page 6242
- [forwarding-class \(Host Outbound Traffic\)](#) on page 6243
- [forwarding-class-set](#) on page 6243
- [forwarding-class-sets](#) on page 6244
- [forwarding-classes](#) on page 6245
- [guaranteed-rate](#) on page 6247
- [host-outbound-traffic](#) on page 6248
- [ieee-802.1](#) on page 6249
- [ieee-802.1 \(Input Congestion Notification\)](#) on page 6250
- [ieee-802.1 \(Output Congestion Notification\)](#) on page 6251
- [import](#) on page 6252
- [ingress \(Buffer Configuration\)](#) on page 6253
- [input \(Congestion Notification\)](#) on page 6254
- [interface \(DCBX\)](#) on page 6255
- [interfaces \(Class of Service\)](#) on page 6256
- [interpolate](#) on page 6257
- [loss-priority \(Classifiers\)](#) on page 6258
- [loss-priority \(Drop Profiles\)](#) on page 6259
- [loss-priority \(Rewrite Rules\)](#) on page 6260
- [multi-destination](#) on page 6261
- [mru](#) on page 6262
- [output \(Congestion Notification\)](#) on page 6263
- [output-traffic-control-profile](#) on page 6264
- [pfc \(Input Congestion Notification\)](#) on page 6265
- [policy-options](#) on page 6266
- [priority \(Schedulers\)](#) on page 6267
- [priority-flow-control](#) on page 6268
- [protocol \(Applications\)](#) on page 6269
- [protocol \(Drop Profile Map\)](#) on page 6270
- [queue-num](#) on page 6271
- [recommendation-tlv](#) on page 6272
- [rewrite-rules](#) on page 6273
- [rx-buffers](#) on page 6274
- [scheduler](#) on page 6275
- [scheduler-map](#) on page 6275

- [scheduler-maps on page 6276](#)
- [schedulers on page 6277](#)
- [shaping-rate on page 6278](#)
- [shared-buffer on page 6280](#)
- [system-defaults on page 6281](#)
- [traceoptions \(Class of Service\) on page 6282](#)
- [traffic-control-profiles on page 6284](#)
- [transmit-rate on page 6285](#)
- [tx-buffers on page 6287](#)
- [unit on page 6288](#)

## application (Application Maps)

---

<b>Syntax</b>	<code>application <i>application-name</i> {     <i>code-points</i> [ <i>aliases</i> ] [ <i>bit-patterns</i> ]; }</code>
<b>Hierarchy Level</b>	[edit policy-options <a href="#">application-maps</a> <i>application-map-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Add an application to an application map and define the application's code points.
<b>Options</b>	<i>application-name</i> —Name of the application.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>

## application (Applications)

---

<b>Syntax</b>	<pre>application <i>application-name</i> {     <i>destination-port</i> <i>port-value</i>;     protocol (tcp   udp);     ether-type <i>type</i>; }</pre>
<b>Hierarchy Level</b>	[edit applications]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Configure properties to define an application.
<b>Options</b>	<p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul>

## application-map

---

<b>Syntax</b>	<code>application-map <i>application-map-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx interface interface-name</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Specify an application map to apply to an interface.
<b>Options</b>	<i>application-map-name</i> —Name of the application map.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 5675</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>

## application-maps

---

<b>Syntax</b>	<pre>application-maps <i>application-map-name</i> {     application <i>application-name</i> {         code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];     } }</pre>
<b>Hierarchy Level</b>	[edit policy-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Define an application map by specifying the applications that belong to the application map.
<b>Options</b>	<p><i>application-map-name</i>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul>



## applications (Applications)

<b>Syntax</b>	<pre> applications {   application application-name {     destination-port port-value;     protocol (tcp   udp);     ether-type type;   } } </pre>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Define applications that DCBX advertises.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>

## applications (DCBX)

---

<b>Syntax</b>	<pre>applications {     no-auto-negotiation; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx interface interface-name</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.1 for the EX Series
<b>Description</b>	Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li></ul>

## buffer-partition (Egress)

**Syntax** `buffer-partition (lossless | lossy | multicast) {  
percent percent;  
}`

**Hierarchy Level** [edit [class-of-service shared-buffer egress](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** The egress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best-effort, lossless, and multicast traffic (multicast also includes broadcast and destination lookup fail traffic). To better support the mix of traffic on your network, you can optimize the allocation of egress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three egress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ports. All ports on the switch receive the same allocation of egress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

**Default** The default egress buffer partition shown in [Table 563 on page 6201](#) supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

**Table 563: Default Egress Shared Buffer Partitioning**

Lossless Partition	Lossy Partition	Multicast Partition
50%	31%	19%

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages must always total 100 percent.

**Options** **lossless**—Shared buffer space reserved for all lossless egress traffic.

**lossy**—Shared buffer space for best-effort unicast egress traffic.

**multicast**—Shared buffer space reserved for all multicast (including broadcast and destination lookup fail) egress traffic.

**percent percent**—The percentage of buffer space to allocate to the specified buffer partition (lossless, lossy, or multicast buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
  - [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
  - [Understanding CoS Buffer Configuration on page 5891](#)

## buffer-partition (Ingress)

**Syntax** `buffer-partition (lossless | lossless-headroom | lossy) {  
percent percentage;  
}`

**Hierarchy Level** [edit [class-of-service shared-buffer ingress](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** The ingress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best effort (best-effort unicast and multicast) and lossless traffic. To better support the mix of traffic on your network, you can optimize the allocation of ingress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three ingress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ingress traffic. All ports on the switch receive the same allocation of ingress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

**Default** The default ingress buffer partition shown in [Table 564 on page 6203](#) supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

**Table 564: Default Ingress Shared Buffer Partitioning**

Lossless Partition	Lossless-Headroom Partition	Lossy Partition
9%	45%	46%

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages always must total 100 percent.

**Options** **lossless**—Shared buffer space reserved for all lossless ingress traffic.

**lossless-headroom**—Shared buffer space reserved to store packets received while either an 802.3x Ethernet PAUSE or a priority-based flow control (PFC) pause is asserted. (When an ingress interface pauses traffic, it must have the buffer space to store all of the packets currently in the buffer, and also all of the packets received before the connected peer stops sending traffic and the wire is cleared of packets.)

**lossy**—Shared buffer space for best-effort ingress traffic.

**percent *percent***—The percentage of buffer space to allocate to the specified buffer partition (lossless, lossless-headroom, or lossy buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

<b>Required Privilege</b>	interfaces—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122</a></li><li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 6179</a></li><li>• <a href="#">Understanding CoS Buffer Configuration on page 5891</a></li></ul>
------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## buffer-size

<b>Syntax</b>	<code>buffer-size (percent <i>percent</i>   remainder);</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service schedulers <i>scheduler-name</i></code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Set the dedicated buffer size of the egress queue that you bind the scheduler to in the scheduler map configuration. The switch allocates space from the global dedicated buffer pool to ports and queues in a hierarchical manner. The switch allocates an equal number of dedicated buffers to each egress port, so each egress port receives the same amount of dedicated buffer space. The amount of dedicated buffer space per port is not configurable.

However, the **buffer-size** statement allows you to control the way each port allocates its share of dedicated buffers to its queues. For example, if a port only uses two queues to forward traffic, you can configure the port to allocate all of its dedicated buffer space to those two ports and avoid wasting buffer space on queues that are not in use. We recommend that the buffer size should be the same size as the minimum guaranteed transmission rate (the **transmit-rate**).

You configure the proportion of port dedicated buffers allocated to a particular output queue using the following process:

1. Configure a scheduler and set the **buffer-size** option to match the scheduler **transmit-rate** value.
2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.

For example, suppose that you want to change the dedicated buffer allocation for FCoE traffic. FCoE traffic is mapped to the `fcoe` forwarding class, and the `fcoe` forwarding class is mapped to queue 3 (this is the default configuration). To use default FCoE traffic mapping, in the scheduler map configuration, map the scheduler to the **fcoe** forwarding class.

3. Associate the scheduler map with the traffic control profile you want to use on the egress ports that carry FCoE traffic.
4. Associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry FCoE traffic.

Queue 3, which is mapped to the `fcoe` forwarding class and therefore to the FCoE traffic, receives the dedicated buffer allocation specified in the **buffer-size** statement.



**NOTE:** The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

**Default** The port allocates dedicated buffers to queues that have an explicitly configured scheduler buffer size. If you do not explicitly configure a scheduler buffer size for a queue, the port serves the explicitly configured queues first. Then the port divides the remaining dedicated buffers equally among the queues that have an explicitly attached scheduler *without* an explicitly configured buffer size configuration. (If you configure a scheduler, but you do not configure the buffer size parameter, the default is equivalent to configuring the buffer size with the **remainder** option.)

If you use the default scheduler and scheduler map on a port (no explicit scheduler configuration), then the port allocates its dedicated buffer pool to queues based on the default scheduling, as shown in [Table 565 on page 6206](#). The default buffer size is the same as the default transmit rate for each default queue:

**Table 565: Default Output Queue Buffer Sizes**

Queue Number	Forwarding Class	Transmit Rate	Buffer Size
0	best-effort	5%	5%
3	fcoe	35%	35%
4	no-loss	35%	35%
7	network-control	5%	5%
8	mcast	20%	20%

Because the default scheduler includes only five forwarding classes, only the queues mapped to those forwarding classes receive dedicated buffers from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

**Options** **percent percent**—Percentage of the port dedicated buffer pool allocated to the queue (or queues) mapped to the scheduler.

**remainder**—Remaining dedicated buffer pool after the port satisfies the needs of the explicitly configured buffers. The port divides the remaining buffers equally among the queues that are explicitly attached to a scheduler but that do not have an explicit buffer size configuration (or are configured with **remainder** as the buffer size).


**Required Privilege** interfaces—To view this statement in the configuration.

**Level** interface-control—To add this statement to the configuration.



- Related Documentation**
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
  - [Understanding CoS Buffer Configuration on page 5891](#)

## cable-length (Congestion Notification)

<b>Syntax</b>	<code>cable-length <i>cable-length-value</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i></a> <a href="#">input</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Specify the length of the cable between the interface and its peer interface in meters. The system uses the cable length and the maximum receive unit (MRU) to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The the shorter the cable length and lower the MRU, the less headroom buffer space is required for PFC.
<div>  <p><b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces <i>interface-name</i></a>] hierarchy level.</p> </div>	
<b>Default</b>	The default cable length value is 100 meters (approximately 328 feet).
<b>Options</b>	<code><i>cable-length-value</i></code> —Length of the cable in meters. (Generally from 1 to 300 meters, but there is no configuration restriction.)
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837</a></li> </ul>

## class-of-service

```

Syntax  class-of-service {
        classifiers {
            (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
                import (classifier-name | default);
                forwarding-class class-name {
                    loss-priority level {
                        code-points [ aliases ] [ bit-patterns ];
                    }
                }
            }
        }
        code-point-aliases {
            (dscp | dscp-ipv6 | ieee-802.1) {
                alias-name bits;
            }
        }
        congestion-notification-profile profile-name {
            input {
                ieee-802.1 {
                    code-point [code-point-bits] {
                        pfc {
                            mru mru-value;
                        }
                    }
                }
                cable-length cable-length-value;
            }
            output {
                ieee-802.1 {
                    code-point [code-point-bits] {
                        flow-control-queue [queue | list-of-queues];
                    }
                }
            }
        }
        drop-profiles {
            profile-name {
                interpolate {
                    fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;
                }
            }
        }
        forwarding-class class-name {
            loss-priority level {
                code-points [ aliases ] [ bit-patterns ];
            }
        }
        forwarding-class class-name {
            scheduler scheduler-name;
        }
        forwarding-class-sets forwarding-class-set-name {
            class class-name;
        }
    }

```

```

}
forwarding-classes {
  class {
    class-name {
      queue-num queue-number <no-loss>;
    }
  }
}
host-outbound-traffic {
  forwarding-class class-name;
  dscp-code-point code-point;
}
interfaces {
  interface-name {
    congestion-notification-profile profile-name {
    }
    forwarding-class lossless-forwarding-class-name;
    forwarding-class-set forwarding-class-set-name {
      output-traffic-control-profile profile-name;
    }
    rewrite-value {
      input {
        ieee-802.1 {
          code-point code-point-bits;
        }
      }
    }
  }
  unit logical-unit-number {
    classifiers {
      (dscp | dscp-ipv6 | ieee-802.1 exp) (classifier-name | default);
    }
    forwarding-class class-name;
    rewrite-rules {
      (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
    }
  }
}
multi-destination {
  classifiers {
    (dscp | ieee-802.1) classifier-name;
  }
}
rewrite-rules {
  (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
    import (rewrite-name | default);
    forwarding-class class-name {
      loss-priority priority code-point (alias | bits);
    }
  }
}
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
}

```

```
schedulers {
  scheduler-name {
    buffer-size (percent percentage | remainder);
    drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile
      drop-profile-name;
    explicit-congestion-notification;
    priority priority;
    shaping-rate (rate | percent percentage);
    transmit-rate (percent percentage);
  }
}
shared-buffer {
  egress {
    percent percent;
    buffer-partition (lossless | lossy | multicast) {
      percent percent
    }
  }
  ingress {
    percent percent;
    buffer-partition (lossless | lossless-headroom | lossy) {
      percent percent
    }
  }
}
system-defaults {
  classifiers exp classifier-name;
}
traffic-control-profiles profile-name {
  guaranteed-rate (rate | percent percentage);
  scheduler-map map-name;
  shaping-rate (rate | percent percentage);
}
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure class-of-service parameters on the switch.

The remaining statements are explained separately.

**Default** If you do not configure any CoS features, the default CoS settings are used.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related  
Documentation**

- [Assigning CoS Components to Interfaces on page 6185](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 6162](#)
- [Configuring a Global MPLS EXP Classifier on page 4479](#)
- [Defining CoS Code-Point Aliases on page 6159](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Example: Configuring ECN on page 6090](#)
- [Configuring CoS Drop Profile Maps on page 6164](#)
- [Defining CoS Forwarding Class Sets on page 6166](#)
- [Defining CoS Forwarding Classes on page 6164](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Defining CoS Queue Schedulers on page 6167](#)
- [Configuring CoS WRED Drop Profiles on page 6163](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 6172](#)
- [Overview of Junos OS CoS for the QFX Series and EX4600 Switch on page 5781](#)

## class (Forwarding Classes)

---

**Syntax**    `class {  
              class-name {  
                  queue-num queue-number <no-loss>;  
              }  
          }`

**Hierarchy Level**    [edit [class-of-service forwarding-classes](#)]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

**Description**    Map one or more forwarding classes to a single queue. You can map unicast forwarding classes to a unicast queue (0 through 7) and multdestination forwarding classes to a multicast queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multicast forwarding class.



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.

---

**Options**    *class-name* —Name of the forwarding class.

The remaining statement is explained separately.

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Forwarding Classes on page 6075](#)
- [Understanding CoS Forwarding Classes on page 5830](#)

---

## class (Forwarding Class Sets)

---

<b>Syntax</b>	<code>class <i>class-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service forwarding-class-sets</a> <i>forwarding-class-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Group forwarding classes into sets of forwarding classes (priority groups). You can group some or all of the configured forwarding classes into up to three unicast forwarding class sets and one multidestination forwarding class set.
<b>Options</b>	<i>class-name</i> —Name of the forwarding class.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6078</a></li><li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5835</a></li></ul>

## classifiers

<b>List of Syntax</b>	<a href="#">Syntax (BA Classifiers) on page 6214</a> <a href="#">Syntax (Multidestination BA Classifiers) on page 6214</a> <a href="#">Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE) on page 6214</a> <a href="#">Syntax (Global EXP Interface Classifier Association with Interfaces) on page 6214</a>
<b>Syntax (BA Classifiers)</b>	<pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1   exp) classifier-name {     import (classifier-name   default);     forwarding-class class-name {       loss-priority level {         code-points [ aliases ] [ bit-patterns ];       }     }   } } </pre>
<b>Syntax (Multidestination BA Classifiers)</b>	<pre> classifiers {   (dscp   ieee-802.1) classifier-name; } </pre>
<b>Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)</b>	<pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1) (default   classifier-name); } </pre>
<b>Syntax (Global EXP Interface Classifier Association with Interfaces)</b>	<pre> classifiers {   exp classifier-name; } </pre>
<b>Hierarchy Level (BA Classifiers)</b>	[edit <a href="#">class-of-service</a> ],
<b>Hierarchy Level (Multidestination BA Classifiers)</b>	[edit <a href="#">class-of-service multi-destination</a> ],
<b>Hierarchy Level (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)</b>	[edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]
<b>Hierarchy Level (Global EXP Classifier)</b>	[edit <a href="#">class-of-service system-defaults</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. EXP statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Define a unicast or multidestination CoS behavior aggregate (BA) classifier.
<b>Options</b>	The statements are explained separately.



<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p) on page 6160</a></li> <li>• <a href="#">Configuring a Global MPLS EXP Classifier on page 4479</a></li> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li> <li>• <a href="#">Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers on page 6069</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419</a></li> </ul>

## code-point (Input Congestion Notification)

<b>Syntax</b>	<pre>code-point [<i>code-point-bits</i>] {     pfc {         mru <i>mru-value</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service congestion-notification-profile</a> <i>profile-name</i> <a href="#">input</a> <a href="#">ieee-802.1</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Enable priority-based flow control (PFC) on an IEEE 802.1p code point (priority).
<b>Options</b>	<p><b><i>code-point-bits</i></b>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5606</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> </ul>

## code-point (Output Congestion Notification)

---

<b>Syntax</b>	<code>code-point [ <i>code-point-bits</i> ] {     <i>flow-control-queue</i> [ <i>queue</i>   <i>list-of-queues</i> ]; }</code>
<b>Hierarchy Level</b>	[edit <i>class-of-service congestion-notification-profile profile-name output ieee-802.1</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Specify the IEEE 802.1p code point bits that identify the traffic you want to enable for priority-based flow control (PFC) pause.
<b>Default</b>	<p>By default, IEEE 802.1p priorities 3 and 4 (code points 011 and 100, respectively) are enabled for PFC pause on all Ethernet interfaces. If you explicitly configure priorities to pause and the output queues on which to enable pause, the explicit configuration overrides the default configuration. When you apply an explicit output congestion notification profile to an interface, only the priorities and queues specified in the output congestion notification profile are enabled for pause on that interface.</p> <p>For example, if you configure an output congestion notification profile that specifies priority 2 (code point 010), then traffic with IEEE 802.1p priority 2 is paused on the configured output queue during periods of congestion. However, traffic with priority 3 and priority 4 is not programmed to pause, because the explicit configuration overwrites the default configuration, and the explicit configuration does not pause priority 3 and priority 4. If you configure an explicit output congestion notification profile, all of the priorities you want to enable for PFC and all of the output queues you want to pause must be explicitly configured.</p>
<b>Options</b>	<p><i>code-point-bits</i>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837</a></li></ul>

## code-point (Rewrite Rules)

<b>Syntax</b>	<code>code-point [ <i>alias</i> ] [ <i>bit-pattern</i> ];</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <a href="#">forwarding-class</a> <a href="#">class-name</a> <a href="#">loss-priority</a> <a href="#">level</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a code-point alias or bit set to apply to a forwarding class for a rewrite rule.
<b>Options</b>	<p><i>alias</i>—Name of the alias.</p> <p><i>bit-pattern</i>—Value of the code-point bits, in decimal form.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li> </ul>

## code-point-aliases

<b>Syntax</b>	<pre>code-point-aliases {   (<a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>) {     <i>alias-name</i> <i>bits</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define an alias for a CoS marker. You can use the alias instead of the bit pattern when you specify the code point during configuration.
<b>Options</b>	<p>(<a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>)—Set the type of classifier for which you are creating an alias.</p> <p><i>alias-name</i>—Name of the code-point alias.</p> <p><i>bits</i> —Value of the code-point bits, in decimal form.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Code-Point Aliases on page 6159</a></li> <li>• <a href="#">Understanding CoS Code-Point Aliases on page 5808</a></li> </ul>

## code-points (Application Maps)

---


Syntax	<code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>
Hierarchy Level	<code>[edit policy-options <b>application-maps</b> <i>application-map-name</i> <b>application</b> <i>application-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
Description	Define one or more code-point aliases or bit sets for an application.
Options	<i>aliases</i> —Name of the alias or aliases.  <i>bit-patterns</i> —Value of the code-point bits, in decimal form.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 5674</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul>

## code-points (CoS)

---

Syntax	<code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>
Hierarchy Level	<code>[edit <b>class-of-service</b> <b>classifiers</b> (<b>dscp</b>   <b>dscp-ipv6</b>   <b>ieee-802.1</b>) <i>classifier-name</i> <b>forwarding-class</b> <i>class-name</i> <b>loss-priority</b> <i>level</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure one or more code-point aliases or bit sets to apply to a forwarding class.
Options	<i>aliases</i> —Name of the alias or aliases.  <i>bit-patterns</i> —Value of the code-point bits, in decimal form.
Required Privilege Level	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li></ul>

## configured-flow-control

<b>Syntax</b>	configured-flow-control { <b>rx-buffers</b> (on   off); <b>tx-buffers</b> (on   off); }
<b>Hierarchy Level</b>	[edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]
<b>Description</b>	<p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the <b>rx-buffers</b> and the <b>tx-buffers</b> values when you configure asymmetric flow control.</p> <p>Use the <b>flow-control</b> and <b>no-flow-control</b> statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <hr/> <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> <hr/>
<b>Default</b>	Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">congestion-notification-profile on page 6220</a></li> <li>• <a href="#">flow-control on page 2659</a></li> </ul>

## congestion-notification-profile

<b>Syntax</b>	<pre> congestion-notification-profile <i>profile-name</i> {   input {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         pfc {           mru <i>mru-value</i>;         }       }     }     cable-length <i>cable-length-value</i>;   }   output {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         flow-control-queue [<i>queue</i>   <i>list-of-queues</i>];       }     }   } } </pre>
<b>Interface Congestion Notification Profile Association</b>	<pre> congestion-notification-profile <i>profile-name</i> { </pre>
<b>Hierarchy Level</b>	<pre> [edit <i>class-of-service</i>], [edit <i>class-of-service interfaces interface-name</i>] </pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a congestion notification profile to enable priority-based flow control (PFC) on traffic specified by an IEEE 802.1 code point, and apply the profile to an interface.



**NOTE:** You must configure PFC for FCoE traffic. Each interface that carries FCoE traffic should be configured for PFC on the FCoE code point (usually 011).

You can attach a maximum of one congestion notification profile to an interface. There is no limit to the total number of congestion notification profiles you can create.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

**Options** *profile-name*—Name of the congestion notification profile.

The remaining statements are explained separately.

**Required Privilege** interface—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 6019](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 6050](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)

## dcbx

---

Syntax	<pre>dcbx {   disable;   interface (<i>interface-name</i>   all) {     disable;     application-map <i>application-map-name</i>;     applications {       no-auto-negotiation;     }     enhanced-transmission-selection {       no-auto-negotiation;       no-recommendation-tlv;       recommendation-tlv {         no-auto-negotiation;       }     }     dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);     priority-flow-control {       no-auto-negotiation;     }   } }</pre>
Hierarchy Level	[edit <a href="#">protocols</a> ]
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p><b>mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	Configure DCBX properties.
Options	The statements are explained separately.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li><li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li><li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li></ul>




## dcbx-version

---

<b>Syntax</b>	<code>dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);</code>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	<p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p>
<b>Default</b>	The default DCBX mode is autonegotiation.
<b>Options</b>	<p><b>auto-negotiate</b>—Automatically negotiate the DCBX version with the connected peer.</p> <p><b>ieee-dcbx</b>—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p><b>dcbx-version-1.01</b>—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <a href="#">Understanding DCBX on page 5580</a></li> </ul>

## destination-port (Applications)

---

<b>Syntax</b>	<code>destination-port <i>port-value</i>;</code>
<b>Hierarchy Level</b>	[edit applications <b>application</b> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with <b>protocol</b> to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at <a href="http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml">http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml</a> for a list of assigned port numbers.</p>
<hr/>	
<div> <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>.</div> <hr/>	
<b>Options</b>	<i>port-value</i> —Identifier for the port.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul>

## disable (DCBX)

---

<b>Syntax</b>	disable
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx</a> ]  [edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 11.3 for EX Series switches.
<b>Description</b>	Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.
<b>Default</b>	DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.  DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> <li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li> </ul>

## drop-probability

---

Syntax	drop-probability 0 drop-probability <i>high-value</i> ;
Hierarchy Level	[edit <a href="#">class-of-service drop-profiles</a> <i>profile-name</i> interpolate]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>When configuring WRED, map the packet <b>drop-probability</b> to the fullness of a queue (<b>fill-level</b>). You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs by specifying a low <b>fill-level</b> value at which packets begin to drop (the drop probability is zero until the queue reaches this level of fullness) and a high <b>fill-level</b> value at which packets drop at the highest drop probability. As the queue fills from the low fill level to the high fill level, the rate of packet drop increases in a linear pattern from zero to the high drop probability.</p>
Options	<p>0—Probability that packets will drop at the lowest <b>fill-level</b> value. This is always zero, because until the queue reaches the specified low <b>fill-level</b> value, no packets are scheduled to drop.</p> <p><b>high-value</b>—The maximum probability that packets will drop before queue fullness exceeds the high value of the queue <b>fill-level</b>, expressed as a percentage. If the queue fills beyond the high <b>fill-level</b> value, all packets drop.</p> <p><b>Range:</b> 0 through 100</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li></ul>

## drop-profile

<b>Syntax</b>	<code>drop-profile <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> <a href="#">drop-profile-map <i>loss-priority</i></a> (low   medium-high   high) <a href="#">protocol <i>protocol</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define drop profiles for random early detection (RED). When a packet arrives, RED checks the queue fill level specified in the drop profile. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.
<b>Options</b>	<i>profile-name</i> —Name of the drop profile.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 6073</a></li> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li> </ul>

## drop-profile-map

<b>Syntax</b>	<code>drop-profile-map <a href="#">loss-priority</a> (low   medium-high   high) <a href="#">protocol <i>protocol</i></a> <a href="#">drop-profile <i>drop-profile-name</i></a>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Map a drop profile to a loss priority and protocol for random early detection (RED). When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 6073</a></li> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li> </ul>

## drop-profiles

---

<b>Syntax</b>	<pre>drop-profiles {     profile-name {         interpolate {             fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;         }     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Define drop profiles for weighted random early detection (WRED).</p> <p>For a packet to be dropped, it must match the drop profile. When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.</p>
<b>Options</b>	<p><b>profile-name</b>—Name of the drop profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li></ul>

## dscp

<b>List of Syntax</b>	<a href="#">Syntax (Classifier) on page 6229</a> <a href="#">Syntax (Code-Point Alias) on page 6229</a> <a href="#">Syntax (Multidestination Classifier) on page 6229</a> <a href="#">Syntax (Interface Classifier Association) on page 6229</a> <a href="#">Syntax (Rewrite Rule) on page 6229</a>
<b>Syntax (Classifier)</b>	<pre>dscp classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>
<b>Syntax (Code-Point Alias)</b>	<pre>dscp alias-name bit-pattern;</pre>
<b>Syntax (Multidestination Classifier)</b>	<pre>dscp classifier-name;</pre>
<b>Syntax (Interface Classifier Association)</b>	<pre>dscp (classifier-name   default);</pre>
<b>Syntax (Rewrite Rule)</b>	<pre>dscp rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>
<b>Hierarchy Level (Classifier)</b>	[edit <a href="#">class-of-service classifiers</a> ],
<b>Hierarchy Level (Code-Point Aliases)</b>	[edit <a href="#">class-of-service code-point-aliases</a> ],
<b>Hierarchy Level (Multidestination Classifier)</b>	[edit <a href="#">class-of-service multi-destination classifiers</a> ],
<b>Hierarchy Level (Interface Classifier Association)</b>	[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ], [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],
<b>Hierarchy Level (Rewrite Rule)</b>	[edit <a href="#">class-of-service rewrite-rules</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.

<b>Description</b>	Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.
<b>Options</b>	<p><i>classifier-name</i>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 6159</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 5914</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li></ul>



## dscp-ipv6

<b>List of Syntax</b>	<a href="#">Syntax (Classifier) on page 6231</a> <a href="#">Syntax (Code-Point Alias) on page 6231</a> <a href="#">Syntax (Interface Classifier Association) on page 6231</a> <a href="#">Syntax (Rewrite Rule) on page 6231</a>
<b>Syntax (Classifier)</b>	<pre>dscp-ipv6 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>
<b>Syntax (Code-Point Alias)</b>	<pre>dscp-ipv6 alias-name bit-pattern;</pre>
<b>Syntax (Interface Classifier Association)</b>	<pre>dscp-ipv6 (classifier-name   default);</pre>
<b>Syntax (Rewrite Rule)</b>	<pre>dscp-ipv6 rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>
<b>Hierarchy (Classifier)</b>	[edit <a href="#">class-of-service classifiers</a> ],
<b>Hierarchy (Code-Point Alias)</b>	[edit <a href="#">class-of-service code-point-aliases</a> ],
<b>Hierarchy (Interface Classifier Association)</b>	[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ], [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],
<b>Hierarchy (Rewrite Rule)</b>	[edit <a href="#">class-of-service rewrite-rules</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.



**NOTE:** There is no DSCP IPv6 classifier for multdestination (multicast, broadcast, and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multidestination DSCP classifier.

<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 6159</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 5914</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li></ul>


---

## dscp-code-point

---

<b>Syntax</b>	<code>dscp-code-point <i>code-point</i>;</code>
<b>Hierarchy Level</b>	[edit class-of-service <a href="#">host-outbound-traffic</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Set the value of the DSCP code point in the type of service (ToS) field of the packet generated by the Routing Engine (host).
<b>Options</b>	<b>code-point</b> —Six-bit DSCP code point value.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 6172</a></li><li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806</a></li></ul>

## egress (Buffer Configuration)

<b>Syntax</b>	<pre>egress {   percent <i>percent</i>;   <b>buffer-partition</b> (lossless   lossy   multicast) {     percent <i>percent</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit <b>class-of-service shared-buffer</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	<p>Configure the global shared buffer pool allocation for egress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for other buffers such as dedicated buffers (buffers allocated permanently to ports).</p> <p>The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports. However, on a port, you can configure the portion of dedicated port buffer space allocated to each queue in the scheduler configuration using the <b>buffer-size</b> option.)</p>
	<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>CAUTION:</b> Changing the buffer configuration is a disruptive event. Traffic stops on <i>all</i> ports until buffer reprogramming is complete.</p> </div> </div>
	<p>You can also partition the shared buffer pool to adjust the egress buffer allocations for different mixes of network traffic using the <b>buffer-partition</b> statement.</p>
<b>Default</b>	The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)
<b>Options</b>	<p><b>percent <i>percent</i></b>—Percentage of available egress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104</a></li> </ul>

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 6179](#)
- [Understanding CoS Buffer Configuration on page 5891](#)

---

## enhanced-transmission-selection

---

<b>Syntax</b>	<pre>enhanced-transmission-selection {     no-auto-negotiation;     no-recommendation-tlv;     recommendation-tlv {         no-auto-negotiation;     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx interface</a> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p>
<b>Options</b>	<p><b>no-auto-negotiation</b>—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p><b>no-recommendation-tlv</b>—Disable automatic negotiation of the ETS Recommendation TLV</p> <p><b>recommendation-tlv</b>—Enable automatic negotiation of ETS Recommendation TLV</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li></ul>

## ether-type

---

<b>Syntax</b>	<code>ether-type <i>ether-type</i>;</code>
<b>Hierarchy Level</b>	[edit applications <a href="#">application</a> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See <a href="http://standards.ieee.org/develop/regauth/ethertype/eth.txt">http://standards.ieee.org/develop/regauth/ethertype/eth.txt</a> for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes.



**NOTE:** To create a FIP application, use the EtherType 0x8914.

---

<b>Options</b>	<i>type</i> —Identifier for the EtherType.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> </ul>

## exp

---

Syntax	<pre>exp classifier-name {     import (classifier-name   default);     forwarding-class class-name {         loss-priority level {             code-points [ aliases ] [ bit-patterns ];         }     } }</pre>
Rewrite Rule Configuration	<pre>exp rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {         loss-priority level {             code-point [ aliases ] [ bit-patterns ];         }     } }</pre>
Global Classifier Association with Interfaces	<pre>exp classifier-name;</pre>
Hierarchy Level	<pre>[edit class-of-service classifiers], [edit class-of-service rewrite-rules] [edit class-of-service system-defaults classifiers]</pre>
Release Information	Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Define the EXP code point mapping that is applied to MPLS packets. EXP classifiers are not applied to any traffic except MPLS traffic. EXP classifiers are applied only to interfaces that are configured as <b>family mpls</b> (for example, <b>set interfaces xe-0/0/35 unit 0 family mpls</b>.)</p> <p>You can configure as many EXP classifiers as you want. However, the switch uses only one EXP classifier as a global MPLS classifier on all interfaces. You specify the global EXP classifier in the <b>[edit class-of-service system-defaults]</b> hierarchy.</p>
Options	<b>classifier-name</b> —Name of the EXP classifier.
Required Privilege Level	<b>interfaces</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Global MPLS EXP Classifier on page 4479</a></li><li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480</a></li><li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li></ul>

## explicit-congestion-notification

<b>Syntax</b>	explicit-congestion-notification;
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service schedulers</a> <i>scheduler-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.
<b>Description</b>	<p>Enable explicit congestion notification (ECN) on the output queue (forwarding class) or output queues (forwarding classes) mapped to the scheduler. ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality.</p> <p>A weighted random early detection (WRED) packet drop profile must be applied to the output queues on which ECN is enabled. ECN uses the WRED drop profile thresholds to mark packets when the output queue experiences congestion.</p> <p>ECN reduces packet loss by forwarding ECN-capable packets during periods of network congestion instead of dropping those packets. (TCP notifies the network about congestion by dropping packets.) During periods of congestion, ECN marks ECN-capable packets that egress from congested queues. When the receiver receives an ECN packet that is marked as experiencing congestion, the receiver echoes the congestion state back to the sender. The sender then reduces its transmission rate to clear the congestion.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring ECN on page 6090</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li> <li>• <a href="#">Understanding CoS Explicit Congestion Notification on page 5926</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li> </ul>


## fill-level

---

Syntax	fill-level <i>low-value</i> fill-level <i>high-value</i> ;
Hierarchy Level	[edit <a href="#">class-of-service drop-profiles profile-name interpolate</a> ]
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<p>When configuring random early detection (RED), map the fullness of a queue to a packet <a href="#">drop-probability</a> value. You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs by specifying a low <b>fill-level</b> value at which packets begin to drop (the drop probability is zero until the queue reaches this level of fullness) and a high <b>fill-level</b> value at which packets drop at the highest drop probability. As the queue fills from the low fill level to the high fill level, the rate of packet drop increases in a linear pattern from zero to the high drop probability.</p>
Options	<p><b>low-value</b>—Fullness of the queue before packets begin to drop, expressed as a percentage. The low value must be less than the high value.</p> <p><b>Range:</b> 0 through 100</p> <p><b>high-value</b>—Fullness of the queue before it reaches the maximum drop probability. If the queue fills beyond the fill level high value, all packets drop. The high value must be greater than the low value.</p> <p><b>Range:</b> 0 through 100</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li></ul>



## flow-control

<b>Syntax</b>	(flow-control   no-flow-control);
<b>Hierarchy Level</b>	[edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p>
	<div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>
	<ul style="list-style-type: none"> <li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li> <li>• <b>no-flow-control</b>—Disable flow control.</li> </ul>
<b>Default</b>	Flow control is disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">configured-flow-control on page 2637</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 2586</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

## flow-control-queue (Output Congestion Notification)

---

<b>Syntax</b>	<code>flow-control-queue [ <i>queue</i>   <i>list-of-queues</i> ];</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> output ieee-802.1 code-point <i>code-point-bits</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	<p>Specify one or more output queues to pause, to support priority-based flow control (PFC). The specified queues pause when the interface receives a PFC frame with a matching IEEE 802.1p code point.</p>
<b>Default</b>	<p>Queue 3 (mapped to the fcoe forwarding class) and queue 4 (mapped to the no-loss forwarding class) are programmed as flow control queues to pause. No other output queues are programmed to pause by default.</p> <p>If you configure flow control queues explicitly, only the queues that you specify are programmed to pause. The explicit flow control queue to pause configuration overrides the default setting, so the queues paused in the default configuration are no longer paused by default.</p> <p>For example, if you configure queue 2 as a flow control queue, then queue 2 pauses when congestion occurs, but queues 3 and 4 do not pause because they were not explicitly specified. To enable pause on output queues 2, 3, and 4, you must explicitly configure all three of the queues as flow control queues.</p> <p>The same behavior applies to the IEEE 802.1p code points (priorities) on which PFC is enabled. By default, priorities 3 (011) and 4 (100) are enabled for PFC pause. If you explicitly configure flow control queues to pause, you must also explicitly configure pause for each priority (code point) that you want to pause, because the explicit configuration overrides the default configuration.</p>
<b>Options</b>	[ <i>queue</i>   <i>list-of-queues</i> ]—The output queue or a list of output queues to pause.
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li></ul>

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837](#)

## forwarding-class

---

<b>List of Syntax</b>	<a href="#">Syntax (Classifier) on page 6242</a> <a href="#">Syntax (Rewrite Rule) on page 6242</a> <a href="#">Syntax (Scheduler Map) on page 6242</a> <a href="#">Syntax (Interface) on page 6242</a>
<b>Syntax (Classifier)</b>	<pre>forwarding-class class-name {     loss-priority level {         code-points [ aliases ] [ bit-patterns ];     } }</pre>
<b>Syntax (Rewrite Rule)</b>	<pre>forwarding-class class-name {     loss-priority level {         code-point [ aliases ] [ bit-patterns ];     } }</pre>
<b>Syntax (Scheduler Map)</b>	<pre>forwarding-class class-name {     scheduler scheduler-name; }</pre>
<b>Syntax (Interface)</b>	<pre>forwarding-class class-name;</pre>
<b>Hierarchy Level (Classifier)</b>	[edit <a href="#">class-of-service classifiers</a> (dscp   dscp-ipv6   ieee-802.1) classifier-name],
<b>Hierarchy Level (Rewrite Rule)</b>	[edit <a href="#">class-of-service rewrite-rules</a> ] (dscp   dscp-ipv6   ieee-802.1) rewrite-name],
<b>Hierarchy Level (Scheduler Map)</b>	[edit <a href="#">class-of-service scheduler-maps</a> map-name],
<b>Hierarchy Level (Interface)</b>	[edit <a href="#">class-of-service interfaces</a> interface-name unit logical-unit-number]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure forwarding class name and option values (classifier configuration), map rewrite rules to forwarding classes (rewrite rules), map forwarding classes to schedulers (scheduler maps), or map forwarding classes to logical interfaces (interfaces).
<b>Options</b>	<b>class-name</b> —Name of the forwarding class.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 6075</a></li><li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li></ul>

- [Defining CoS Rewrite Rules on page 6182](#)
- [Understanding CoS Forwarding Classes on page 5830](#)
- [Understanding CoS Rewrite Rules on page 5914](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)

## forwarding-class (Host Outbound Traffic)

<b>Syntax</b>	<code>forwarding-class <i>class-name</i>;</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service</code> <code>host-outbound-traffic</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Define forwarding class name for outbound host traffic (traffic generated by the Routing Engine).
<b>Options</b>	<i>class-name</i> —Name of the forwarding class.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 6172</a></li> <li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806</a></li> </ul>

## forwarding-class-set

<b>Syntax</b>	<code>forwarding-class-set <i>forwarding-class-set-name</i> {     <code>output-traffic-control-profile</code> <i>profile-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service</code> <code>interfaces</code> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Apply a previously defined forwarding class set to an output traffic control profile.
<b>Options</b>	<i>forwarding-class-set-name</i> —Name of the forwarding class set.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li> <li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5835</a></li> </ul>

## forwarding-class-sets

---

<b>Syntax</b>	<code>forwarding-class-sets <i>forwarding-class-set-name</i> {     class <i>class-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Assign forwarding classes to forwarding class sets (priority groups).
<b>Options</b>	<p><i>forwarding-class-set-name</i>—Name of the forwarding class set.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6078</a></li><li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5835</a></li></ul>

## forwarding-classes

**Syntax**

```
forwarding-classes {
  class {
    class-name {
      queue-num queue-number <no-loss>;
    }
  }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Map one or more forwarding classes to a single queue. You can configure up to 12 forwarding classes (8 unicast forwarding classes on queues 0 through 7 and 4 multidestination forwarding classes on queues 8 through 11) and map them to queues. You can map multiple forwarding classes to a single queue using the **class** statement. All forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if



you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

---

<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<code>interfaces</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 6075</a></li><li>• <a href="#">Understanding CoS Forwarding Classes on page 5830</a></li></ul>



## guaranteed-rate

<b>Syntax</b>	<code>guaranteed-rate (rate  percent <i>percentage</i>);</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service traffic-control-profiles</code> <i>traffic-control-profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a guaranteed minimum rate of transmission for a traffic control profile. The sum of the guaranteed rates of all of the forwarding class sets (priority groups) on a port should not exceed the total port bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group (forwarding class set) can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.
	<div>  <p><b>NOTE:</b> You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.</p> </div>
<b>Default</b>	If you do not specify a guaranteed rate, the guaranteed rate is zero (0) and there is no minimum guaranteed bandwidth.
	<div>  <p><b>NOTE:</b> If you do not configure a guaranteed rate for a traffic control profile, the queues that belong to any forwarding class set (priority group) that uses that traffic control profile cannot have a configured transmit rate. The result is that there is no minimum guaranteed bandwidth for those queues and that those queues can be starved during periods of congestion.</p> </div>
<b>Options</b>	<p><b>percent <i>percentage</i></b>—Minimum percentage of transmission capacity allocated to the forwarding class set or logical interface.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Minimum transmission rate allocated to the forwarding class set or logical interface, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 10,000,000,000 bps</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
  - [Understanding CoS Traffic Control Profiles on page 5880](#)
  - [output-traffic-control-profile on page 6264](#)

---

## host-outbound-traffic

---

<b>Syntax</b>	<pre>host-outbound-traffic {     forwarding-class <i>class-name</i>;     dscp-code-point <i>code-point</i>; }</pre>
<b>Hierarchy Level</b>	[edit class-of-service]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Allow queue selection for traffic generated by the Routing Engine (host). The selected queue must be configured properly. You can also configure specific DSCP code point bits for the type of service (ToS) field of the generated packets. This configuration does not affect transit packets or incoming packets. This is a global configuration that only affects packets originating on the Routing Engine. If you do not configure an output queue for host outbound traffic, the switch uses the default queue mapping.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 6172</a></li><li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 5806</a></li></ul>

## ieee-802.1

<b>List of Syntax</b>	<a href="#">Syntax (Classifier) on page 6249</a> <a href="#">Syntax (Code-Point Alias) on page 6249</a> <a href="#">Syntax (Multidestination Classifier) on page 6249</a> <a href="#">Syntax (Interface Classifier Association) on page 6249</a> <a href="#">Syntax (Rewrite Rule) on page 6249</a>
<b>Syntax (Classifier)</b>	<pre> ieee-802.1 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>
<b>Syntax (Code-Point Alias)</b>	<pre> ieee-802.1 alias-name bit-pattern;</pre>
<b>Syntax (Multidestination Classifier)</b>	<pre> ieee-802.1 classifier-name;</pre>
<b>Syntax (Interface Classifier Association)</b>	<pre> ieee-802.1 (classifier-name   default);</pre>
<b>Syntax (Rewrite Rule)</b>	<pre> ieee-802.1 rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>
<b>Hierarchy Level (Classifier)</b>	[edit <a href="#">class-of-service classifiers</a> ],
<b>Hierarchy Level (Code-Point Alias)</b>	[edit <a href="#">class-of-service code-point-aliases</a> ],
<b>Hierarchy Level (Multidestination Classifier)</b>	[edit <a href="#">class-of-service multi-destination classifiers</a> ],
<b>Hierarchy Level (Interface Classifier Association)</b>	[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ], [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],
<b>Hierarchy Level (Rewrite Rule)</b>	[edit <a href="#">class-of-service rewrite-rules</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.

<b>Description</b>	Configure an IEEE 802.1 classifier, configure an IEEE 802.1 code-point alias, apply a fixed IEEE 802.1 classifier to an interface, or apply an IEEE-802.1 rewrite rule.
<b>Options</b>	<i>classifier-name</i> —Name of the classifier.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 6159</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 5914</a></li></ul>

---

## ieee-802.1 (Input Congestion Notification)

---

<b>Syntax</b>	<pre>ieee-802.1 {     code-point [<i>code-point-bits</i>] {         pfc {             mru <i>mru-value</i>;         }     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service congestion-notification-profile</a> <i>profile-name</i> <a href="#">input</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point.
<b>Options</b>	The statements are described separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5606</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li></ul>

## ieee-802.1 (Output Congestion Notification)


<b>Syntax</b>	<pre> ieee-802.1 {   code-point [ code-point-bits ] {     flow-control-queue [ queue   list-of-queues ];   } } </pre>
<b>Hierarchy Level</b>	[edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>output</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point on output queues.
<b>Options</b>	The statements are described separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837</a></li> </ul>

## import

---

<b>Syntax</b>	<code>import (<i>import</i>   default);</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> ], [edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a default or previously defined classifier.
<b>Options</b>	<p><b><i>import</i></b>—Name of the classifier mapping configured at the <a href="#">[edit class-of-service classifiers]</a> hierarchy level.</p> <p><b>default</b>—Default classifier mapping.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p><a href="#">interfaces</a>—To view this statement in the configuration.</p> <p><a href="#">interface-control</a>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 5914</a></li></ul>

## ingress (Buffer Configuration)

<b>Syntax</b>	<pre>ingress {   <b>buffer-partition</b> (lossless   lossless-headroom   lossy) {     percent <i>percent</i>;   }   percent <i>percent</i>; }</pre>
<b>Hierarchy Level</b>	[edit <b>class-of-service shared-buffer</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	<p>Configure the global shared buffer pool allocation for ingress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for buffers such as dedicated buffers (buffers allocated permanently to ports) and headroom buffers (buffers that help prevent packet loss on lossless flows).</p> <p>The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared ingress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)</p>
	<div>  <p><b>CAUTION:</b> Changing the buffer configuration is a disruptive event. Traffic stops on <i>all</i> ports until buffer reprogramming is complete.</p> </div>
	<p>You can also partition the shared buffer pool to adjust the ingress buffer allocations for different mixes of network traffic using the <b>buffer-partition</b> statement.</p>
<b>Default</b>	The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)
<b>Options</b>	<p><b>percent <i>percent</i></b>—Percentage of available ingress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116</a></li> </ul>

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic](#) on page 6122
- [Configuring Global Ingress and Egress Shared Buffers](#) on page 6179
- [Understanding CoS Buffer Configuration](#) on page 5891

## input (Congestion Notification)

---

**Syntax**

```
input {  
  ieee-802.1 {  
    code-point [code-point-bits] {  
      pfc {  
        mru mru-value;  
      }  
    }  
  }  
  cable-length cable-length-value;  
}
```

**Hierarchy Level** [edit [class-of-service congestion-notification-profile](#) *profile-name*]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure priority-based flow control (PFC) on incoming traffic.

**Options** The remaining statements are explained separately.

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS PFC for FCoE Traffic](#) on page 5606
- [Configuring CoS PFC \(Congestion Notification Profiles\)](#) on page 6174
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\)](#) on page 5559



## interface (DCBX)

<b>Syntax</b>	<pre> interface (<i>interface-name</i>   all) {   disable;   application-map <i>application-map-name</i>;   applications {     no-auto-negotiation;   }   enhanced-transmission-selection {     no-auto-negotiation;     no-recommendation-tlv;     recommendation-tlv {       no-auto-negotiation;     }   }   dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);   priority-flow-control {     no-auto-negotiation;   } } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p><b>Mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>
<b>Description</b>	Configure DCBX properties on an interface.
<b>Options</b>	<p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>

## interfaces (Class of Service)

---

**Syntax**

```
interfaces {  
  interface-name {  
    congestion-notification-profile profile-name {  
    }  
    forwarding-class lossless-forwarding-class-name;  
    forwarding-class-set forwarding-class-set-name {  
      output-traffic-control-profile profile-name;  
    }  
    rewrite-value {  
      input {  
        ieee-802.1 {  
          code-point code-point-bits;  
        }  
      }  
    }  
  }  
  unit logical-unit-number {  
    classifiers {  
      (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);  
    }  
    forwarding-class class-name;  
    rewrite-rules {  
      (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);  
    }  
  }  
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure interface-specific CoS properties for incoming packets.

**Options** *interface-name*—Name of the interface.  
  
The statements are explained separately.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Unicast Classifiers on page 6066](#)
- [Example: Configuring Forwarding Classes on page 6075](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Assigning CoS Components to Interfaces on page 6185](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Defining CoS Rewrite Rules on page 6182](#)

- [Interfaces Overview on page 2389](#)

## interpolate

---

<b>Syntax</b>	<pre> interpolate {   fill-level <i>low-value</i> fill-level <i>high-value</i>;   drop-probability 0 drop-probability <i>high-value</i>; } </pre>
<b>Hierarchy Level</b>	[edit class-of-service <b>drop-profiles</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Specify values for interpolating the relationship between queue fill level and drop probability.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li> </ul>

## loss-priority (Classifiers)

---

<b>Syntax</b>	<code>loss-priority <i>level</i> {     <i>code-points</i> [ <i>aliases</i> ] [ <i>bit-patterns</i> ]; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> <a href="#">forwarding-class</a> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure packet loss priority value for a specific set of code-point aliases and bit patterns.
<b>Options</b>	<p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p><code>interfaces</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 6066</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li></ul>

## loss-priority (Drop Profiles)

---

<b>Syntax</b>	<code>loss-priority <i>level</i> <i>protocol protocol</i> <i>drop-profile profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service schedulers scheduler-name drop-profile-map</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure packet loss priority value for a drop profile mapped to a system drop profile.
<b>Options</b>	<p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Packet has low loss priority.</li> <li>• <b>medium-high</b>—Packet has medium-high loss priority.</li> <li>• <b>high</b>—Packet has high loss priority.</li> </ul> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 6073</a></li> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li> </ul>

## loss-priority (Rewrite Rules)

---

<b>Syntax</b>	<code>loss-priority <i>level</i> {     <i>code-point</i> (<i>alias</i>   <i>bit-pattern</i>); }</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service rewrite-rules</code> ( <code>dscp</code>   <code>dscp-ipv6</code>   <code>ieee-802.1</code> ) <i>rewrite-name</i> <i>forwarding-class</i> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and loss priority. Packets that match the forwarding class and loss priority are rewritten with the rewrite code-point alias or bit pattern.
<b>Options</b>	<p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p><code>interfaces</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 5914</a></li></ul>


---

## multi-destination

---

<b>Syntax</b>	<pre>multi-destination {   classifiers {     (dscp   ieee-802.1) classifier-name;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define a multicast CoS behavior aggregate (BA) classifier.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers on page 6069</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li><li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li></ul>

## mrp

<b>Syntax</b>	<code>mrp <i>mrp-value</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> input ieee-802.1 code-point <i>code-point-bits</i> pfc</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure the maximum receive unit (MRU) of the interface in bytes (incoming packet sizes must be less than or equal to the MRU, or the packets are dropped). The system uses the MRU and the cable length to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The lower the MRU and the shorter the cable length, the less headroom buffer space is required for PFC.
<div>  <p><b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces <i>interface-name</i></a>] hierarchy level.</p> </div>	
<b>Default</b>	For priority 3 traffic, the default MRU value is 2500 bytes.  For priority 4 traffic, the default MRU value is 9612 bytes.
<b>Options</b>	<b><i>mrp-value</i></b> —Value of the maximum packet receive unit size in bytes (generally from 1500 to 9216 bytes, but there is no configuration restriction).
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837</a></li> </ul>



## output (Congestion Notification)

<b>Syntax</b>	<pre>output {   ieee-802.1 {     code-point [code-point-bits] {       flow-control-queue [queue   list-of-queues];     }   } }</pre>
<b>Hierarchy Level</b>	[edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Configure priority-based flow control (PFC) on output queues.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837</a></li> </ul>

## output-traffic-control-profile

---

<b>Syntax</b>	<code>output-traffic-control-profile <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">forwarding-class-set</a> <i>forwarding-class-set-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Apply an output traffic scheduling and shaping profile to a forwarding class set (priority group).
<b>Options</b>	<i>profile-name</i> —Name of the traffic-control profile to apply to the specified forwarding class set.
<b>Required Privilege Level</b>	<code>interfaces</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6094</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 5880</a></li></ul>

## pfc (Input Congestion Notification)

<b>Syntax</b>	<code>pfc {     <b>mru</b> <i>mru-value</i>; }</code>
<b>Hierarchy Level</b>	[edit <b>class-of-service</b> <b>congestion-notification-profile</b> <i>profile-name</i> <b>input</b> <b>ieee-802.1</b> <b>code-point</b> <i>code-point-bits</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Enable and configure ingress interface priority-based flow control (PFC).
<b>Options</b>	The remaining statement is explained separately.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 5837</a></li> </ul>

## policy-options

---

**Syntax**    `policy-options`

```
  application-maps application-map-name {
    application application-name {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name match-type <actions>;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
    }
  }
```

**Hierarchy Level**    [edit]

**Release Information**    Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.1 for the EX Series.


**Description**    Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level**    storage—To view this statement in the configuration.  
storage-control—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 5672](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 5595](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 5589](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

## priority (Schedulers)


<b>Syntax</b>	<code>priority <i>priority</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service schedulers</a> <i>scheduler-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the packet-scheduling drop priority value.
<b>Options</b>	<p><b><i>priority</i></b>—It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Scheduler has low priority.</li> <li>• <b>strict-high</b>—Scheduler has strict high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.</li> </ul>
<div>  <p><b>NOTE:</b> We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.</p> </div>	
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li> </ul>

## priority-flow-control

---

<b>Syntax</b>	<pre>priority-flow-control {     no-auto-negotiation; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx interface</a> (all   <i>interface-name</i> )]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 11.3 for EX Series switches.
<b>Description</b>	Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.
<b>Options</b>	<b>no-auto-negotiation</b> —Disable automatic negotiation of PFC.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li><li>• <a href="#">Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5606</a></li><li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li><li>• <a href="#">Understanding Priority-Based Flow Control</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li></ul>

## protocol (Applications)

<b>Syntax</b>	<code>protocol (tcp   udp);</code>
<b>Hierarchy Level</b>	[edit applications <b>application</b> <i>application-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	Networking protocol type, which combines with <b>destination-port</b> to identify an application type.
<div>  <b>NOTE:</b> To create an application for iSCSI, use the protocol <b>tcp</b> with the destination port number <b>3260</b>. </div>	
<b>Options</b>	<b>tcp</b> —Transmission Control Protocol  <b>udp</b> —User Datagram Protocol
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 5672</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 5589</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>

## protocol (Drop Profile Map)

---

<b>Syntax</b>	<code>protocol protocol <b>drop-profile</b> profile-name;</code>
<b>Hierarchy Level</b>	[edit <b>class-of-service schedulers scheduler-name drop-profile-map loss-priority</b> (low   medium-high   high)]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the protocol type for the specified drop profile.
<b>Options</b>	<p><b>protocol</b>—Type of protocol. The protocol can be:</p> <ul style="list-style-type: none"><li>• <b>any</b>—Accept any protocol type.</li></ul> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Example: Configuring Drop Profile Maps on page 6073</a></li><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 6071</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 5909</a></li></ul>



## queue-num

<b>Syntax</b>	<code>queue-num <i>queue-number</i> &lt;no-loss&gt;;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service forwarding-classes class</a> <i>class-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. No-loss option introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Map a forwarding class to an output queue number. Optionally, configure the forwarding class as a lossless forwarding class.

You can map some or all of the eight unicast forwarding classes to a unicast queue (0 through 7) or some or all of the four multdestination (multicast, broadcast, destination lookup fail) forwarding classes to the same multdestination queue (8 through 11), providing that you do not map one forwarding class to more than one queue. The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.

<b>Options</b>	<b><i>queue-number</i></b> —Number of the CoS unicast queue (0 through 7) or the CoS multidestination queue (8 through 11).  <b><i>no-loss</i></b> —Optional packet drop attribute keyword to configure the forwarding class as lossless.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 6075</a></li><li>• <a href="#">Understanding CoS Forwarding Classes on page 5830</a></li></ul>

---

## recommendation-tlv


---

<b>Syntax</b>	<pre>recommendation-tlv {     no-auto-negotiation; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">protocols dcbx interface interface-name enhanced-transmission-selection</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	Enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.)
<b>Default</b>	DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.
<b>Options</b>	<b><i>no-auto-negotiation</i></b> —Disable sending of the ETS recommendation TLV.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 5724</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li></ul>

## rewrite-rules

<b>List of Syntax</b>	<a href="#">Syntax (Rewrite Rule Configuration) on page 6273</a> <a href="#">Syntax (Rewrite Rule Association with Interface) on page 6273</a>
<b>Syntax (Rewrite Rule Configuration)</b>	<pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1   exp) rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {       loss-priority priority code-point (alias   bits);     }   } }</pre>
<b>Syntax (Rewrite Rule Association with Interface)</b>	<pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1   exp) rewrite-name; }</pre>
<b>Hierarchy Level (Rewrite Rule Configuration)</b>	[edit <a href="#">class-of-service</a> ],
<b>Hierarchy Level (Rewrite Rule Association with Interface)</b>	[edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. EXP statement introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	<p>Configure rewrite rules that map traffic to code points when traffic exits the system, and apply the rewrite rules to a specific interface.</p> <p>MPLS EXP rewrite rules can only be bound to logical interfaces, not to physical interfaces. You can configure as many EXP rewrite rules as you want, but you can use only 16 EXP rewrite rules on switch interfaces at any given time.</p>
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 6182</a></li> <li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 5914</a></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419</a></li> </ul>

## rx-buffers

<b>Syntax</b>	rx-buffers (on   off);
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p>
	<p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p>
<b>Default</b>	Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.
<b>Options</b>	<b>on   off</b> —Enable or disable an interface to generate and send Ethernet PAUSE messages.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2659</a></li> <li>• <a href="#">tx-buffers on page 2742</a></li> </ul>

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## scheduler

---

<b>Syntax</b>	<code>scheduler <i>scheduler-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Map a scheduler to a forwarding class using a scheduler map.
<b>Options</b>	<i>scheduler-name</i> —Name of the scheduler to map to the forwarding class.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li> </ul>

## scheduler-map

---

<b>Syntax</b>	<code>scheduler-map <i>map-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service traffic-control-profiles <i>traffic-control-profile-name</i></a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Associate a scheduler map with a traffic control profile.
<b>Options</b>	<i>map-name</i> —Name of the scheduler map.
<b>Required Privilege Level</b>	interfaces—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6094</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li> <li>• <a href="#">Understanding CoS Traffic Control Profiles on page 5880</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li> </ul>

## scheduler-maps

---



<b>Syntax</b>	<pre>scheduler-maps {   map-name {     forwarding-class class-name scheduler scheduler-name;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a scheduler map name to map a scheduler configuration to a forwarding class.
<b>Options</b>	<p><i>map-name</i>—Name of the scheduler map.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li><li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li></ul>

## schedulers

<b>Syntax</b>	<pre> schedulers {   scheduler-name {     buffer-size (percent <i>percentage</i>   remainder);     drop-profile-map loss-priority (low   medium-high   high) protocol <i>protocol</i> drop-profile       drop-profile-name;     explicit-congestion-notification;     priority <i>priority</i>;     shaping-rate (<i>rate</i>   percent <i>percentage</i>);     transmit-rate (percent <i>percentage</i>);   } } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify scheduler name and parameter values such minimum bandwidth ( <b>transmit-rate</b> ), maximum bandwidth ( <b>shaping-rate</b> ), and priority ( <b>priority</b> ).
<b>Options</b>	<p><b>scheduler-name</b> —Name of the scheduler.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 6073</a></li> <li>• <a href="#">Example: Configuring ECN on page 6090</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li> </ul>

## shaping-rate

---

<b>Syntax</b>	<code>shaping-rate (rate   percent <i>percentage</i>);</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service schedulers <i>scheduler-name</i></code> ], [edit <code>class-of-service traffic-control-profiles <i>profile-name</i></code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure the shaping rate. The shaping rate throttles the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class set. You specify the maximum bandwidth for a queue by using a scheduler map to associate a forwarding class (queue) with a scheduler that has a configured shaping rate. You specify the maximum bandwidth for a forwarding class set by setting the shaping rate for a traffic control profile, and then applying the traffic control profile and a forwarding class set to an interface.</p> <p>We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate (the minimum guaranteed bandwidth for a queue) or the configured guaranteed rate (the minimum guaranteed bandwidth for a forwarding class set).</p> <div><div></div><div><p><b>NOTE:</b> When you set the maximum bandwidth (shaping-rate value) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured shaping-rate value.</p></div></div> <div><div></div><div><p><b>NOTE:</b> We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.</p></div></div>
<b>Default</b>	If you do not configure a shaping rate, the default shaping rate is 100 percent (all of the available bandwidth), which is the equivalent of no rate shaping.
<b>Options</b>	<p><b>percent <i>percentage</i></b>—Shaping rate as a percentage of the available interface bandwidth. <b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Peak (maximum) rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). <b>Range:</b> 1000 through 10,000,000,000 bps</p>



**Required Privilege** interfaces—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Understanding CoS Traffic Control Profiles on page 5880](#)

## shared-buffer

```
Syntax  shared-buffer {
        egress {
            buffer-partition (lossless | lossy | multicast) {
                percent percent
            }
            percent percent;
        }
        ingress {
            percent percent;
            buffer-partition (lossless | lossless-headroom | lossy) {
                percent percent
            }
        }
    }
```

Hierarchy Level [edit [class-of-service](#)]

Release Information Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Configure the global shared buffer pool allocation to ports. Shared buffers are a pool of buffer space that the system can allocate dynamically across all of its ports as memory space is needed. Some buffer space is reserved for dedicated buffers (buffers allocated permanently to ports), headroom buffers (buffers that help prevent packet loss on lossless flows), and other buffers.

Configure the way the system uses the available (user-configurable) buffer space by setting the **shared-buffer** percentage for the ingress buffer pool and for the egress buffer pool.

The percentage you specify is the percentage of available buffer space allocated to the global shared ingress buffer pool or to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until the buffer reprogramming is complete.

You can also partition the ingress shared buffer pool and the egress shared buffer pool to adjust the buffer allocations for different mixes of network traffic (best-effort, lossless, multicast) using the **buffer-partition** statement.

**Options** The statements are explained separately.

<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122</a></li> <li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 6179</a></li> <li>• <a href="#">Understanding CoS Buffer Configuration on page 5891</a></li> </ul>

## system-defaults

---

<b>Syntax</b>	<pre>system-defaults {   classifiers exp classifier-name; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Configure the global EXP classifier used on all interfaces to classify MPLS traffic.</p> <p>Although you can configure as many EXP classifiers as you want, the switch uses only one EXP classifier as a global MPLS classifier on all interfaces. All switch interfaces use the EXP classifier specified as the system default to classify MPLS traffic.</p>
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Global MPLS EXP Classifier on page 4479</a></li> <li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers on page 4480</a></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 4419</a></li> <li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li> </ul>

## traceoptions (Class of Service)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt;     &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;flag-modifier&gt;;     no-remote-trace }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">class-of-service</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set class-of-service (CoS) tracing options.



**NOTE:** The `traceoptions` statement is not supported on the QFabric system.

**Default** Traceoptions is disabled.

**Options** **file *filename***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***. The traceoption output continues in a second trace file named ***trace-file.1***. When ***trace-file.1*** reaches its maximum size, output continues in a third file named ***trace-file.2***, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

**Range:** 2 through 1000 files

**Default:** 1 trace file

**flag**—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Trace all operations.
- **asynch**—Trace asynchronous configuration processing.
- **chassis-scheduler**—Trace chassis stream scheduler processing.
- **cos-adjustment**—Trace CoS rate adjustments.
- **dynamic**—Trace dynamic CoS functions.
- **hardware-database**—Trace the chassis hardware database related processing.
- **init**—Trace initialization events.

- **performance-monitor**—Trace performance monitor counters.
- **process**—Trace configuration processing.
- **restart**—Trace restart processing.
- **route-socket**—Trace route-socket events.
- **show**—Trace show command servicing.
- **snmp**—Trace SNMP-related processing.
- **util**—Trace utilities.

The following are the global tracing options:

- **all**—Perform all tracing operations
- **parse**—Trace parser processing.

**no-remote-trace**—(Optional) Disable remote tracing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.




<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

## traffic-control-profiles

---

<b>Syntax</b>	<pre>traffic-control-profiles <i>profile-name</i> {     <b>guaranteed-rate</b> (<i>rate</i>  percent <i>percentage</i>);     <b>scheduler-map</b> <i>map-name</i>;     <b>shaping-rate</b> (<i>rate</i>  percent <i>percentage</i>); }</pre>
<b>Hierarchy Level</b>	[edit <b>class-of-service</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.
<b>Options</b>	<p><b>profile-name</b>—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6094</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6078</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 6185</a></li><li>• <a href="#">output-traffic-control-profile on page 6264</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 5880</a></li></ul>

## transmit-rate

<b>Syntax</b>	<code>transmit-rate (rate   percent <i>percentage</i>);</code>
<b>Hierarchy Level</b>	[edit <code>class-of-service schedulers <i>scheduler-name</i></code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the minimum transmission rate or percentage for a queue (forwarding class) scheduler. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.
	<p> <b>NOTE:</b> The <code>transmit-rate</code> setting works only if you also configure the <code>guaranteed-rate</code> in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the guaranteed rate, the minimum guaranteed rate for individual queues that you set using the <code>transmit-rate</code> statement does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate.</p> <p> <b>NOTE:</b> You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.</p> <p> <b>NOTE:</b> For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.</p>
<b>Default</b>	If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 11 are:

Queue Number	Default Minimum Guaranteed Bandwidth
0 (best-effort)	5 %
1	0
2	0
3 (fcoe)	35 %
4 (no-loss)	35 %
5	0
6	0
7 (network control)	5 %
8 (mcast)	20 %
9	0
10	0
11	0

Configure schedulers if you want to change the minimum guaranteed bandwidth and other queue characteristics.

**Options** *rate*—Minimum transmission rate for the queue, in bps. You can specify a value in bits-per-second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

**Range:** 1000 through 10,000,000,000 bps

**percent** *percentage*—Minimum percentage of transmission capacity allocated to the queue. A percentage of zero means that there is no minimum bandwidth guarantee for the queue.

**Range:** 0 through 100 percent


**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)



## tx-buffers

<b>Syntax</b>	tx-buffers (on   off);
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p>
<div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>	
<b>Default</b>	Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.
<b>Options</b>	on   off—Enable or disable an interface to respond to an Ethernet PAUSE message.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 2659</a></li> <li>• <a href="#">rx-buffers on page 2728</a></li> </ul>

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 6177](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 6178](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

---

## unit

---

**Syntax**     `unit logical-unit-number {  
                  classifiers {  
                    (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);  
                  }  
                  forwarding-class class-name;  
                  rewrite-rules {  
                    (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);  
                  }  
                  }`

**Hierarchy Level**     [edit **class-of-service interfaces** *interface-name*]

**Release Information**     Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**     Configure a logical interface on the physical device. You must configure a logical interface to use the physical device.

**Options**     *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,385

The remaining statements are explained separately.

**Required Privilege Level**     interfaces—To view this statement in the configuration.  
                                         interface-control—To add this statement to the configuration.

**Related Documentation**     • [Assigning CoS Components to Interfaces on page 6185](#)

# Administration

- [Routine Monitoring on page 6289](#)
- [Operational Commands on page 6295](#)

## Routine Monitoring

---

- [Monitoring CoS Classifiers on page 6289](#)
- [Monitoring CoS Forwarding Classes on page 6290](#)
- [Monitoring Interfaces That Have CoS Components on page 6291](#)
- [Monitoring CoS Rewrite Rules on page 6292](#)
- [Monitoring CoS Scheduler Maps on page 6293](#)
- [Monitoring CoS Value Aliases on page 6294](#)

## Monitoring CoS Classifiers

**Purpose** Display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

**Action** To monitor CoS classifiers in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier
```

To monitor a particular classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier name classifier-name
```

To monitor a particular type of classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier type classifier-type
```

**Meaning** [Table 566 on page 6289](#) summarizes key output fields for CoS classifiers.

**Table 566: Summary of Key CoS Classifier Output Fields**

Field	Values
Classifier	Name of a classifier.

Table 566: Summary of Key CoS Classifier Output Fields (*continued*)

Field	Values
Code point type	Type of classifier: <ul style="list-style-type: none"> <li>• <b>dscp</b>—All classifiers of the DSCP type.</li> <li>• <b>ieee-802.1</b>—All classifiers of the IEEE 802.1 type.</li> <li>• <b>ieee-mcast</b>—All classifiers of the IEEE 802.1 multicast type.</li> </ul>
Index	Internal index of the classifier.
Code point	DSCP or IEEE 802.1 code point value of the incoming packets, in bits. These values are used for classification.
Forwarding Class	Name of the forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch.
Loss Priority	Loss priority value that the classifier assigns to the incoming packet based on its code point value.

- Related Documentation**
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 6160](#)
  - [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 6162](#)

## Monitoring CoS Forwarding Classes

**Purpose** Use the monitoring functionality to view the current assignment of CoS forwarding classes to queue numbers on the system.

**Action** To monitor CoS forwarding classes in the CLI, enter the following CLI command:

```
user@switch> show class-of-service forwarding-class
```

**Meaning** [Table 567 on page 6291](#) summarizes key output fields for CoS forwarding classes.

Table 567: Summary of Key CoS Forwarding Class Output Fields

Field	Values
Forwarding Class	<p>Names of forwarding classes assigned to queue numbers. By default, the following unicast forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.</li> <li>• <b>fcoe</b>—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.</li> <li>• <b>no-loss</b>—Provides guaranteed delivery for TCP lossless traffic</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> <p>By default, the following multideestination forwarding class is assigned to queue 8:</p> <ul style="list-style-type: none"> <li>• <b>mcast</b>—Provides no special CoS handling of packets.</li> </ul>
Queue	<p>Queue number corresponding to the forwarding class name.</p> <p>By default, four queues (0, 3, 4, and 7) are assigned to unicast forwarding classes and one queue (8) is assigned to a multideestination forwarding class.</p>
No-Loss	<p>Packet drop attribute associated with each forwarding class:</p> <ul style="list-style-type: none"> <li>• Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion)</li> <li>• Enabled—The forwarding class is configured for lossless transport</li> </ul> <p><b>NOTE:</b> To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p>

- Related Documentation**
- [Defining CoS Forwarding Classes on page 6164](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)

## Monitoring Interfaces That Have CoS Components

**Purpose** Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

**Action** To monitor interfaces that have CoS components in the CLI, enter the command:

```
user@switch> show class-of-service interface
```

To monitor a specific interface in the CLI, enter the command:

```
user@switch> show class-of-service interface interface-name
```

**Meaning** [Table 568 on page 6292](#) summarizes key output fields for CoS interfaces.

Table 568: Summary of Key CoS Interfaces Output Fields

Field	Values
Physical interface	Name of a physical interface to which CoS components are assigned.
Index	Index of this interface or the internal index of a specific object.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Scheduler map	Name of the scheduler map associated with this interface.
Congestion-notification	Status of congestion notification (enabled or disabled).
Rewrite Input IEEE Code-point	(Fibre Channel NP_Port interfaces only) IEEE 802.1p code point (priority) the interface assigns to incoming Fibre Channel (FC) traffic when the interface encapsulates the FC traffic in Ethernet before forwarding it onto the FCoE network.
Logical Interface	Name of a logical interface on the physical interface to which CoS components are assigned.
Object	Category of an object—for example, <b>classifier</b> , <b>scheduler-map</b> , or <b>rewrite</b> .
Name	Name of the object—for example, <b>ba-classifier</b> .
Type	Type of the object—for example, <b>ieee8021p</b> for a classifier.

**Related Documentation** • [Assigning CoS Components to Interfaces on page 6185](#)

## Monitoring CoS Rewrite Rules

**Purpose** Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

**Action** To monitor CoS rewrite rules in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule
```

To monitor a particular rewrite rule in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule name rewrite-rule-name
```

To monitor a particular type of rewrite rule (for example, DSCP, DSCP IPv6, or IEEE-802.1) in the CLI, enter the CLI command:

```
user@switch> show class-of-service rewrite-rule type rewrite-rule-type
```

**Meaning** [Table 569 on page 6293](#) summarizes key output fields for CoS rewrite rules.

**Table 569: Summary of Key CoS Rewrite Rule Output Fields**

Field	Values
Rewrite rule	Name of the rewrite rule.
Code point type	Rewrite rule type: <ul style="list-style-type: none"> <li>• <b>dscp</b>—For IPv4 DiffServ traffic.</li> <li>• <b>dscp-ipv6</b>—For IPv6 Diffserv traffic.</li> <li>• <b>ieee-802.1</b>—For Layer 2 traffic.</li> </ul>
Index	Internal index for the rewrite rule.
Forwarding class	Name of the forwarding class that is used to determine CoS values for rewriting in combination with loss priority.  Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.
Loss priority	Level of loss priority that is used to determine CoS values for rewriting in combination with forwarding class.
Code point	Rewrite code point value.

**Related Documentation** • [Defining CoS Rewrite Rules on page 6182](#)

## Monitoring CoS Scheduler Maps

**Purpose** Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

**Action** To monitor CoS scheduler maps in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map
```

To monitor a specific scheduler map in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map scheduler-map-name
```

**Meaning** [Table 570 on page 6293](#) summarizes key output fields for CoS scheduler maps.

**Table 570: Summary of Key CoS Scheduler Maps Output Fields**

Field	Values
Scheduler map	Name of the scheduler map.
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.

Table 570: Summary of Key CoS Scheduler Maps Output Fields (*continued*)

Field	Values
Scheduler	Name of the scheduler.
Forwarding class	Names of the forwarding classes to which the scheduler is assigned.
Transmit rate	Configured transmit rate of the scheduler as a percentage of the total interface bandwidth.
Priority	<p>Scheduling priority of a queue:</p> <ul style="list-style-type: none"> <li>• <b>strict-high</b> or <b>high</b>—Packets in this queue are transmitted first. Only one queue can be configured as <b>strict-high</b> or <b>high</b>.</li> <li>• <b>low</b>—Packets in this queue are transmitted after packets in the <b>strict-high</b> queue.</li> </ul>
Drop Profiles	Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.
Loss Priority	Drop profile associated with each packet loss priority. You can configure different drop profiles for <b>low</b> , <b>medium-high</b> , and <b>high</b> loss priority traffic.
Protocol	Transport protocol of the drop profile for the particular priority.
Name	Name of the drop profile.

**Related Documentation** • [Defining CoS Queue Schedulers on page 6167](#)

## Monitoring CoS Value Aliases

**Purpose** Use the monitoring functionality to display information about the CoS value aliases that the system is currently using to represent DSCP and IEEE 802.1p code point bits.

**Action** To monitor CoS value aliases in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases
```

To monitor a specific type of code-point alias (for example, DSCP or IEEE 802.1) in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases ieee-802.1
```

**Meaning** [Table 571 on page 6295](#) summarizes key output fields for CoS value aliases.



Table 571: Summary of Key CoS Value Alias Output Fields

Field	Values
Code point type	Type of the CoS value: <ul style="list-style-type: none"> <li>• <b>dscp</b>—Examines Layer 3 packet headers for IP packet classification.</li> <li>• <b>ieee-802.1</b>—Examines Layer 2 packet headers for packet classification.</li> </ul>
Alias	Name given to a set of bits—for example, <b>af11</b> is a name for bits <b>001010</b> .
Bit pattern	Set of bits associated with the alias.

**Related Documentation** • [Defining CoS Code-Point Aliases on page 6159](#)

## Operational Commands

- [show class-of-service](#)
- [show class-of-service classifier](#)
- [show class-of-service code-point-aliases](#)
- [show class-of-service congestion-notification](#)
- [show class-of-service drop-profile](#)
- [show class-of-service forwarding-class](#)
- [show class-of-service forwarding-class-set](#)
- [show class-of-service forwarding-table](#)
- [show class-of-service forwarding-table classifier](#)
- [show class-of-service forwarding-table classifier mapping](#)
- [show class-of-service forwarding-table drop-profile](#)
- [show class-of-service forwarding-table rewrite-rule](#)
- [show class-of-service forwarding-table rewrite-rule mapping](#)
- [show class-of-service forwarding-table scheduler-map](#)
- [show class-of-service interface](#)
- [show class-of-service multi-destination](#)
- [show class-of-service rewrite-rule](#)
- [show class-of-service scheduler-map](#)
- [show class-of-service shared-buffer](#)
- [show class-of-service traffic-control-profile](#)
- [show dcbx](#)
- [show dcbx neighbors](#)
- [show interfaces queue](#)
- [show pfe filter hw summary](#)

- `show pfe next-hop`
- `show pfe route`
- `show pfe terse`
- `show pfe version`

## show class-of-service

<b>Syntax</b>	<b>show class-of-service</b>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the class-of-service (CoS) information.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring CoS Value Aliases on page 6294</a></li> <li>• <a href="#">Monitoring CoS Classifiers on page 6289</a></li> <li>• <a href="#">Monitoring CoS Forwarding Classes on page 6290</a></li> <li>• <a href="#">Monitoring Interfaces That Have CoS Components on page 6291</a></li> <li>• <a href="#">Monitoring CoS Scheduler Maps on page 6293</a></li> <li>• <a href="#">Monitoring CoS Rewrite Rules on page 6292</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of- service on page 6298</a>
<b>Output Fields</b>	Table 572 on page 6297 lists the output fields for the <b>show class-of-service</b> command. Output fields are listed in the approximate order in which they appear.

**Table 572: show class-of-service Output Fields**

Field Name	Field Description	Level of Output
<b>Forwarding class</b>	The forwarding class configuration: <ul style="list-style-type: none"> <li>• <b>Forwarding class</b>—Name of the forwarding class.</li> <li>• <b>ID</b>—Forwarding class ID.</li> <li>• <b>Queue</b>—Queue number.</li> </ul>	All levels
<b>Code point type</b>	The type of code-point alias: <ul style="list-style-type: none"> <li>• <b>dscp</b>—Aliases for DiffServ code point (DSCP) values.</li> <li>• <b>ieee-802.1</b>—Aliases for IEEE 802.1p values.</li> </ul>	All levels
<b>Alias</b>	Names given to CoS values.	All levels
<b>Bit pattern</b>	Set of bits associated with an alias.	All levels
<b>Classifier</b>	Name of the classifier.	All levels
<b>Code point</b>	Code-point values.	All levels
<b>Loss priority</b>	Loss priority assigned to specific CoS values and aliases of the classifier.	All levels
<b>Rewrite rule</b>	Name of the rewrite rule if one has been configured.	All levels

Table 572: show class-of-service Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Drop profile</b>	Name of the drop profile.	All levels
<b>Type</b>	Type of drop profile. QFX Series supports only the <b>discrete</b> type of drop-profile.	All levels
<b>Fill level</b>	Percentage of queue buffer fullness in a drop profile at which packets begin to drop during periods of congestion.	All levels
<b>Scheduler map</b>	Name of the scheduler map.	All levels
<b>Scheduler</b>	Name of the scheduler.	All levels
<b>Transmit rate</b>	Transmission rate of the scheduler.	All levels
<b>Buffer size</b>	Delay buffer size in the queue.	All levels
<b>Drop profiles</b>	Drop profiles configured for the specified scheduler.	All levels
<b>Protocol</b>	Transport protocol corresponding to the drop profile.	All levels
<b>Name</b>	Name of the drop profile.	All levels
<b>Queues supported</b>	Number of queues that can be configured on the interface.	All levels
<b>Queues in use</b>	Number of queues currently configured.	All levels
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Scheduler map</b>	Name of the scheduler map.	All levels
<b>Congestion-notification</b>	Enabled if a congestion notification profile is applied to the interface; disabled if no congestion notification profile is applied to the interface.	All levels
<b>Forwarding class set</b>	Name of the forwarding class set (priority group).	
<b>Index</b>	Internal index of an object.	All levels

## Sample Output

### show class-of- service

```

user@switch> show class-of-service
Forwarding class      ID      Queue
  best-effort         0        0
    fcoe              1        3
   no-loss            2        4
network-control      3        7
      mcast           8        8

Code point type: dscp

```

```

Alias          Bit pattern
af11           001010
af12           001100
...           ...

Code point type: ieee-802.1
Alias          Bit pattern
af11           100
...           ...

Classifier: dscp-default, Code point type: dscp, Index: 7
Code point    Forwarding class    Loss priority
000000        best-effort         low
000001        best-effort         low
...           ...                 ...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point    Forwarding class    Loss priority
000           best-effort         low
001           best-effort         low
010           best-effort         low
011           fcoe                low
100           no-loss             low
101           best-effort         low
110           network-control     low
111           network-control     low

Drop profile:<default-drop-profile>, Type: discrete, Index: 1
Fill level
100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 21
Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer
Limit: none,
Priority: low
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any      1      <default-drop-profile>
  Medium high   any      1      <default-drop-profile>
  High          any      1      <default-drop-profile>

Scheduler: <default-fcoe>, Forwarding class: fcoe, Index: 50
Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer
Limit: none,
Priority: low
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any      1      <default-drop-profile>
  Medium high   any      1      <default-drop-profile>
  High          any      1      <default-drop-profile>

Scheduler: <default-noloss>, Forwarding class: no-loss, Index: 51
Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer
Limit: none,
Priority: low

```

```
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       1      <default-drop-profile>
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>

Scheduler: <default-nc>, Forwarding class: network-control, Index: 23
  Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer
Limit: none,
  Priority: low
  Excess Priority: low
  drop-profile-map-set-type: mark
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      <default-drop-profile>
    Medium high   any       1      <default-drop-profile>
    High          any       1      <default-drop-profile>

Scheduler: <default-mcast>, Forwarding class: mcast, Index: 49
  Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent, Buffer
Limit: none,
  Priority: low
  Excess Priority: low
  drop-profile-map-set-type: mark
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      <default-drop-profile>
    Medium high   any       1      <default-drop-profile>
    High          any       1      <default-drop-profile>

Physical interface: xe-0/0/0, Index: 129
Queues supported: 12, Queues in use: 12
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

Physical interface: xe-0/0/1, Index: 130
Queues supported: 12, Queues in use: 12
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

...           ...           ...

Forwarding class set: lan-fcset, Type: normal-type, Forwarding class set index:
7
  Forwarding class                Index
  best-effort                     0
```

## show class-of-service classifier

<b>Syntax</b>	show class-of-service classifier <name <i>name</i> > <type dscp   type dscp-ipv6   type exp   type ieee-802.1   type inet-precedence>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.
<b>Options</b>	<p><b>none</b>—Display all classifiers.</p> <p><b>name <i>name</i></b>—(Optional) Display named classifier.</p> <p><b>type dscp</b>—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.</p> <p><b>type dscp-ipv6</b>—(Optional) Display all classifiers of the DSCP for IPv6 type.</p> <p><b>type exp</b>—(Optional) Display all classifiers of the MPLS experimental (EXP) type.</p> <p><b>type ieee-802.1</b>—(Optional) Display all classifiers of the ieee-802.1 type.</p> <p><b>type inet-precedence</b>—(Optional) Display all classifiers of the inet-precedence type.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service classifier type ieee-802.1 on page 6302</a> <a href="#">show class-of-service classifier type ieee-802.1 (QFX Series) on page 6302</a>
<b>Output Fields</b>	<a href="#">Table 573 on page 6301</a> describes the output fields for the <b>show class-of-service classifier</b> command. Output fields are listed in the approximate order in which they appear.

**Table 573: show class-of-service classifier Output Fields**

Field Name	Field Description
<b>Classifier</b>	Name of the classifier.
<b>Code point type</b>	Type of the classifier: <b>exp</b> (not on EX Series switch), <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>ieee-802.1</b> , or <b>inet-precedence</b> .
<b>Index</b>	Internal index of the classifier.
<b>Code point</b>	Code point value used for classification
<b>Forwarding class</b>	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Table 573: show class-of-service classifier Output Fields (*continued*)

Field Name	Field Description
Loss priority	Loss priority value used for classification. For most platforms, the value is <b>high</b> or <b>low</b> . For some platforms, the value is <b>high</b> , <b>medium-high</b> , <b>medium-low</b> , or <b>low</b> .

## Sample Output

### show class-of-service classifier type ieee-802.1

```

user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point      Forwarding Class      Loss priority
000             best-effort           low
001             best-effort           high
010             expedited-forwarding  low
011             expedited-forwarding  high
100             assured-forwarding    low
101             assured-forwarding    medium-high
110             network-control       low
111             network-control       high

Classifier: users-ieee802.1, Code point type: ieee-802.1
Code point      Forwarding class      Loss priority
100             expedited-forwarding  low

```

### show class-of-service classifier type ieee-802.1 (QFX Series)

```

user@switch> show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point      Forwarding class      Loss priority
000             best-effort           low
001             best-effort           low
010             best-effort           low
011             fcoe                  low
100             no-loss               low
101             best-effort           low
110             network-control       low
111             network-control       low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
Code point      Forwarding class      Loss priority
000             mcast                 low
001             mcast                 low
010             mcast                 low
011             mcast                 low
100             mcast                 low
101             mcast                 low
110             mcast                 low
111             mcast                 low

```



## show class-of-service code-point-aliases

<b>Syntax</b>	<code>show class-of-service code-point-aliases</code> <code>&lt;dscp   dscp-ipv6   exp   ieee-802.1   inet-precedence&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.
<b>Options</b>	<p><b>none</b>—Display code point aliases of all code point types.</p> <p><b>dscp</b>—(Optional) Display Differentiated Services code point (DSCP) aliases.</p> <p><b>dscp-ipv6</b>—(Optional) Display IPv6 DSCP aliases.</p> <p><b>exp</b>—(Optional) Display MPLS EXP code point aliases.</p> <p><b>ieee-802.1</b>—(Optional) Display IEEE-802.1 code point aliases.</p> <p><b>inet-precedence</b>—(Optional) Display IPv4 precedence code point aliases.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service code-point-aliases exp on page 6304</a>
<b>Output Fields</b>	<a href="#">Table 574 on page 6303</a> describes the output fields for the <b>show class-of-service code-point-aliases</b> command. Output fields are listed in the approximate order in which they appear.

**Table 574: show class-of-service code-point-aliases Output Fields**

Field Name	Field Description
<b>Code point type</b>	Type of the code points displayed: <b>dscp</b> , <b>dscp-ipv6</b> (not on EX Series switch), <b>exp</b> (not on EX Series switch or the QFX Series), <b>ieee-802.1</b> , or <b>inet-precedence</b> (not on the QFX Series).
<b>Alias</b>	Alias for a bit pattern.
<b>Bit pattern</b>	Bit pattern for which the alias is displayed.

## Sample Output

`show class-of-service code-point-aliases exp`

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
Alias      Bit pattern
af11      100
af12      101
be        000
be1       001
cs6       110
cs7       111
ef        010
ef1       011
nc1       110
nc2       111
```

## show class-of-service congestion-notification

<b>Syntax</b>	show class-of-service congestion-notification
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display whether priority-based flow control (PFC) is enabled for each IEEE 802.1p code point.
<b>Options</b>	<b>none</b> —Display the PFC state for all IEEE 802.1p code points.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 6174</a></li> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 5606</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 6019</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 6028</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 6036</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 6050</a></li> <li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway</a></li> <li>• <a href="#">Example: Configuring PFC Across Layer 3 Interfaces on page 6138</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 5559</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 575 on page 6305</a> describes the output fields for the <b>show class-of-service congestion-notification</b> command. Output fields are listed in the approximate order in which they appear.

**Table 575: show class-of-service congestion-notification Output Fields**

Field Name	Field Description
<b>Type</b>	Type of interfaces on which congestion notification is applied. Congestion notification is applied on input interfaces.
<b>Index</b>	Index of this congestion notification profile.
<b>Name</b>	Name of the congestion notification profile.
<b>Cable Length</b>	Length of the attached physical cable in meters. The default value is 100 meters.

Table 575: show class-of-service congestion-notification Output Fields (*continued*)

Field Name	Field Description
Priority	IEEE 802.1p code point.
PFC	State of PFC for the corresponding code point, either <b>enabled</b> or <b>disabled</b> .
MRU	<p>Maximum receive unit of the interface in bytes. (Incoming traffic that exceeds the MRU size of an interface is dropped.) The default values are:</p> <ul style="list-style-type: none"> <li>2500 bytes for priority 3 traffic</li> <li>9216 bytes for priority 4 traffic</li> </ul> <p><b>NOTE:</b> If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.</p>
Flow-Control-Queues	Output queue mapping to IEEE 802.1p code points (priorities). Explicit output queue to priority mapping overwrites the default configuration, and only explicitly mapped queues are displayed in the output. Flow control is only enabled on a queue when you enable PFC on the corresponding priority in the input stanza of the congestion notification profile.

## Sample Output

### show class-of-service congestion-notification

```

user@switch> show class-of-service congestion-notification
Name: fcoe_p3_cnp, Index: 12037
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Enabled   2500
  100      Enabled   9216
  101      Disabled
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  000
  001      0
  010      1
  011      2
  100      3
  101      4
  110      5
  111      6
          7

```

Name: fcoe\_p3\_p5\_cnp, Index: 46484

Type: Input

Cable Length: 100 m

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

Type: Output

Priority	Flow-Control-Queues
011	
	3
101	
	5

## show class-of-service drop-profile

<b>Syntax</b>	<code>show class-of-service drop-profile</code> <code>&lt;profile-name profile-name&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display data points for each class-of-service (CoS) random early detection (RED) drop profile.
<b>Options</b>	<b>none</b> —Display all drop profiles. <b>profile-name profile-name</b> —(Optional) Display the specified profile only.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service drop-profile on page 6309</a> <a href="#">show class-of-service drop-profile (EX4200 Switch) on page 6309</a> <a href="#">show class-of-service drop-profile (EX8200 Switch) on page 6309</a>
<b>Output Fields</b>	<a href="#">Table 576 on page 6308</a> describes the output fields for the <b>show class-of-service drop-profile</b> command. Output fields are listed in the approximate order in which they appear.

**Table 576: show class-of-service drop-profile Output Fields**

Field Name	Field Description
<b>Drop profile</b>	Name of a drop profile.
<b>Type</b>	Type of drop profile: <ul style="list-style-type: none"> <li><b>discrete</b> (default)</li> <li><b>interpolated</b> (EX8200 switches only)</li> </ul>
<b>Index</b>	Internal index of this drop profile.
<b>Fill Level</b>	Percentage fullness of a queue.
<b>Drop probability</b>	Drop probability at this fill level.

## Sample Output

### show class-of-service drop-profile

```

user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100         100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
     0           0
     1           1
     2           2
     4           4
     5           5
     6           6
     8           8
    10          10
    12          15
    14          20
    15          23
... 64 entries total
    90          96
    92          96
    94          97
    95          98
    96          98
    98          99
    99          99
   100         100

```

### show class-of-service drop-profile (EX4200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level
    100
Drop profile: dp1, Type: discrete, Index: 40496
  Fill level
    10

```

### show class-of-service drop-profile (EX8200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100         100
Drop profile: dp1, Type: interpolated, Index: 40496
  Fill level    Drop probability
     0           0
     1          80
     2          90
     4          90
     5          90
     6          90
     8          90
    10          90
    12          91
    14          91
    15          91
    16          91

```

18	91
20	91
22	92
24	92
25	92
26	92
28	92
30	92
32	93
34	93
35	93
36	93
38	93
40	93
42	94
44	94
45	94
46	94
48	94
49	94
51	95
52	95
54	95
55	95
56	95
58	95
60	95
62	96
64	96
65	96
66	96
68	96
70	96
72	97
74	97
75	97
76	97
78	97
80	97
82	98
84	98
85	98
86	98
88	98
90	98
92	99
94	99
95	99
96	99
98	99
99	99
100	100
Drop profile: dp2, Type: discrete, Index: 40499	
Fill level	Drop probability
10	5
50	50



## show class-of-service forwarding-class

<b>Syntax</b>	show class-of-service forwarding-class
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS on EX Series Switches</a></li> <li>• <a href="#">Example: Configuring Forwarding Classes on page 6075</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> <li>• <a href="#">Monitoring CoS Forwarding Classes</a></li> <li>• <a href="#">Defining CoS Forwarding Classes (CLI Procedure)</a></li> <li>• <a href="#">Defining CoS Forwarding Class Sets on page 6166</a></li> <li>• <a href="#">Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-class on page 6312</a> <a href="#">show class-of-service forwarding-class (EX8200 Switch) on page 6312</a> <a href="#">show class-of-service forwarding-class (QFX Series) on page 6312</a>
<b>Output Fields</b>	Table 577 on page 6311 describes the output fields for the <b>show class-of-service forwarding-class</b> command. Output fields are listed in the approximate order in which they appear.

**Table 577: show class-of-service forwarding-class Output Fields**

Field Name	Field Description
<b>Forwarding class</b>	Name of the forwarding class.
<b>ID</b>	Forwarding class identifier.
<b>Queue</b>	CoS queue mapped to the forwarding class.
<b>Policing priority</b>	Not supported on EX Series switches or the QFX Series and can be ignored.
<b>Fabric priority</b>	(EX8200 switches only) Fabric priority for the forwarding class, either <b>high</b> or <b>low</b> . Determines the priority of packets entering the switch fabric.

Table 577: show class-of-service forwarding-class Output Fields (*continued*)

Field Name	Field Description
<b>No-Loss</b>	<p>(QFX Series only) Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes:</p> <ul style="list-style-type: none"> <li>Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is <b>drop</b>).</li> <li>Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is <b>no-loss</b>).</li> </ul>

## Sample Output

### show class-of-service forwarding-class

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Policing priority
best-effort           0        0      normal
expedited-forwarding  1        5      normal
assured-forwarding    2        1      normal
network-control       3        7      normal

```

## Sample Output

### show class-of-service forwarding-class (EX8200 Switch)

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Fabric priority
best-effort           0        0      low
expedited-forwarding  1        5      low
assured-forwarding    2        1      low
network-control       3        7      low
mcast-be              4        2      low
mcast-ef              5        4      low
mcast-af              6        6      low

```

## Sample Output

### show class-of-service forwarding-class (QFX Series)

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Policing priority No-Loss
best-effort           0        0      normal      Disabled
fcoe                  1        3      normal      Enabled
no-loss               2        4      normal      Enabled
network-control       3        7      normal      Disabled
mcast                 8        8      normal      Disabled

```

## show class-of-service forwarding-class-set

<b>Syntax</b>	<code>show class-of-service forwarding-class-set</code> <code>&lt;forwarding-class-set-name&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the forwarding classes associated with each forwarding class set.
<b>Options</b>	<p><b>none</b>—Display all forwarding class sets.</p> <p><b>forwarding-class-set-name</b>—(Optional) Display the forwarding classes associated with the specified forwarding class set.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 5835</a></li> <li>• <a href="#">Defining CoS Forwarding Class Sets on page 6166</a></li> <li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6078</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 578 on page 6313</a> describes the output fields for the <b>show class-of-service forwarding-class-set</b> command. Output fields are listed in the approximate order in which they appear.

**Table 578: show class-of-service forwarding-class-set Output Fields**

Field Name	Field Description
Forwarding class set	Name of the forwarding class set.
Type	Internal Junos OS type.
Forwarding class set index	Index of this forwarding class set.
Forwarding class	Name of a forwarding class.
Index	Index of this forwarding class.

## Sample Output

### show class-of-service forwarding-class-set

```

user@switch> show class-of-service forwarding-class-set
Forwarding class set: san_fcset, Type: normal-type, Forwarding class set index:
37839
  Forwarding class      Index
  fcoe                  1

Forwarding class set: lan_fcset, Type: normal-type, Forwarding class set index:

```

37840

Forwarding class  
best-effort

Index  
0

Forwarding class set: multicast\_fcset, Type: normal-type, Forwarding class set  
index: 37841

Forwarding class  
mcast

Index  
8

## show class-of-service forwarding-table

<b>List of Syntax</b>	<a href="#">Syntax on page 6315</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 6315</a>
<b>Syntax</b>	show class-of-service forwarding-table
<b>Syntax (TX Matrix and TX Matrix Plus Router)</b>	show class-of-service forwarding-table <lcc <i>number</i> >   <sfc <i>number</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all <b>show class-of-service forwarding-table</b> commands in succession.
<b>Options</b>	<p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> <li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li> <li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li> <li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li> <li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li> </ul> <p><b>sfc <i>number</i></b>—(TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace <i>number</i> with 0.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table on page 6316</a> <a href="#">show class-of-service forwarding-table lcc (TX Matrix Plus Router) on page 6317</a>
<b>Output Fields</b>	See the output field descriptions for <b>show class-of-service forwarding-table</b> commands: <ul style="list-style-type: none"> <li>• <a href="#">show class-of-service forwarding-table classifier</a></li> <li>• <a href="#">show class-of-service forwarding-table classifier mapping</a></li> <li>• <a href="#">show class-of-service forwarding-table drop-profile</a></li> </ul>

- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table loss-priority-map*
- *show class-of-service forwarding-table loss-priority-map mapping*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

## Sample Output

### show class-of-service forwarding-table

```

user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry #   Code point   Forwarding-class #   PLP
  0         000         0                   0
  1         001         0                   1
  2         010         1                   0
  3         011         1                   1
  4         100         2                   0
  5         101         2                   1
  6         110         3                   0
  7         111         3                   1

Interface      Index      Table Index/      Q num      Table type
sp-0/0/0.1001   66         11                11         IPv4 precedence
sp-0/0/0.2001   67         11                11         IPv4 precedence
sp-0/0/0.16383  68         11                11         IPv4 precedence
fe-0/0/0.0      69         11                11         IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
  Priority low

```

```

    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1
      Drop
Entry  Fullness(%)  Probability(%)
   0           100           100

```

### show class-of-service forwarding-table lcc (TX Matrix Plus Router)

```

user@host> show class-of-service forwarding-table lcc 0
lcc0-re0:

```

```

-----
Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP
Entry #   Code point   Forwarding-class #   PLP
   0       000000         0         0
   1       000001         0         0
   2       000010         0         0
   3       000011         0         0
   4       000100         0         0
   5       000101         0         0
   6       000110         0         0
   7       000111         0         0
   8       001000         0         0
   9       001001         0         0
  10       001010         0         0
  11       001011         0         0
  12       001100         0         0
  13       001101         0         0
  14       001110         0         0
  15       001111         0         0
  16       010000         0         0
  17       010001         0         0
  18       010010         0         0
  19       010011         0         0
  20       010100         0         0
  21       010101         0         0
  22       010110         0         0
  23       010111         0         0
  24       011000         0         0
  25       011001         0         0
  26       011010         0         0
  27       011011         0         0
  28       011100         0         0
  29       011101         0         0
  30       011110         0         0
  31       011111         0         0
  32       100000         0         0
  33       100001         0         0
  34       100010         0         0
  35       100011         0         0
  36       100100         0         0
  37       100101         0         0
  38       100110         0         0
  39       100111         0         0

```

40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
...			



## show class-of-service forwarding-table classifier

<b>Syntax</b>	show class-of-service forwarding-table classifier
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table classifier on page 6319</a>
<b>Output Fields</b>	<a href="#">Table 579 on page 6319</a> describes the output fields for the <b>show class-of-service forwarding-table classifier</b> command. Output fields are listed in the approximate order in which they appear.

**Table 579: show class-of-service forwarding-table classifier Output Fields**

Field Name	Field Description
<b>Classifier table index</b>	Index of the classifier table.
<b>entries</b>	Total number of entries.
<b>Table type</b>	Type of code points in the table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), or <b>IPv6 DSCP</b> .
<b>Entry #</b>	Entry number.
<b>Code point</b>	Code point value used for classification.
<b>Forwarding-class #</b>	Forwarding class to which the code point is assigned.
<b>PLP</b>	Packet loss priority value set by classification. For most platforms, the value can be <b>0</b> or <b>1</b> . For some platforms, the value is <b>0</b> , <b>1</b> , <b>2</b> , or <b>3</b> . The value <b>0</b> represents low PLP. The value <b>1</b> represents <b>high</b> PLP. The value <b>2</b> represents medium-low PLP. The value <b>3</b> represents medium-high PLP.

## Sample Output

### show class-of-service forwarding-table classifier

```

user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP

Entry #   Code point   Forwarding-class #   PLP

```

0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	1	1
11	001011	0	0
...			
60	111100	0	0
61	111101	0	0
62	111110	0	0
63	111111	0	0

## show class-of-service forwarding-table classifier mapping

<b>Syntax</b>	show class-of-service forwarding-table classifier mapping
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table classifier mapping on page 6321</a>
<b>Output Fields</b>	<a href="#">Table 580 on page 6321</a> describes the output fields for the <b>show class-of-service forwarding-table classifier mapping</b> command. Output fields are listed in the approximate order in which they appear.

**Table 580: show class-of-service forwarding-table classifier mapping Output Fields**

Field Name	Field Description
<b>Table index/ Q num</b>	If the table type is <b>Fixed</b> , the number of the queue to which the interface is mapped. For all other types, this value is the classifier index number.
<b>Interface</b>	Name of the logical interface. This field can also show the physical interface (QFX Series).
<b>Index</b>	Logical interface index.
<b>Table type</b>	Type of code points in the table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>Fixed</b> , <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), or <b>IPv6 DSCP</b> .

## Sample Output

### show class-of-service forwarding-table classifier mapping

```

user@host> show class-of-service forwarding-table classifier mapping
Table index/
Interface      Index  Q num  Table type
so-5/0/0.0     10    62436  DSCP
so-0/1/0.0     11    62436  DSCP
so-0/2/0.0     12      1  Fixed
so-0/2/1.0     13    62436  DSCP
so-0/2/1.0     13    62437  IEEE 802.1
so-0/2/2.0     14    62436  DSCP
so-0/2/2.0     14    62438  IPv4 precedence

```



## show class-of-service forwarding-table drop-profile

<b>Syntax</b>	show class-of-service forwarding-table drop-profile
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the data points of all random early detection (RED) drop profiles as they exist in the forwarding table.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table drop-profile on page 6323</a>
<b>Output Fields</b>	<a href="#">Table 581 on page 6323</a> describes the output fields for the <b>show class-of-service forwarding-table drop-profile</b> command. Output fields are listed in the approximate order in which they appear.

**Table 581: show class-of-service forwarding-table drop-profile Output Fields**

Field Name	Field Description
RED drop profile index	Index of this drop profile.
# entries	Number of entries in a particular RED drop profile index.
Entry	Drop profile entry number.
Fullness(%)	Percentage fullness of a queue.
Drop probability(%)	Drop probability at this fill level.

## Sample Output

### show class-of-service forwarding-table drop-profile

```

user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
      Drop
Entry    Fullness(%)  Probability(%)
  0         100           100

RED drop profile index: 8742, # entries: 3
      Drop
Entry    Fullness(%)  Probability(%)
  0         10           10
  1         20           20
  2         30           30

```

RED drop profile index: 24627, # entries: 64

Entry	Fullness(%)	Drop	
		Probability(%)	
0	0	0	
1	1	1	
2	2	2	
3	4	4	
...			
61	98	99	
62	99	99	
63	100	100	

RED drop profile index: 25393, # entries: 64

Entry	Fullness(%)	Drop	
		Probability(%)	
0	0	0	
1	1	1	
2	2	2	
3	4	4	
...			
61	98	98	
62	99	99	
63	100	100	

## show class-of-service forwarding-table rewrite-rule

<b>Syntax</b>	show class-of-service forwarding-table rewrite-rule
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display mapping of queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table rewrite-rule on page 6325</a>
<b>Output Fields</b>	<a href="#">Table 582 on page 6325</a> describes the output fields for the <b>show class-of-service forwarding-table rewrite-rule</b> command. Output fields are listed in the approximate order in which they appear.

**Table 582: show class-of-service forwarding-table rewrite-rule Output Fields**

Field Name	Field Description
Rewrite table index	Index for this rewrite rule.
# entries	Number of entries in this rewrite rule.
Table type	Type of table: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>EXP-PUSH-3</b> (not on the QFX Series), <b>EXP-SWAP-PUSH-2</b> , (J Series routers only), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), <b>IPv6 DSCP</b> , or <b>Fixed</b> .
Q#	Queue number to which this entry is assigned.
Low bits	Code point value for low-priority loss profile.
State	State of this code point: <b>enabled</b> , <b>rewritten</b> , or <b>disabled</b> .
High bits	Code point value for high-priority loss profile.

## Sample Output

### show class-of-service forwarding-table rewrite-rule

```

user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q#      Low bits  State      High bits  State
0       000111  Enabled    001010    Enabled
2       000000  Disabled   001100    Enabled

```

1	101110	Enabled	110111	Enabled
3	110000	Enabled	111000	Enabled



## show class-of-service forwarding-table rewrite-rule mapping

<b>Syntax</b>	show class-of-service forwarding-table rewrite-rule mapping
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For each logical interface, display the table identifier of the rewrite rule map for each code point type.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table rewrite-rule mapping on page 6327</a>
<b>Output Fields</b>	<a href="#">Table 583 on page 6327</a> describes the output fields for the <b>show class-of-service forwarding-table rewrite-rule mapping</b> command. Output fields are listed in the approximate order in which they appear.

**Table 583: show class-of-service forwarding-table rewrite-rule mapping Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the logical interface. This field can also show the physical interface (QFX Series).
<b>Index</b>	Logical interface index.
<b>Table index</b>	Rewrite table index.
<b>Type</b>	Type of classifier: <b>DSCP</b> , <b>EXP</b> (not on the QFX Series), <b>EXP-PUSH-3</b> (not on the QFX Series), <b>EXP-SWAP-PUSH-2</b> (not on the QFX Series), <b>Frame-Relay DE</b> (J Series routers only), <b>IEEE 802.1</b> , <b>IPv4 precedence</b> (not on the QFX Series), <b>IPv6 DSCP</b> , or <b>Fixed</b> .

## Sample Output

### show class-of-service forwarding-table rewrite-rule mapping

```

user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface      Index  Table index  Type
so-5/0/0.0     10     3753        DSCP
so-0/1/0.0     11     3753        DSCP
so-0/2/0.0     12     3753        DSCP
so-0/2/1.0     13     3753        DSCP
so-0/2/2.0     14     3753        DSCP
so-0/2/3.0     15     3753        DSCP

```

## show class-of-service forwarding-table scheduler-map

<b>Syntax</b>	show class-of-service forwarding-table scheduler-map
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For each physical interface, display the scheduler map information as it exists in the forwarding table.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service forwarding-table scheduler-map on page 6329</a>
<b>Output Fields</b>	<a href="#">Table 584 on page 6328</a> describes the output fields for the <b>show class-of-service forwarding-table scheduler-map</b> command. Output fields are listed in the approximate order in which they appear.

**Table 584: show class-of-service forwarding-table scheduler-map Output Fields**

Field Name	Field Description
Interface	Name of the physical interface.
Index	Physical interface index.
Map index	Scheduler map index.
Num of queues	Number of queues defined in this scheduler map.
Entry	Number of this entry in the scheduler map.
Scheduler index	Scheduler policy index.
Forwarding-class #	Forwarding class number to which this entry is applied.
Tx rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.
Max buffer delay	Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or the keyword <b>remainder</b> , which indicates that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	<ul style="list-style-type: none"> <li><b>high</b>—Queue priority is high.</li> <li><b>low</b>—Queue priority is low.</li> </ul>
PLP high	Drop profile index for a high packet loss priority profile.

Table 584: show class-of-service forwarding-table scheduler-map Output Fields (*continued*)

Field Name	Field Description
PLP low	Drop profile index for a low packet loss priority profile.
PLP medium-high	Drop profile index for a medium-high packet loss priority profile.
PLP medium-low	Drop profile index for a medium-low packet loss priority profile.
TCP PLP high	Drop profile index for a high TCP packet loss priority profile.
TCP PLP low	Drop profile index for a low TCP packet loss priority profile.
Policy is exact	If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled.

## Sample Output

### show class-of-service forwarding-table scheduler-map

```

user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
    Policy is exact
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

```

## show class-of-service interface

---

<b>Syntax</b>	<code>show class-of-service interface</code> <code>&lt;comprehensive   detail&gt; &lt;interface-name&gt;</code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Forwarding class map information added in Junos OS Release 9.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.</p> <p>Options <b>detail</b> and <b>comprehensive</b> introduced in Junos OS Release 11.4.</p>
<b>Description</b>	Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.
<b>Options</b>	<p><b>none</b>—Display CoS associations for all physical and logical interfaces.</p> <p><b>comprehensive</b>—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.</p> <p><b>detail</b>—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.</p> <p>If the <b>interface</b> <i>interface-name</i> is a physical interface, the output includes:</p> <ul style="list-style-type: none"><li>• Brief QoS information about the physical interface</li><li>• Brief QoS information about the logical interface</li><li>• CoS information about the physical interface</li><li>• Brief information about filters or policers of the logical interface</li><li>• Brief CoS information about the logical interface</li></ul> <p>If the <b>interface</b> <i>interface-name</i> is a logical interface, the output includes:</p> <ul style="list-style-type: none"><li>• Brief QoS information about the logical interface</li><li>• Information about filters or policers for the logical interface</li><li>• CoS information about the logical interface</li></ul> <p><b>interface-name</b>—(Optional) Display class-of-service (CoS) associations for the specified interface.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service interface (Physical) on page 6341</a>

[show class-of-service interface \(Logical\) on page 6342](#)  
[show class-of-service interface \(Gigabit Ethernet\) on page 6342](#)  
[show class-of-service interface \(PPPoE Interface\) on page 6342](#)  
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 6342](#)  
[show class-of-service interface detail on page 6343](#)  
[show class-of-service interface comprehensive on page 6343](#)  
[show class-of-service interface \(ACX Series Routers\) on page 6354](#)

**Output Fields** [Table 585 on page 6331](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

**Table 585: show class-of-service interface Output Fields**

Field Name	Field Description
<b>Physical interface</b>	Name of a physical interface.
<b>Index</b>	Index of this interface or the internal index of this object.
<b>Dedicated Queues</b>	Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.
<b>Queues supported</b>	Number of queues you can configure on the interface.
<b>Queues in use</b>	Number of queues currently configured.
<b>Total non-default queues created</b>	Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.
<b>Rewrite Input IEEE Code-point</b>	(QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
<b>Shaping rate</b>	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the <b>Shaping rate</b> field is displayed for either the physical interface or the logical interface.
<b>Scheduler map</b>	Name of the output scheduler map associated with this interface.
<b>Scheduler map forwarding class sets</b>	(QFX Series only) Name of the fabric forwarding class set scheduler map associated with a QFabric system Interconnect device interface.
<b>Input shaping rate</b>	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
<b>Input scheduler map</b>	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
<b>Chassis scheduler map</b>	Name of the scheduler map associated with the packet forwarding component queues.
<b>Rewrite</b>	Name and type of the rewrite rules associated with this interface.
<b>Classifier</b>	Name and type of classifiers associated with this interface.

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Forwarding-class-map</b>	Name of the forwarding map associated with this interface.
<b>Congestion-notification</b>	(QFX Series only) Congestion notification state, <b>enabled</b> or <b>disabled</b> .
<b>Logical interface</b>	Name of a logical interface.
<b>Object</b>	Category of an object: <b>Classifier</b> , <b>Fragmentation-map</b> (for LSQ interfaces only), <b>Scheduler-map</b> , <b>Rewrite</b> , or <b>Translation Table</b> (for IQE PICs only).
<b>Name</b>	Name of an object.
<b>Type</b>	Type of an object: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>ieee-802.1</b> , <b>ip</b> , or <b>inet-precedence</b> .
<b>Link-level type</b>	Encapsulation on the physical interface.
<b>MTU</b>	MTU size on the physical interface.
<b>Speed</b>	Speed at which the interface is running.
<b>Loopback</b>	Whether loopback is enabled and the type of loopback.
<b>Source filtering</b>	Whether source filtering is enabled or disabled.
<b>Flow control</b>	Whether flow control is enabled or disabled.
<b>Auto-negotiation</b>	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.
<b>Remote-fault</b>	(Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Device flags</b>	<p>The <b>Device flags</b> field provides information about the physical device and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Down</b>—Device has been administratively disabled.</li> <li>• <b>Hear-Own-Xmit</b>—Device receives its own transmissions.</li> <li>• <b>Link-Layer-Down</b>—The link-layer protocol has failed to connect with the remote endpoint.</li> <li>• <b>Loopback</b>—Device is in physical loopback.</li> <li>• <b>Loop-Detected</b>—The link layer has received frames that it sent, thereby detecting a physical loopback.</li> <li>• <b>No-Carrier</b>—On media that support carrier recognition, no carrier is currently detected.</li> <li>• <b>No-Multicast</b>—Device does not support multicast traffic.</li> <li>• <b>Present</b>—Device is physically present and recognized.</li> <li>• <b>Promiscuous</b>—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.</li> <li>• <b>Quench</b>—Transmission on the device is quenched because the output buffer is overflowing.</li> <li>• <b>Recv-All-Multicasts</b>—Device is in multicast promiscuous mode and therefore provides no multicast filtering.</li> <li>• <b>Running</b>—Device is active and enabled.</li> </ul>
<b>Interface flags</b>	<p>The <b>Interface flags</b> field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>Admin-Test</b>—Interface is in test mode and some sanity checking, such as loop detection, is disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Hardware-Down</b>—Interface is nonfunctional or incorrectly connected.</li> <li>• <b>Link-Layer-Down</b>—Interface keepalives have indicated that the link is incomplete.</li> <li>• <b>No-Multicast</b>—Interface does not support multicast traffic.</li> <li>• <b>No-receive No-transmit</b>—Passive monitor mode is configured on the interface.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>Pop all MPLS labels from packets of depth</b>—MPLS labels are removed as packets arrive on an interface that has the <b>pop-all-labels</b> statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> <li>• <b>1</b>—Takes effect for incoming packets with one label only.</li> <li>• <b>2</b>—Takes effect for incoming packets with two labels only.</li> <li>• <b>[ 1 2 ]</b>—Takes effect for incoming packets with either one or two labels.</li> </ul> </li> <li>• <b>Promiscuous</b>—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.</li> <li>• <b>Recv-All-Multicasts</b>—Interface is in multicast promiscuous mode and provides no multicast filtering.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Flags</b>	<p>The <b>Logical interface flags</b> field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC Encapsulation</b>—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).</li> <li>• <b>Device-down</b>—Device has been administratively disabled.</li> <li>• <b>Disabled</b>—Interface is administratively disabled.</li> <li>• <b>Down</b>—A hardware failure has occurred.</li> <li>• <b>Clear-DF-Bit</b>—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.</li> <li>• <b>Hardware-Down</b>—Interface protocol initialization failed to complete successfully.</li> <li>• <b>PFC</b>—Protocol field compression is enabled for the PPP session.</li> <li>• <b>Point-To-Point</b>—Interface is point-to-point.</li> <li>• <b>SNMP-Traps</b>—SNMP trap notifications are enabled.</li> <li>• <b>Up</b>—Interface is enabled and operational.</li> </ul>
<b>Encapsulation</b>	Encapsulation on the logical interface.
<b>Admin</b>	Administrative state of the interface ( <b>Up</b> or <b>Down</b> )
<b>Link</b>	Status of physical link ( <b>Up</b> or <b>Down</b> ).
<b>Proto</b>	Protocol configured on the interface.
<b>Input Filter</b>	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
<b>Output Filter</b>	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.
<b>Link flags</b>	<p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <b>ACFC</b>—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.</li> <li>• <b>Give-Up</b>—Link protocol does not continue connection attempts after repeated failures.</li> <li>• <b>Loose-LCP</b>—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-LMI</b>—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.</li> <li>• <b>Loose-NCP</b>—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.</li> <li>• <b>Keepalives</b>—Link protocol keepalives are enabled.</li> <li>• <b>No-Keepalives</b>—Link protocol keepalives are disabled.</li> <li>• <b>PFC</b>—Protocol field compression is configured. The PPP session negotiates the PFC option.</li> </ul>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.
<b>CoS queues</b>	Number of CoS queues configured.



Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .
<b>Statistics last cleared</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>
<b>IPv6 transit statistics</b>	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.
<b>Input errors</b>	Input errors on the interface. The labels are explained in the following list: <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Bucket Drops</b>—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>HS link FIFO overflows</b>—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> </ul>

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Output errors</b>	<p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Drops</b> field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>HS link FIFO underflows</b>—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeds the MTU of the interface.</li> </ul>
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.
<b>Queue counters</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the <b>Dropped packets</b> field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
<b>SONET alarms</b> <b>SONET defects</b>	<p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: <b>SONET PHY</b>, <b>SONET section</b>, <b>SONET line</b>, and <b>SONET path</b>.</p>
<b>SONET PHY</b>	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET PHY</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>PLL Lock</b>—Phase-locked loop</li> <li>• <b>PHY Light</b>—Loss of optical signal</li> </ul>

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>SONET section</b>	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET section</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B1</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>SEF</b>—Severely errored framing</li> <li>• <b>LOS</b>—Loss of signal</li> <li>• <b>LOF</b>—Loss of frame</li> <li>• <b>ES-S</b>—Errored seconds (section)</li> <li>• <b>SES-S</b>—Severely errored seconds (section)</li> <li>• <b>SEFS-S</b>—Severely errored framing seconds (section)</li> </ul>
<b>SONET line</b>	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET line</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B2</b>—Bit interleaved parity for SONET line overhead</li> <li>• <b>REI-L</b>—Remote error indication (near-end line)</li> <li>• <b>RDI-L</b>—Remote defect indication (near-end line)</li> <li>• <b>AIS-L</b>—Alarm indication signal (near-end line)</li> <li>• <b>BERR-SF</b>—Bit error rate fault (signal failure)</li> <li>• <b>BERR-SD</b>—Bit error rate defect (signal degradation)</li> <li>• <b>ES-L</b>—Errored seconds (near-end line)</li> <li>• <b>SES-L</b>—Severely errored seconds (near-end line)</li> <li>• <b>UAS-L</b>—Unavailable seconds (near-end line)</li> <li>• <b>ES-LFE</b>—Errored seconds (far-end line)</li> <li>• <b>SES-LFE</b>—Severely errored seconds (far-end line)</li> <li>• <b>UAS-LFE</b>—Unavailable seconds (far-end line)</li> </ul>

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>SONET path</b>	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> <li>• <b>Seconds</b>—Number of seconds the defect has been active.</li> <li>• <b>Count</b>—Number of times that the defect has gone from inactive to active.</li> <li>• <b>State</b>—State of the error. A state other than <b>OK</b> indicates a problem.</li> </ul> <p>The <b>SONET path</b> field has the following subfields:</p> <ul style="list-style-type: none"> <li>• <b>BIP-B3</b>—Bit interleaved parity for SONET section overhead</li> <li>• <b>REI-P</b>—Remote error indication</li> <li>• <b>LOP-P</b>—Loss of pointer (path)</li> <li>• <b>AIS-P</b>—Path alarm indication signal</li> <li>• <b>RDI-P</b>—Path remote defect indication</li> <li>• <b>UNEQ-P</b>—Path unequipped</li> <li>• <b>PLM-P</b>—Path payload (signal) label mismatch</li> <li>• <b>ES-P</b>—Errored seconds (near-end STS path)</li> <li>• <b>SES-P</b>—Severely errored seconds (near-end STS path)</li> <li>• <b>UAS-P</b>—Unavailable seconds (near-end STS path)</li> <li>• <b>ES-PFE</b>—Errored seconds (far-end STS path)</li> <li>• <b>SES-PFE</b>—Severely errored seconds (far-end STS path)</li> <li>• <b>UAS-PFE</b>—Unavailable seconds (far-end STS path)</li> </ul>
<b>Received SONET overhead</b>  <b>Transmitted SONET overhead</b>	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> <li>• <b>C2</b>—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P.</li> <li>• <b>F1</b>—Section user channel byte. This byte is set aside for the purposes of users.</li> <li>• <b>K1</b> and <b>K2</b>—These bytes are allocated for APS signaling for the protection of the multiplex section.</li> <li>• <b>J0</b>—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter.</li> <li>• <b>S1</b>—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal.</li> <li>• <b>Z3</b> and <b>Z4</b>—Allocated for future use.</li> </ul>
<b>Received path trace</b>  <b>Transmitted path trace</b>	<p>SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>
<b>HDLC configuration</b>	<p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> <li>• <b>Policing bucket</b>—Configured state of the receiving policer.</li> <li>• <b>Shaping bucket</b>—Configured state of the transmitting shaper.</li> <li>• <b>Giant threshold</b>—Giant threshold programmed into the hardware.</li> <li>• <b>Runt threshold</b>—Runt threshold programmed into the hardware.</li> </ul>

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
<b>Packet Forwarding Engine configuration</b>	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> <li>• <b>PLP byte</b>—Packet Level Protocol byte.</li> </ul>
<b>CoS information</b>	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> <li>• <b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li>• <b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li>• <b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li>• <b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li>• <b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li>• <b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li>• <b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul>
<b>Forwarding classes</b>	Total number of forwarding classes supported on the specified interface.
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.
<b>Queue</b>	Queue number.
<b>Forwarding classes</b>	Forwarding class name.
<b>Queued Packets</b>	Number of packets queued to this queue.
<b>Queued Bytes</b>	Number of bytes queued to this queue. The byte counts vary by PIC type.
<b>Transmitted Packets</b>	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.
<b>Transmitted Bytes</b>	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
<b>Tail-dropped packets</b>	Number of packets dropped because of tail drop.

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
Transmit rate	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
Rate Limit	<p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> <li><b>None</b>—No rate limit.</li> <li><b>exact</b>—Queue transmits at the configured rate.</li> </ul>
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as <b>low</b> or <b>high</b> .
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .

Table 585: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> <li>• <b>Loss priority</b>—Packet loss priority for drop profile assignment.</li> <li>• <b>Protocol</b>—Transport protocol for drop profile assignment.</li> <li>• <b>Index</b>—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles.</li> <li>• <b>Name</b>—Name of the drop profile.</li> <li>• <b>Type</b>—Type of the drop profile: <b>discrete</b> or <b>interpolated</b>.</li> <li>• <b>Fill Level</b>—Percentage fullness of a queue.</li> <li>• <b>Drop probability</b>—Drop probability at this fill level.</li> </ul>
Adjustment information	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> <li>• <b>Adjusting application</b>—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> <li>• The adjusting application can appear as <b>anclp LS-0</b>, which is the Junos OS Access Node Control Profile process (<b>anclpd</b>) that performs shaping-rate adjustments on schedule nodes.</li> <li>• The adjusting application can also appear as <b>pppoe</b>, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).</li> </ul> </li> <li>• <b>Adjustment type</b>—Type of adjustment: <b>absolute</b> or <b>delta</b>.</li> <li>• <b>Configured shaping rate</b>—Shaping rate configured for the scheduler node or queue.</li> <li>• <b>Adjustment value</b>—Value of adjusted shaping rate.</li> <li>• <b>Adjustment target</b>—Level of shaping-rate adjustment performed: <b>node</b> or <b>queue</b>.</li> <li>• <b>Adjustment overhead-accounting mode</b>—Configured shaping mode: <b>frame</b> or <b>cell</b>.</li> </ul>

## Sample Output

### show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4

```

Total non-default queues created: 4  
 Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no  
 Shaping rate: 32000

Object	Name	Type	Index
Scheduler-map	<default>		27
Rewrite	exp-default	exp	21
Classifier	exp-default	exp	5
Classifier	ipprec-compatibility	ip	8
Forwarding-class-map	exp-default	exp	5

### show class-of-service interface (Logical)

user@host> show class-of-service interface so-0/2/3.0

Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no  
 Shaping rate: 32000

Object	Name	Type	Index
Scheduler-map	<default>		27
Rewrite	exp-default	exp	21
Classifier	exp-default	exp	5
Classifier	ipprec-compatibility	ip	8
Forwarding-class-map	exp-default	exp	5

### show class-of-service interface (Gigabit Ethernet)

user@host> show class-of-service interface ge-6/2/0

Physical interface: ge-6/2/0, Index: 175  
 Queues supported: 4, Queues in use: 4  
 Scheduler map: <default>, Index: 2  
 Input scheduler map: <default>, Index: 3  
 Chassis scheduler map: <default-chassis>, Index: 4

### show class-of-service interface (PPPoE Interface)

user@host> show class-of-service interface pp0.1

Logical interface: pp0.1, Index: 85

Object	Name	Type	Index
Traffic-control-profile	tcp-pppoe.o.pp0.1	Output	2726446535
Classifier	ipprec-compatibility	ip	13

Adjusting application: PPPoE  
 Adjustment type: absolute  
 Adjustment value: 5000000  
 Adjustment overhead-accounting mode: cell  
 Adjustment target: node

### show class-of-service interface (T4000 Routers with Type 5 FPCs)

user@host> show class-of-service interface xe-4/0/0

Physical interface: xe-4/0/0, Index: 153  
 Queues supported: 8, Queues in use: 4  
 Shaping rate: 5000000000 bps  
 Scheduler map: <default>, Index: 2  
 Congestion-notification: Disabled

Logical interface: xe-4/0/0.0, Index: 77

Index	Object	Name	Type
13	Classifier	ipprec-compatibility	ip



## show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
```

```
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
```

```
Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps
Scheduler map: interface-scheduler-map, Index: 58414
Input shaping rate: 10000 bps
878674 Input scheduler map: scheduler-map, Index: 15103
Chassis scheduler map: <default-chassis>, Index: 4
Congestion-notification: Disabled
```

```
Logical interface ge-0/3/0.0
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
inet
mpls
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.0	up	up	inet		
			mpls		
Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.0	up	up	inet		
			mpls		

```
Logical interface: ge-0/3/0.0, Index: 68
```

Object	Name	Type	Index
Rewrite	exp-default	exp (mpls-any)	33
Classifier	exp-default	exp	10
Classifier	ipprec-compatibility	ip	13

```
Logical interface ge-0/3/0.1
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2
inet
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.1	up	up	inet		
Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.1	up	up	inet		

```
Logical interface: ge-0/3/0.1, Index: 69
```

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

## show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
Interface index: 138, SNMP ifIndex: 601, Generation: 141
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Enabled,
Auto-negotiation: Enabled, Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
```

```

CoS queues      : 4 supported, 4 maximum usable queues
Schedulers     : 256
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
Last flapped   : 2010-09-07 06:35:22 PDT (15:14:42 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 total statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Label-switched interface (LSI) traffic statistics:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 af3                0                0                0
1 af2                0                0                0
2 ef2                0                0                0
3 ef1                0                0                0

Egress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 af3                0                0                0
1 af2                0                0                0
2 ef2                0                0                0
3 ef1                0                0                0

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets     0            0
Unicast packets   0            0
Broadcast packets 0            0
Multicast packets 0            0

```

```

CRC/Align errors                0                0
FIFO errors                     0                0
MAC control frames              0                0
MAC pause frames                0                0
Oversized frames                0
Jabber frames                   0
Fragment frames                 0
VLAN tagged frames              0
Code violations                  0
Filter statistics:
  Input packet count             0
  Input packet rejects           0
  Input DA rejects               0
  Input SA rejects               0
  Output packet count            0
  Output packet pad count        0
  Output packet error count      0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue            Bandwidth          Buffer Priority
Limit                           %          bps      %          usec
  2 ef2                         39          19500   0          120    high
none
  Direction : Input
  CoS transmit queue            Bandwidth          Buffer Priority
Limit                           %          bps      %          usec
  0 af3                         30          3000    45          0      low
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes   :                0                0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets : 0 0 pps
    RL-dropped bytes : 0 0 bps
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets : 0 0 pps
    RL-dropped bytes : 0 0 bps
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets : 0 0 pps
    RL-dropped bytes : 0 0 bps
    RED-dropped packets : 0 0 pps

```

```

    RED-dropped bytes      :                0          0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0          0 pps
    RL-dropped bytes      :                0          0 bps
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available
Queue: 1, Forwarding classes: af2
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets                :             108546          0 pps
    Bytes                  :          12754752        376 bps
  Transmitted:
    Packets                :             108546          0 pps
    Bytes                  :          12754752        376 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps

```

Scheduler map: interface-scheduler-map, Index: 58414

Scheduler: ef2, Forwarding class: ef2, Index: 39155

Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer Limit: none, Priority: high

Excess Priority: unspecified

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Input shaping rate: 10000 bps

Input scheduler map: scheduler-map

Scheduler map: scheduler-map, Index: 15103

Scheduler: af3, Forwarding class: af3, Index: 35058

Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer Limit: none, Priority: low

Excess Priority: unspecified

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	40582	green
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	18928	yellow

Drop profile: green, Type: discrete, Index: 40582

Fill level	Drop probability
50	0
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: yellow, Type: discrete, Index: 18928

Fill level	Drop probability
50	0
100	100

Chassis scheduler map: < default-drop-profile>

Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25

Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low

Excess Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25  
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low  
 Excess Priority: low  
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

```

Fill level      Drop probability
    100          100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
Excess Priority: low
Drop profiles:
  Loss priority  Protocol    Index    Name
  Low            any         1        < default-drop-profile>
  Medium low     any         1        < default-drop-profile>
  Medium high    any         1        < default-drop-profile>
  High           any         1        < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Congestion-notification: Disabled
Forwarding class
priority Policing priority          ID      Queue  Restricted queue  Fabric
af3      normal                    0       0          0             low
af2      normal                    1       1          1             low
ef2      normal                    2       2          2             high
ef1      normal                    3       3          3             high
af1      normal                    4       4          0             low

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 172, Route table: 0
Flags: Sendbcst-pkt-to-re
Input Filters: filter-in-ge-0/3/0.0-i,
Policer: Input: p1-ge-0/3/0.0-inet-i
Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0

```



Flags: Is-Primary  
Output Filters: exp-filter,,,,,

Logical interface ge-1/2/0.0 (Index 347) (SNMP ifIndex 638) (Generation 156)

Forwarding class ID	Queue	Restricted queue	Fabric priority	Policing priority
SPU priority				
best-effort	0	0	low	normal
low				

Aggregate Forwarding-class statistics per forwarding-class

Aggregate Forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

Input unicast bytes: 0  
Output unicast bytes: 0  
Input unicast packets: 0  
Output unicast packets: 0

Input multicast bytes: 0  
Output multicast bytes: 0  
Input multicast packets: 0  
Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:

Input unicast bytes: 0  
Output unicast bytes: 0  
Input unicast packets: 0  
Output unicast packets: 0

Input multicast bytes: 0  
Output multicast bytes: 0  
Input multicast packets: 0  
Output multicast packets: 0

IPv4 protocol forwarding-class statistics:

Forwarding-class statistics:

Forwarding-class best-effort statistics:

Input unicast bytes: 0  
Output unicast bytes: 0  
Input unicast packets: 0  
Output unicast packets: 0

Input multicast bytes: 0  
Output multicast bytes: 0  
Input multicast packets: 0  
Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:

Input unicast bytes: 0  
Output unicast bytes: 0  
Input unicast packets: 0  
Output unicast packets: 0

Input multicast bytes: 0  
Output multicast bytes: 0  
Input multicast packets: 0  
Output multicast packets: 0

IPv6 protocol forwarding-class statistics:  
 Forwarding-class statistics:  
 Forwarding-class best-effort statistics:

Input unicast bytes: 0  
 Output unicast bytes: 0  
 Input unicast packets: 0  
 Output unicast packets: 0

Input multicast bytes: 0  
 Output multicast bytes: 0  
 Input multicast packets: 0  
 Output multicast packets: 0

Forwarding-class expedited-forwarding statistics:

Input unicast bytes: 0  
 Output unicast bytes: 0  
 Input unicast packets: 0  
 Output unicast packets: 0

Input multicast bytes: 0  
 Output multicast bytes: 0  
 Input multicast packets: 0  
 Output multicast packets: 0

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2

Input packets : 0

Output packets: 0

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.0	up	up	inet	filter-in-ge-0/3/0.0-i	
			mpls		exp-filter
Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.0	up	up	inet	p1-ge-0/3/0.0-inet-i	
			mpls		

Filter: filter-in-ge-0/3/0.0-i

Counters:

Name	Bytes	Packets
count-filter-in-ge-0/3/0.0-i	0	0

Filter: exp-filter

Counters:

Name	Bytes	Packets
count-exp-seven-match	0	0
count-exp-zero-match	0	0

Policers:

Name	Packets
p1-ge-0/3/0.0-inet-i	0

Logical interface: ge-0/3/0.0, Index: 68

Object	Name	Type	Index
Rewrite	exp-default	exp (mpls-any)	33

Rewrite rule: exp-default, Code point type: exp, Index: 33

Forwarding class	Loss priority	Code point	
af3	low	000	
af3	high	001	
af2	low	010	
af2	high	011	
ef2	low	100	
ef2	high	101	
ef1	low	110	
ef1	high	111	
Object	Name	Type	Index
Classifier	exp-default	exp	10

Classifier: exp-default, Code point type: exp, Index: 10

Code point	Forwarding class	Loss priority	
000	af3	low	
001	af3	high	
010	af2	low	
011	af2	high	
100	ef2	low	
101	ef2	high	
110	ef1	low	
111	ef1	high	
Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

Code point	Forwarding class	Loss priority		
000	af3	low		
001	af3	high		
010	af3	low		
011	af3	high		
100	af3	low		
101	af3	high		
110	ef1	low		
111	ef1	high		
Forwarding class	ID	Queue	Restricted queue	Fabric
priority Policing priority				
af3	0	0	0	low
af2	1	1	1	low
ef2	2	2	2	high
ef1	3	3	3	high
af1	4	4	0	low

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2

Traffic statistics:

Input bytes : 0  
Output bytes : 0  
Input packets: 0  
Output packets: 0

Local statistics:

Input bytes : 0  
Output bytes : 0  
Input packets: 0

```

Output packets:          0
Transit statistics:
Input bytes  :          0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:          0          0 pps
Protocol inet, MTU: 1500, Generation: 174, Route table: 0
Flags: Sendbroadcast-pkt-to-re

```

```

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2
Input packets : 0
Output packets: 0

```

```

Interface      Admin Link Proto Input Filter      Output Filter
ge-0/3/0.1     up   up   mpls
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/3/0.1     up   up

```

```

Logical interface: ge-0/3/0.1, Index: 69
Object          Name          Type          Index
Classifier       ipprec-compatibility  ip          13

```

```
Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13
```

```

Code point      Forwarding class      Loss priority
000             af3                   low
001             af3                   high
010             af3                   low
011             af3                   high
100             af3                   low
101             af3                   high
110             ef1                   low
111             ef1                   high

```

```

Forwarding class      ID      Queue  Restricted queue  Fabric
priority Policing priority
af3                   0        0        0                low
normal
af2                   1        1        1                low
normal
ef2                   2        2        2                high
normal
ef1                   3        3        3                high
normal
af1                   4        4        0                low
normal

```

### show class-of-service interface (ACX Series Routers)

```

user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

Logical interface: at-0/0/0.0, Index: 69

```

Logical interface: at-0/0/0.32767, Index: 70

Physical interface: at-0/0/1, Index: 133

Queues supported: 4, Queues in use: 4

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Logical interface: at-0/0/1.0, Index: 71

Logical interface: at-0/0/1.32767, Index: 72

Physical interface: ge-0/1/0, Index: 146

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Rewrite	dscp-default	dscp	31
Classifier	d1	dscp	11331
Classifier	ci	ieee8021p	583

Logical interface: ge-0/1/0.0, Index: 73

Object	Name	Type	Index
Rewrite	custom-exp	exp (mpls-any)	46413

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Rewrite	ri	ieee8021p (outer)	35392
Classifier	ci	ieee8021p	583

Physical interface: ge-0/1/3, Index: 149

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/3.0, Index: 77

Object	Name	Type	Index
Rewrite	custom-exp2	exp (mpls-any)	53581

Physical interface: ge-0/1/4, Index: 150

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
--------	------	------	-------

```

Classifier                ipprec-compatibility  ip                                13

Physical interface: ge-0/1/5, Index: 151
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip        13

Physical interface: ge-0/1/6, Index: 152
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip        13

Physical interface: ge-0/1/7, Index: 153
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  d1          dscp      11331

Physical interface: ge-0/2/0, Index: 154
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip        13

Physical interface: ge-0/2/1, Index: 155
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip        13

Logical interface: ge-0/2/1.0, Index: 78

Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip        13

Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157
Queues supported: 8, Queues in use: 5
  Scheduler map: <default>, Index: 2
  Congestion-notification: Disabled
Object      Name      Type      Index
Classifier  ipprec-compatibility  ip        13

Logical interface: xe-0/3/1.0, Index: 81

[edit]
user@host-g11#

```



## show class-of-service multi-destination

<b>Syntax</b>	show class-of-service multi-destination
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For each class-of-service (CoS) multideestination classifier, display the classifier type.
<b>Options</b>	<b>none</b> —Display all multideestination classifiers.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Multideestination (Multicast, Broadcast, DLF) BA Classifiers on page 6162</a></li> <li>• <a href="#">Example: Configuring Multideestination (Multicast, Broadcast, DLF) Classifiers on page 6069</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 5810</a></li> <li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 5820</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 586 on page 6358</a> describes the output fields for the <b>show class-of-service multi-destination</b> command. Output fields are listed in the approximate order in which they appear.

**Table 586: show class-of-service multi-destination Output Fields**

Field Name	Field Description
Family ethernet	Family to which the classifier belongs.
Classifier Name	Name of the classifier.
Classifier Type	Type of the classifier: <b>dscp</b> or <b>ieee-802.1</b> .
Classifier Index	Internal index of the classifier.

## Sample Output

### show class-of-service multi-destination

```
user@switch> show class-of-service multi-destination
```

```

Family ethernet:
Classifier Name      Classifier Type      Classifier Index
ba-mcast-classifier  ieee-802.1          62376

```



## show class-of-service rewrite-rule

<b>Syntax</b>	show class-of-service rewrite-rule <name <i>name</i> > <type <i>type</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the mapping of forwarding classes and loss priority to code point values.
<b>Options</b>	<p><b>none</b>—Display all rewrite rules.</p> <p><b>name <i>name</i></b>—(Optional) Display the specified rewrite rule.</p> <p><b>type <i>type</i></b>—(Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>dscp</b>—For IPv4 traffic.</li> <li>• <b>dscp-ipv6</b>—For IPv6 traffic.</li> <li>• <b>exp</b>—For MPLS traffic.</li> <li>• <b>frame-relay-de</b>—(J Series routers only) For Frame Relay traffic.</li> <li>• <b>ieee-802.1</b>—For Layer 2 traffic.</li> <li>• <b>inet-precedence</b>—For IPv4 traffic.</li> </ul>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service rewrite-rule type dscp on page 6360</a> <a href="#">show class-of-service rewrite-rule type dscp (QFX Series) on page 6360</a>
<b>Output Fields</b>	<a href="#">Table 587 on page 6359</a> describes the output fields for the <b>show class-of-service rewrite-rule</b> command. Output fields are listed in the approximate order in which they appear.

**Table 587: show class-of-service rewrite-rule Output Fields**

Field Name	Field Description
<b>Rewrite rule</b>	Name of the rewrite rule.
<b>Code point type</b>	Type of rewrite rule: <b>dscp</b> , <b>dscp-ipv6</b> , <b>exp</b> , <b>frame-relay-de</b> , or <b>inet-precedence</b> .
<b>Forwarding class</b>	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch.
<b>Index</b>	Internal index for this particular rewrite rule.
<b>Loss priority</b>	Loss priority for rewriting.

Table 587: show class-of-service rewrite-rule Output Fields (*continued*)

Field Name	Field Description
Code point	Code point value to rewrite.

## Sample Output

### show class-of-service rewrite-rule type dscp

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
  Forwarding class      Loss priority      Code point
  gold                  high              000000
  silver                low               110000
  silver                high              111000
  bronze                low               001010
  bronze                high              001100
  lead                  high              101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
  Forwarding class      Loss priority      Code point
  gold                  low               000111
  gold                  high              001010
  silver                low               110000
  silver                high              111000
  bronze                high              001100
  lead                  low               101110
  lead                  high              110111

```

## Sample Output

### show class-of-service rewrite-rule type dscp (QFX Series)

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp, Index: 31
  Forwarding class      Loss priority      Code point
  best-effort           low               000000
  best-effort           high              000000
  fcoe                  low               101110
  fcoe                  high              101110
  no-loss               low               001010
  no-loss               high              001100
  network-control       low               110000
  network-control       high              111000

```

## show class-of-service scheduler-map

<b>Syntax</b>	show class-of-service scheduler-map <name>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.
<b>Options</b>	<b>none</b> —Display all scheduler maps.  <b>name</b> —(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service scheduler-map on page 6362</a>
<b>Output Fields</b>	<a href="#">Table 588 on page 6361</a> describes the output fields for the <b>show class-of-service scheduler-map</b> command. Output fields are listed in the approximate order in which they appear.

**Table 588: show class-of-service scheduler-map Output Fields**

Field Name	Field Description
<b>Scheduler map</b>	Name of the scheduler map.
<b>Index</b>	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.
<b>Scheduler</b>	Name of the scheduler.
<b>Forwarding class</b>	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
<b>Transmit rate</b>	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword <b>remainder</b> , which indicates that the scheduler receives the remaining bandwidth of the interface.
<b>Rate Limit</b>	Rate limiting configuration of the queue. Possible values are <b>none</b> , meaning no rate limiting, and <b>exact</b> , meaning the queue only transmits at the configured rate.
<b>Maximum buffer delay</b>	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword <b>remainder</b> to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.
<b>Priority</b>	Scheduling priority: <b>low</b> or <b>high</b> .

Table 588: show class-of-service scheduler-map Output Fields (*continued*)

Field Name	Field Description
Excess priority	Priority of excess bandwidth: <b>low</b> , <b>medium-low</b> , <b>medium-high</b> , <b>high</b> , or <b>none</b> .
Explicit Congestion Notification	(QFX Series only) Explicit congestion notification (ECN) state: <ul style="list-style-type: none"> <li>Disable—ECN is disabled on the specified scheduler</li> <li>Enable—ECN is enabled on the specified scheduler</li> </ul> ECN is disabled by default.
Adjust minimum	Minimum shaping rate for an adjusted queue, in bps.
Adjust percent	Bandwidth adjustment applied to a queue, in percent.
Drop profiles	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
Loss priority	Packet loss priority for drop profile assignment.
Protocol	Transport protocol for drop profile assignment.
Name	Name of the drop profile.

## Sample Output

### show class-of-service scheduler-map

```

user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   8724   aa-drop-profile
  Low           TCP       9874   bb-drop-profile
  High          non-TCP   8833   cc-drop-profile
  High          TCP       8484   dd-drop-profile

Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   8724   aa-drop-profile
  Low           TCP       9874   bb-drop-profile
  High          non-TCP   8833   cc-drop-profile
  High          TCP       8484   dd-drop-profile

```

## show class-of-service shared-buffer

<b>Syntax</b>	show class-of-service shared-buffer <egress   ingress>
<b>Release Information</b>	Command introduced in Junos OS Release 12.3 for the QFX Series.
<b>Description</b>	Display the shared buffer allocation and partitioning configuration.
<b>Options</b>	<p><b>none</b>—Display ingress and egress shared buffer settings.</p> <p><b>egress</b>—(Optional) Display the egress shared buffer settings.</p> <p><b>ingress</b>—(Optional) Display the ingress shared buffer settings.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 6104</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 6116</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 6122</a></li> <li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 6179</a></li> <li>• <a href="#">Understanding CoS Buffer Configuration on page 5891</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show class-of-service shared-buffer on page 6364</a>
<b>Output Fields</b>	Table 589 on page 6363 describes the output fields for the <b>show class-of-service shared-buffer</b> command. Output fields are listed in the approximate order in which they appear.

**Table 589: show class-of-service shared-buffer Output Fields**

Field Name	Field Description
Ingress	Ingress shared buffer configuration.
Total Buffer	Total buffer space available to the ports in KB. This is the combined dedicated buffer pool and shared buffer pool.
Dedicated Buffer	Buffer space allocated to the dedicated buffer pool in KB.
Shared Buffer	Buffer space allocated to the shared buffer pool in KB.
Lossless	Buffer space allocated to the lossless traffic buffer pool in KB.

Table 589: show class-of-service shared-buffer Output Fields (*continued*)

Field Name	Field Description
<b>Lossless Headroom</b>	Buffer space allocated to the lossless headroom traffic buffer pool to support priority-based flow control (PFC) and Ethernet PAUSE in KB. (Ingress ports only.)
<b>Lossy</b>	Buffer space allocated to the lossy (best-effort) traffic buffer pool in KB.
<b>Lossless Headroom Utilization</b>	Utilization of the ingress lossless headroom buffer pool. (These fields can help you to determine how much headroom buffer space you need to reserve to support PFC and Ethernet PAUSE for lossless flows.)
<b>Node Device</b>	Index number that identifies the switch. On a QFX3500 switch, this field always has a value of zero (0).
<b>Total</b>	Size of the lossless headroom ingress buffer pool in KB.
<b>Used</b>	Amount in KB of lossless headroom ingress buffer used.
<b>Free</b>	Amount in KB of lossless headroom ingress buffer free (unused).
<b>Egress</b>	Egress shared buffer configuration.
<b>Multicast</b>	Buffer space allocated to the multicast traffic buffer pool in KB. (Egress ports only.)

## Sample Output

### show class-of-service shared-buffer

```
user@switch> show class-of-service shared-buffer
```

```
Ingress:
```

```
Total Buffer      : 9360.00 KB
Dedicated Buffer   : 2158.00 KB
Shared Buffer      : 7202.00 KB
  Lossless        : 648.18 KB
  Lossless Headroom : 3240.90 KB
  Lossy           : 3312.92 KB
```

```
Lossless Headroom Utilization:
```

```
Node Device      Total      Used      Free
0                3240.90 KB  0.00 KB  3240.90 KB
```

```
Egress:
```

```
Total Buffer      : 9360.00 KB
Dedicated Buffer   : 2704.00 KB
Shared Buffer      : 6656.00 KB
  Lossless        : 3328.00 KB
  Multicast       : 1264.64 KB
  Lossy           : 2063.36 KB
```

## show class-of-service traffic-control-profile

<b>Syntax</b>	<code>show class-of-service traffic-control-profile</code> <code>&lt;profile-name&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Routers.
<b>Description</b>	For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and Trio MPC/MIC interfaces only, display traffic shaping and scheduling profiles.  (ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.
<b>Options</b>	<b>none</b> —Display all profiles.  <b>profile-name</b> —(Optional) Display information about a single profile.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show class-of-service traffic-control-profile on page 6367</a> <a href="#">show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC) on page 6367</a> <a href="#">show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces) on page 6367</a>
<b>Output Fields</b>	<a href="#">Table 590 on page 6365</a> describes the output fields for the <b>show class-of-service traffic-control-profile</b> command. Output fields are listed in the approximate order in which they appear.

**Table 590: show class-of-service traffic-control-profile Output Fields**

Field Name	Field Description
<b>Traffic control profile</b>	Name of the traffic control profile.
<b>Index</b>	Index number of the traffic control profile.
<b>ATM Service</b>	(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: <ul style="list-style-type: none"> <li>cbr—Constant bit rate.</li> <li>rtvbr—Real time variable bit rate.</li> <li>nrtvbr—Non real time variable bit rate.</li> <li>ubr—Unspecified bit rate.</li> </ul>
<b>Maximum Burst Size</b>	Configured maximum burst size, in cells.
<b>Peak rate</b>	Configured peak rate, in cps.

Table 590: show class-of-service traffic-control-profile Output Fields (*continued*)

Field Name	Field Description
<b>Sustained rate</b>	Configured sustained rate, in cps.
<b>Shaping rate</b>	Configured shaping rate, in bps.  <b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
<b>Shaping rate burst</b>	Configured burst size for the shaping rate, in bytes.  <b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.
<b>Shaping rate priority high</b>	Configured shaping rate for high-priority traffic, in bps.
<b>Shaping rate priority medium</b>	Configured shaping rate for medium-priority traffic, in bps.
<b>Shaping rate priority low</b>	Configured shaping rate for low-priority traffic, in bps.
<b>Shaping rate excess high</b>	Configured shaping rate for high-priority excess traffic, in bps.
<b>Shaping rate excess low</b>	Configured shaping rate for low-priority excess traffic, in bps.
<b>Scheduler map</b>	Name of the associated scheduler map.
<b>Delay Buffer rate</b>	Configured delay buffer rate, in bps.
<b>Excess rate</b>	Configured excess rate, in percent or proportion.
<b>Excess rate high</b>	Configured excess rate for high priority traffic, in percent or proportion.
<b>Excess rate low</b>	Configured excess rate for low priority traffic, in percent or proportion.
<b>Guaranteed rate</b>	Configured guaranteed rate, in bps or cps.  <b>NOTE:</b> (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values: <ul style="list-style-type: none"> <li>• <b>cbr</b>—Guaranteed rate is equal to the configured peak rate in cps.</li> <li>• <b>rtvbr</b>—Guaranteed rate is equal to the configured sustained rate in cps.</li> <li>• <b>nrtvbr</b>—Guaranteed rate is equal to the configured sustained rate in cps.</li> </ul>
<b>Guaranteed rate burst</b>	Configured burst size for the guaranteed rate, in bytes.
<b>adjust-minimum</b>	Configured minimum shaping rate for an adjusted queue, in bps.



Table 590: show class-of-service traffic-control-profile Output Fields (*continued*)

Field Name	Field Description
overhead accounting mode	Configured shaping mode: <b>Frame Mode</b> or <b>Cell Mode</b> .
Overhead bytes	Configured byte adjustment value.

## Sample Output

### show class-of-service traffic-control-profile

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
  Scheduler map: m1
  Delay Buffer rate: 500000
  Guaranteed rate: 1000000

Traffic control profile: Profile2, Index: 57624
  Scheduler map: m2
  Delay Buffer rate: 600000
  Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
  Scheduler map: m3
  Delay Buffer rate: 800000
  Guaranteed rate: 3000000
  .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
  Scheduler map: m4
  Delay Buffer rate: 750000
  Guaranteed rate: 4000000
  ..adjust-minimum 20000000

```

### show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
  ATM Service: RTVBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
  Shaping rate: 1000 cps
  Shaping rate burst: 500 cells
  Delay Buffer rate: 2000 cps
  Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
  ATM Service: UBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode

```

### show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
  ATM Service: RTVBR
  Shaping rate: 2000 cps

```

Shaping rate burst: 200 cells  
Scheduler map: <default>  
Delay Buffer rate: 1000 cps  
Guaranteed rate: 1700 cps

## show dcbx

<b>Syntax</b>	show dcbx
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 5724</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> </ul>
<b>Output Fields</b>	<a href="#">Table 463 on page 5723</a> lists the output fields for the <b>show dcbx</b> command. Output fields are listed in the approximate order in which they appear.

Table 591: show dcbx output fields

Field Name	Field Description
DCBX	Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"> <li>• Enabled—DCBX is enabled on the switch or on the specified interface</li> <li>• Disabled—DCBX is disabled on the switch or on the specified interface</li> </ul>
Interface	Name of the interface

## Sample Output

### show dcbx

```

user@switch> show dcbx
DCBX                : Enabled
Interface           DCBX
xe-0/0/9.0          enabled
xe-0/0/32.0         enabled
xe-0/0/36.0         enabled

```

## show dcbx neighbors

<b>Syntax</b>	<b>show dcbx neighbors</b> <interface <i>interface-name</i> > <terse>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.3 for EX Series switches.
<b>Description</b>	Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.
<b>Options</b>	<b>none</b> —Display information about all DCBX neighbor interfaces.  <b>interface-name</b> —(Optional) Display information for the specified interface.  <b>terse</b> —Display the specified level of output.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DCBX Autonegotiation on page 5669</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 5595</a></li> <li>• <a href="#">Example: Configuring an FCoE Transit Switch</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 5515</a></li> <li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li> <li>• <a href="#">dcbx on page 5685</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 6383</a> <a href="#">show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 6385</a> <a href="#">show dcbx neighbors terse (QFX Series) on page 6387</a> <a href="#">show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 6387</a> <a href="#">show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 6388</a> <a href="#">show dcbx neighbors (EX4500 Switch: Includes ETS) on page 6389</a>
<b>Output Fields</b>	<a href="#">Table 464 on page 5724</a> lists the output fields for the <b>show dcbx neighbors</b> command. Output fields are listed in the approximate order in which they appear.

Table 592: show dcbx neighbors Output Fields

Field Name	Field Description
Interface	Name of the interface.

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Parent Interface	Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.
Active-application-map	Name of the application map applied to the interface.
Protocol-Mode	<p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> <li>IEEE DCBX Version—The interface uses IEEE DCBX mode.</li> <li>DCBX Version 1.01—The interface uses DCBX version 1.01.</li> </ul> <p><b>NOTE:</b> On interfaces that use the IEEE DCBX mode, the <b>show dcbx neighbors interface <i>interface-name</i></b> operational command does not include application, PFC, or ETS operational state in the output.</p>
Protocol-State	<p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li><b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> <li><b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.</li> </ul>
Local-Advertisement	<p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>
Operational version	Version of the DCBX standard used.
sequence-number	<p>Number of state change messages sent to the peer.</p> <p>If the interface <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p> <p>If the interface <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>acknowledge-id</b> number in the <b>Peer-Advertisement</b> section.</p>
acknowledge-id	<p>Number of acknowledge messages received from the peer.</p> <p>If the <b>Protocol-State</b> value is <b>in-sync</b>, this number should match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p> <p>If the <b>Protocol-State</b> value is <b>ack-pending</b>, this number does not match the <b>sequence-number</b> value in the <b>Peer-Advertisement</b> section.</p>

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Peer-Advertisement</b>	(DCBX Version 1.01 only)  Status of advertisements that the peer sends to the local interface.
<b>Operational version</b>	Version of the DCBX standard used.
<b>sequence-number</b>	<p>Number of state change messages the peer sent to the local interface.</p> <p>If this number matches the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized.</p> <p>If this number does not match the <b>acknowledge-id</b> number in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.</p>
<b>acknowledge-id</b>	<p>Number of acknowledge messages the peer has received from the local interface.</p> <p>If this number matches the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization.</p> <p>If this number does not match the <b>sequence-number</b> value in the <b>Local-Advertisement</b> field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.</p>

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Feature: PFC</b>	Priority-based flow control (PFC) feature DCBX state information.
<b>Protocol-State</b>	(DCBX Version 1.01 only)  DCBX protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—PFC autonegotiation is disabled.</li> </ul>
<b>Operational State</b>	(DCBX Version 1.01 only)  Operational state of the feature: <b>enabled</b> or <b>disabled</b> .
<b>Local-Advertisement</b>	Status of advertisements that the local interface sends to the peer.
<b>Enable</b>	(DCBX Version 1.01 only)  State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>
<b>Willing</b>	Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the PFC configuration from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the PFC configuration from the peer.</li> </ul>
<b>Mac auth Bypass Capability</b>	(IEEE DCBX only)  (QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is <b>no</b> .
<b>Error</b>	(DCBX Version 1.01 only)  Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Operational State</b>	<p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>
<b>Maximum Traffic Classes capable to support PFC</b>	<p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>
<b>Code Point</b>	<p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>
<b>Admin Mode</b>	<p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>
<b>Operational Mode</b>	<p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> <li>• <b>Enable</b>—PFC is enabled on the code point.</li> <li>• <b>Disable</b>—PFC is disabled on the code point.</li> </ul>
<b>Peer-Advertisement</b>	<p>Status of advertisements that the peer sends to the local interface.</p>
<b>Enable</b>	<p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>
<b>Willing</b>	<p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the PFC configuration from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the PFC configuration from the local interface.</li> </ul>
<b>Error</b>	<p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>



Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Operational State</b>	<p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled on the interface</li> <li>• <b>Disabled</b>—PFC is disabled on the interface</li> </ul>
<b>Mac auth Bypass Capability</b>	<p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The connected peer supports MAC authentication bypass.</li> <li>• <b>No</b>—The connected peer does not support MAC authentication bypass.</li> </ul>
<b>Maximum Traffic Classes capable to support PFC</b>	<p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> <li>• <b>6</b> (EX Series switches)</li> <li>• <b>8</b> (QFX Series)</li> </ul>
<b>Code Point</b>	<p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>
<b>Admin Mode</b>	<p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—PFC is enabled for the code point.</li> <li>• <b>Disabled</b>—PFC is disabled for the code point.</li> </ul>

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Feature: Application</b>	State information for the DCBX application.
<b>Protocol-State</b>	<p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface.</li> <li>• <b>not-applicable</b>—The local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.</li> </ul>
<b>Local-Advertisement</b>	<p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to <b>no-auto-negotiation</b> (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>
<b>Enable</b>	<p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>
<b>Willing</b>	<p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The local interface is willing to learn the FCoE interface state from the peer.</li> <li>• <b>No</b>—The local interface is not willing to learn the FCoE interface state from the peer.</li> </ul>
<b>Error</b>	<p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. The local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. The local and peer configuration are not compatible.</li> </ul>
<b>Appl-Name</b>	Name of the application:

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Ethernet-Type</b>	<p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>
<b>Socket-Number</b>	<p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>
<b>Priority-Field or Priority-Map</b>	<p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>
<b>Status</b>	<p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul> <p><b>NOTE:</b> If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p>
<b>Peer-Advertisement</b>	<p>Status of advertisements that the peer sends to the local interface.</p>
<b>Enable</b>	<p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>
<b>Willing</b>	<p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> <li>• <b>Yes</b>—The peer is willing to learn the FCoE interface state from the local interface.</li> <li>• <b>No</b>—The peer is not willing to learn the FCoE interface state from the local interface.</li> </ul>

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Error</b>	(DCBX Version 1.01 only)  Configuration compatibility error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error detected. Local and peer configuration are compatible.</li> <li>• <b>Yes</b>—Error detected. Local and peer configuration are not compatible.</li> </ul>
<b>Appl-Name</b>	Name of the application: <ul style="list-style-type: none"> <li>• <b>FCoE</b>—Fibre Channel over Ethernet</li> </ul>
<b>Ethernet-Type</b>	Ethernet type (EtherType) of the application. For example, <b>0x8906</b> indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.
<b>Socket-Number</b>	Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.
<b>Priority-Field or Priority-Map</b>	Priority assigned to the application.
<b>Status</b>	(DCBX Version 1.01 only)  Peer interface status: <ul style="list-style-type: none"> <li>• <b>Enabled</b>—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.)</li> <li>• <b>Disabled</b>—The local configuration and the peer configuration do not match.</li> </ul>

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
<b>Feature: ETS</b>	Enhanced Transmission Selection (ETS) DCBX state information.
<b>Protocol-State</b>	(DCBX Version 1.01 only)  ETS protocol state synchronization status: <ul style="list-style-type: none"> <li>• <b>in-sync</b>—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> <li>• <b>ack-pending</b>—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.</li> </ul>
<b>Operational State</b>	(DCBX Version 1.01 only)  Operational state of the feature, <b>enabled</b> or <b>disabled</b> .
<b>Local-Advertisement</b>	Status of advertisements that the local interface sends to the peer.
<b>Enable</b>	(DCBX Version 1.01 only)  State that the local interface advertises to the peer: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>
<b>TLV Type</b>	(IEEE DCBX only)  Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Recommendation-or-Configuration</b>—Advertises both TLVs.</li> </ul>
<b>Willing</b>	Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise <b>No</b> for this field): <ul style="list-style-type: none"> <li>• <b>Yes</b>—Local interface is willing to learn the ETS state from the peer.</li> <li>• <b>No</b>—Local interface is not willing to learn the ETS state from the peer.</li> </ul>
<b>Credit Based Shaper</b>	

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	(IEEE DCBX only)
	Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .
<b>Error</b>	(DCBX Version 1.01 only)  Configuration error status: <ul style="list-style-type: none"> <li>• <b>No</b>—No error. This should always be the switch ETS error state.</li> <li>• <b>Yes</b>—Error detected.</li> </ul>
<b>Maximum Traffic Classes capable to support PFC</b>	(DCBX Version 1.01 only)  Largest number of traffic classes the local interface supports for PFC.
<b>Maximum Traffic Classes supported</b>	(IEEE DCBX only)  Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
<b>Code Point</b>	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
<b>Priority-Group</b>	Class-of-service (CoS) priority group (forwarding class set) identification number.
<b>Percentage B/W</b>	Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
<b>Transmission Selection Algorithm</b>	(IEEE DCBX only)  The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .
<b>Peer-Advertisement</b>	Status of advertisements that the peer sends to the local interface.
<b>Enable</b>	

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	(DCBX Version 1.01 only)  State that the peer advertises to the local interface: <ul style="list-style-type: none"> <li>• <b>Yes</b>—The feature is enabled.</li> <li>• <b>No</b>—The feature is disabled.</li> </ul>
<b>TLV Type</b>	(IEEE DCBX only)  Type of ETS TLV: <ul style="list-style-type: none"> <li>• <b>Configuration</b>—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration.</li> <li>• <b>Recommendation</b>—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration.</li> <li>• <b>Configuration/Recommendation</b>—Advertises both TLVs.</li> </ul>
<b>Willing</b>	Willingness of the peer to learn the ETS state from the local interface using DCBX: <ul style="list-style-type: none"> <li>• <b>Yes</b>—Peer is willing to learn the ETS state from the local interface.</li> <li>• <b>No</b>—Peer is not willing to learn the ETS state from the local interface.</li> </ul>
<b>Credit Based Shaper</b>	(IEEE DCBX only)  Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always <b>No</b> .
<b>Error</b>	(DCBX Version 1.01 only)  Configuration error status of the peer: <ul style="list-style-type: none"> <li>• <b>No</b>—No error in peer ETS TLV.</li> <li>• <b>Yes</b>—Error in peer ETS TLV.</li> </ul>
<b>Maximum Traffic Classes capable to support PFC</b>	(DCBX Version 1.01 only)  Largest number of traffic classes the local interface supports for PFC.
<b>Maximum Traffic Classes supported</b>	(IEEE DCBX only)  Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
<b>Code Point</b>	

Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
<b>Priority-Group</b>	CoS priority group (forwarding class set) identification number.
<b>Percentage B/W</b>	Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
<b>Transmission Selection Algorithm</b>	(IEEE DCBX only)  Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is <b>ETS</b> .
<b>PFC</b>	(QFX Series, <b>terse</b> option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> <li>• Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled</li> <li>• Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled</li> <li>• Not Advt—Interface does not advertise PFC to the connected peer</li> </ul>
<b>ETS</b>	( <b>terse</b> option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> <li>• Advt—Interface advertises ETS TLVs</li> <li>• Disabled—ETS is disabled on the interface (interface does not advertise ETS)</li> </ul>
<b>ETS Rec</b>	( <b>terse</b> option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> <li>• Advt—Peer interface advertises ETS TLVs</li> <li>• Not Advt—Peer interface does not advertise ETS</li> </ul> <p><b>NOTE:</b> When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>



Table 592: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Version	<p>(<b>terse</b> option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> <li>• <b>IEEE</b>—The interface uses IEEE DCBX.</li> <li>• <b>1.01</b>—The interface uses DCBX version 1.01.</li> </ul> <p>When the DCBX version used is the result of autonegotiation, the term (<b>Auto</b>) appears next to the version. For example, <b>IEEE (Auto)</b> indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p>

## Sample Output

### show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
  Operational version: 1
  sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
  Operational version: 1
  sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode      Operational Mode
000             Disabled       Disable
001             Disabled       Disable
010             Disabled       Disable
011             Enabled        Enable
100             Enabled        Enable
101             Disabled       Disable
110             Disabled       Disable
111             Disabled       Disable

Peer-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode
000             Disabled

```

001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001110	Enabled
iSCSI		3260	10000000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906	N/A	00001110	Enabled

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1

111	7
Priority-Group	Percentage B/W
0	40%
1	5%

### show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

```
user@switch> show dcbx neighbors interface xe-0/0/0
```

```
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
```

```
Active-application-map: app-map-1
```

```
Protocol-Mode: IEEE-DCBX Version
```

```
Feature: PFC
```

```
Local-Advertisement:
```

```
Willing: No
```

```
Mac auth Bypass Capability: No
```

```
Operational State: Enabled
```

```
Maximum Traffic Classes capable to support PFC: 8
```

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

```
Peer-Advertisement:
```

```
Willing: No
```

```
Mac auth Bypass Capability: No
```

```
Operational State: Enabled
```

```
Maximum Traffic Classes capable to support PFC: 8
```

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

```
Feature: Application
```

```
Local-Advertisement:
```

Appl-Name	Ethernet-Type	Socket-Number	Priority-field
FCoE	0x8906		00001110
iSCSI		3260	10000000

```
Peer-Advertisement:
```

Appl-Name	Ethernet-Type	Socket-Number	Priority-field
-----------	---------------	---------------	----------------

FCoE	0x8906	N/A	00001110
------	--------	-----	----------

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

Peer-Advertisement:

TLV Type: Recommendation

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0

101	1
110	1
111	7
Priority-Group	Percentage B/W
0	40%
1	5%
Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

### show dcbx neighbors terse (QFX Series)

```

user@switch> show dcbx neighbors terse
Interface Parent PFC ETS ETS Version
Interface
xe-0/0/8.0 - Enabled Advt Advt IEEE (Auto)
xe-0/0/9.0 - Disabled Disabled 1.01
xe-0/0/11.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/12.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/32.0 - Enabled Advt Not Advt IEEE
xe-0/0/36.0 - Not Advt Advt Advt IEEE

```

### show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)

```

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
  Operational version: 0
  sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
  Operational version: 0
  sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 6

Code Point      Admin Mode
000             Disabled
001             Disabled
010             Disabled
011             Enabled
100             Disabled
101             Disabled
110             Disabled
111             Disabled

```

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Status	Appl-Name	Ethernet-Type	Socket-Number	Priority-Map
Enabled	FCoE	0x8906		00001000

**show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)**

user@switch&gt; show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

## Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

## Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

## Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled
iscsi		3260	00100000	Enabled

## Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled
iscsi		3260	00100000	Enabled

**show dcbx neighbors (EX4500 Switch: Includes ETS)**

user@switch&gt; show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0  
 Protocol-State: in-sync  
 Active-application-map: map\_iscsi

## Local-Advertisement:

Operational version: 0

sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Enabled
001	Enabled
010	Disabled
011	Disabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Enabled
001	Disabled
010	Disabled
011	Disabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00000001	Enabled
iscsi		3260	00000010	Enabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00010000	Enabled
iscsi		3260	00010000	Enabled

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled



## Local-Advertisement:

Enable: Yes, Willing: No, Error: No  
Maximum Traffic Classes supported : 3

Code Point	Priority-Group
000	7
001	7
010	7
011	7
100	7
101	7
110	7
111	7
Priority-Group	Percentage B/W
7	100%

## Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No  
Maximum Traffic Classes supported : 8

Code Point	Priority-Group
000	0
001	1
010	0
011	0
100	2
101	0
110	0
111	0
Priority-Group	Percentage B/W
0	30%
1	40%
2	30%

## show interfaces queue

---

**Syntax**    show interfaces queue  
              <aggregate | remaining-traffic>  
              <both-ingress-egress>  
              <egress>  
              <forwarding-class *forwarding-class*>  
              <ingress>  
              <interface-name *interface-name*>  
              <l2-statistics>

**Release Information**    Command introduced before Junos OS Release 7.4.  
                              **both-ingress-egress**, **egress**, and **ingress** options introduced in Junos OS Release 7.6.  
                              Command introduced in Junos OS Release 11.1 for the QFX Series.  
                              **l2-statistics** option introduced in Junos OS Release 12.1.

**Description**    Display class-of-service (CoS) queue information for physical interfaces.

**Options**    **none**—Show detailed CoS queue statistics for all physical interfaces.

**aggregate**—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)

**both-ingress-egress**—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)

**egress**—(Optional) Display egress queue statistics.

**forwarding-class *forwarding-class***—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.

**ingress**—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)

**interface-name *interface-name***—(Optional) Show detailed CoS queue statistics for the specified interface.

**l2-statistics**—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles

**remaining-traffic**—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 257 on page 2813](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 257 on page 2813](#).

Table 593: Layer 2 Overhead, Transmitted Packets/Bytes

Protocol	Fragmentation		LFI
	First fragmentation	Second to <i>n</i> fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

## Layer 2 Statistics—Fragmentation Overhead Calculation

## MLPPP/MC-MLPPP Overhead details:

=====

## Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

## Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

## MLFR (FRF15) Overhead details:

=====

## Fragment 1:

```

Framereley header         : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

## Fragments 2 ...n :

```

Framereley header         : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

## MFR (FRF16) Overhead details:

=====

Fragment 1:  
Fragmentation header : 2 bytes  
Framereplay header : 2 bytes  
Inner proto : 2 bytes  
HDLC flag and FCS : 4 bytes

Fragments 2 ...n :  
Fragmentation header : 2 bytes  
Framereplay header : 2 bytes  
HDLC flag and FCS : 4 bytes

## Overhead with LFI

MLPPP(Long & short sequence):  
=====

Outer PPP header	: 4 bytes
HDLC flag and FCS	: 4 bytes

MLFR (FRF15):  
=====

Framereplay header	: 2 bytes
Control,NLPID	: 2 bytes
HDLC flag and FCS	: 4 bytes

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

## Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

**Required Privilege Level**

view

**List of Sample Output**

[show interfaces queue \(Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC\) on page 6400](#)  
[show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 6401](#)  
[show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 6403](#)  
[show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 6403](#)  
[show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 6404](#)  
[show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 6408](#)  
[show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 6411](#)  
[show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 6413](#)  
[show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 6414](#)  
[show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 6415](#)  
[show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 6418](#)  
[show interfaces queue \(QFX Series\) on page 6428](#)  
[show interfaces queue l2-statistics \(lsq interface\) on page 6429](#)  
[show interfaces queue lsq \(lsq-ifd\) on page 6429](#)

**Output Fields** Table 258 on page 2816 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

**Table 594: show interfaces queue Output Fields**

Field Name	Field Description
<b>Physical interface</b>	Name of the physical interface.
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.
<b>SNMP ifindex</b>	SNMP index number for the interface.
<b>Forwarding classes supported</b>	Total number of forwarding classes supported on the specified interface.
<b>Forwarding classes in use</b>	Total number of forwarding classes in use on the specified interface.
<b>Ingress queues supported</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
<b>Ingress queues in use</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
<b>Output queues supported</b>	Total number of output queues supported on the specified interface.
<b>Output queues in use</b>	Total number of output queues in use on the specified interface.
<b>Egress queues supported</b>	Total number of egress queues supported on the specified interface.
<b>Egress queues in use</b>	Total number of egress queues in use on the specified interface.
<b>Queue counters (Ingress)</b>	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>
<b>Burst size</b>	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command:	
<b>Queue</b>	Queue number.
<b>Forwarding classes</b>	Forwarding class name.

Table 594: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
<b>Queued Packets</b>	<p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>
<b>Queued Bytes</b>	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 259 on page 2819</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>
<b>Transmitted Packets</b>	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 2812</a></p>
<b>Transmitted Bytes</b>	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 259 on page 2819</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 2812</a></p>
<b>Tail-dropped packets</b>	Number of packets dropped because of tail drop.
<b>RL-dropped packets</b>	<p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>
<b>RL-dropped bytes</b>	<p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 2814</a>.</p>

Table 594: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 259 on page 2819</a>.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> <li>(J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority bytes dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority bytes dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Byte counts vary by interface hardware. [Table 259 on page 2819](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware.

[Table 259 on page 2819](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.



Table 595: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet.</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>• Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>• Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>• RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted.</li> </ul>	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 595: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul>	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
  Transmitted:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
    Tail-dropped packets :              0              0 pps
    RL-dropped packets :          205050862          602295 pps
    RL-dropped bytes   :          13595326612       327648832 bps
    RED-dropped packets :              0              0 pps
      Low              :              0              0 pps
      Medium-low       :              0              0 pps
      Medium-high      :              0              0 pps
      High             :              0              0 pps
    RED-dropped bytes   :              0              0 bps
      Low              :              0              0 bps
      Medium-low       :              0              0 bps
      Medium-high      :              0              0 bps
      High             :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps
```

#### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :              5              0 pps
    Bytes            :              242              0 bps
  Transmitted:
    Packets          :              5              0 pps
    Bytes            :              242              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets :              0              0 pps
    RED-dropped bytes   :              0              0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets          :          42603765          595484 pps
```

```

Bytes          :          5453281920          609776496 bps
Transmitted:
Packets        :          42603765          595484 pps
Bytes          :          5453281920          609776496 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets        :          45          0 pps
Bytes          :          3930          0 bps
Transmitted:
Packets        :          45          0 pps
Bytes          :          3930          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 4, Forwarding classes: af11
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 5, Forwarding classes: ef11
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets        :          31296413          437436 pps
Bytes          :          4005940864          447935200 bps
Transmitted:
Packets        :          31296413          437436 pps
Bytes          :          4005940864          447935200 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
RED-dropped bytes  :          0          0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps

```

```

Transmitted:
Packets      :                0                0 pps
Bytes        :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps

```

#### show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :                5240762                3404 pps
    Bytes        :            3020710354            15934544 bps
  Transmitted:
    Packets      :                5240762                3404 pps
    Bytes        :            3020710354            15934544 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                0                0 pps
    RED-dropped bytes :                0                0 bps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                0                0 pps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :                2480391                1650 pps
    Bytes        :            1304685666            6945704 bps
  Transmitted:
    Packets      :                2478740                1650 pps
    Bytes        :            1303817240            6945704 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                1651                0 pps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                1651                0 pps
    RED-dropped bytes :                868426                0 bps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                868426                0 pps

```

#### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:

```

```

Packets      :      13      0 pps
Bytes        :      622      0 bps
Transmitted:
Packets      :      13      0 pps
Bytes        :      622      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: af1
Queued:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Transmitted:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Transmitted:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Transmitted:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps

```

#### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets      :      148450735      947295 pps
Bytes        :      8016344944      409228848 bps
Transmitted:
Packets      :      76397439      487512 pps
Bytes        :      4125461868      210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets :      72053285      459783 pps
Low          :      72053285      459783 pps
Medium-low   :      0      0 pps
Medium-high  :      0      0 pps
High         :      0      0 pps
RED-dropped bytes  :      3890877444      198626472 bps

```

```

Low : 3890877444 198626472 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 410278257 473940 pps
Bytes : 22156199518 204742296 bps
Transmitted:
Packets : 4850003 4033 pps
Bytes : 261900162 1742256 bps
Tail-dropped packets : Not Available
RED-dropped packets : 405425693 469907 pps
Low : 405425693 469907 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 21892988124 203000040 bps
Low : 21892988124 203000040 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```

```

Queued:
  Packets      :          76605230          485376 pps
  Bytes       :          5209211400        264044560 bps
Transmitted:
  Packets      :          76444631          484336 pps
  Bytes       :          5198235612        263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets :          160475          1040 pps
  Low         :          160475          1040 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :          10912300        565760 bps
  Low         :          10912300        565760 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Transmitted:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Tail-dropped packets : Not Available
RED-dropped packets :              0              0 pps
  Low         :              0              0 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :              0              0 bps
  Low         :              0              0 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          4836136          3912 pps
  Bytes       :          333402032        2139056 bps
Transmitted:
  Packets      :          3600866          1459 pps
  Bytes       :          244858888        793696 bps
Tail-dropped packets : Not Available
RED-dropped packets :          1225034          2450 pps
  Low         :          1225034          2450 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :          83302312        1333072 bps
  Low         :          83302312        1333072 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Transmitted:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Tail-dropped packets : Not Available

```



RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

#### Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

##### Queued:

Packets	:	77059796	486384 pps
Bytes	:	3544750624	178989576 bps

##### Transmitted:

Packets	:	77059797	486381 pps
Bytes	:	3544750670	178988248 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

##### Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

##### Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

##### Queued:

Packets	:	4846580	3934 pps
Bytes	:	222942680	1447768 bps

##### Transmitted:

Packets	:	4846580	3934 pps
Bytes	:	222942680	1447768 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps

```

#### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
  Transmitted:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps

```

```

    RED-dropped bytes      :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                :                7055              1 pps
    Bytes                  :            451552              512 bps
  Transmitted:
    Packets                :                7055              1 pps
    Bytes                  :            451552              512 bps
    Tail-dropped packets : Not Available
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :                1031              0 pps
    Bytes                  :            143292              0 bps
  Transmitted:
    Packets                :                1031              0 pps
    Bytes                  :            143292              0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes     :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes     :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes     :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                :                77009             11 pps
    Bytes                  :            6894286             7888 bps
  Transmitted:
    Packets                :                77009             11 pps
    Bytes                  :            6894286             7888 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0                0 pps
    RL-dropped bytes     :                0                0 bps
    RED-dropped packets   :                0                0 pps
    RED-dropped bytes     :                0                0 bps

```

## Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

## Queued:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

## Transmitted:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 3, Forwarding classes: network-control

## Queued:

Packets	:	94386	12 pps
Bytes	:	13756799	9568 bps

## Transmitted:

Packets	:	94386	12 pps
Bytes	:	13756799	9568 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

#### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                254                0 pps
    Bytes        :            16274                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps
    RED-dropped bytes  :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      : Not Available
    Bytes        :                0                0 bps
  Transmitted:
    Packets      :                0                0 pps
    Bytes        :                0                0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :                0                0 pps

```

```

    RED-dropped bytes      :                0          0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                3          0 pps
    Bytes                  :               126          0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0          0 pps
    RED-dropped bytes      :                0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0          0 pps
    RED-dropped bytes      :                0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0          0 pps
    RED-dropped bytes      :                0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                : Not Available
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0          0 pps
    RED-dropped bytes      :                0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :             80564692          0 pps
    Bytes                  :          3383717100          0 bps
  Transmitted:
    Packets                :             80564692          0 pps
    Bytes                  :          3383717100          0 bps
    Tail-dropped packets   :                0          0 pps
    RED-dropped packets    :                0          0 pps
    RED-dropped bytes      :                0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :             80564685          0 pps
    Bytes                  :          3383716770          0 bps
  Transmitted:
    Packets                :             80564685          0 pps

```

```

Bytes : 3383716770 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Transmitted:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 288 0 pps
Bytes : 18450 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```



```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
  Transmitted:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

#### show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
  Transmitted:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :      0      0 pps
    Low          :      0      0 pps

```

```

Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps

```

```

      High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 109355853 471736 pps
    Bytes : 7436199152 256627968 bps
  Transmitted:
    Packets : 109355852 471736 pps
    Bytes : 7436198640 256627968 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps

```

```
Transmitted:
Packets      :                0                0 pps
Bytes        :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low        :                0                0 pps
  Medium-low :                0                0 pps
  Medium-high:                0                0 pps
  High       :                0                0 pps
RED-dropped bytes :                0                0 bps
  Low        :                0                0 bps
  Medium-low :                0                0 bps
  Medium-high:                0                0 bps
  High       :                0                0 bps
```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```
user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

  Interface index: 192, SNMP ifIndex: 1948

  Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
  Lam

  Forwarding classes: 16 supported, 9 in use

  Egress queues: 8 supported, 8 in use

  Queue: 0, Forwarding classes: DEFAULT

  Queued:

    Packets      :                214886                13449 pps
    Bytes        :                9884756            5164536 bps

  Transmitted:

    Packets      :                214886                13449 pps
    Bytes        :                9884756            5164536 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
      Low        :                0                0 pps
      Medium-low :                0                0 pps
      Medium-high:                0                0 pps
      High       :                0                0 pps
    RED-dropped bytes :                0                0 bps
      Low        :                0                0 bps
      Medium-low :                0                0 bps
```

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Low	:	0	0 bps
-----	---	---	-------

Medium-low	:	0	0 bps
------------	---	---	-------

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High	:	0	0 bps
------	---	---	-------

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps



High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

#### Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

##### Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

##### Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

##### Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps
Queue: 2, Forwarding classes: PRIVATE			
Queued:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Transmitted:			
Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	32843	0 pps
---------	---	-------	-------

Bytes	:	2641754	56 bps
-------	---	---------	--------

Transmitted:

Packets	:	32843	0 pps
---------	---	-------	-------

Bytes	:	2641754	56 bps
-------	---	---------	--------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Low	:	0	0 bps
-----	---	---	-------

Medium-low	:	0	0 bps
------------	---	---	-------

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High : 0 0 bps

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available

```

Total-dropped packets:	0	0 pps
Total-dropped bytes :	0	0 bps

#### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           1           0 pps
    Bytes        :          1001          0 bps
  Transmitted:
    Packets      :           5           0 pps
    Bytes        :          1062          0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :           1           0 pps
    Bytes        :          1500          0 bps
  Transmitted:
    Packets      :           6           0 pps
    Bytes        :          1573          0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :           1           0 pps
    Bytes        :           512          0 bps
  Transmitted:
    Packets      :           3           0 pps
    Bytes        :           549          0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
=====

```

#### show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0

```

## Queue: 0, Forwarding classes: be

## Queued:

Packets	:	55576	1206 pps
Bytes	:	29622008	5145472 bps

## Transmitted:

Packets	:	55576	1206 pps
Bytes	:	29622008	5145472 bps
Tail-dropped packets	:	0	0 pps
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 1, Forwarding classes: ef

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 2, Forwarding classes: af

## Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

## Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## Queue: 3, Forwarding classes: nc



Queued:			
Packets	:	22231	482 pps
Bytes	:	11849123	2057600 bps
Transmitted:			
Packets	:	22231	482 pps
Bytes	:	11849123	2057600 bps
Tail-dropped packets	:	0	0 pps
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## show pfe filter hw summary

<b>Syntax</b>	show pfe filter hw summary
<b>Release Information</b>	Command introduced in Junos OS Release 14.1X53-D10 for the QFX Series.
<b>Description</b>	<p>Display a summary of the access control list (ACL; also known as firewall filter) ternary content-addressable memory (TCAM) hardware utilization to show the allocated, used, and free TCAM entry space.</p> <p>Command supported on standalone QFX Series switches, QFX5100-only (pure QFX5100) Virtual Chassis Fabric (VCF), QFX5100-only (pure QFX5100) Virtual Chassis (VC), and QFX3500-only (pure QFX3500) VC.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Planning the Number of Firewall Filters to Create on page 5236</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show pfe summary on page 6433</a>
<b>Output Fields</b>	<a href="#">Table 596 on page 6432</a> lists the output fields for the <b>show pfe filter</b> command. Output fields are listed in the approximate order in which they appear.

**Table 596: show pfe filter Output Fields**

Field Name	Field Description
<b>Group</b>	<p>ACL ingress and egress filter groups:</p> <ul style="list-style-type: none"> <li>• iRACL group—ingress routing ACL filter group</li> <li>• iVACL group—ingress VLAN ACL filter group</li> <li>• iPACL group—ingress port ACL filter group</li> <li>• ePACL group—egress port ACL filter group</li> <li>• eVACL group—egress VLAN ACL filter group</li> <li>• eRACL group—egress routing ACL filter group</li> <li>• eRACL IPv6 group—egress IPv6 routing ACL filter group</li> </ul>
<b>Group-ID</b>	Internal identification number of the filter group.
<b>Allocated</b>	Number of TCAM filter entries allocated to the filter group.
<b>Used</b>	Number of TCAM filter entries used by the filter group.
<b>Free</b>	Number of TCAM filter entries available for use by the filter group.

## Sample Output

### show pfe summary

```
user@switch> show pfe summary
```

Group	Group-ID	Allocated	Used	Free
-----				
> Ingress filter groups:				
iRACL group	14	512	4	508
iVACL group	13	512	2	510
iPACL group	12	256	2	254
> Egress filter groups:				
ePACL group	20	256	3	253
eVACL group	21	256	4	252
eRACL group	22	256	245	11
eRACL IPV6 group	24	256	3	253

## show pfe next-hop

---

<b>List of Syntax</b>	<a href="#">Syntax on page 6434</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 6434</a>
<b>Syntax</b>	<code>show pfe next-hop</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<code>show pfe next-hop</code> <code>&lt;fpc <i>slot</i>&gt;</code> <code>&lt;interface <i>interface-name</i>&gt;</code> <code>&lt;lcc <i>number</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Packet Forwarding Engine next-hop information.
<b>Options</b>	<p><b>none</b>—Display all Packet Forwarding Engine next-hop information.</p> <p><b>fpc <i>slot</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"><li>On a TX Matrix router, if you specify the number of a T640 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 31.</li><li>On a TX Matrix Plus router, if you specify the number of a T1600 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 31.</li><li>On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the <b>lcc <i>number</i></b> option (the recommended method), replace <b><i>slot</i></b> with a value from 0 through 7. Otherwise, replace <b><i>slot</i></b> with a value from 0 through 63.</li></ul> <p>For example, the following commands have the same result:</p> <pre>user@host&gt; show pfe next-hop fpc 1 lcc 1 user@host&gt; show pfe next-hop fpc 9</pre> <p><b>interface <i>interface-name</i></b>—(Optional) Display the Packet Forwarding Engine next-hop interface.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine next-hop interface for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine next-hop interface for the router (or line-card chassis) that is connected to a TX Matrix Plus router.</p>

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**Required Privilege Level** admin

**Related Documentation**

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**

[show pfe next-hop on page 6436](#)  
[show pfe next-hop fpc \(TX Matrix Router\) on page 6436](#)  
[show pfe next-hop fpc \(TX Matrix Plus Router\) on page 6437](#)

**Output Fields** Table 597 on page 6436 lists the output fields for the **show pfe next-hop** command. Output fields are listed in the approximate order in which they appear.

**Table 597: show pfe next-hop Output Fields**

Field Name	Field Description
ID	The next-hop ID for the entry.
Type	The next-hop type for the entry.
Interface	The interface to which the next-hop entry is assigned.
Protocol	The protocol type for the next-hop entry.
Encap	Encapsulation type for the next-hop entry.
Next Hop Addr	Next-hop address for the next-hop entry.
MTU	MTU value for the nexthop entry.

## Sample Output

### show pfe next-hop

```

user@host> show pfe next-hop
Nexthop Info:
  ID      Type      Interface      Protocol      Encap      Next Hop Addr      MTU
  ----      -      -      -      -      -      -
  4         Mcast      -              IPv4          -          0.0.0.0             0
  5         Bcast      -              IPv4          -          -                   0
  7         Discard     -              IPv4          -          -                   0
  8         MDiscard    -              IPv4          -          -                   0
  9         Reject      -              IPv4          -          -                   0
  13        Local      -              IPv4          -          192.168.4.60        0
  14        Resolve    fxp0.0         IPv4          Unspecified   -                   0
  17        Local      -              IPv4          -          127.0.0.1           0
  18        Unicast     fxp0.0         IPv4          Unspecified   192.168.4.254       0
  21        Local      -              IPv4          -          11.1.0.1            0
  22        Unicast     at-0/1/0.0     IPv4          ATM SNAP      11.1.0.2            4482
  ...

```

### show pfe next-hop fpc (TX Matrix Router)

```

user@host> show pfe next-hop fpc 1
Slot 1
Nexthop Info:
  ID      Type      Interface      Next Hop Addr      Protocol      Encap      MTU
  ----      -      -      -      -      -      -
  5         Mcast      -              default            IPv4          -          0
  6         Bcast      -              -                  IPv4          -          0
  8         Discard     -              -                  IPv4          -          0
  9         MDiscard    -              -                  IPv4          -          0
  13        Mcast      -              default            IPV6          -          0
  17        MDiscard    -              -                  IPV6          -          0
  18        Reject      -              -                  IPV6          -          0
  24        Discard     -              -                  None          -          0

```

```

68      Local -          192.168.66.113      IPv4      -      0
69      Resolve fxp0.0    -          IPv4      Unspecified 0
70      Unicast fxp0.0    192.168.71.254 IPv4      Unspecified 0
256     Local -          10.71.71.1        IPv4      -      0
257     Local -          127.0.0.1        IPv4      -      0
258     Mcast.local..1   default    IPv4      Unspecified 0
259     Bcast.local..1   -          IPv4      Unspecified 0
261     Discard.local..1 -          IPv4      Unspecified 0
262     MDiscard.local..1 -          IPv4      Unspecified 0
269     Mcast.local..1   default    IPV6      Unspecified 0
271     Discard.local..1 -          IPV6      Unspecified 0
...

```

### show pfe next-hop fpc (TX Matrix Plus Router)

```
user@host> show pfe next-hop fpc 0
```

Slot 0

ID	Type	Interface	Next Hop Addr	Protocol	Encap	MTU
31	Mcast	-	default	IPv4	-	0
32	Bcast	-	-	IPv4	-	0
34	Discard	-	-	IPv4	-	0
35	MDiscard	-	-	IPv4	-	0
36	Reject	-	-	IPv4	-	0
39	Mcast	-	default	IPv6	-	0
42	Discard	-	-	IPv6	-	0
43	MDiscard	-	-	IPv6	-	0
44	Reject	-	-	IPv6	-	0
49	Receive	-	-	MPLS	-	0
50	Discard	-	-	MPLS	-	0
111	Mcast	.local..1	default	IPv4	Unspecified	0
112	Bcast	.local..1	-	IPv4	Unspecified	0
114	Discard	.local..1	-	IPv4	Unspecified	0
115	MDiscard	.local..1	-	IPv4	Unspecified	0
116	Reject	.local..1	-	IPv4	Unspecified	0
119	Mcast	.local..1	default	IPv6	Unspecified	0
122	Discard	.local..1	-	IPv6	Unspecified	0
123	MDiscard	.local..1	-	IPv6	Unspecified	0
124	Reject	.local..1	-	IPv6	Unspecified	0
191	Mcast	.local..2	default	IPv4	Unspecified	0
192	Bcast	.local..2	-	IPv4	Unspecified	0
194	Discard	.local..2	-	IPv4	Unspecified	0
195	MDiscard	.local..2	-	IPv4	Unspecified	0
196	Reject	.local..2	-	IPv4	Unspecified	0
322	Local	-	10.1.0.5	IPv4	-	0
323	Resolve	bcm0.0	-	IPv4	Unspecified	0
326	Local	-	129.0.0.5	IPv4	-	0
327	Resolve	bcm0.0	-	IPv4	Unspecified	0
328	Local	-	fe80::201:ff:fe01:5	IPv6	-	0
329	Receive	bcm0.0	ff02::1:ff01:5	IPv6	Unspecified	0
330	Receive	bcm0.0	fe80::	IPv6	Unspecified	0
331	Resolve	bcm0.0	-	IPv6	Unspecified	0
332	Local	-	fec0::a:1:0:5	IPv6	-	0
333	Receive	bcm0.0	ff02::1:ff00:5	IPv6	Unspecified	0
334	Receive	bcm0.0	fec0::	IPv6	Unspecified	0
335	Resolve	bcm0.0	-	IPv6	Unspecified	0
348	Local	-	192.168.178.4	IPv4	-	0
349	Resolve	em0.0	-	IPv4	Unspecified	0

350	Unicast	em0.0	192.168.178.126	IPv4	Unspecified	0
357	Local	-	fe80::201:1ff:fe01:5	IPv6	-	0
512	Local	-	10.255.178.11	IPv4	-	0
513	Local	-	127.0.0.1	IPv4	-	0
515	Local	-	abcd::10:255:178:11	IPv6	-	0
516	Local	-	fe80::200:ff:fe00:0	IPv6	-	0
517	Local	-	127.0.0.1	IPv4	-	0
518	Mcast	.local..3	default	IPv4	Unspecified	0
519	Bcast	.local..3	-	IPv4	Unspecified	0
521	Discard	.local..3	-	IPv4	Unspecified	0
522	MDiscard	.local..3	-	IPv4	Unspecified	0
523	Reject	.local..3	-	IPv4	Unspecified	0
531	Mcast	.local..3	default	IPv6	Unspecified	0
533	Discard	.local..3	-	IPv6	Unspecified	0
534	MDiscard	.local..3	-	IPv6	Unspecified	0
535	Reject	.local..3	-	IPv6	Unspecified	0
539	Mgroup	-	-	IPv4	-	0
540	Bcast	ge-15/0/3.0	-	IPv4	Ethernet	0
541	Receive	ge-15/0/3.0	14.2.1.0	IPv4	Ethernet	0
542	Local	-	14.2.1.1	IPv4	-	0
543	Resolve	ge-15/0/3.0	-	IPv4	Ethernet	0
544	Bcast	ge-31/0/4.0	-	IPv4	Ethernet	0
545	Receive	ge-31/0/4.0	14.1.1.0	IPv4	Ethernet	0
546	Local	-	14.1.1.1	IPv4	-	0
547	Resolve	ge-31/0/4.0	-	IPv4	Ethernet	0
548	Unicast	ge-31/0/4.0	14.1.1.2	IPv4	Ethernet	0
549	Unicast	ge-15/0/3.0	14.2.1.2	IPv4	Ethernet	0
550	Bcast	ae1.0	-	IPv4	Ethernet	0
551	Receive	ae1.0	11.1.1.0	IPv4	Ethernet	0
552	Local	-	11.1.1.1	IPv4	-	0
553	Resolve	ae1.0	-	IPv4	Ethernet	0
554	Aggreg.	ae1.0	-	IPv4	Ethernet	0
555	Unicast	ge-23/0/8.0	11.1.1.2	IPv4	Ethernet	0
556	Unicast	ge-7/0/9.0	11.1.1.2	IPv4	Ethernet	0
557	Aggreg.	ae1.0	-	MPLS	Ethernet	0
558	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
559	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0
560	Aggreg.	ae1.0	-	MPLS	Ethernet	0
561	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
562	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0



## show pfe route

<b>List of Syntax</b>	<a href="#">Syntax on page 6439</a> <a href="#">Syntax (EX Series Switches) on page 6439</a> <a href="#">Syntax (QFX Series) on page 6439</a> <a href="#">Syntax (MX Series) on page 6439</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 6439</a>
<b>Syntax</b>	<pre>show pfe route &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Syntax (QFX Series)</b>	<pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt; &lt;hw (host   lpm   multicast)&gt;&gt; &lt;&lt;clnp&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;mpls&gt; &lt;summary&gt; &lt;hw&gt;</pre>
<b>Syntax (MX Series)</b>	<pre>show pfe route &lt;&lt;inet6   ip&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;dhcp&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<pre>show pfe route &lt;fpc slot&gt; &lt;&lt;inet6   ip   iso&gt; &lt;prefix prefix&gt;   &lt;table &lt;table-name&gt; &lt;index index&gt; &lt;prefix prefix&gt;&gt;&gt; &lt;lcc number&gt; &lt;mpls&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.3 for the MX Series.</p> <p>Command option <b>hw</b> introduced in Junos OS Release 14.1X53-D10 for the QFX Series.</p>
<b>Description</b>	<p>Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.</p>



**NOTE:** The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the **show route forwarding table** command. For more information, see the [CLI Explorer](#).

**Options** **none**—Display all Packet Forwarding Engine forwarding table information.

**clnp**—(Optional) Show International Standards Organization (ISO) connectionless-mode network protocol (CLNP) route table information.

**dhcp**—(Optional) Display Packet Forwarding Engine DHCP-Snooping route table information.

**fpc slot**—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.

- On a TX Matrix router, if you specify the number of a T640 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **31**.
- On a TX Matrix Plus router, if you specify the number of a T1600 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **31**.
- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot** with a value from **0** through **7**. Otherwise, replace **slot** with a value from **0** through **63**.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

**host**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display host routes installed in the on-chip hardware table.

**hw**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display routes installed in the on-chip hardware table (as opposed to displaying routes from the routing table and the PFE forwarding table before they are installed in the hardware).

**index index**—(Optional) Display table index.

**inet6**—(Optional) Display Packet Forwarding Engine IPv6 routes.

**ip**—(Optional) Display Packet Forwarding Engine IPv4 routes.

**iso**—(Optional) Display ISO version routing tables.

**lcc *number***—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the router (line-card chassis) that houses the FPC.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**lpm**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display longest prefix match (LPM) routes installed in the on-chip hardware table.

**mpls**—(Optional) Display Packet Forwarding Engine MPLS information.

**multicast**—(QFX standalone switches, pure mode QFX5100-only VCF and VC, and pure mode QFX3500-only VC) (Optional) Display multicast routes installed in the on-chip hardware table.

**prefix *prefix***—(Optional) IPv4 or IPv6 prefix for which to show table entries.

**summary**—(Optional) Display summary of Packet Forwarding Engine information.

**table <*table-name*>**—(Optional) Display table information.

**Required Privilege Level**

admin

**Related Documentation**

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Routing Matrix with a TXP-Mixed-LCC-3D Configuration*

**List of Sample Output**

[show pfe route ip on page 6443](#)  
[show pfe route iso on page 6443](#)  
[show pfe route lcc summary \(TX Matrix Router\) on page 6443](#)  
[show pfe route lcc summary \(TX Matrix Plus Router\) on page 6445](#)  
[show pfe route summary \(MX Series Router\) on page 6446](#)  
[show pfe route summary hw \(QFX Series\) on page 6447](#)

[show pfe route ip hw host \(QFX Series\) on page 6447](#)

**Output Fields** [Table 598 on page 6442](#) lists the output fields for the **show pfe route** command. Output fields are listed in the approximate order in which they appear.

**Table 598: show pfe route Output Fields**

Field Name	Field Description
<b>Destination</b>	Destination address for the entry.
<b>NH IP Addr</b>	Next-hop IP address for the entry.
<b>Type</b>	Next-hop type for the entry
<b>NH ID</b>	Next-hop ID for the entry
<b>Encap</b>	Encapsulation type for the next-hop entry.
<b>Interface</b>	Interface to which the next-hop entry is assigned.

[Table 599 on page 6442](#) lists the output fields for the QFX Series **show pfe route** hardware table (**hw**) commands. Output fields are listed in the approximate order in which they appear.

**Table 599: QFX Series show pfe route Hardware Table Output Fields**

Field Name	Field Description
<b>Max</b>	Maximum routing entries per route type.
<b>Used</b>	Number of routing entries consumed per route type.
<b>Free</b>	Number of unused routing entries per route type.
<b>% Free</b>	Percentage of unused routing entries per route type.
<b>Rtt</b>	Internal routing engine index number of the route table.
<b>VRF</b>	Internal hardware index number for the corresponding route table.
<b>Destination</b>	Destination address for the entry.
<b>Type</b>	( <b>show pfe route summary hw</b> )—Route type for the entry: IPv4 or IPv6 route, and host, LPM, or multicast route.  ( <b>show pfe route (ip   inet6) hw</b> )—Next-hop type for the entry.
<b>NH ID</b>	Next-hop ID for the entry
<b>Interface</b>	Interface to which the next-hop entry is assigned.

Table 599: QFX Series show pfe route Hardware Table Output Fields (*continued*)

Field Name	Field Description
HW NH-ID	Internal hardware index number of the next-hop.
Src-MAC-Address	Source MAC address.
Port	Port number.
Dst-MAC-Address	Destination MAC address.
VLAN	ID of the multicast group VLAN.
GROUP	Internal hardware index number of the multicast group next-hop.
CLASS	Internal class number of the multicast group.

## Sample Output

### show pfe route ip

```
user@host> show pfe route ip
```

```
IPv4 Route Table 0, default.0, 0x0:
```

Destination	NH IP Addr	Type	NH ID	Interface
default		Discard	8	
127.0.0.1	127.0.0.1	Local	256	
172.16/12	192.168.71.254	Unicast	68	fxp0.0
192.168.0/18	192.168.71.254	Unicast	68	fxp0.0
192.168.40/22	192.168.71.254	Unicast	68	fxp0.0
192.168.64/18	192.168.71.254	Unicast	68	fxp0.0
192.168.64/21		Resolve	67	fxp0.0
192.168.71.249	192.168.71.249	Local	66	
192.168.220.0/30		Resolve	303	fe-0/0/0.0
192.168.220.0	192.168.220.0	Receive	301	fe-0/0/0.0
224.0.0.1		Mcast	5	
255.255.255.255		Bcast	6	

```
...
```

### show pfe route iso

```
user@host# show pfe route iso
```

```
CLNS Route Table 0, CLNP.0, 0x0:
```

Destination	Type	NH ID	Interface
default	Reject	60	
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152	Local	514	Local
49.0001.00a0.c96b.c491/72	Local	536	

### show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary
```

## Slot 0

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	43	3081
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## Slot 1

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	43	3081
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## Slot 16

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	41	2938
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## Slot 17

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	41	2938
1	4	281

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	68

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	9	717
1	5	389

## show pfe route lcc summary (TX Matrix Plus Router)

user@host> show pfe route lcc 2 summary

Slot 0

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833

Slot 6

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815

2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833

...

**show pfe route summary (MX Series Router)**

user@host&gt; show pfe route summary

Slot 0

## DHCP-Snooping Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	144

## IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

## MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

## IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833



...

**show pfe route summary hw (QFX Series)**

```

user@switch> show pfe route summary hw
Slot 0
Unit: 0
Profile active: l2-profile-three
Type          Max      Used      Free      % free
-----
IPv4 Host      8192     103      8073     98.55
IPv4 LPM       16384     9      16369     99.91
IPv4 Mcast     4096      2      4037     98.56

IPv6 Host      4096      6      4037     98.56
IPv6 LPM(< 64) 8192      3      8185     99.91
IPv6 LPM(> 64) 256      1    255     99.61
IPv6 Mcast     2048      0      2019     98.58

```

**show pfe route ip hw host (QFX Series)**

```

user@switch> show pfe route ip host hw
Slot 0
Unit: 0
IPv4 Host entries present: 103
Rtt  VRF  Destination                                     Type  NH-ID  Interface
      HW NH-ID  Src-MAC-Address  Port Dst-MAC-Address
-----
4    3    255.255.255.255                                     Bcast  1695   .local.    .4
ifl 550 100003 00:00:00:01:02:03 127 00:00:00:01:02:03
0    1    200.1.1.42   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.56   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.61   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    11.1.1.2   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.73   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.76   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.18   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.5   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.23   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    101.1.1.255   Bcast  1664   ae0        .0
ifl 544 100003 00:00:00:01:02:03 127 00:00:00:01:02:03
0    1    200.1.1.40   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23
0    1    200.1.1.58   Unicast 1743   et-0/1/1   .0
ifl 559 100268 84:18:88:de:96:fd 53 00:00:00:21:12:23. . .
. . .

```

## show pfe terse

---

<b>List of Syntax</b>	<a href="#">Syntax on page 6448</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Router) on page 6448</a> <a href="#">Syntax (MX Series Router) on page 6448</a>
<b>Syntax</b>	show pfe terse
<b>Syntax (TX Matrix and TX Matrix Plus Router)</b>	show pfe terse <lcc <i>number</i>   scc> <sfc <i>number</i> >
<b>Syntax (MX Series Router)</b>	show pfe terse <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Packet Forwarding Engine status information.
<b>Options</b>	<b>none</b> —Display brief information about the Packet Forwarding Engine.  <b>all-members</b> —(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.  <b>lcc <i>number</i></b> —(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for the router (or line-card chassis) that is connected to a TX Matrix Plus router. Replace <i>number</i> with the following values depending on the LCC configuration: <ul style="list-style-type: none"><li>• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.</li><li>• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.</li><li>• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li><li>• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.</li></ul> <b>local</b> —(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member.

**member *member-id***—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

**scc**—(TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).

**sfc**—(TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).

**Required Privilege Level** admin

**List of Sample Output** [show pfe terse \(TX Matrix Router\) on page 6449](#)  
[show pfe terse \(TX Matrix Plus Router\) on page 6449](#)  
[show pfe terse sfc \(TX Matrix Plus Router\) on page 6449](#)

## Sample Output

### show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot State Flags Uptime
0 SFM Present Online 0x0bf 01:25:42
2 SFM Present Online 0x0bf 01:25:40
0 FPC Present Online 0x102 01:25:57
1 FPC Present Online 0x102 01:25:55
2 FPC Present Online 0x102 01:25:53
```

### show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:
-----
Slot Type Slot State Uptime
0 LCC Present Online 2d 05:26

lcc0-re0:
-----
Slot Type Slot State Uptime
0 GFPC Present Online 2d 05:25
1 GFPC Present Online 2d 05:25
```

### show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:
-----
Slot Type Slot State Uptime
0 LCC Present Online 2d 05:25
```

## show pfe version

---

<b>Syntax</b>	show pfe version <brief   detail>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display Packet Forwarding Engine version information.
<b>Options</b>	brief   detail—Display the specified level of output.
<b>Required Privilege Level</b>	admin
<b>List of Sample Output</b>	<a href="#">show pfe version brief on page 6450</a> <a href="#">show pfe version detail on page 6450</a>

### Sample Output

#### show pfe version brief

```
user@host> show pfe version brief
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
```

#### show pfe version detail

```
user@host> show pfe version detail
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC

junos-core01.juniper.net:/volume/build/junos/rpd_feb11/11.1/development/20101111.0/obj-i386/
junos/usr.sbin/pfed
```

## CHAPTER 75

# Troubleshooting

- [Troubleshooting Procedures on page 6451](#)

## Troubleshooting Procedures

---

- [Troubleshooting Dropped FCoE Traffic on page 6451](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth on page 6454](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth on page 6455](#)
- [Troubleshooting Egress Queue Bandwidth Impacted by Congestion on page 6456](#)
- [Troubleshooting an Unexpected Rewrite Value on page 6457](#)
- [Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic on page 6458](#)

## Troubleshooting Dropped FCoE Traffic

**Problem**    **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

**Cause**    There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.

5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.
6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

---

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

**Solution** The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 5606 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

#### Related Documentation

- [show class-of-service congestion-notification on page 6305](#)
- [show class-of-service forwarding-class-set on page 6313](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth

**Problem**     **Description:** The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth (shaping rate) configured for the queue.

**Cause**        When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.



The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

**Solution** When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

**Related  
Documentation**

- [shaping-rate on page 6278](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

**Problem** **Description:** The minimum bandwidth of a queue or a priority group when measured at the egress port exceeds the minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).

**Cause** When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.



**NOTE:** The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

**Solution** When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit

rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

- Related Documentation**
- [guaranteed-rate on page 6247](#)
  - [transmit-rate on page 6285](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
  - [Example: Configuring Queue Schedulers on page 6081](#)
  - [Understanding CoS Output Queue Schedulers on page 5868](#)

## Troubleshooting Egress Queue Bandwidth Impacted by Congestion

**Problem**    **Description:** Congestion on an egress port causes egress queues to receive less bandwidth than expected. Egress port congestion can impact the amount of bandwidth allocated to queues on the congested port and, in some cases, on ports that are not congested.

**Cause**       Egress queue congestion can cause the ingress port buffer to fill above a certain threshold and affect the flow to the queues on the egress port. One queue receives its configured bandwidth, but the other queues on the egress port are affected and do not receive their configured share of bandwidth.

**Solution**    The solution is to configure a drop profile to apply weighted random early detection (WRED) to the queue or queues on the congested ports.

Configure a drop profile on the queue that is receiving its configured bandwidth. This queue is preventing the other queues from receiving their expected bandwidth. The drop profile prevents the queue from affecting the other queues on the port.

To configure a WRED profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

- Related Documentation**
- [drop-profile on page 6227](#)
  - [Example: Configuring WRED Drop Profiles on page 6071](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Understanding CoS WRED Drop Profiles on page 5909](#)

## Troubleshooting an Unexpected Rewrite Value

**Problem**    **Description:** Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

**Cause**    If you configure a rewrite rule for a forwarding class on an egress port but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

**Solution**    If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.



**TIP:** If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.



**NOTE:** There are no default rewrite rules. You can bind one rewrite rule for each type (DSCP and IEEE 802.1) to a given interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value associations.

1. Assign a rewrite value to a forwarding class. Add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:

```
[edit class-of-service rewrite-rules]
```

```
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **fcoe** is being randomly rewritten, and you want to use IEEE 802.1 code point **011** for the **fcoe** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class fcoe loss-priority high code-point
011
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1
custom-rw
```

#### Related Documentation

- [interfaces on page 6256](#)
- [rewrite-rules on page 6273](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Monitoring CoS Rewrite Rules on page 6292](#)

## Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic

**Problem**    **Description:** In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets.

**Cause**      Failure of a queue to transmit packets for 12 consecutive seconds may be due to:

- A strict-high priority queue consuming all of the port bandwidth
- Several queues consuming all of the port bandwidth
- Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
- Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

**Solution**    If the cause is a strict-high priority queue or other queues consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the queues that are using all of the port bandwidth and preventing other queues from obtaining bandwidth on the port. You configure a maximum rate by creating a scheduler, using a scheduler map to apply it to a forwarding class (which maps to an output queue), and applying the scheduler map to the port using a forwarding class set and a traffic control profile.

To configure rate shaping using the CLI:

1. Name the existing scheduler or create a scheduler and define the maximum bandwidth as a rate or as a percentage:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name shaping-rate (rate | percent percentage)
```

2. Configure a scheduler map to associate the scheduler with the forwarding class (queue) that is consuming all of the port bandwidth:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

3. Associate the scheduler map with a traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name scheduler-map
scheduler-map-name
```

4. Associate the traffic control profile (and thus the scheduler map that contains the rate shaping queue scheduler) with a forwarding class set and apply them to the interface that is being reset:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set fc-set-name
output-traffic-control-profile traffic-control-profile-name
```

For example, a strict-high priority queue is using all of the bandwidth on interface **shpnode:xe-0/0/10** and preventing other queues from transmitting for 12 consecutive seconds. You decide to set a maximum rate of 7 Gbps on the strict-high priority queue to ensure that at least 3 Gbps of the port bandwidth is available to service other queues.

[Table 600 on page 6459](#) shows the topology for this example:

**Table 600: Components of the Rate Shaping Troubleshooting Example**

Component	Settings
Affected interface	<b>shpnode:xe-0/0/10</b>
Scheduler (strict-high priority scheduler)	Name: <b>shp-sched</b> Shaping rate: <b>7g</b> Priority: <b>strict-high</b>  <b>NOTE:</b> This example assumes that the scheduler already exists and has been configured as <b>strict-high</b> priority, but that rate shaping to prevent the strict-high priority traffic from using all of the port bandwidth has not been applied.
Scheduler map	Name: <b>shp-map</b> Forwarding class to associate with the <b>shp-sched</b> scheduler: <b>strict-high</b>  <b>NOTE:</b> This example assumes that a strict-high priority forwarding class has been configured and assigned the name <b>strict-high</b> .
Traffic control profile	Name: <b>shp-tcp</b>  <b>NOTE:</b> This example does not describe how to define a complete traffic control profile.

Table 600: Components of the Rate Shaping Troubleshooting Example (*continued*)

Component	Settings
Forwarding class set	Name: <b>shp-pg</b>

To configure the scheduler, map it to the strict-high priority forwarding class, and apply it to interface **shpnode:xe-0/0/10** using the CLI:

1. Specify the scheduler for the strict-high priority queue (**shp-sched**) with a maximum bandwidth of 7 Gbps:

```
[edit class-of-service schedulers]
user@switch# set shp-sched shaping-rate 7g
```

2. Configure a scheduler map (**shp-map**) that associates the scheduler (**shp-sched**) with the forwarding class (**strict-high**):

```
[edit class-of-service scheduler-maps]
user@switch# set shp-map forwarding-class strict-high scheduler shp-sched
```

3. Associate the scheduler map **shp-map** with a traffic control profile (**shp-tcp**):

```
[edit class-of-service traffic-control-profiles]
user@switch# set shp-tcp scheduler-map shp-map
```

4. Associate the traffic control profile **shp-tcp** with a forwarding class set (**shp-pg**) and the affected interface (**shpnode:xe-0/0/10**):

```
[edit class-of-service]
user@switch# set interfaces shpnode:xe-0/0/10 forwarding-class-set shp-pg
output-traffic-control-profile shp-tcp
```

#### Related Documentation

- [Understanding CoS Output Queue Schedulers on page 5868](#)
- [Defining CoS Queue Scheduling Priority on page 6171](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 6094](#)
- [Example: Configuring Forwarding Class Sets on page 6078](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)

## PART 21

# Network Management and Monitoring

- [Overview on page 6463](#)
- [Configuration on page 6565](#)
- [Administration on page 6799](#)
- [Troubleshooting on page 6893](#)





CHAPTER 76

# Overview

- [Network Management on page 6463](#)
- [Automation on page 6469](#)
- [Junos Space on page 6489](#)
- [Network Analytics on page 6490](#)
- [sFlow Technology on page 6509](#)
- [SNMP on page 6513](#)
- [System Logging on page 6560](#)

## Network Management

- [Understanding Device and Network Management Features on page 6463](#)
- [Understanding Network Management Implementation on the QFabric System on page 6466](#)
- [Understanding Telnet on the QFabric System on page 6467](#)
- [Understanding Tracing and Logging Operations on page 6468](#)

## Understanding Device and Network Management Features

After you install a QFX Series product or EX4600 switch in your network, you need to manage the device. The products support features that you use to manage the device within the network, including the management of configuration, system performance, fault monitoring, and remote access.

[Table 601 on page 6463](#) lists the device and network management features on the QFX Series and EX4600.

Table 601: Device and Network Management Features on the QFX Series and EX4600

Feature	Typical Uses	Documentation
AI-Scripts and Advanced Insight Manager (AIM)—Automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems, and submit problem reports to Juniper Support Systems.	Fault management	<a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a>

**Table 601: Device and Network Management Features on the QFX Series and EX4600 (*continued*)**

Feature	Typical Uses	Documentation
Alarms and LEDs on the switch—Show status of hardware components and indicate warning or error conditions.	Fault management	<a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 7192</a>
Firewall filters—Control the packets that are sent to and from the network, balance network traffic, and optimize performance.	Performance management	<ul style="list-style-type: none"> <li>• <a href="#">Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</a></li> <li>• <a href="#">Overview of Firewall Filters on page 5209</a></li> </ul>
In-band management—Enables connection to the switch using the same interfaces through which customer traffic flows. Communication between the switch and a remote console is typically enabled using SSH and Telnet services. SSH provides secure encrypted communications, whereas Telnet provides unencrypted, and therefore less secure, access to the switch.	Remote access management	<ul style="list-style-type: none"> <li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1361</a></li> <li>• <a href="#">Configuring Telnet Service for Remote Access to a Router or Switch</a></li> </ul>
Juniper Networks Junos OS automation scripts—Configuration and operations automation tools provided by Junos OS. These tools include commit scripts, operation scripts, event scripts, and event policies. Commit scripts enforce custom configuration rules, whereas operation scripts, event policies, and event scripts automate network troubleshooting and management.	<ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>	<a href="#">Junos OS Automation Library</a>
Junos OS command-line interface (CLI)—CLI configuration statements that enable you to configure the switch based on your networking requirements, such as security, service, and performance.	<ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• User access management</li> <li>• Remote access management</li> </ul>	<a href="#">CLI User Guide</a>
Junos Space software—Multipurpose GUI-based network management system that includes a base platform, the Network Application Platform, and other optional applications such as Ethernet Design, Service Now, Service Insight, and Virtual Control.	<ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Junos Space Support on page 6489</a></li> <li>• <a href="#">Junos Space Network Application Platform User Guide</a></li> </ul>

**Table 601: Device and Network Management Features on the QFX Series and EX4600 (*continued*)**

Feature	Typical Uses	Documentation
<p>Junos XML API—XML representation of Junos OS configuration statements and operational mode commands. Junos XML configuration tag elements are the content to which the Junos XML protocol operations apply. Junos XML operational tag elements are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device. The Junos XML API also includes tag elements that are the counterpart to Junos CLI configuration statements.</p>	<ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Junos XML API Configuration Developer Reference</a></li> <li>• <a href="#">Junos XML API Operational Developer Reference</a></li> </ul>
<p>NETCONF XML management protocol—XML-based management protocol that client applications use to request and change configuration information on routing, switching, and security platforms running Junos OS. The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b>, <b>set</b>, and <b>commit</b> to perform those operations.</p>	<ul style="list-style-type: none"> <li>• Configuration management</li> <li>• Performance management</li> <li>• Fault management</li> </ul>	<p><a href="#">NETCONF XML Management Protocol Developer Guide</a></p>
<p>Operational mode commands—May be used to do the following:</p> <ul style="list-style-type: none"> <li>• Monitor switch performance. For example, the <b>show chassis routing-engine</b> command shows the CPU utilization of the Routing Engine. High CPU utilization of the Routing Engine can affect performance of the switch.</li> <li>• View current activity and status of the device or network. For example, you can use the <b>ping</b> command to monitor and diagnose connectivity problems, and the <b>traceroute</b> command to locate points of failure on the network.</li> </ul>	<ul style="list-style-type: none"> <li>• Performance management</li> <li>• Fault management</li> </ul>	<p><a href="#">CLI Explorer</a></p>

**Table 601: Device and Network Management Features on the QFX Series and EX4600 (continued)**

Feature	Typical Uses	Documentation
Out-of-band management—Enables connection to the switch through a management interface. Out-of-band management is supported on two dedicated management Ethernet interfaces as well as on the console and auxiliary ports. The management Ethernet interfaces connect directly to the Routing Engine. No transit traffic is allowed through the interfaces, separating customer and management traffic and ensuring that congestion or failures in the transit network do not affect the management of the switch.	Remote access management	<ul style="list-style-type: none"> <li>• <a href="#">Connecting a QFX3500 Device to a Network for Out-of-Band Management</a></li> <li>• <a href="#">Connecting a QFX Series Device to a Management Console</a></li> <li>• <a href="#">Configuring Console and Auxiliary Port Properties on page 6583</a></li> </ul>
SNMP Configuration Management MIB—Provides notification for configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made. A history of the last 32 configuration changes is kept in jnxCmChgEventTable.	Configuration management	<a href="#">SNMP MIBs and Traps Reference</a>
SNMP MIBs and traps—Enable the monitoring of network devices from a central location. Use SNMP requests such as <b>get</b> and <b>walk</b> to monitor and view system activity.  The QFX3500 switch supports SNMP Version 1 (v1), v2, and v3, and both standard and Juniper Networks enterprise-specific MIBs and traps.	Fault management	<ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs and Traps Reference</a></li> <li>• <a href="#">Understanding the Implementation of SNMP on page 6513</a></li> </ul>
System log messages—Log details of system and user events, including errors. You can specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> <li>• Fault management</li> <li>• User access management</li> </ul>	<ul style="list-style-type: none"> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> <li>• <a href="#">Overview of Junos OS System Log Messages on page 6560</a></li> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6561</a></li> </ul>

## Understanding Network Management Implementation on the QFabric System

This topic describes network management features on the QFabric system that are implemented differently than on other devices running Junos OS.

The following network management features are supported on the QFabric system:

- **System log messages**—The QFabric system monitors events that occur on its component devices, distributes system log messages about those events to all external system log message servers (hosts) that are configured, and archives the messages. Component devices include Node devices, Interconnect devices, Director devices, and

the Virtual Chassis. You configure system log messages at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view messages.

- **Simple Network Management Protocol (SNMP) Version 1 (v1) and v2c**—SNMP monitors network devices from a central location. The SNMP implementation on the QFabric system supports the basic SNMP architecture of Junos OS with some limitations, including a reduced set of MIB objects, read-only access for SNMP communities, and limited support for SNMP requests. You configure SNMP at the **[edit snmp]** hierarchy level. Only the **show snmp statistics** operational mode command is supported, but you can issue SNMP requests using external SNMP client applications.
- **Advanced Insight Solutions (AIS)**—AIS provides tools and processes to automate the delivery of support services for the QFabric system. AIS components include Advanced Insight Scripts (AI-Scripts) and Advanced Insight Manager (AIM). You install AI-Scripts using the **request system scripts add** operational mode command. However, the **jais-activate-scripts.slax** file used during installation is preconfigured for the QFabric system and cannot be changed.



**NOTE:** Do not install Junos Space and AIS on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

#### Related Documentation

- [Advanced Insight Scripts \(AI-Scripts\) Release Notes](#)
- [Understanding Device and Network Management Features on page 6463](#)
- [Overview of Junos OS System Log Messages on page 6560](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
- [SNMP MIBs Support on page 6530](#)

## Understanding Telnet on the QFabric System

This topic describes the support for the Telnet protocol on QFabric systems.

Telnet service is available for devices running Junos OS, including QFX Series devices. However, on QFabric systems, Telnet support is limited and the following conditions apply:

- You can telnet from a QFabric system to external devices that are connected to the QFabric system by way of the network Node group. To connect to these external devices, issue the **telnet** command from the QFabric default partition CLI.
- You cannot use the Telnet protocol to connect from the QFabric system default partition CLI to individual components. To access system components, you must issue the **request component login** command instead.

#### Related Documentation

- [request component login on page 1480](#)
- [telnet](#)

## Understanding Tracing and Logging Operations

Tracing and logging operations enable you to track events that occur in the switch—both normal operations and error conditions—and to track the packets that are generated by or passed through the switch. The results of tracing and logging operations are placed in files in the `/var/log` directory on the switch.

The Junos OS supports remote tracing for the following processes:

- **chassisd**—Chassis-control process
- **eventd**—Event-processing process
- **cosd**—Class-of-service process

You configure remote tracing by using the **tracing** statement at the **[edit system]** hierarchy level.



**NOTE:** The **tracing** statement is not supported on the QFX3000 QFabric system.

If you enabled remote tracing but wish to disable it for specific processes on the switch, use the **no-remote-trace** statement at the **[edit process-name traceoptions]** hierarchy level. This feature does not alter local tracing functionality in any way, and logging files are stored on the switch.

Logging operations use a system logging mechanism similar to the UNIX **syslogd** utility to record systemwide, high-level operations, such as interfaces going up or down and users logging in to or out of the switch. You configure these operations by using the **syslog** statement at the **[edit system]** hierarchy level and by using the **options** statement at the **[edit ethernet-switching-options]** hierarchy level.

Tracing operations record more detailed information about the operations of the switch, including packet forwarding and routing information. To configure tracing operations, use the **traceoptions** statement.



**NOTE:** The **traceoptions** statement is not supported on the QFX3000 QFabric system.

You can define tracing operations in different portions of the switch configuration:

- **SNMP agent activity tracing operations**—Define tracing of the activities of SNMP agents on the switch. You configure SNMP agent activity tracing operations at the **[edit snmp]** hierarchy level.
- **Global switching tracing operations**—Define tracing for all switching operations. You configure global switching tracing operations at the **[edit ethernet-switching-options]** hierarchy level of the configuration.

- Protocol-specific tracing operations—Define tracing for a specific routing protocol. You configure protocol-specific tracing operations in the **[edit protocols]** hierarchy when configuring the individual routing protocol. Protocol-specific tracing operations override any equivalent operations that you specify in the global **traceoptions** statement. If there are no equivalent operations, they supplement the global tracing options. If you do not specify any protocol-specific tracing, the routing protocol inherits all the global tracing operations.
- Tracing operations within individual routing protocol entities—Some protocols allow you to define more granular tracing operations. For example, in Border Gateway Protocol (BGP), you can configure peer-specific tracing operations. These operations override any equivalent BGP-wide operations or, if there are no equivalents, supplement them. If you do not specify any peer-specific tracing operations, the peers inherit, first, all the BGP-wide tracing operations and, second, the global tracing operations.
- Interface tracing operations—Define tracing for individual interfaces and for the interface process itself. You define interface tracing operations at the **[edit interfaces]** hierarchy level of the configuration.
- Remote tracing—To enable system-wide remote tracing, configure the **destination-override syslog host** statement at the **[edit system tracing]** hierarchy level. This specifies the remote host running the system log process (syslogd), which collects the traces. Traces are written to files on the remote host in accordance with the syslogd configuration in **/etc/syslog.conf**. By default, remote tracing is not configured.

To override the system-wide remote tracing configuration for a particular process, include the **no-remote-trace** statement at the **[edit process-name traceoptions]** hierarchy. When **no-remote-trace** is enabled, the process does local tracing.

To collect traces, use the **local0** facility as the selector in the **/etc/syslog.conf** file on the remote host. To separate traces from various processes into different files, include the process name or trace-file name (if it is specified at the **[edit process-name traceoptions file]** hierarchy level) in the Program field in the **/etc/syslog.conf** file. If your system log server supports parsing hostname and program name, then you can separate traces from the various processes.



**NOTE:** During a commit check, warnings about the **traceoptions** configuration (for example, mismatch in trace file sizes or number of trace files) are not displayed on the console. However, these warnings are logged in the system log messages when the new configuration is committed.

#### Related Documentation

- [Overview of Junos OS System Log Messages on page 6560](#)

## Automation

- [Overview of QFX5100 Switch Automation Enhancements on page 6470](#)
- [Overview of Python with QFX5100 Switch Automation Enhancements on page 6471](#)
- [Understanding Automation Scripts Support on page 6473](#)

- [How Commit Scripts Work on page 6474](#)
- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 6480](#)
- [Overview of Generating Persistent or Transient Configuration Changes on page 6481](#)
- [Required Boilerplate for Commit Scripts on page 6485](#)
- [How Op Scripts Work on page 6486](#)
- [Required Boilerplate for Op Scripts on page 6487](#)

## Overview of QFX5100 Switch Automation Enhancements

The QFX5100 switch automation enhancements are designed to support the increasing needs of large data centers for more automation and programmability.

- [Features of the QFX5100 Switch Automation Enhancements on page 6470](#)

### Features of the QFX5100 Switch Automation Enhancements

---

To use the QFX5100 switch automation enhancements, you must install the `jinstall-qfx-5-flex-x.tgz` software bundle. This software bundle is identical to the other QFX5100 switch software bundle except that Veriexec is disabled, which enables you to run unsigned programs, such as programs that you develop with Python, Chef, and Puppet. The QFX5100 switch automation enhancements include the following features:

- The factory default configuration is a Layer 3 configuration. (The standard default factory configuration is Layer 2.)
- Safeguards ensure that you cannot overwrite essential Junos OS files, including system log notifications.
- Zero Touch Provisioning (ZTP) allows you to provision new switches in your network automatically, without manual intervention. See [“Understanding Zero Touch Provisioning” on page 32](#).
- The installation automatically sets up and reserves a 1-gigabit user partition on your system. You can use this partition to store your binaries and additional packages.
- The user partition is not overwritten when you upgrade or downgrade the software to a QFX5100 switch Junos OS image that does not contain the automation enhancements.



**NOTE:** If you make changes to the user partition while performing a unified in-service software upgrade (unified ISSU), the changes might be lost.

---

- The Python interpreter is included by default.
  - You can invoke Python directly from the shell. See [“Invoking the Python Interpreter” on page 6587](#).



- Starting with Junos OS Release 14.1X53-D10, three Open Source Python modules are pre-installed in the `jinstall-qfx-5-flex-x.tgz` software bundle. See “[Overview of Python with QFX5100 Switch Automation Enhancements](#)” on page 6471 for details.
- Chef for Junos OS and Puppet for Junos OS automation tools for provisioning and managing computer networking and storage resources are included.
- For further information on Chef, see [Chef for Junos Getting Started Guide](#).
- For further information on Puppet, see [Puppet for Junos OS Documentation](#).



**NOTE:** For full compatibility, you must use only Chef for Junos OS and Puppet for Junos OS rather than the standard FreeBSD versions of Chef and Puppet software.



**CAUTION:** Download additional third party packages at your own risk.

#### Related Documentation

- *Installing Junos OS Software with QFX5100 Switch Automation Enhancements*
- [Invoking the Python Interpreter on page 6587](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions on page 6897](#)

## Overview of Python with QFX5100 Switch Automation Enhancements

Python is a programming language that lets you work more quickly and integrate your systems more effectively. The Python interpreter is included within the Junos operating system (Junos OS) `jinstall-qfx-5-flex-x.tgz` software bundle.

Python is also suitable as an extension language for customizable applications.

Starting with Junos OS Release 14.1X53-D10, these Open Source Python modules are pre-installed in the `jinstall-qfx-5-flex-x.tgz` software bundle:

- **ncclient**—Facilitates client scripting and application development through the NETCONF protocol. See <http://ncclient.grnet.gr/0.3.2/> for documentation of some of the external APIs of the ncclient Python module. At the bottom of this list, see examples of usage of some of these APIs with sample scripts.
- **lxml**—Combines the speed and XML feature completeness of the C libraries libxml2 and libxslt with the simplicity of a native Python API. See <http://lxml.de/tutorial.html/> for documentation of some of the external APIs of the lxml Python module.
- **jinja2**—Serves as a fast, secure, designer-friendly templating language. See <http://jinja.pocoo.org/docs/api/> for documentation of some of the external APIs of the jinja2 Python module.

Example usage of some of the APIs of the ncclient Python module follows:

**\* Example of "connect" and "command" API:**

```
from ncclient import manager

def connect(host, port, user, password):
    conn = manager.connect(host=host,
                           port=port,
                           username=user,
                           password=password,
                           timeout=10,
                           device_params = {'name': 'junos'},
                           hostkey_verify=False)

    print 'show version'
    print '*' * 30
    result = conn.command('show version', format='text')
    print result.xpath('output')[0].text

if __name__ == '__main__':
    connect('router', '22', 'netconf', 'juniper!')
```

**\* Example of "compare\_configuration" API:**

```
from ncclient import manager
from ncclient.xml_ import *

import time

def connect(host, port, user, password, source):
    conn = manager.connect(host=host,
                           port=port,
                           username=user,
                           password=password,
                           timeout=10,
                           device_params = {'name': 'junos'},
                           hostkey_verify=False)

    compare_config = conn.compare_configuration(rollback=3)
    print compare_config.tostring

if __name__ == '__main__':
    connect('router', 830, 'netconf', 'juniper!', 'candidate')
```

**\* Example of "lock", "load\_configuration", "validate", "commit", "discard\_changes", "unlock" APIs:**

```
from ncclient import manager
from ncclient.xml_ import *

import time

def connect(host, port, user, password, source):
    conn = manager.connect(host=host,
                           port=port,
                           username=user,
                           password=password,
                           timeout=10,
                           device_params = {'name': 'junos'},
                           hostkey_verify=False)

    print 'locking configuration'
    lock = conn.lock()
```

```

# build configuration element
config = new_ele('system')
sub_ele(config, 'host-name').text = 'foo'
sub_ele(config, 'domain-name').text = 'bar'

send_config = conn.load_configuration(config=config)
print send_config.tostring

check_config = conn.validate()
print check_config.tostring

compare_config = conn.compare_configuration()
print compare_config.tostring

print 'commit confirmed 300'
#commit_config = conn.commit(confirmed=True, timeout='300')
commit_config = conn.commit()
print commit_config.tostring

print 'sleeping for 5 sec...'
time.sleep(5)

discard_changes = conn.discard_changes()
print discard_changes.tostring

print 'unlocking configuration'
unlock = conn.unlock()
print unlock.tostring

if __name__ == '__main__':
    connect('router', 830, 'netconf', 'juniper!', 'candidate')

```



**NOTE:** For information on using Python, refer to your Python documentation.

#### Related Documentation

- [Installing Junos OS Software with QFX5100 Switch Automation Enhancements](#)
- [Invoking the Python Interpreter on page 6587](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions on page 6897](#)

## Understanding Automation Scripts Support

This document describes the support for the Junos OS automation scripts on the QFabric system Director devices.

Junos OS automation consists of a suite of tools used to automate operational and configuration tasks on network devices running Junos OS. The automation tools, which leverage the native XML capabilities of the Junos OS, include commit scripts, operation (op) scripts, event policies and event scripts, and macros.



**NOTE:** Event policies and event scripts are not supported on the QFabric system at this time.

The QFabric system supports Junos OS automation scripts that are written in Stylesheet Language Alternative Syntax (SLAX) version 1.0.

Commit scripts automate the commit process and enforce custom configuration rules. You can use commit scripts to generate specific errors and warnings, and customize configurations and configuration templates. When a candidate configuration is committed, it is inspected by each active commit script. If a configuration violates your custom rules and the scripts generate an error, the commit fails. If the commit is successful, any configuration changes (both transient and permanent) are incorporated into the active configuration before it is passed to the Director software, which distributes the configuration to all applicable QFabric system components, including Node devices and Node servers.

Op scripts automate operational and troubleshooting tasks. Op scripts can be executed manually from the Junos OS CLI or NETCONF XML management protocol, or they can be called from another script.

The QFabric system supports the following automation script features:

- Commit scripts and op scripts are supported.
- Scripts written in SLAX version 1 are supported.
- Scripts are configured and deployed from the Director group. Since there is more than one Director device in a Director group, scripts must be deployed by each Director device or deployed in the shared media space.
- Scripts are stored in the shared media at this location:  
`/pbdata/mgd_shared/partition-ip/var/db/scripts`. Under this directory, commit scripts are stored in the **commit** subdirectory, and op scripts are stored in the **op** subdirectory.
- Scripts are not stored in flash memory.

#### Related Documentation

- [How Commit Scripts Work on page 6474](#)
- [How Op Scripts Work on page 6486](#)
- [Required Boilerplate for Commit Scripts on page 6485](#)
- [Required Boilerplate for Op Scripts on page 6487](#)
- [Controlling the Execution of Commit Scripts on page 6588](#)

## How Commit Scripts Work


You enable commit scripts by listing the names of one or more commit script files at the **[edit system scripts commit]** hierarchy level. These scripts contain instructions that enforce custom configuration rules. Commit scripts are invoked during the commit process before the standard Junos OS validity checks are performed.

When you perform a commit operation, Junos OS executes each script in turn, passing the information in the candidate configuration to the scripts. The script inspects the configuration, performs the necessary tests and validations, and generates a set of instructions for performing certain actions. These actions include generating error, warning, and system log messages. If errors are generated, the commit operation fails and the candidate configuration remains unchanged. This is the same behavior that occurs with standard commit errors.

Commit scripts can also generate changes to the system configuration. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Figure 219 on page 6475 shows the flow of commit script input and output.

**Figure 219: Commit Script Input and Output**



Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

The following sections discuss several important concepts related to the commit script input and output:

- [Commit Script Input on page 6475](#)
- [Commit Script Output on page 6476](#)
- [Commit Scripts and the Junos OS Commit Model on page 6477](#)

### Commit Script Input

The input for a commit script is the postinheritance candidate configuration in Junos XML API format. The term *postinheritance* means that all configuration group values have been inherited by their targets in the candidate configuration and the inactive portions of the configuration have been removed. For more information about configuration groups, see the *CLI User Guide*.

When you issue the **commit** command, Junos OS automatically generates the candidate configuration in XML format and reads it into the management (mgd) process, at which time the input is evaluated by any commit scripts.

To display the XML format of the postinheritance configuration, issue the **show | display commit-scripts view** command:

```
[edit]
user@host# show | display commit-scripts view
```

To display all configuration groups data, including script-generated changes to the groups, issue the **show groups | display commit-scripts** command:

```
[edit]
user@host# show groups | display commit-scripts
```

To save the commit script input to a file, add the **save** command to the command line:

```
[edit]
user@host# show | display commit-scripts view | save filename.xml
```

By default, the file is placed in your home directory on the switch, router, or security device.

### Commit Script Output

---

To specify the desired commit script output—including warning, error, and system log messages, persistent changes, and transient changes—the script can contain tags that appear in any order, in any number. The tags for specifying output are as follows:

- **<xnm:warning>**—Generates a warning message
- **<xnm:error>**—Generates an error message.
- **<syslog><message>**—Generates a system log message.
- **<change>**—Generates a persistent change to the configuration.
- **<transient-change>**—Generates a transient change to the configuration.
- **<xsl:call-template name="jcs:emit-change">**
  - <xsl:with-param name="content">**—Generates a persistent change relative to the current context node as defined by an XPath expression.
- **<xsl:call-template name="jcs:emit-change">**
  - <xsl:with-param name="tag" select="transient-change"/>**
    - <xsl:with-param name="content">**—Generates a transient change relative to the current context node as defined by an XPath expression.
- **<xsl:call-template name="jcs:emit-change">**
  - <xsl:with-param name="message">**
    - <xsl:text>**—Generates a warning message in conjunction with a configuration change. You can use this set of tags to generate a notification that the configuration has been changed.

Junos OS processes this output and performs the appropriate actions. Errors and warnings are passed back to the Junos OS CLI or to a Junos XML protocol client application. The presence of an error automatically causes the commit operation to fail. Persistent and transient changes are loaded into the appropriate configuration database.

To test the output of error, warning, and system log messages from commit scripts, issue the **commit check | display xml** command:

```
[edit]
user@host# commit check | display xml
```

To display a detailed trace of commit script processing, issue the **commit check | display detail** command:

```
[edit]
user@host# commit check | display detail
```



**NOTE:** System log messages do not appear in the trace output, so you cannot use the commit check operation to test script-generated system log messages. Furthermore, system log messages are written to the system log during a commit operation, but not during a commit check operation.

#### Related Documentation

- *Example: Protecting the Junos OS Configuration from Modification or Deletion.*
- *jcs:emit-change Template*

### Commit Scripts and the Junos OS Commit Model

Junos OS uses a commit model to update the device's configuration. This model allows you to make a series of changes to a candidate configuration without affecting the operation of the device. When the changes are complete, you can commit the configuration. The commit operation saves the candidate configuration changes into the current configuration.

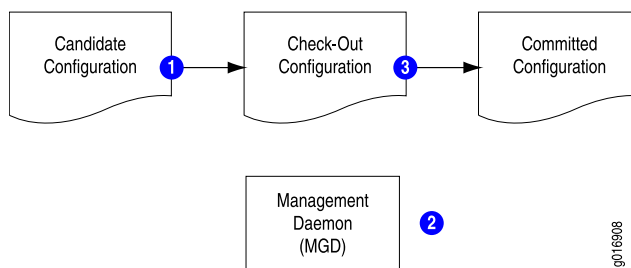
When you commit a set of changes in the candidate configuration, two methods are used to forward these changes to the current configuration:

- Standard commit model—Used when no commit scripts are active on the device.
- Commit script model—Incorporates commit scripts into the commit model.

#### Standard Commit Model

In the standard commit model, the management (mgd) process validates the candidate configuration based on standard Junos validation rules. If the configuration file is valid, it becomes the current active configuration. [Figure 220 on page 6478](#) and the accompanying discussion explain how the standard commit model works:

Figure 220: Standard Commit Model



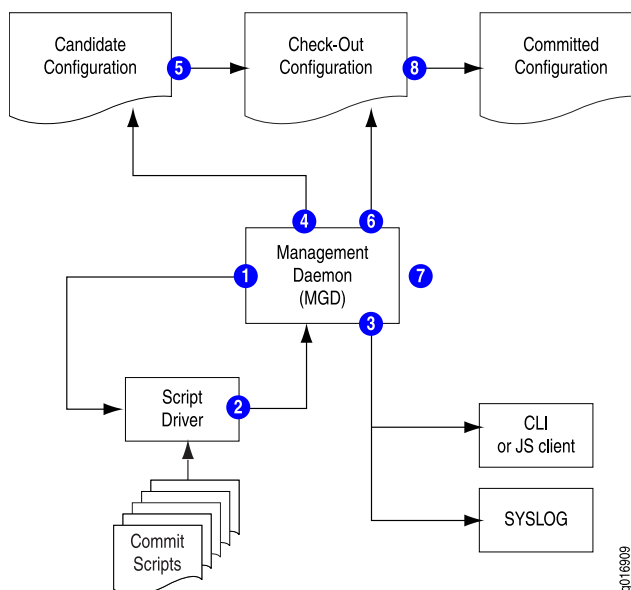
In the standard commit model, the software performs the following steps:

1. When the candidate configuration is committed, it is copied to become the checkout configuration.
2. The mgd process validates the checkout configuration.
3. If no error occurs, the checkout configuration is copied as the current active configuration.

### Commit Model with Commit Scripts

When commit scripts are added to the standard commit model, the process becomes more complex. The mgd process first passes an XML-formatted checkout configuration to a script driver, which handles the verification of the checkout configuration by the commit scripts. When verification is complete, the script driver returns an XML *action file* to the mgd process. The mgd process follows the instructions in the action file to update the candidate and checkout configurations, issue messages to the CLI, and write information to the system log as required. After processing the action file, the mgd process performs the standard Junos OS validation. [Figure 221 on page 6478](#) and the accompanying discussion explain this process.

Figure 221: Commit Model with Commit Scripts Added





In the commit script model, Junos OS performs the following steps:

1. When the candidate configuration is committed, the mgd process sends the XML-formatted candidate configuration to the script driver.
2. Each enabled commit script is invoked against the candidate configuration, and each script can generate a set of actions for the mgd process to perform. The actions are collected in an XML action file.
3. The mgd process performs the following actions in response to **<error>**, **<warning>**, and **<syslog>** tag elements in the action file:
  - **<error>**—The mgd process halts the commit process (that is, the commit operation fails), returns an error message to the CLI or Junos XML protocol client, and takes no further action.
  - **<warning>**—The mgd process forwards the message to the CLI or the Junos XML protocol client.
  - **<syslog>**—The mgd process forwards the message to the system log process.
4. If the action file includes any **<change>** tag elements, the mgd process loads the requested changes into the candidate configuration.
5. The candidate configuration is copied to become the checkout configuration.
6. If the action file includes any **<transient-change>** tag elements, the mgd process loads the requested changes into the checkout configuration.
7. The mgd process validates the checkout configuration.
8. If there are no validation errors, the checkout configuration is copied to become the current active configuration.



**NOTE:** Commit scripts cannot make configuration changes to protected statements or within protected hierarchies. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made. Failure to modify a protected configuration element does not halt the commit script or the commit process.

Changes that are made to the candidate configuration during the commit operation are not evaluated by the custom rules during that commit operation. However, persistent changes are maintained in the candidate configuration and are evaluated by the custom rules during subsequent commit operations. For more information about how commit scripts change the candidate configuration, see [“Avoiding Potential Conflicts When Using Multiple Commit Scripts” on page 6480](#).

Transient changes are never evaluated by the custom rules in commit scripts, because they are made to the checkout configuration only after the commit scripts have evaluated the candidate configuration and the candidate is copied to become the checkout configuration. To remove a transient change from the configuration, remove, disable, or deactivate the commit script (as discussed in *Controlling Execution of Commit Scripts During Commit Operations*), or comment out the code that generates the transient change.

For more information about differences between persistent and transient changes, see [“Overview of Generating Persistent or Transient Configuration Changes”](#) on page 6481.

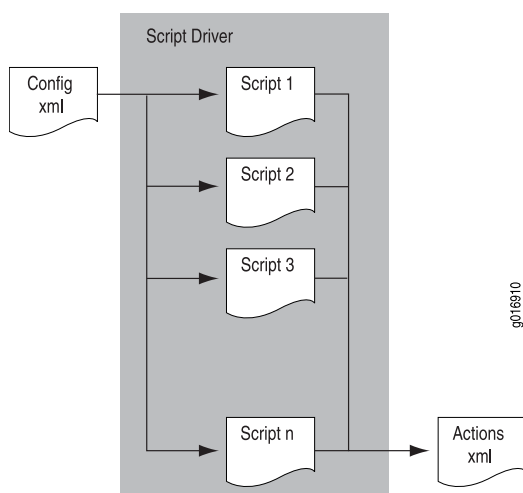
**Related Documentation**

- [Avoiding Potential Conflicts When Using Multiple Commit Scripts on page 6480](#)

## Avoiding Potential Conflicts When Using Multiple Commit Scripts

When you use multiple commit scripts, each script evaluates the original candidate configuration file. Changes made by one script are not evaluated by the other scripts. This means that conflicts between scripts might not be resolved when the scripts are first applied to the configuration. The commit scripts are executed in the order they are listed at the `[edit system scripts commit]` hierarchy level, as illustrated in [Figure 222 on page 6480](#).

**Figure 222: Configuration Evaluation by Multiple Commit Scripts**



As an example of a conflict between commit scripts, suppose that commit script **A.xsl** is created to ensure that the device uses the domain name server with IP address 192.168.0.255. Later, the DNS server's address is changed to 192.168.255.255 and a second script, **B.xsl**, is added to check that the device uses the DNS server with that address. However, script **A.xsl** is not removed or disabled.

Because each commit script evaluates the original candidate configuration, the final result of executing both scripts **A.xsl** and **B.xsl** depends on which DNS server address is configured in the original candidate configuration. If the now outdated address of 192.168.0.255 is configured, script **B.xsl** changes it to 192.168.255.255. However, if the correct address of 192.168.255.255 is configured, script **A.xsl** changes it to the incorrect value 192.168.0.255.

As another example of a potential conflict between commit scripts, suppose that a commit script protects a hierarchy using the **protect** attribute. If a second commit script attempts to modify or delete the hierarchy or the statements within the hierarchy, Junos OS issues a warning during the commit process and prevents the configuration change.

Exercise care to ensure that you do not introduce conflicts between scripts like those described in the examples. As a method of checking for conflicts with persistent changes, you can issue two separate **commit** commands.

**Related Documentation**

- [How Commit Scripts Work on page 6474](#)

## Overview of Generating Persistent or Transient Configuration Changes

Junos OS commit scripts enforce custom configuration rules. When a candidate configuration includes statements that you have decided must not be included in your configuration, or when the candidate configuration omits statements that you have decided are required, commit scripts can automatically change the configuration and thereby correct the problem.

- [Differences Between Persistent and Transient Changes on page 6481](#)
- [Interaction of Configuration Changes and Configuration Groups on page 6484](#)
- [Tag Elements and Templates for Generating Changes on page 6484](#)

### Differences Between Persistent and Transient Changes

Configuration changes made by commit scripts can be *persistent* or *transient*.

A persistent change remains in the candidate configuration and affects routing operations until you explicitly delete it, even if you subsequently remove or disable the commit script that generated the change and reissue the **commit** command. In other words, removing the commit script does not cause a persistent change to be removed from the configuration.

A transient change, in contrast, is made in the *checkout configuration* but not in the candidate configuration. The checkout configuration is the configuration database that is inspected for standard Junos OS syntax just before it is copied to become the active configuration on the device. If you subsequently remove or disable the commit script that made the change and reissue the **commit** command, the change is no longer made to the checkout configuration and so does not affect the active configuration. In other words, removing the commit script effectively removes a transient change from the configuration.

A common use for transient changes is to eliminate the need to repeatedly configure and display well-known policies, thus allowing these policies to be enforced implicitly. For example, if MPLS must be enabled on every interface with an International Organization for Standardization (ISO) protocol enabled, the change can be transient, so that the repetitive or redundant configuration data need not be carried or displayed in the candidate configuration. Furthermore, transient changes allow you to write script instructions that apply the change only if a set of conditions is met.

Persistent and transient changes are loaded into the configuration in the same manner that the **load replace** configuration mode command loads an incoming configuration. When generating a persistent or transient change, adding the **replace="replace"** attribute to a configuration element produces the same behavior as a **replace:** tag in a **load replace** operation.

By default, Junos OS merges the incoming configuration and the candidate configuration. New statements and hierarchies are added, and conflicting statements are overridden. When generating a persistent or transient change, if you add the **replace="replace"** attribute to a configuration element, Junos OS replaces the existing configuration element with the incoming configuration element. If the **replace="replace"** attribute is added to a configuration element, but there is no existing element of the same name in the current configuration, the incoming configuration element is added into the configuration. Elements that do not have the **replace** attribute are merged into the configuration.

Persistent and transient changes are loaded before the standard Junos validation checks are performed. This means any configuration changes introduced by a commit script are validated for correct syntax. If the syntax is correct, the new configuration becomes the active, operational device configuration.

Protected elements in the configuration hierarchy cannot be modified or deleted by either a persistent or a transient change. If a commit script attempts to modify or delete a protected statement or hierarchy, Junos OS issues a warning that the change cannot be made, and proceeds with the commit.

Persistent and transient changes have several important differences, as described in [Table 602 on page 6482](#).

**Table 602: Differences Between Persistent and Transient Changes**

Persistent Changes	Transient Changes
A persistent change is represented in a commit script by the <b>&lt;change&gt;</b> tag.	A transient change is represented in a commit script by the <b>&lt;transient-change&gt;</b> tag.
Another way to represent a persistent change is with the <b>content</b> parameter inside a call to the <b>jcs:emit-change</b> template.	Another way to represent a transient change is to use the <b>content</b> parameter and the <b>tag transient</b> parameter inside a call to the <b>jcs:emit-change</b> template.
The <b>jcs:emit-change</b> template is a helper template contained in the <b>junos.xsl</b> import file.	
You can use persistent changes to perform any Junos XML protocol operation, such as activate, deactivate, delete, insert (reorder), comment (annotate), and replace sections of the configuration.	Like persistent changes, you can use transient changes to perform any Junos XML protocol operation. However, some Junos XML protocol operations do not make sense to use with transient changes, such as generating comments and inactive settings.
Persistent changes are always loaded during the commit process if no errors are generated by any commit scripts or by the standard Junos OS validity check.	For transient changes to be loaded, you must include the <b>allow-transients</b> statement at the <b>[edit system scripts commit]</b> hierarchy level. If you enable a commit script that generates transient changes and you do not include the <b>allow-transients</b> statement in the configuration, the CLI generates an error message and the commit operation fails.
	Like persistent changes, transient changes must pass the standard Junos OS validity check.
	You cannot use a commit script to generate the <b>allow-transients</b> statement at the <b>[edit system scripts commit]</b> hierarchy level. Rather, you must include this statement directly by using the CLI.

Table 602: Differences Between Persistent and Transient Changes (*continued*)

Persistent Changes	Transient Changes
<p>Persistent changes work like the <b>load replace</b> configuration mode command, and the change is added to the candidate configuration.</p> <p>When generating a persistent change, if you add the <b>replace="replace"</b> attribute to a configuration element, Junos OS replaces the existing element in the candidate configuration with the incoming configuration element. If there is no existing element of the same name in the candidate configuration, the incoming configuration element is added into the configuration. Elements that do not have the <b>replace</b> attribute are merged into the configuration.</p>	<p>Transient changes work like the <b>load replace</b> configuration mode command, and the change is added to the checkout configuration.</p> <p>When generating a transient change, if you add the <b>replace="replace"</b> attribute to a configuration element, Junos OS replaces the existing element in the checkout configuration with the incoming configuration element. If there is no existing element of the same name in the checkout configuration, the incoming configuration element is added into the configuration. Elements that do not have the <b>replace</b> attribute are merged into the configuration.</p> <p>Transient changes are not copied to the candidate configuration. For this reason, transient changes are not saved in the configuration if the associated commit script is deleted or deactivated.</p>
<p>After a persistent change is committed, the software treats it like a change you make by directly editing and committing the candidate configuration.</p> <p>After the persistent changes are copied to the candidate configuration, they are copied to the checkout configuration. If the changes pass the standard Junos OS validity checks, the changes are propagated to the switch, router, or security device components.</p>	<p>Each time a transient change is committed, the software updates the checkout configuration database. After the transient changes pass the standard Junos OS validity checks, the changes are propagated to the device components.</p>
<p>After committing a script that causes a persistent change to be generated, you can view the persistent change by issuing the <b>show</b> configuration mode command:</p> <pre>user@host# show</pre> <p>This command displays persistent changes only, not transient changes.</p>	<p>After committing a script that causes a transient change to be generated, you can view the transient change by issuing the <b>show   display commit-scripts</b> configuration mode command:</p> <pre>user@host# show   display commit-scripts</pre> <p>This command displays both persistent and transient changes.</p>
<p>Persistent changes must conform to your custom configuration design rules as dictated by commit scripts.</p> <p>This does not become apparent until after a second commit operation because persistent changes are not evaluated by commit script rules on the current commit operation. The subsequent commit operation fails if the persistent changes do not conform to the rules imposed by the commit scripts configured during the first commit operation.</p>	<p>Transient changes are never tested by and do not need to conform to your custom rules. This is caused by the order of operations in the Junos OS commit model, which is explained in detail in <a href="#">“Commit Scripts and the Junos OS Commit Model” on page 6477</a>.</p>
<p>A persistent change remains in the configuration even if you delete, disable, or deactivate the commit script instructions that generated the change.</p>	<p>If you delete, disable, or deactivate the commit script instructions that generate a transient change, the change is removed from the configuration after the next commit operation. In short, if the associated instructions or the entire commit script is removed, the transient change is also removed.</p>

Table 602: Differences Between Persistent and Transient Changes (*continued*)

Persistent Changes	Transient Changes
As with direct CLI configuration, you can remove a persistent change by rolling back to a previous configuration that did not include the change and issuing the <b>commit</b> command. However, if you do not disable or deactivate the associated commit script, and the problem that originally caused the change to be generated still exists, the change is automatically regenerated when you issue another <b>commit</b> command.	You cannot remove a transient change by rolling back to a previous configuration.
You can alter persistent changes directly by editing the configuration using the CLI.	<p>You cannot directly alter or delete a transient change by using the Junos OS CLI, because the change is not in the candidate configuration.</p> <p>To alter the contents of a transient change, you must alter the statements in the commit script that generates the transient change.</p>

### Interaction of Configuration Changes and Configuration Groups

Any configuration change you can make by directly editing the configuration using the Junos OS command-line interface (CLI) can also be generated by a commit script as a persistent or transient change. This includes values specified at a specific hierarchy level or in configuration groups. As with direct CLI configuration, values specified in the *target* override values inherited from a configuration group. The target is the statement to which you apply a configuration group by including the **apply-groups** statement.

If you define persistent or transient changes as belonging to a configuration group, the configuration groups are applied in the order you specify in the **apply-groups** statements, which you can include at any hierarchy level except the top level. You can also disable inheritance of a configuration group by including the **apply-groups-except** statement at any hierarchy level except the top level.



**CAUTION:** Each commit script inspects the postinheritance view of the configuration. If a candidate configuration contains a configuration group, be careful when using a commit script to change the related target configuration, because doing so might alter the intended inheritance from the configuration group.

Also be careful when using a commit script to change a configuration group, because the configuration group might be generated by an application that performs a load replace operation on the group during each commit operation.

For more information about configuration groups, see the *CLI User Guide*.

### Tag Elements and Templates for Generating Changes

To generate changes, you can use the **jcs:emit-change** template, which implicitly includes **<change>** and **<transient-change>** XML elements; or you can explicitly include **<change>**

and `<transient-change>` XML elements. Using the `jcs:emit-change` template allows you to set the hierarchical context of the change once rather than multiple times.

The `<change>` and `<transient-change>` elements are similar to the `<load-configuration>` operation defined by the Junos XML management protocol. The possible contents of the `<change>` and `<transient-change>` elements are the same as the contents of the `<configuration>` tag element used in the Junos XML protocol operation `<load-configuration>`. For complete details about the `<load-configuration>` element, see the *Junos XML Management Protocol Developer Guide*.

## Required Boilerplate for Commit Scripts

When you write commit scripts, you use Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) tools provided with Junos OS. These tools include basic boilerplate that you must include in all commit scripts, optional extension functions that accomplish scripting tasks more easily, and named templates that make commit scripts easier to read and write, which you import from a file called `junos.xsl`. For more information about the extension functions and templates, see *Junos Script Automation: Understanding Extension Functions in the jcs and slax Namespaces* and *Junos Script Automation: Named Templates in the jcs Namespace Overview*.

Commit scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

You must include either XSLT or SLAX boilerplate as the starting point for all commit scripts that you create. The XSLT boilerplate follows:

### XSLT Boilerplate for Commit Scripts

```

1  <?xml version="1.0" standalone="yes"?>
2  <xsl:stylesheet version="1.0"
3    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4    xmlns:junos="http://xml.juniper.net/junos/*/junos"
5    xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6    xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7    <xsl:import href="../../../import/junos.xsl"/>

8    <xsl:template match="configuration">
9      <!-- ... Insert your code here ... -->
10   </xsl:template>
11 </xsl:stylesheet>
```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI). This PI specifies that the code is written in XML using version 1.0. The XML PI, if present, must be the first noncomment token in the script file.

```
1  <?xml version="1.0"?>
```

Lines 2 through 6 set the style sheet element and the associated namespaces. Line 2 sets the style sheet version as 1.0. Lines 3 through 6 list all the namespace mappings commonly used in commit scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```
2 <xsl:stylesheet version="1.0"
3   xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4   xmlns:junos="http://xml.juniper.net/junos/*/junos"
5   xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6   xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of the Junos OS. The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in *Junos Script Automation: Named Templates in the jcs Namespace Overview* and *Junos Named Templates in the jcs Namespace Summary*.

```
7 <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `<configuration>` element, which is the node selected by the `<xsl:template match="/">` template, contained in the `junos.xsl` import file. The `<xsl:template match="configuration">` element allows you to exclude the `/configuration/` root element from all XML Path Language (XPath) expressions in the script and begin XPath expressions with the top Junos OS hierarchy level. For more information, see *XPath Overview*.

```
8 <xsl:template match="configuration">
```

Add your code between Lines 8 and 9.

Line 9 closes the template.

```
9 </xsl:template>
```

Line 10 closes the style sheet and the commit script.

```
10 </xsl:stylesheet>
```

### SLAX Boilerplate for Commit Scripts

The corresponding SLAX boilerplate is as follows:

```
version 1.0;
ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xsl";

match configuration {
/*
* Insert your code here
*/
}
```

## How Op Scripts Work

Op scripts execute Junos OS operational commands and inspect the resulting output. After inspection, op scripts can automatically correct errors within the device running Junos OS based on this output.




You add op scripts to device operations by listing the filenames of one or more op script files within the **[edit system scripts op]** hierarchy level. These files must be added to the appropriate op script file directory. For more information about op script file directories, see *Storing Scripts in Flash Memory*. Once added to the device, op scripts are invoked from the command line, using the **op filename** command.

You can use op scripts to generate changes to the device configuration by including the **<load-configuration>** tag element. Because the changes are loaded before the standard validation checks are performed, they are validated for correct syntax, just like statements already present in the configuration before the script is applied. If the syntax is correct, the configuration is activated and becomes the active, operational device configuration.

Figure 223 on page 6487 shows a high-level view of the flow of op script input and output.

**Figure 223: Op Script Input and Output**



## Required Boilerplate for Op Scripts

When you write operation (op) scripts, you use Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) tools provided with Junos OS. These tools include basic boilerplate that you must include in all op scripts, optional extension functions that accomplish scripting tasks more easily, and named templates that make scripts easier to read and write, which you import from a file called **junos.xml**. For more information about the extension functions and templates, see *Junos Script Automation: Understanding Extension Functions in the jcs and slax Namespaces* and *Junos Script Automation: Named Templates in the jcs Namespace Overview*.

Op scripts are based on Junos XML and Junos XML protocol tag elements. Like all XML elements, angle brackets enclose the name of a Junos XML or Junos XML protocol tag element in its opening and closing tags. This is an XML convention, and the brackets are a required part of the complete tag element name. They are not to be confused with the angle brackets used in the documentation to indicate optional parts of Junos OS CLI command strings.

You must include either XSLT or SLAX boilerplate as the starting point for all op scripts that you create. The XSLT boilerplate follows:

### XSLT Boilerplate for Op Scripts

```

1 <?xml version="1.0" standalone="yes"?>
2 <xsl:stylesheet version="1.0"
3   xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4   xmlns:junos="http://xml.juniper.net/junos/*/junos"
5   xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6   xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
7   <xsl:import href="../import/junos.xml"/>
8   <xsl:template match="/">

```

```
9      <op-script-results>
      <!-- ... insert your code here ... -->
10    </op-script-results>
11  </xsl:template>
      <!-- ... insert additional template definitions here ... -->
12 </xsl:stylesheet>
```

Line 1 is the Extensible Markup Language (XML) processing instruction (PI), which marks this file as XML and specifies the version of XML as 1.0. The XML PI, if present, must be the first non-comment token in the script file.

```
1 <?xml version="1.0"?>
```

Line 2 opens the style sheet and specifies the XSLT version as 1.0.

```
2 <xsl:stylesheet version="1.0"
```

Lines 3 through 6 list all the namespace mappings commonly used in operation scripts. Not all of these prefixes are used in this example, but it is not an error to list namespace mappings that are not referenced. Listing all namespace mappings prevents errors if the mappings are used in later versions of the script.

```
3   xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
4   xmlns:junos="http://xml.juniper.net/junos/*/junos"
5   xmlns:xnm="http://xml.juniper.net/xnm/1.1/xnm"
6   xmlns:jcs="http://xml.juniper.net/junos/commit-scripts/1.0">
```

Line 7 is an XSLT import statement. It loads the templates and variables from the file referenced as `../import/junos.xsl`, which ships as part of Junos OS (in the file `/usr/libdata/cscript/import/junos.xsl`). The `junos.xsl` file contains a set of named templates you can call in your scripts. These named templates are discussed in *Junos Script Automation: Named Templates in the jcs Namespace Overview* and *Junos Named Templates in the jcs Namespace Summary*.

```
7   <xsl:import href="../import/junos.xsl"/>
```

Line 8 defines a template that matches the `</>` element. The `<xsl:template match="/">` element is the root element and represents the top level of the XML hierarchy. All XML Path Language (XPath) expressions in the script must start at the top level. This allows the script to access all possible Junos XML and Junos XML protocol remote procedure calls (RPCs). For more information, see *XPath Overview*.

```
8   <xsl:template match="/">
```

After the `<xsl:template match="/">` tag element, the `<op-script-results>` and `</op-script-results>` container tags must be the top-level child tags, as shown in Lines 9 and 10.

```
9      <op-script-results>
      <!-- ... insert your code here ... -->
10    </op-script-results>
```

Line 11 closes the template.

```
11  </xsl:template>
```

Between Line 11 and Line 12, you can define additional XSLT templates that are called from within the `<xsl:template match="/">` template.

Line 12 closes the style sheet and the op script.

```
12 </xsl:stylesheet>
```

### SLAX Boilerplate for Op Scripts

The corresponding SLAX boilerplate is as follows:

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
import "../import/junos.xml";

match / {
  <op-script-results> {
    /*
     * Insert your code here
     */
  }
}
```

## Junos Space

- [Understanding Junos Space Support on page 6489](#)

### Understanding Junos Space Support

The Juniper Networks Junos Space application, running on a JA1500 appliance or a Junos Space Virtual Appliance, is a comprehensive platform for building and deploying applications for collaboration, productivity, and network infrastructure and operations management. Junos Space provides a runtime environment implemented as a fabric of virtual and physical appliances.

The Junos Space Network Management Platform software comprises various applications for network management and configuration, including:

- Junos Space Administration—Provides management of Junos Space fabric, databases, licenses, applications, authentication servers, tags, permission labels, DMI schemas, and troubleshooting.
- Network Director—Provides unified management of supported Juniper Networks devices in your network. By providing full network life cycle management, Network Director simplifies the discovery, configuration, visualization, monitoring, and administration of large networks.
- Service Automation—Provides an end-to-end solution designed to streamline operations and enable proactive network management for Junos OS devices. The solution consists of Advanced Insight Scripts (AI-Scripts), Junos Space Service Now and Service Insight applications, and Juniper Support Systems (JSS).



**NOTE:** Do not install Junos Space and AI-Scripts on the control plane network EX4200 switches or EX4200 Virtual Chassis in a QFX3000 QFabric system

Before you can use Junos Space Network Director to manage the QFX Series device, you must ensure that the configuration on the device meets the requirements for all managed devices. For example:

- The device configuration has a static management IP address that is reachable from the Junos Space server.
- There is a user with full administrative privileges for Junos Space administration.
- SNMP is enabled (only if you plan on using SNMP as part of the device discovery).
- In Junos Space, set up a default device management interface (DMI) schema for the QFX Series device.

For more information about Network Director requirements, see the *Network Director Quick Start Guide* at:

[http://www.juniper.net/techpubs/en\\_US/network-director1.5/information-products/pathway-pages/index.html](http://www.juniper.net/techpubs/en_US/network-director1.5/information-products/pathway-pages/index.html)

For more information about Junos Space, go to:

[http://www.juniper.net/techpubs/en\\_US/release-independent/junos-space/index.html](http://www.juniper.net/techpubs/en_US/release-independent/junos-space/index.html)

#### Related Documentation

- [Configuring SNMP on page 1356](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 1361](#)

## Network Analytics

---

- [Network Analytics Overview on page 6490](#)
- [Understanding Network Analytics Configuration and Status on page 6497](#)
- [Understanding Network Analytics Streaming Data on page 6499](#)
- [Understanding Enhanced Network Analytics Streaming Data on page 6501](#)
- [Understanding Enhanced Analytics Local File Output on page 6506](#)
- [Prototype File for the Google Protocol Buffer Stream Format on page 6508](#)

### Network Analytics Overview

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed. The analytics manager (analyticsm) in the Packet Forwarding Engine collects traffic and queue statistics, and the analytics daemon (analyticd) in the Routing Engine analyzes the data and generates reports. You can enable network analytics by configuring microburst monitoring and high-frequency traffic statistics monitoring.



**NOTE:** In Junos OS Release 13.2X51-D15, the network analytics feature was enhanced, and extensive changes were made to the CLI statements and hierarchies. If you upgrade to Junos OS Release 13.2X51-D15 or later from a release prior to 13.2X51-D15, network analytics configurations committed in previous releases will appear on your device, but the feature is disabled. To enable this feature, you must reconfigure it using the new CLI statements and hierarchies.

For more information, see:

- [Analytics Feature Overview on page 6491](#)
- [Network Analytics Enhancements Overview on page 6492](#)
- [Summary of CLI Changes on page 6493](#)

### Analytics Feature Overview

You enable network analytics by configuring queue (microburst) monitoring and high-frequency traffic statistics monitoring. You use microburst monitoring to look at traffic queue conditions in the network. A microburst occurrence indicates to the Packet Forwarding Engine that a user-specified queue depth or latency threshold is reached. The queue depth is the buffer (in bytes) containing the data, and latency is the time (in nanoseconds or microseconds) the data stays in the queue.

You can configure queue monitoring based on either queue depth or latency (but not both), and configure the frequency (polling interval) at which the Packet Forwarding Engine checks for microbursts and sends the data to the Routing Engine for processing. You may configure queue monitoring globally for all physical interfaces on the system, or for a specific interface on the switch. However, the specified queue monitoring interval applies either to all interfaces, or none; you cannot configure the interval for each interface.

You use high-frequency traffic statistics monitoring to collect traffic statistics at specified polling intervals. Similar to the queue monitoring interval, the traffic monitoring interval applies either to all interfaces, or none; you cannot configure the interval for each interface.

Both traffic and queue monitoring are disabled by default. You must configure each type of monitoring using the CLI. In each case, the configuration for an interface always takes precedence over the global configuration.



**NOTE:** You can configure traffic and queue monitoring for physical interfaces only; logical interfaces and Virtual Chassis port (VCP) interfaces are not supported.

The analyticsd daemon in the Routing Engine generates local log files containing queue and traffic statistics records. You can specify the log filename and size, and the number of log files. If you do not configure a filename, the data is not saved.

You can display the local log file or specify a server to receive the streaming data containing the queue and traffic statistics.

For each port, information for the last 10 records of traffic statistics and 100 records of queue statistics is cached. You may view this information by using the **show analytics** commands.

To store traceoptions data, you configure the **traceoptions** statement at the **[edit services analytics]** hierarchy level.

### Network Analytics Enhancements Overview

---

Beginning in Junos OS Release 13.2X51-D15, the network analytics feature provides the following enhancements:

- **Resources**—Consist of interfaces and system. The interfaces resource allows you to configure an interface name and an associated resource profile name for each interface. With the system resource, you can configure the polling intervals for queue monitoring and traffic monitoring, and an associated resource profile for the system.
- **Resource profile**—A template that contains the configurations for queue and traffic monitoring, such as depth threshold and latency threshold values, and whether each type of monitoring is enabled or disabled. Once a resource profile is configured, you apply it to a system or interfaces resource.
- **Collector**—A server for collecting queue and traffic monitoring statistics, and can be a local or remote server. You can configure a local server to store monitoring statistics in a log file, or a remote server to receive streamed statistics data.
- **Export profile**—You must configure an export profile if you wish to send streaming data to a remote collector. In the export profile, you define the category of streamed data (system-wide or interface-specific) to determine stream type the collector will receive. You can specify both system and interface stream categories. System data includes system information and status of queue and traffic monitoring. Interface-specific data includes interface information, queue and traffic statistics, and link, queue, and traffic status.
- **Google Protocol Buffer (GBP) stream format**—A new streaming format for monitoring statistics data that is sent to a remote collector in a single AnRecord message. This stream format provides nine types of information, including:
  - **System information**—General system information, including boot time, model information, serial number, number of ports, and so on.
  - **System queue status**—Queue status for the system in general.
  - **System traffic status**—Traffic status for the system in general.
  - **Interface information**—Includes SNMP index, slot, port, and other information.
  - **Queue statistics for interfaces**—Queue statistics for specific interfaces.
  - **Traffic statistics for interfaces**—Traffic statistics for specific interfaces.
  - **Link status for interfaces**—Includes link speed, state, and so on.

- Queue status for interfaces—Queue status for specific interfaces.
- Traffic status for interfaces—Traffic status for specific interfaces.
- The **analytics.proto** file—Provides a template for the GBP stream format. This file can be used for writing your analytics server application. To download the file, go to:  
[http://www.juniper.net/techpubs/en\\_US/junos13.2/topics/reference/proto-files/analytics-proto.txt](http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/proto-files/analytics-proto.txt)
- Use of threshold values—The Analytics Manager (analyticsm) will generate a queue statistics record when the lower queue depth or latency threshold value is exceeded.
- User Datagram Protocol (UDP)—Additional transport protocol you can configure, in addition to Transmission Control Protocol (TCP), for the remote streaming server port.
- Single file for local logging—Replaces the separate log files for queue and traffic statistics.
- Change in latency measurement—Configuration and reporting of latency values have changed from microseconds to nanoseconds.
- Change in reporting of the collection time in UTC format—Statistics collection time is reported in microseconds instead of milliseconds.
- New operational mode command **show analytics collector**—Replaces the **show analytics streaming-server** command.
- Changes in command output format—Include the following changes:
  - Addition of unicast, multicast, and broadcast packet counters in queue and traffic statistics.
  - Reversal of the sequence of statistics information in the output. The most recent record is displayed at the beginning, and the oldest record at the end of the output.
  - Removal of traffic or queue monitoring status information from the global portion of the **show analytics configuration** and **show analytics status** command output if there is no global configuration.
  - Addition of **n/a** to the interface-specific portion of the **show analytics configuration** and **show analytics status** command output if a parameter is not configured (for example, depth threshold or latency threshold).

### Summary of CLI Changes

Beginning in Junos OS Release 13.2X51-D15, enhancements to the network analytics feature result in changes in the CLI when you configure the feature. See [Table 603 on page 6494](#) for a summary of CLI changes.

Table 603: Network Analytics CLI Changes

Task	CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10	CLI for Junos OS Release 13.2X51-D15 and later
Configuring global queue and traffic monitoring polling interval	<pre>[edit services analytics]  traffic-statistics {   interval <i>interval</i>; } queue-statistics {   interval <i>interval</i>; }</pre>	<pre>[edit services analytics]  resource {   system {     polling-interval {       queue-monitoring <i>interval</i>;       traffic-monitoring <i>interval</i>;     }   } }</pre>
Configuring local files for traffic and queue statistics reporting	<pre>[edit services analytics]  traffic-statistics {   file <i>filename</i>;   size <i>size</i>;   files <i>number</i>; } queue-statistics {   file <i>filename</i>;   size <i>size</i>;   files <i>number</i>; }</pre>	<pre>[edit services analytics]  collector {   local {     file <i>filename</i> {       files <i>number</i>;       size <i>size</i>;     }   } }</pre>
Enabling queue statistics and traffic monitoring, and specifying the depth threshold for all interfaces (globally)	<pre>[edit services analytics]  interfaces {   all {     queue-statistics;     traffic-statistics;     depth-threshold {       high <i>number</i>;       low <i>number</i>;     }   } }</pre>	<p>Requires defining a resource profile and applying it to the system:</p> <ol style="list-style-type: none"> <li>To define a resource profile: <pre>[edit services analytics]  resource-profiles {   <i>profile-name</i> {     queue-monitoring;     traffic-monitoring;     depth-threshold {       high <i>number</i>;       low <i>number</i>;     }   } }</pre> </li> <li>To apply a profile to the system: <pre>[edit services analytics]  resource {   system {     resource-profile <i>profile-name</i>;   } }</pre> </li> </ol>



Table 603: Network Analytics CLI Changes (*continued*)

Task	CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10	CLI for Junos OS Release 13.2X51-D15 and later
Enabling queue statistics and traffic monitoring, and specifying the latency threshold for one interface	<pre>[edit services analytics] interfaces {   interface {     queue-statistics;     traffic-statistics;     latency-threshold       high <i>number</i>;       low <i>number</i>;   } }</pre>	<p>Requires defining a resource profile and applying it to the interface:</p> <ol style="list-style-type: none"> <li>To define a resource profile: <pre>[edit services analytics] resource-profiles {   profile-name {     queue-monitoring;     traffic-monitoring;     latency-threshold {       high <i>number</i>;       low <i>number</i>;     }   } }</pre> </li> <li>To apply a profile to the interface: <pre>[edit services analytics] resource {   interfaces {     interface-name {       resource-profile <i>profile-name</i>;     }   } }</pre> </li> </ol>
<p>Configuring the streaming data format (JSON, CSV, or TSV) to send to a remote server</p> <p><b>NOTE:</b> Junos OS Release 13.2X51-D15 added support for the GPB stream format and configuration of the transport protocols (TCP or UDP).</p>	<pre>[edit services analytics] streaming-servers {   address <i>ip-address</i> {     port <i>number</i> {       stream-format <i>format</i>;     }   } }</pre>	<p>Requires defining the stream format in an export profile and applying the profile to the collector.</p> <ol style="list-style-type: none"> <li>To configure the stream format: <pre>[edit services analytics] export-profiles {   profile-name {     stream-format <i>format</i>;   } }</pre> </li> <li>To apply an export profile to the collector: <pre>[edit services analytics] collector {   address <i>ip-address</i> {     port <i>number</i> {       transport <i>protocol</i> {         export-profile <i>profile-name</i>;       }     }   } }</pre> </li> </ol>

Table 603: Network Analytics CLI Changes (*continued*)

Task	CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10	CLI for Junos OS Release 13.2X51-D15 and later
Configuring the streaming message types (queue or traffic statistics) to send to a remote server	<pre> [edit services analytics] streaming-servers {   address <i>ip-address</i> {     port <i>number</i> {       stream-type <i>type</i>;       stream-type <i>type</i>;     }   } } </pre>	<p>Requires defining an export profile and applying it to the collector:</p> <ol style="list-style-type: none"> <li>To define an export profile: <pre> [edit services analytics] export-profiles {   <i>profile-name</i> {     interface {       information;       statistics {         queue;         traffic;       }       status {         link;         queue;         traffic;       }     }   }   system {     information;     status {       queue;       traffic;     }   } } </pre> </li> <li>To apply an export profile to the collector: <pre> [edit services analytics] collector {   address <i>ip-address</i> {     port <i>number</i> {       export-profile <i>profile-name</i>;     }   } } </pre> </li> </ol>

Table 603: Network Analytics CLI Changes (*continued*)

Task	CLI for Junos OS Release 13.2X50-D15 and 13.2X51-D10	CLI for Junos OS Release 13.2X51-D15 and later
Configuring the transport protocol for sending streaming data to an external server	No configuration is available. Only the TCP protocol is supported.	Configuration is available. Both TCP and UDP protocols are supported, and can be configured for the same port.  [edit services analytics]  collector { address <i>ip-address</i> { port <i>number1</i> { transport tcp; transport udp; } port <i>number2</i> { transport udp; } } }
Show information about remote streaming server or collector	Issue the <b>show analytics streaming-sever</b> command.	Issue the <b>show analytics collector</b> command.

Related Documentation • [analytics on page 6667](#)

## Understanding Network Analytics Configuration and Status

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. You can enable network analytics by configuring traffic and queue statistics monitoring.



**NOTE:** This topic describes the configuration and status output from Junos OS Release 13.2X50-D15 and 13.2X51-D10 only.

If you had enabled traffic or queue monitoring, you can issue the **show analytics configuration** and **show analytics status** commands to view the global interface configuration and status and that of specific interfaces. The output that is displayed depends on your configuration at the global interface and specific interface levels. For example:

- A global interface configuration (for all interfaces) to disable monitoring supersedes the configuration to enable it on an interface.
- The interface configuration to enable or disable monitoring supersedes the global interface configuration, unless monitoring had been disabled globally for all interfaces.
- If there is no configuration, whether for all interfaces or a specific interface, monitoring is disabled by default (see [Table 604 on page 6498](#)).

Table 604 on page 6498 describes the correlation between the user configuration and the settings that are displayed.

**Table 604: Configuration and Status Output in Junos OS Release 13.2X51-D10 and 13.2X50-D15**

User Configuration	Global or System Settings		Specific Interface Settings	
	Configuration	Status	Configuration	Status
No global or specific interface configuration. This is the default setting.	Auto	Auto	Auto	Disabled
No global interface configuration but the specific interface monitoring is disabled.	Auto	Auto	Disabled	Disabled
No global interface configuration but the specific interface monitoring is enabled.	Auto	Auto	Enabled	Enabled
Monitoring is disabled globally and there is no interface configuration.	Disabled	Disabled	Auto	Disabled
Monitoring is disabled at both the global and specific interface levels.	Disabled	Disabled	Disabled	Disabled
Monitoring is disabled at the global interface level but is enabled at the specific interface level. The global interface <i>Disabled</i> setting supersedes the <i>Enabled</i> setting for a specific interface.	Disabled	Disabled	Enabled	Disabled
Monitoring is enabled for all interfaces but there is no configuration for the specific interface .	Enabled	Enabled	Auto	Enabled
Monitoring is enabled at both the global and specific interface levels.	Enabled	Enabled	Enabled	Enabled
Monitoring is enabled for all interfaces but is disabled for the specific interface.	Enabled	Enabled	Disabled	Disabled

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [analytics on page 6667](#)
  - [queue-statistics on page 6676](#)
  - [traffic-statistics on page 6680](#)
  - [show analytics configuration on page 6835](#)
  - [show analytics status on page 6841](#)

## Understanding Network Analytics Streaming Data

This topic describes the network analytics queue and traffic statistics that are streamed to remote servers.

You can configure one or more remote servers to receive streamed data containing queue and traffic statistics. The format of the streamed data can be Javascript Object Notation (JSON), Comma-separated Values (CSV), or Tab-separated Values (TSV).



**NOTE:** The output shown in this topic applies to Junos OS Release 13.2X51-D10 only. The time is displayed in the Unix epoch format (also known as Unix time or POSIX time).

The following examples show the streamed queue statistics data output in different formats.

- JSON format:

```
{"record-type":"queue-stats","time":1383453988263,"router-id":"qfx5100-switch",
"port":"xe-0/0/18","latency":0,"queue-depth":208}
```

- CSV format:

```
q,1383454067604,qfx5100-switch,xe-0/0/18,0,208
```

- TSV format:

```
q      585870192561703872      qfx5100-switch      xe-0/0/18      (null)
208    2
```

[Table 605 on page 6499](#) describes the output fields for streamed queue statistics data in the order they appear.

**Table 605: Streamed Queue Statistics Data Output Fields**

Field	Description
record-type	Type of statistics. Displayed as: <ul style="list-style-type: none"> <li>• <b>queue-stats</b> (JSON format)</li> <li>• <b>q</b> (CSV or TSV format)</li> </ul>
time	Time (in Unix epoch format) at which the statistics were captured.
router-id	ID of the network analytics host device.
port	Name of the physical port configured for network analytics.
latency	Traffic queue latency in milliseconds.
queue depth	Depth of the traffic queue in bytes.

The following examples show the streamed traffic statistics data output in different formats.

- JSON format:

```
{"record-type":"traffic-stats","time":1383453986763,"router-id":"qfx5100-switch",
"port":"xe-0/0/16","rxpkt":26524223621,"rxpps":8399588,"rxbyte":3395100629632,
"rxbps":423997832,"rxdrop":0,"rxerr":0,"txpkt":795746503,"txpps":0,"txbyte":101855533467,
"txbps":0,"txdrop":0,"txerr":0}
```

- CSV format:

```
t,1383454072924,qfx5100-switch,xe-0/0/19,1274299748,82950,163110341556,85603312,0,0,
27254178291,8300088,3488534810679,600002408,27268587050,3490379142400
```

- TSV format:

```
t      1383454139025    qfx5100-switch    xe-0/0/19      1279874033      82022
163823850036    84801488      0      0      27811618258    8199630
3559887126455    919998736      27827356915    3561901685120
```

[Table 606 on page 6500](#) describes the output fields for streamed traffic statistics data in the order they appear.

**Table 606: Streamed Traffic Statistics Data Output Fields**

Field	Description
record-type	Type of statistics. Displayed as: <ul style="list-style-type: none"> <li><b>traffic-stats</b> (JSON format)</li> <li><b>t</b> (CSV or TSV format)</li> </ul>
time	Time (in Unix epoch format) at which the statistics were captured.
router-id	ID of the network analytics host device.
port	Name of the physical port configured for network analytics.
rxpkt	Total packets received.
rxpps	Total packets received per second.
rxbyte	Total bytes received.
rxbps	Total bytes received per second.
rxdrop	Total incoming packets dropped.
rxerr	Total packets with errors.
txpkt	Total packets transmitted.
txpps	Total packets transmitted per second.
txbyte	Total bytes transmitted.

Table 606: Streamed Traffic Statistics Data Output Fields (*continued*)

Field	Description
txbps	Total bytes transmitted per second.
txdrop	Total transmitted bytes dropped.
txerr	Total transmitted packets with errors (dropped).

**Related  
Documentation**

- [Network Analytics Overview on page 6490](#)
- [show analytics streaming-servers on page 6845](#)
- [streaming-servers on page 6678](#)

## Understanding Enhanced Network Analytics Streaming Data

Network analytics monitoring data can be streamed to remote servers called collectors. You can configure one or more collectors to receive streamed data containing queue and traffic statistics. This topic describes the streamed data output.



**NOTE:** This topic applies to Junos OS Release 13.2X51-D15 or later.

Starting in Junos OS Release 13.2X51-D15, network analytics supports the following streaming data formats and output:

- [Google Protocol Buffer \(GPB\) on page 6501](#)
- [JavaScript Object Notation \(JSON\) on page 6504](#)
- [Comma-separated Values \(CSV\) on page 6504](#)
- [Tab-separated Values \(TSV\) on page 6504](#)
- [Queue Statistics Output for JSON, CSV, and TSV on page 6505](#)
- [Traffic Statistics Output for JSON, CSV, and TSV on page 6505](#)

### Google Protocol Buffer (GPB)

Support for the Google Protocol Buffer (GPB) streaming format has been added in Junos OS Release 13.2X51-D15. This streaming format provides:

- Support for nine types of messages, based on resource type (system-wide or interface-specific).
- Sends messages in a hierarchical format.
- You can generate other stream format messages (JSON, CSV, TSV) from GPB formatted messages.
- Includes a 8-byte message header. See [Table 607 on page 6502](#) for more information.

[Table 607 on page 6502](#) describes the GPB stream format message header.

Table 607: GPB Stream Format Message Header Information

Byte Position	Field
0 to 3	Length of message
4	Message version
5 to 7	Reserved for future use

The following GPB prototype file (**analytics.proto**) provides details about the streamed data:

```
package analytics;

// Traffic statistics related info
message TrafficStatus {
    optional uint32          status          = 1;
    optional uint32          poll_interval   = 2;
}

// Queue statistics related info
message QueueStatus {
    optional uint32          status          = 1;
    optional uint32          poll_interval   = 2;
    optional uint64          lt_high         = 3;
    optional uint64          lt_low          = 4;
    optional uint64          dt_high         = 5;
    optional uint64          dt_low          = 6;
}

message LinkStatus {
    optional uint64          speed           = 1;
    optional uint32          duplex          = 2;
    optional uint32          mtu             = 3;
    optional bool            state           = 4;
    optional bool            auto_negotiation= 5;
}

message InterfaceInfo {
    optional uint32          snmp_index      = 1;
    optional uint32          index           = 2;
    optional uint32          slot           = 3;
    optional uint32          port           = 4;
    optional uint32          media_type      = 5;
    optional uint32          capability      = 6;
    optional uint32          porttype       = 7;
}

message InterfaceStatus {
    optional LinkStatus      link            = 1;
    optional QueueStatus     queue_status    = 2;
    optional TrafficStatus   traffic_status  = 3;
}

message QueueStats {
    optional uint64          timestamp       = 1;
    optional uint64          queue_depth     = 2;
    optional uint64          latency         = 3;
}
```



```

}

message TrafficStats {
    optional uint64      timestamp      = 1;
    optional uint64      rxpkt          = 2;
    optional uint64      rxucpkt       = 3;
    optional uint64      rxmcpkt      = 4;
    optional uint64      rxbcpkt      = 5;
    optional uint64      rxpps        = 6;
    optional uint64      rxbyte       = 7;
    optional uint64      rxbps        = 8;
    optional uint64      rxrcerr      = 9;
    optional uint64      rxdropkt     = 10;
    optional uint64      txpkt        = 11;
    optional uint64      txucpkt      = 12;
    optional uint64      txmcpkt     = 13;
    optional uint64      txbcpkt     = 14;
    optional uint64      txpps        = 15;
    optional uint64      txbyte       = 16;
    optional uint64      txbps        = 17;
    optional uint64      txrcerr      = 18;
    optional uint64      txdropkt     = 19;
}

message InterfaceStats {
    optional TrafficStats traffic_stats = 1;
    optional QueueStats  queue_stats  = 2;
}

//Interface message
message Interface {
    required string      name          = 1;
    optional bool        deleted       = 2;
    optional InterfaceInfo information = 3;
    optional InterfaceStats stats      = 4;
    optional InterfaceStatus status    = 5;
}

message SystemInfo {
    optional uint64      boot_time     = 1;
    optional string      model_info    = 2;
    optional string      serial_no     = 3;
    optional uint32      max_ports     = 4;
    optional string      collector     = 5;
    repeated string      interface_list = 6;
}

message SystemStatus {
    optional QueueStatus queue_status = 1;
    optional TrafficStatus traffic_status = 2;
}

//System message
message System {
    required string      name          = 1;
    optional bool        deleted       = 2;
    optional SystemInfo  information = 3;
    optional SystemStatus status      = 4;
}

message AnRecord {

```

```

optional uint64      timestamp      = 1;
optional System      system         = 2;
repeated Interface   interface      = 3;
}

```

### JavaScript Object Notation (JSON)

The JavaScript Object Notation (JSON) streaming format supports the following data:

- Queue statistics data. For example:

```

{"record-type":"queue-stats","time":1383453988263,"router-id":"qfx5100-switch",
"port":"xe-0/0/18","latency":0,"queue-depth":208}

```

See [Table 605 on page 6499](#) for more information about queue statistics output fields.

- Traffic statistics. For example:

```

{"record-type":"traffic-stats","time":1383453986763,"router-id":"qfx5100-switch",
"port":"xe-0/0/16","rxpkt":26524223621,"rxpps":8399588,"rxbyte":3395100629632,
"rxbps":423997832,"rxdrop":0,"rxerr":0,"txpkt":795746503,"txpps":0,"txbyte":101855533467,
"txbps":0,"txdrop":0,"txerr":0}

```

See [Table 606 on page 6500](#) for more information about traffic statistics output fields.

### Comma-separated Values (CSV)

The Comma-separated Values (CSV) streaming format supports the following data:

- Queue statistics. For example:

```

q,1383454067604,qfx5100-switch,xe-0/0/18,0,208

```

See [Table 605 on page 6499](#) for more information about queue statistics output fields.

- Traffic statistics. For example:

```

t,1383454072924,qfx5100-switch,xe-0/0/19,1274299748,82950,163110341556,85603312,0,0,
27254178291,8300088,3488534810679,600002408,27268587050,3490379142400

```

See [Table 606 on page 6500](#) for more information about traffic statistics output fields.

### Tab-separated Values (TSV)

The Tab-separated Values (TSV) streaming format supports the following data:

- Queue statistics. For example:

```

q      585870192561703872      qfx5100-switch      xe-0/0/18      (null)
208      2

```

See [Table 605 on page 6499](#) for more information about queue statistics output fields.

- Traffic statistics. For example:

```

t      1383454139025      qfx5100-switch      xe-0/0/19      1279874033      82022
163823850036      84801488      0      0      27811618258      8199630
3559887126455      919998736      27827356915      3561901685120

```

See [Table 606 on page 6500](#) for more information about traffic statistics output fields.

### Queue Statistics Output for JSON, CSV, and TSV

Table 605 on page 6499 describes the output fields for streamed queue statistics data in the order they appear.

**Table 608: Streamed Queue Statistics Data Output Fields**

Field	Description
record-type	Type of statistics. Displayed as: <ul style="list-style-type: none"> <li>• <b>queue-stats</b> (JSON format)</li> <li>• <b>q</b> (CSV or TSV format)</li> </ul>
time	Time (in Unix epoch format) at which the statistics were captured.
router-id	ID of the network analytics host device.
port	Name of the physical port configured for network analytics.
latency	Traffic queue latency in milliseconds.
queue depth	Depth of the traffic queue in bytes.

### Traffic Statistics Output for JSON, CSV, and TSV

Table 606 on page 6500 describes the output fields for streamed traffic statistics data in the order they appear.

**Table 609: Streamed Traffic Statistics Data Output Fields**

Field	Description
record-type	Type of statistics. Displayed as: <ul style="list-style-type: none"> <li>• <b>traffic-stats</b> (JSON format)</li> <li>• <b>t</b> (CSV or TSV format)</li> </ul>
time	Time (in Unix epoch format) at which the statistics were captured.
router-id	ID of the network analytics host device.
port	Name of the physical port configured for network analytics.
rxpkt	Total packets received.
rxpps	Total packets received per second.
rxbyte	Total bytes received.
rxbps	Total bytes received per second.

Table 609: Streamed Traffic Statistics Data Output Fields (*continued*)

Field	Description
rxdrop	Total incoming packets dropped.
rxerr	Total packets with errors.
txpkt	Total packets transmitted.
txpps	Total packets transmitted per second.
txbyte	Total bytes transmitted.
txbps	Total bytes transmitted per second.
txdrop	Total transmitted bytes dropped.
txerr	Total transmitted packets with errors (dropped).

#### Related Documentation

- [Network Analytics Overview on page 6490](#)
- [Prototype File for the Google Protocol Buffer Stream Format on page 6508](#)
- [address \(Analytics Collector\)](#)
- [collector \(Analytics\)](#)
- [show analytics collector on page 6833](#)

## Understanding Enhanced Analytics Local File Output

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. You enable network analytics by configuring queue or traffic statistics monitoring, or both. In addition, you can configure a local file for storing the traffic and queue statistics records.



**NOTE:** This topic describes the local file output in Junos OS Release 13.2X51-D15 and later. For information about local file output from earlier releases, see the [monitor start \(Analytics\)](#) topic.

Beginning in Junos OS Release 13.2X51-D15, the traffic and queue monitoring statistics can be stored locally in a single file. The following example shows the output from the **monitor start** command.

```
root@qfx5100-33> monitor start an
root@qfx5100-33>
*** an ***
q,1393947567698432,qfx5100-33,xe-0/0/19,1098572,1373216
q,1393947568702418,qfx5100-33,xe-0/0/19,1094912,1368640
q,1393947569703415,qfx5100-33,xe-0/0/19,1103065,1378832
t,1393947569874528,qfx5100-33,xe-0/0/16,12603371884,12603371884,0,0,
```

```

8426023,1613231610488,8628248712,0,3,5916761,5916761,0,0,0,757345408,0,0,0
t,1393947569874528,qfx5100-33,xe-0/0/18,12601953614,12601953614,0,0,
8446737,1613050071660,8649421552,0,5,131761619,131761619,0,0,84468,
16865487232,86495888,0,0
t,1393947569874528,qfx5100-33,xe-0/0/19,126009250,126009250,0,0,84469,
16129184128,86496392,0,0,12584980342,12584980342,0,0,8446866,1610877487744,
8649588432,12593703960,0
q,1393947575698402,qfx5100-33,xe-0/0/19,1102233,1377792
q,1393947576701398,qfx5100-33,xe-0/0/19,1107724,1384656

```

See [Table 610 on page 6507](#) for queue statistics output, and [Table 611 on page 6507](#) for traffic statistics output. The fields in the tables are listed in the order they appear in the output example.

**Table 610: Output Fields for Queue Statistics in Local Analytics File**

Field	Description	Example in Output
Record type	Type of statistics (queue or traffic monitoring)	<b>q</b>
Time (microseconds)	Unix epoch (or Unix time) in microseconds at which the statistics were captured.	<b>1393947567698432</b>
Router ID	ID of the network analytics host device.	<b>qfx5100-33</b>
Port	Name of the physical port configured for network analytics.	<b>xe-0/0/19</b>
Latency (nanoseconds)	Traffic queue latency in nanoseconds.	<b>1098572</b>
Queue depth (bytes)	Depth of the traffic queue in bytes.	<b>1373216</b>

**Table 611: Output Fields for Traffic Statistics in Local Analytics File**

Field	Description	Example in Output
Record type	Type of statistics (queue or traffic monitoring)	<b>t</b>
Time (microseconds)	Unix epoch (or Unix time) in microseconds at which the statistics were captured.	<b>1393947569874528</b>
Router ID	ID of the network analytics host device.	<b>qfx5100-33</b>
Port	Name of the physical port configured for network analytics.	<b>xe-0/0/16</b>
rxpkt	Total packets received.	<b>12603371884</b>
rxucpkt	Total unicast packets received.	<b>12603371884</b>
rxmcpkt	Total multicast packets received.	<b>0</b>
rxbcpkt	Total broadcast packets received.	<b>0</b>
rxpps	Total packets received per second.	<b>8426023</b>

Table 611: Output Fields for Traffic Statistics in Local Analytics File (*continued*)

Field	Description	Example in Output
rxbyte	Total octets received.	1613231610488
rxbps	Total bytes received per second.	8628248712
rxdroppkt	Total incoming packets dropped.	0
rxrcerr	CRC/Align errors received.	3
txpkt	Total packets transmitted.	5916761
txucpkt	Total unicast packets transmitted.	5916761
txmcpkt	Total multicast packets transmitted.	0
txbcpkt	Total broadcast packets transmitted.	0
txpps	Total packets transmitted per second.	0
txbyte	Total octets transmitted.	757345408
txbps	Bytes per second transmitted.	0
txdroppkt	Total transmitted packets dropped.	0
txrcerr	CRC/Align errors transmitted.	0

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [analytics on page 6667](#)

## Prototype File for the Google Protocol Buffer Stream Format

The Google Protocol Buffer (GBP) stream format is used for streaming monitoring statistics data to a remote collector in a single AnRecord message.

The **analytics.proto** file provides a template for the GBP stream format. This file can be used for writing your analytics server application.

To download the GPB prototype file, go to:

[http://www.juniper.net/techpubs/en\\_US/junos13.2/topics/reference/proto-files/analytics-proto.txt](http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/proto-files/analytics-proto.txt)

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [analytics on page 6667](#)
  - [export-profiles](#)

---

## sFlow Technology

---

- [Understanding How to Use sFlow Technology for Network Monitoring on a Switch on page 6509](#)

### Understanding How to Use sFlow Technology for Network Monitoring on a Switch

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station called a *collector*. You can configure sFlow technology on a Juniper Networks switch to continuously monitor traffic at wire speed on all interfaces simultaneously.

This topic describes:

- [Sampling Mechanism and Architecture of sFlow Technology on Switches on page 6509](#)
- [Adaptive Sampling on page 6511](#)
- [sFlow Agent Address Assignment on page 6511](#)
- [sFlow Limitations on Switches on page 6512](#)

---

#### Sampling Mechanism and Architecture of sFlow Technology on Switches

---

sFlow technology uses the following two sampling mechanisms:

- **Packet-based sampling**—Samples one packet out of a specified number of packets from an interface enabled for sFlow technology. Only the first 128 bytes of each packet are sent to the collector. Data collected include the Ethernet, IP, and TCP headers, along with other application-level headers (if present). Although this type of sampling might not capture infrequent packet flows, the majority of flows are reported over time, allowing the collector to generate a reasonably accurate representation of network activity. To configure packet-based sampling, you must specify a sample rate.
- **Time-based sampling**—Samples interface statistics at a specified interval from an interface enabled for sFlow technology. Statistics such as Ethernet interface errors are captured. To configure time-based sampling, you must specify a polling interval.

The sampling information is used to create a network traffic visibility picture. The Juniper Networks Junos operating system (Junos OS) fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks* (see <http://faqs.org/rfcs/rfc3176.html>).



**NOTE:** sFlow technology on the switches samples only raw packet headers. A raw Ethernet packet is the complete Layer 2 network frame.

---

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent's two main activities are random sampling and statistics gathering. It combines interface counters and flow samples and sends them across the network to the sFlow collector as UDP datagrams, directing those datagrams

to the IP address and UDP destination port of the collector. Each datagram contains the following information:

- The IP address of the sFlow agent
- The number of samples
- The interface through which the packets entered the agent
- The interface through which the packets exited the agent
- The source and destination interface for the packets
- The source and destination VLAN for the packets

EX Series switches, QFX Series switches, and the QFabric systems adopt the distributed sFlow architecture. The sFlow agent has two separate sampling entities that are associated with each Packet Forwarding Engine in case of switches and nodes in case of a QFabric system. These sampling entities are known as subagents. Each subagent has a unique ID that is used by the collector to identify the data source. A subagent has its own independent state and forwards its own sample messages to the sFlow agent. The sFlow agent is responsible for packaging the samples into datagrams and sending them to the sFlow collector. Because sampling is distributed across subagents, the protocol overhead associated with sFlow technology is significantly reduced at the collector.



**NOTE:** On the QFabric system, an sFlow collector must be reachable through the data network. Because each Node device has all routes stored in the default routing instance, the collector IP address should be included in the default routing instance to ensure the collector's reachability from the Node device.

---



**NOTE:** You cannot configure sFlow monitoring on a link aggregation group (LAG), but you can configure it individually on a LAG member interface.

---

Infrequent sampling flows might not be reported in the sFlow information, but over time the majority of flows are reported. Based on a configured sampling rate  $N$ , 1 out of  $N$  packets is captured and sent to the collector. This type of sampling does not provide a 100 percent accurate result in the analysis, but it does provide a result with quantifiable accuracy. A user-configured polling interval defines how often the sFlow data for a specific interface are sent to the collector, but an sFlow agent can also schedule polling.



**NOTE:** We recommend that you configure the same sample rate for all the ports in a line card. If you configure different sample rates, the lowest value is used for all ports on the line card..

---





**NOTE:** If the mastership assignment changes in a Virtual Chassis setup, sFlow technology continues to function.

### Adaptive Sampling

To ensure sampling accuracy and efficiency, EX Series switches and QFX Series devices use adaptive sFlow sampling. Adaptive sampling monitors the overall incoming traffic rate on the device and provides feedback to the interfaces to dynamically adapt their sampling rate to traffic conditions. The sFlow agent reads the statistics on the interfaces every few seconds (12 seconds for EX Series switches and 5 seconds for QFX Series devices) and identifies five interfaces with the highest number of samples.

On a Flexible PIC Concentrator (FPC), when the CPU processing limit is reached because of sflow sample processing, a binary backoff algorithm is initiated. This reduces the sampling load, arriving through the top five sample-producing interfaces on that FPC by half. The backoff algorithm achieves this by doubling the sampling rate on these five earmarked interfaces. This process is repeated until the CPU-load due to sflow on the given FPC comes down to an acceptable level.

On a QFabric system, sFlow technology monitors the interfaces on each node device as a group, and implements the binary backoff algorithm based on the traffic on that group of interfaces.



**NOTE:** On the QFX Series standalone switches, if you configure sFlow technology monitoring on multiple interfaces and with a high sampling rate, we recommend that you specify a collector that is on the data network instead of on the management network. Having a high volume of sFlow technology monitoring traffic on the management network might interfere with other management interface traffic.

Using adaptive sampling prevents overloading of the CPU and keeps the device operating at its optimum level even when there is a change in traffic patterns on the interfaces. The reduced sampling rate is used until the device is rebooted or when a new sampling rate is configured.



**NOTE:** sFlow technology on EX Series switches does not support graceful restart. When a graceful restart occurs, the adaptive sampling rate is set to the user-configured sampling rate.

### sFlow Agent Address Assignment

The sFlow collector uses the sFlow agent's IP address to determine the source of the sFlow data. You can configure the IP address of the sFlow agent to ensure that the agent ID of the sFlow agent remains constant. If you do not specify the IP address to be assigned

to the agent, an IP address is automatically assigned to the agent based on the following order of priority of interfaces configured on the device:

EX Series Devices	QFX Series Devices
<ol style="list-style-type: none"> <li>1. Virtual Management Ethernet (VME) interface</li> <li>2. Management Ethernet interface</li> </ol>	<ol style="list-style-type: none"> <li>1. Management Ethernet interface me0 IP address</li> <li>2. Any Layer 3 interface if the me0 IP address is not available</li> </ol>

If a particular interface is not configured, the IP address of the next interface in the priority list is used as the IP address for the agent. Once an IP address is assigned to the agent, the agent ID is not modified until the sFlow service is restarted. At least one interface has to be configured for an IP address to be assigned to the agent. When the agent's IP address is assigned automatically, the IP address is dynamic and changes when the switch reboots.

On the QFabric system, the following default values are used if the optional parameters are not configured:

- Agent ID is the management IP address of the default partition.
- Source IP is the management IP address of the default partition.

In addition, the QFabric system subagent ID (which is included in the sFlow datagrams) is the ID of the node group from which the datagram is sent to the collector.

sFlow data can be used to provide network traffic visibility information. You can explicitly configure the source IP address to be assigned to the sFlow datagrams. If you do not explicitly configure the IP address, the IP address of any of the configured Layer 3 network interfaces is used as the source IP address. If a Layer 3 IP address is not configured, then the agent IP address is used as the source IP address.

### sFlow Limitations on Switches

On the QFX Series, limitations of sFlow traffic sampling include the following:

- sFlow sampling on ingress interfaces does not capture CPU-bound traffic.
- sFlow sampling on egress interfaces does not support broadcast and multicast packets.
- Egress samples do not contain modifications made to the packet in the egress pipeline.
- If a packet is discarded because of a firewall filter, the reason code for discarding the packet is not sent to the collector.
- The out-priority field for a VLAN is always set to 0 (zero) on ingress and egress samples.
- On QFX5100 standalone switches and the QFX Series Virtual Chassis (including mixed QFX Series Virtual Chassis), egress firewall filters are not applied to sFlow sampling packets. On these platforms, the software architecture is different from that on other QFX Series devices—sFlow packets are sent by the Routing Engine (not the line card on the host) and do not transit the switch. Egress firewall filters affect data packets that are transiting a switch, but do not affect packets sent by the Routing Engine. As a result, sFlow sampling packets are always sent to the sFlow collector.

EX9200 switches support configuration of only one sampling rate (inclusive of ingress and egress rates) on an FPC. To support compatibility with the sflow configuration of other Juniper Networks products, EX9200 switches still accept multiple rate configuration on different interfaces of the same FPC. However, the switch programs the lowest rate as the sampling rate for all the interfaces of that FPC. The sFlow show command (**show sflow interfaces**) displays the configured rate and the actual (effective) rate. However, different rates on different FPCs is still supported on EX9200 switches.

#### Related Documentation

- [Example: Monitoring Network Traffic Using sFlow Technology on page 6571](#)
- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches](#)
- [Configuring sFlow Technology on page 6596](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\)](#)
- [Monitoring Interface Status and Traffic](#)

## SNMP

- [Understanding the Implementation of SNMP on page 6513](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
- [Fabric Chassis MIB on page 6518](#)
- [Utility MIB on page 6522](#)
- [SNMPv3 Overview on page 6523](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524](#)
- [Understanding RMON on page 6525](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6527](#)
- [Understanding Health Monitoring on page 6529](#)
- [SNMP MIBs Support on page 6530](#)
- [SNMP Traps Support on page 6546](#)
- [MIB Objects for the QFX Series on page 6558](#)

## Understanding the Implementation of SNMP

The QFX Series products support the Simple Network Management Protocol (SNMP) that is implemented in the Junos OS software.



**NOTE:** By default, SNMP is not enabled on devices running Junos OS. For information on enabling SNMP on a device running Junos OS, see [“Configuring SNMP” on page 1356](#).

A typical SNMP implementation includes the following components:

- Network management system (NMS)—The NMS is a combination of hardware and software that is used to monitor and administer a network. Software running on the

NMS includes the SNMP manager, which collects information about network connectivity, activity, and events by polling the managed devices.

- **Managed device**—A managed device (also called a network element) is any device managed by the NMS. Routers and switches are common examples of managed devices. The SNMP agent is the SNMP process that resides on the managed device and communicates with the NMS.
- **SNMP agent**—The SNMP agent exchanges network management information with SNMP manager software running on an NMS, or host. The agent responds to requests for information and actions from the manager. The agent also controls access to the agent's MIB, the collection of objects that can be viewed or changed by the SNMP manager.

SNMP data is stored in a highly structured, hierarchical format known as a management information base (MIB). The MIB structure is based on a tree structure, which defines a grouping of objects into related sets. Each object in the MIB is associated with an object identifier (OID), which names the object. The “leaf” in the tree structure is the actual managed object instance, which represents a resource, event, or activity that occurs in your network device. The SNMP implementation in Junos OS uses both standard (developed by IETF and documented in RFCs) and Juniper Networks enterprise-specific MIBs.

Communication between the agent and the manager occurs in one of the following forms:

- **Get, GetBulk, and GetNext requests**—The manager requests information from the agent; the agent returns the information in a **Get** response message.
- **Set requests**—The manager changes the value of a MIB object controlled by the agent; the agent indicates status in a **Set** response message.
- **Traps notification**—The agent sends traps to notify the manager of significant events that occur on the network device.

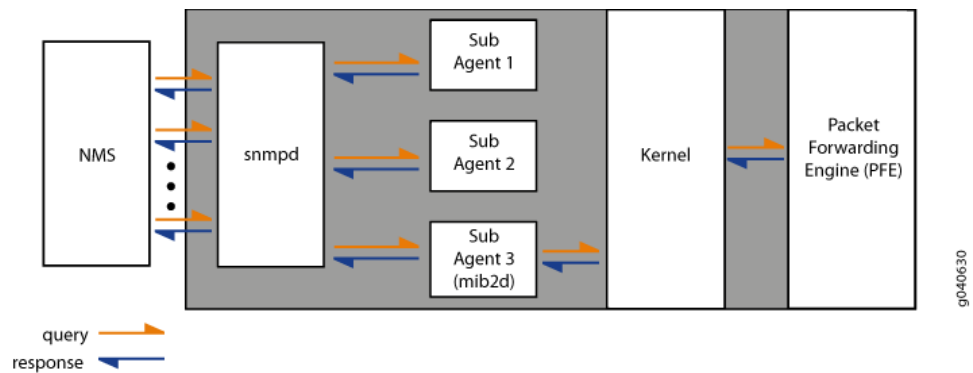
The processes maintaining the SNMP management data include:

- A master SNMP agent (known as SNMP process, or `snmpd`) that resides on the managed device and is managed by the NMS or host.
- Various subagents that reside on different modules of Junos OS, such as the Routing Engine, and are managed by the master SNMP agent.
- Junos OS processes that share data with the subagents when polled for SNMP data (for example, interface-related MIBs).

When an NMS polls the master agent for data, the master agent immediately shares the data with the NMS if the requested data is available from the master agent or one of the subagents. However, if the requested data is not maintained by the master agent or subagents, the subagent polls the Junos OS kernel or the process that maintains that data. The Junos OS kernel may need to get the data from the Packet Forwarding Engine. On receiving the required data, the subagent passes the response back on to the master agent, which in turn passes it on to the NMS.

Figure 224 on page 6515 shows the communication flow among the NMS, SNMP master agent (snmpd), SNMP subagents, Junos OS kernel, and Packet Forwarding Engine.

Figure 224: SNMP Communication Flow



When a significant event, most often an error or a failure, occurs on a network device, the SNMP agent sends notifications to the SNMP manager. SNMP notifications can be sent as traps (unconfirmed notifications) or inform requests (confirmed notifications).

Junos OS supports trap queuing to ensure that traps are not lost because of temporary unavailability of routes. Two types of queues, destination queues and a throttle queue, are formed to ensure delivery of traps and control the trap traffic. On QFX Series products, the maximum size of trap queues (throttle queue plus destination queue) is 40,960 traps. The maximum size of any one queue is 20,480 traps.

Junos OS forms a destination queue when a trap to a particular destination is returned because the host is not reachable, and it adds the subsequent traps to the same destination to the queue. Junos OS checks for availability of routes every 30 seconds, and sends the traps from the destination queue in a round-robin fashion.

If the trap delivery fails, the trap is added back to the queue, and the delivery attempt counter and the next delivery attempt timer for the queue are reset. Subsequent attempts occur at progressive intervals of 1 minute, 2 minutes, 4 minutes, and 8 minutes. The maximum delay between the attempts is 8 minutes, and the maximum number of attempts is ten. After ten unsuccessful attempts, the destination queue and all the traps in the queue are deleted.

Junos OS also has a throttle mechanism to control the number of traps (throttle threshold) sent during a particular time period (throttle interval). The throttle mechanism ensures consistency in trap traffic, especially when large numbers of traps are generated because of interface status changes. The throttle interval period begins when the first trap arrives at the throttle. All traps within the trap threshold are processed, and the traps beyond the threshold limit are queued. The default throttle threshold is 500 traps, and the throttle interval default is 5 seconds.



**NOTE:** You cannot configure trap queueing in Junos OS. You cannot view information about trap queues except for what is provided in the system logs.

- Related Documentation**
- [Configuring SNMP on page 1356](#)
  - [SNMP MIBs Support on page 6530](#)
  - [SNMP Traps Support on page 6546](#)

## Understanding the Implementation of SNMP on the QFabric System

SNMP monitors network devices from a central location. The QFabric system supports the basic SNMP architecture of Junos OS, but its implementation of SNMP differs from that of other devices running Junos OS. This topic provides an overview of the SNMP implementation on the QFabric system.

As in other SNMP systems, the SNMP manager resides on the network management system (NMS) of the network to which the QFabric system belongs. The SNMP agent resides in the QFabric Director software and is responsible for receiving and distributing all traps as well as responding to all the queries of the SNMP manager. For example, traps that are generated by a Node device are sent to the SNMP agent in the Director software, which in turn processes and sends them to the target IP addresses that are defined in the SNMP configuration.



**NOTE:** In its SNMP implementation, the QFabric system acts as an SNMP proxy server, and requires more time to process SNMP requests than a typical Junos OS device does. The default timeout setting on most SNMP client applications is 3 seconds, which is not enough time for the QFabric system to respond to SNMP requests, so the results of your `mibwalk` command may be incomplete. For this reason, we recommend that you change the SNMP timeout setting to 5 seconds or longer for the QFabric system to complete the responses to your requests.

Support for SNMP on the QFabric system includes:

- Support for the SNMP Version 1 (v1) and v2.



**NOTE:** Only SNMPv2 traps are supported on the QFabric system.

- Support for the following standard MIBs:
  - RFC 1155, *Structure and Identification of Management Information for TCP/IP-based Internets*
  - RFC 1157, *A Simple Network Management Protocol (SNMP)*
  - RFC 1212, *Concise MIB Definitions*
  - RFC 1213, *Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II* (partial support, including the system group and interfaces group)
  - RFC 1215, *A Convention for Defining Traps for use with the SNMP*
  - RFC 1901, *Introduction to Community-based SNMPv2*

- RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC 1907, *Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol Using SMIv2*
- RFC 2012, *SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2*
- RFC 2013, *SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2*
- RFC 2233, *The Interfaces Group MIB Using SMIv2*
- RFC 2571, *An Architecture for Describing SNMP Management Frameworks* (read-only access) (excluding SNMPv3)
- RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (read-only access) (excluding SNMPv3)
- RFC 2576, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 2578, *Structure of Management Information Version 2 (SMIv2)*
- RFC 2579, *Textual Conventions for SMIv2*
- RFC 2580, *Conformance Statements for SMIv2*
- RFC 2665, *Definitions of Managed Objects for the Ethernet-like Interface Types*
- RFC 2863, *The Interfaces Group MIB*
- RFC 3410, *Introduction and Applicability Statements for Internet Standard Management Framework* (excluding SNMPv3)
- RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Framework* (excluding SNMPv3)
- RFC 3412, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)* (excluding SNMPv3)
- RFC 3413, *Simple Network Management Protocol (SNMP) Applications* (excluding SNMPv3)
- RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)*
- RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework* (excluding SNMPv3)
- RFC 4188, *Definitions of Managed Objects for Bridges*

- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- RFC 4363b, *Q-Bridge VLAN MIB*
- Support for the following Juniper Networks enterprise-specific MIBs:
  - Chassis MIB (mib-jnx-chassis.txt)
  - Class-of-Service MIB (mib-jnx-cos.txt)
  - Configuration Management MIB (mib-jnx-cfgmngmt.txt)
  - Fabric Chassis MIB (mib-jnx-fabric-chassis.txt)
  - Interface MIB Extensions (mib-jnx-if-extensions.txt)
  - Power Supply Unit MIB (mib-jnx-power-supply-unit.txt)
  - QFabric MIB (mib-jnx-qf-smi.txt)
  - Utility MIB (mib-jnx-util.txt)
- Support for operational mode commands—Limited to the **show snmp statistics** command. You may issue other SNMP requests, including **get**, **get next**, and **walk** requests, by using external SNMP client applications.

**Related  
Documentation**

- [SNMP MIBs Support on page 6530](#)
- [SNMP Traps Support on page 6546](#)

## Fabric Chassis MIB

The Juniper Networks enterprise-specific SNMP Fabric Chassis MIB (mib-jnx-fabric-chassis) provides hardware information about the QFabric system and its component devices in a single MIB. The Fabric Chassis MIB is based on the Juniper Networks enterprise-specific Chassis MIB that provides information for individual devices. Unlike the Chassis MIB, the Fabric Chassis MIB represents the QFabric system component devices as part of the QFabric system. Only the information from the Fabric Chassis MIB (and not from individual Chassis MIBs) is available to SNMP management clients of the QFabric system.

The Fabric Chassis MIB uses the basic information structure of the Chassis MIB, but adds another level of indexing that provides detailed information about QFabric system devices. Each physical device in a QFabric system (such as a Node device or an Interconnect device) is represented with its hardware components, including the power supply, fans, and front and rear cards.

As in other SNMP systems, the SNMP manager resides on the network management system (NMS) of the network to which the QFabric system belongs. The SNMP agent (snmpd) resides in the QFabric system Director software and is responsible for receiving and distributing all traps as well as responding to all queries from the SNMP manager. In addition, there is an SNMP subagent running in the Routing Engine of each Node group and Interconnect device. The SNMP subagent manages the information about the component device, and that information is communicated to the SNMP agent in the Director software as needed. Traps that are generated by a Node device are sent to the



SNMP agent in the Director software, which in turn processes and sends them to the target IP addresses that are defined in the SNMP configuration.

Table 612 on page 6519 describes the tables and objects in the Fabric Chassis MIB.

**Table 612: Fabric Chassis MIB Tables and Objects**

Table or Object Name	Root OID	Description
<b>Tables with Counterparts in the Chassis MIB</b>		
jnxFabricContainersTable	1.3.6.1.4.1.2636.3.42.2.2.2	<p>Provides information about different types of containers in QFabric system devices.</p> <ul style="list-style-type: none"> <li>Containers for Interconnect devices include fan trays, power supply units, control boards, and so on.</li> <li>Containers for Node devices include fan trays, power supply units, Flexible PIC Concentrator (FPC), PICs, and so on.</li> <li>Containers for the Director devices include CPU, memory, fan trays, power supply units, and hard disks. The containers have a non-hierarchical or flat structure, and components in them are organized as siblings to each other.</li> </ul>
jnxFabricContentsTable	1.3.6.1.4.1.2636.3.42.2.2.3	<p>Contains contents that are present across all devices represented in the jnxFabricDeviceTable object. This table includes all field replaceable units (FRUs) and non-FRUs for QFabric system devices.</p> <ul style="list-style-type: none"> <li>Contents in the Interconnect devices include fan trays and control boards.</li> <li>Contents in the Node devices include fan trays and power supply units.</li> <li>Contents in the Director devices include CPUs, memory, fan trays, power supply units, and hard disks, but do not include network interface cards (NICs).</li> </ul>
jnxFabricFilledTable	1.3.6.1.4.1.2636.3.42.2.2.4	<p>Shows the status of containers in QFabric devices. The jnxFabricFilledState object represents the state of the component: (1) unknown, (2) empty, or (3) filled.</p> <p><b>NOTE:</b> The jnxFabricFilledTable object does not contain information about the Director group.</p>
jnxFabricOperatingTable	1.3.6.1.4.1.2636.3.42.2.2.5	<p>Represents different operating parameters for the contents that are populated in the jnxFabricContentsTable object.</p> <ul style="list-style-type: none"> <li>Contents in each Node device and Interconnect device include fan trays, power supply units, FPC, PIC, and Routing Engine.</li> <li>Contents in the Director device include CPUs, memory, fan trays, power supply units, and hard disks, but do not include network interface cards (NICs).</li> </ul> <p>The jnxFabricOperatingState object provides the state of the device: (1) unknown, (2) running, (3) ready, (4) reset, (5) runningAtFullSpeed (for fans only), (6) down, (6) off (for power supply units), or (7) standby.</p>

Table 612: Fabric Chassis MIB Tables and Objects (*continued*)

Table or Object Name	Root OID	Description
jnxFabricRedundancyTable	1.3.6.1.4.1.2636.3.42.2.2.6	<p>Represents the redundancy information that is available at different subsystem levels across the QFabric system. Information about the Routing Engines in Node devices is included, but there are no corresponding entries for Interconnect devices in this table. The jnxFabricRedundancyState object indicates the state of the subsystem: (1) unknown, (2) master, (3) backup, or (4) disabled.</p> <p><b>NOTE:</b> Information about redundant Director devices, virtual machines (VMs) within Director groups, and Virtual Chassis devices is not available at this time.</p>
jnxFabricFruTable	1.3.6.1.4.1.2636.3.42.2.2.7	<p>Contains all FRUs for the QFabric system in the jnxFabricDeviceTable table. The FRUs are listed regardless of whether or not they are installed or online. The jnxFabricFruState object represents the state of the FRU, including online, offline, or empty, and so on. This table also contains information about each FRU, such as name, type, temperature, time last powered on, and time last powered off.</p> <p><b>NOTE:</b> The jnxFabricFruTable table does not include network interface cards (NICs) on Director devices.</p>
<b>Table Specific to the Fabric Chassis MIB</b>		
jnxFabricDeviceTable	1.3.6.1.4.1.2636.3.42.2.2.1	<p>Contains information about all devices in the QFabric system. This table organizes scalar variables represented in the Chassis MIB into a table format for the QFabric system component devices. Columns in this table include device information such as model, device alias, and serial number. The jnxFabricDeviceIndex identifies each QFabric system device (Node device, Interconnect device, and Director device).</p> <p><b>NOTE:</b> At this time, information about the Virtual Chassis is not available.</p> <p><b>NOTE:</b> The following objects are not supported:</p> <ul style="list-style-type: none"> <li>jnxFabricDeviceEntryRevision</li> <li>jnxFabricDeviceEntryFirmwareRevision</li> <li>jnxFabricDeviceEntryKernelMemoryUsedPercent</li> </ul>

#### Scalar Variables

Table 612: Fabric Chassis MIB Tables and Objects (*continued*)

Table or Object Name	Root OID	Description
<p>The following scalar variables are supported:</p> <ul style="list-style-type: none"> <li>• jnxFabricClass</li> <li>• jnxFabricDescr</li> <li>• jnxFabricSerialNo</li> <li>• jnxFabricRevision</li> <li>• jnxFabricLastInstalled</li> <li>• jnxFabricContentsLastChange</li> <li>• jnxFabricFilledLastChange</li> </ul>	1.3.6.1.4.1.2636.3.42.2.1	<p>Describe the QFabric system as a whole.</p> <p><b>NOTE:</b> The jnxFabricFirmwareRevision scalar variable is not supported at this time.</p>

Table 613 on page 6521 describes the SNMPv2 traps that are defined in the Fabric Chassis MIB.



**NOTE:** Only SNMPv2 traps are supported on the QFabric system.

Table 613: Fabric Chassis MIB SNMPv2 Traps

Trap Group and Name	Root OID	Description
<p>jnxFabricChassisTraps group—Includes the following traps:</p> <ul style="list-style-type: none"> <li>• jnxFabricPowerSupplyFailure</li> <li>• jnxFabricFanFailure</li> <li>• jnxFabricOverTemperature</li> <li>• jnxFabricRedundancySwitchover</li> <li>• jnxFabricFruRemoval</li> <li>• jnxFabricFruInsertion</li> <li>• jnxFabricFruPowerOff</li> <li>• jnxFabricFruPowerOn</li> <li>• jnxFabricFruFailed</li> <li>• jnxFabricFruOffline</li> <li>• jnxFabricFruOnline</li> <li>• jnxFabricFruCheck</li> <li>• jnxFabricFEBSwitchover</li> <li>• jnxFabricHardDiskFailed</li> <li>• jnxFabricHardDiskMissing</li> <li>• jnxFabricBootFromBackup</li> </ul>	1.3.6.1.4.1.2636.4.19	<p>Indicates an alarm condition.</p> <p><b>NOTE:</b> Hardware events on the Director group are detected by scanning. As a result, a trap may not be generated until up to 30 seconds after the event has occurred.</p> <p><b>NOTE:</b> The software does not distinguish between the fan removal and fan failure events on the Director group. In each case, both the jnxFabricFanFailure and jnxFabricFruFailed traps are generated.</p>

Table 613: Fabric Chassis MIB SNMPv2 Traps (*continued*)

Trap Group and Name	Root OID	Description
<p>jnxFabricChassisOKTraps group—Includes the following traps:</p> <ul style="list-style-type: none"> <li>jnxFabricPowerSupplyOK</li> <li>jnxFabricFanOK</li> <li>jnxFabricTemperatureOK</li> <li>jnxFabricFruOK</li> </ul>	1.3.6.1.4.1.2636.4.20	Indicates an alarm cleared condition.

For more information, see the Fabric Chassis MIB at:

[http://www.juniper.net/techpubs/en\\_US/junos13.1/topics/reference/mibs/mib-jnx-fabric-chassis.txt](http://www.juniper.net/techpubs/en_US/junos13.1/topics/reference/mibs/mib-jnx-fabric-chassis.txt)

**Related  
Documentation**

- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
- [Chassis MIBs](#)

## Utility MIB

The Juniper Networks enterprise-specific Utility MIB, whose object ID is {jnxUtilMibRoot 1}, defines objects for counters, integers, and strings. The Utility MIB contains one table for each of the following five data types:

- 32-bit counters
- 64-bit counters
- Signed integers
- Unsigned integers
- Octet strings

Each data type has an arbitrary ASCII name, which is defined when the data is populated, and a timestamp that shows the last time when the data instance was modified. For a downloadable version of this MIB, see

[http://www.juniper.net/techpubs/en\\_US/junos14.1/topics/reference/mibs/mib-jnx-util.txt](http://www.juniper.net/techpubs/en_US/junos14.1/topics/reference/mibs/mib-jnx-util.txt).

For information about the enterprise-specific Utility MIB objects, see the following topics:

- [jnxUtilCounter32Table](#)
- [jnxUtilCounter64Table](#)
- [jnxUtilIntegerTable](#)
- [jnxUtilUintTable](#)
- [jnxUtilStringTable](#)

**Related  
Documentation**

- [Juniper Networks Enterprise-Specific MIBs](#)
- [Juniper Networks Enterprise-Specific MIBs](#)

- [Standard SNMP MIBs Supported by Junos OS](#)
- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)

## SNMPv3 Overview

The QFX3500 switch supports SNMP version 3 (SNMPv3). SNMPv3 enhances the functionality of SNMPv1 and SNMPv2c by supporting user authentication and data encryption. SNMPv3 uses the user-based security model (USM) to provide security for SNMP messages, and the view-based access control model (VACM) for user access control.

SNMPv3 features include:

- With USM, the SNMP messages between the SNMP manager and the agent can have the message source authenticated and the data integrity checked. USM reduces messaging delays and message replays by enforcing timeout limits and by checking for duplicate message request IDs.
- VACM complements USM by providing user access control for SNMP queries to the agent. You define access privileges that you wish to extend to a group of one or more users. Access privileges are determined by the security model parameters (**usm**, **v1**, or **v2**) and security level parameters (**authentication**, **privacy**, or **none**). For each security level, you must associate one MIB view for the group. Associating a MIB view with a group grants the read, write, or notify permission to a set of MIB objects for the group.
- You configure security parameters for each user, including the username, authentication type and authentication password, and privacy type and privacy password. The username given to each user is in a format that is dependent on the security model configured for that user.
- To ensure messaging security, another type of username, called the security name, is included in the messaging data that is sent between the local SNMP server and the destination SNMP server. Each user name is mapped to a security name, but the security name is in a format that is independent of the security model.
- Trap entries in SNMPv3 are created by configuring the notify, notify filter, target address, and target parameters. The **notify** statement specifies the type of notification (trap) and contains a single tag that defines a set of target addresses to receive a trap. The notify filter defines access to a collection of trap object identifiers (OIDs). The target address defines the address of an SNMP management application and other attributes used in sending notifications. Target parameters define the message processing and security parameters used in sending notifications to a particular target.

### Related Documentation

- [Assigning a Security Name to a Group on page 6612](#)
- [Configuring Access Privileges for a Group on page 6611](#)
- [Configuring SNMP Informs on page 6614](#)
- [Creating SNMPv3 Users on page 6609](#)

## Minimum SNMPv3 Configuration on a Device Running Junos OS

To configure the minimum requirements for SNMPv3, include the following statements at the `[edit snmp v3]` and `[edit snmp]` hierarchy levels:



**NOTE:** You must configure at least one view (notify, read, or write) at the `[edit snmp view-name]` hierarchy level.

```
[edit snmp]
view view-name {
  oid object-identifier (include | exclude);
}
[edit snmp v3]
notify name {
  tag tag-name;
}
notify-filter profile-name {
  oid object-identifier (include | exclude);
}
snmp-community community-index {
  security-name security-name;
}
target-address target-address-name {
  address address;
  target-parameters target-parameters-name;
}
target-parameters target-parameters-name {
  notify-filter profile-name;
  parameters {
    message-processing-model (v1 | v2c | v3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
    }
  }
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix){
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
```

```

    }
  }
  security-to-group {
    security-model (usm | v1 | v2c) {
      security-name security-name {
        group group-name;
      }
    }
  }
}

```

#### Related Documentation

- [Creating SNMPv3 Users on page 6609](#)
- [Configuring MIB Views on page 6605](#)
- [Defining Access Privileges for an SNMP Group](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 6613](#)
- [Configuring SNMP Informs on page 6614](#)
- [Complete SNMPv3 Configuration Statements](#)
- [Example: SNMPv3 Configuration](#)

## Understanding RMON

- [RMON Overview on page 6525](#)
- [Alarm Thresholds and Events on page 6526](#)

### RMON Overview

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

An operational support system (OSS) or a fault-monitoring system can be used to automatically monitor events that track many different metrics, including performance, availability, faults, and environmental data. For example, an administrator might want to know when the internal temperature of a chassis has risen above a configured threshold, which might indicate that a chassis fan tray is faulty, the chassis air flow is impeded, or the facility cooling system in the vicinity of the chassis is not operating normally.

The RMON MIB also defines tables that store various statistics for Ethernet interfaces, including the **etherStatsTable** and the **etherHistoryTable**. The **etherStatsTable** contains cumulative real-time statistics for Ethernet interfaces, such as the number of unicast, multicast, and broadcast packets received on an interface. The **etherHistoryTable** maintains a historical sample of statistics for Ethernet interfaces. The control of the **etherHistoryTable**, including the interfaces to track and the sampling interval, is defined by the RMON **historyControlTable**.

To enable RMON alarms, you perform the following steps:


1. Configure SNMP, including trap groups. You configure SNMP at the `[edit snmp]` hierarchy level.
2. Configure rising and falling events in the `eventTable`, including the event types and trap groups. You can also configure events using the CLI at the `[edit snmp rmon event]` hierarchy level.
3. Configure alarms in the `alarmTable`, including the variables to monitor, rising and falling thresholds, the sampling types and intervals, and the corresponding events to generate when alarms occur. You can also configure alarms using the CLI at the `[edit snmp rmon alarm]` hierarchy level.

Extensions to the `alarmTable` are defined in the Juniper Networks enterprise-specific MIB `jnxRmon` (`mib-jnx-rmon.txt`).

### Alarm Thresholds and Events

By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range (see [Figure 225 on page 6526](#)).

Figure 225: Setting Thresholds



Events are only generated when the alarm threshold is first crossed in any one direction rather than after each sample interval. For example, if a rising threshold alarm, along with its corresponding event, is raised, no more threshold crossing events occur until a corresponding falling alarm occurs. This considerably reduces the quantity of events that are produced by the system, making it easier for operations staff to react when events do occur.

Before you configure remote monitoring, you should identify what variables need to be monitored and their allowable operational range. This requires some period of baselining to determine the allowable operational ranges. An initial baseline period of at least



3 months is not unusual when you first identify the operational ranges and define thresholds, but baseline monitoring should continue over the life span of each monitored variable.

**Related Documentation**

- [Configuring RMON Alarms and Events on page 6606](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6527](#)

## RMON MIB Event, Alarm, Log, and History Control Tables

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and its corresponding event are generated. The event can be logged and can generate an SNMP trap.

[Table 614 on page 6527](#) provides each field in the RMON eventTable, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the `[edit snmp rmon]` hierarchy level.

**Table 614: RMON Event Table**

Field	Description	Statement [edit snmp rmon]
eventDescription	Text description of this event.	<b>description</b>
eventType	Type of event (for example, log, trap, or log and trap).	<b>type</b>
eventCommunity	Trap group to which to send this event, as defined in the Junos OS configuration. (This is not the same as the SNMP community.)	<b>community</b>
eventOwner	Entity (for example, manager) that created this event.	—
eventStatus	Status of this row (for example, valid, invalid, or createRequest).	—

[Table 615 on page 6527](#) provides each field in the RMON alarmTable, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the `[edit snmp rmon]` hierarchy level.

**Table 615: RMON Alarm Table**

Field	Description	Statement [edit snmp rmon]
alarmStatus	Status of this row (for example, valid, invalid, or createRequest)	—
alarmInterval	Sampling period (in seconds) of the monitored variable	<b>interval</b>
alarmVariable	Object identifier (OID) and instance of the variable to be monitored	—

Table 615: RMON Alarm Table (*continued*)

Field	Description	Statement [edit snmp rmon]
alarmValue	Actual value of the sampled variable	—
alarmSampleType	Sample type (absolute or delta changes)	<b>sample-type</b>
alarmStartupAlarm	Initial alarm (rising, falling, or either)	<b>startup-alarm</b>
alarmRisingThreshold	Rising threshold against which to compare the value	<b>rising-threshold</b>
alarmFallingThreshold	Falling threshold against which to compare the value	<b>falling-threshold</b>
alarmRisingEventIndex	Index (row) of the rising event in the event table	<b>rising-event-index</b>
alarmFallingEventIndex	Index (row) of the falling event in the event table	<b>falling-event-index</b>

Table 616 on page 6528 provides each field in the `jnxRmon jnxRmonAlarmTable`, which is an extension to the RMON `alarmTable`. You can troubleshoot the RMON agent, `rmopd`, that runs on a switch by inspecting the contents of the `jnxRmonAlarmTable` object.

Table 616: jnxRmon Alarm Table

Field	Description
<code>jnxRmonAlarmGetFailCnt</code>	Number of times the internal <b>Get</b> request for the variable failed
<code>jnxRmonAlarmGetFailTime</code>	Value of the <code>sysUpTime</code> object when the last failure occurred
<code>jnxRmonAlarmGetFailReason</code>	Reason why the <b>Get</b> request failed
<code>jnxRmonAlarmGetOkTime</code>	Value of the <code>sysUpTime</code> object when the variable moved out of failure state
<code>jnxRmonAlarmState</code>	Status of this alarm entry

Table 617 on page 6528 provides each field in the RMON `historyControlTable`, the description of the field, and the corresponding Junos OS statement that you can use to configure the field. The Junos OS statements reside at the **[edit snmp rmon history]** hierarchy level. The `historyControlTable` controls the RMON `etherHistoryTable`.

Table 617: RMON History Control Table

Field	Description	Statement [edit snmp rmon history]
<code>historyControlDataSource</code>	Identifies the source of the data for which historical data was collected.	<b>interface</b>

Table 617: RMON History Control Table (*continued*)

Field	Description	Statement [edit snmp rmon history]
historyControlBucketsRequested	Requested number of discrete time intervals over which data is to be saved.	<b>bucket-size</b>
historyControlBucketsGranted	Number of discrete sampling intervals over which data is to be saved.	—
historyControlInterval	Interval, in seconds, over which the data is sampled for each bucket.	<b>interval</b>
historyControlOwner	Entity that configured this entry.	<b>owner</b>
historyControlStatus	Status of this entry.	—

**Related Documentation**

- [Configuring RMON Alarms and Events on page 6606](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [Understanding RMON on page 6525](#)

## Understanding Health Monitoring

Health monitoring is an SNMP feature that extends the RMON alarm infrastructure to provide monitoring for a predefined set of objects (such as file system usage, CPU usage, and memory usage), and for Junos OS processes.

You enable the health monitor feature using the **health-monitor** statement at the **[edit snmp]** hierarchy level. You can also configure health monitor parameters such as a falling threshold, rising threshold, and interval. If the value of a monitored object exceeds the rising or falling threshold, an alarm is triggered and an event may be logged.

The falling threshold is the lower threshold for the monitored object instance. The rising threshold is the upper threshold for the monitored object instance. Each threshold is expressed as a percentage of the maximum possible value. The interval represents the period of time, in seconds, over which the object instance is sampled and compared with the rising and falling thresholds.

Events are only generated when a threshold is first crossed in any one direction, rather than after each sample interval. For example, if a rising threshold alarm, along with its corresponding event, is raised, no more threshold crossing events occur until a corresponding falling alarm occurs.

System log entries for health monitor events have a corresponding HEALTHMONITOR tag and not a generic SNMPD\_RMON\_EVENTLOG tag. However, the health monitor sends generic RMON risingThreshold and fallingThreshold traps. You can use the **show snmp health-monitor** operational command to view information about health monitor alarms and logs.

When you configure the health monitor, monitoring information for certain object instances is available, as shown in [Table 618 on page 6530](#).

**Table 618: Monitored Object Instances**

Object	Description
jnxHrStoragePercentUsed.1	Monitors the <b>/dev/ad0s1a</b> : file system on the switch. This is the root file system mounted on <b>/</b> .
jnxHrStoragePercentUsed.2	Monitors the <b>/dev/ad0s1e</b> : file system on the switch. This is the configuration file system mounted on <b>/config</b> .
jnxOperatingCPU (RE0)	Monitors CPU usage by the Routing Engine (RE0).
jnxOperatingBuffer (RE0)	Monitors the amount of memory available on the Routing Engine (RE0).
sysApplElmtRunCPU	Monitors the CPU usage for each Junos OS process (also called daemon). Multiple instances of the same process are monitored and indexed separately.
sysApplElmtRunMemory	Monitors the memory usage for each Junos OS process. Multiple instances of the same process are monitored and indexed separately.

- Related Documentation**
- [Configuring Health Monitoring on page 6609](#)
  - [falling-threshold \(Health Monitor\) on page 1415](#)
  - [interval \(Health Monitor\) on page 1420](#)
  - [rising-threshold \(Health Monitor\) on page 1451](#)
  - [show snmp health-monitor on page 6867](#)

## SNMP MIBs Support

The QFX Series standalone switches, QFX Series Virtual Chassis, and QFabric systems support standard MIBs and Juniper Networks enterprise-specific MIBs.

For more information, see:

- [MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis on page 6530](#)
- [MIBs Supported on QFabric Systems on page 6539](#)

### MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

The QFX Series standalone switches and QFX Series Virtual Chassis support both standard MIBs and Juniper Networks enterprise-specific MIBs. For more information, see:

- [Table 619 on page 6531](#) for standard MIBs.
- [Table 620 on page 6536](#) for Juniper Networks enterprise-specific MIBs.

Table 619: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

RFC	Additional Information
IEEE 802.1ab section 12.1, <i>Link Layer Discovery Protocol (LLDP) MIB</i>	Supported tables and objects: <ul style="list-style-type: none"> <li>• lldpRemManAddrOID</li> <li>• lldpLocManAddrOID</li> <li>• lldpReinitDelay</li> <li>• lldpNotificationInterval</li> <li>• lldpStatsRxPortFramesDiscardedTotal</li> <li>• lldpStatsRxPortFramesError</li> <li>• lldpStatsRxPortTLVsDiscardedTotal</li> <li>• lldpStatsRxPortTLVsUnrecognizedTotal</li> <li>• lldpStatsRxPortAgeoutsTotal</li> </ul>
IEEE 802.3ad, <i>Aggregation of Multiple Link Segments</i>	The following tables and objects are supported: <ul style="list-style-type: none"> <li>• dot3adAggPortTable, dot3adAggPortListTable, dot3adAggTable, and dot3adAggPortStatsTable</li> <li>• dot3adAggPortDebugTable (only dot3adAggPortDebugRxState, dot3adAggPortDebugMuxState, dot3adAggPortDebugActorSyncTransitionCount, dot3adAggPortDebugPartnerSyncTransitionCount, dot3adAggPortDebugActorChangeCount, and dot3adAggPortDebugPartnerChangeCount)</li> <li>• dot3adTablesLastChanged</li> </ul>
RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i>	—
RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>	—
RFC 1212, <i>Concise MIB Definitions</i>	—
RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i>	The following areas are supported: <ul style="list-style-type: none"> <li>• MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> <li>• Statistics counters</li> <li>• IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>)</li> <li>• ipAddrTable</li> <li>• SNMP management</li> <li>• Interface management</li> </ul> </li> <li>• SNMPv1 <b>Get</b>, <b>GetNext</b> requests, and SNMPv2 <b>GetBulk</b> request</li> <li>• Junos OS-specific secured access list</li> <li>• Master configuration keywords</li> <li>• Reconfigurations upon SIGHUP</li> </ul>

**Table 619: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

RFC	Additional Information
RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>	Support is limited to MIB II SNMP version 1 traps and version 2 notifications.
RFC 1286, <i>Definitions of Managed Objects for Bridges</i>	—
RFC 1657, <i>Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2</i>	—
RFC 1850, <i>OSPF Version 2 Management Information Base</i>	The following table, objects, and traps are not supported: <ul style="list-style-type: none"> <li>• Host Table</li> <li>• ospfOriginateNewLsas and ospfRxNewLsas objects</li> <li>• ospfOriginateLSA, ospfLsdbOverflow, and ospfLsdbApproachingOverflow traps</li> </ul>
RFC 1901, <i>Introduction to Community-based SNMPv2</i>	—
RFC 1905, <i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>	—
RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>	—
RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>	—
RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i>	<b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.
RFC 2287, <i>Definitions of System-Level Managed Objects for Applications</i>	The following objects are supported: <ul style="list-style-type: none"> <li>• sysApplInstallPkgTable</li> <li>• sysApplInstallElmtTable</li> <li>• sysApplElmtRunTable</li> <li>• sysApplMapTable</li> </ul>

**Table 619: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

RFC	Additional Information
RFC 2570, <i>Introduction to Version 3 of the Internet-standard Network Management Framework</i>	—
RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access)	<b>NOTE:</b> RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.
RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access)	<b>NOTE:</b> RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.
RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	<b>NOTE:</b> RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.
RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>	—
RFC 2579, <i>Textual Conventions for SMIv2</i>	—
RFC 2580, <i>Conformance Statements for SMIv2</i>	—
RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>	—
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	Support does not include row creation, the Set operation, and the vrrpStatsPacketLengthErrors object.
RFC 2790, <i>Host Resources MIB</i>	Support is limited to the following objects: <ul style="list-style-type: none"> <li>Only hrStorageTable. The file systems <code>/</code>, <code>/config</code>, <code>/var</code>, and <code>/tmp</code> always return the same index number. When SNMP restarts, the index numbers for the remaining file systems might change.</li> <li>Only the objects of the hrSystem and hrSWInstalled groups.</li> </ul>
RFC 2819, <i>Remote Network Monitoring Management Information Base</i>	The following objects are supported: <ul style="list-style-type: none"> <li>etherStatsTable (for Ethernet interfaces only), alarmTable, eventTable, and logTable.</li> <li>historyControlTable and etherHistoryTable (except the etherHistoryUtilization object).</li> </ul>
RFC 2863, <i>The Interfaces Group MIB</i>	<b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.
RFC 2932, <i>IPv4 Multicast Routing MIB</i>	—

**Table 619: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

RFC	Additional Information
RFC 2933, <i>Internet Group Management Protocol (IGMP) MIB</i>	—
RFC 2934, <i>Protocol Independent Multicast MIB for IPv4</i>	In Junos OS, RFC 2934 is implemented based on a draft version, <i>pimmib.mib</i> , of the now standard RFC.
RFC 3410, <i>Introduction and Applicability Statements for Internet Standard Management Framework</i>	—
RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i>	<b>NOTE:</b> RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571.
RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>	<b>NOTE:</b> RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572.
RFC 3413, <i>Simple Network Management Protocol (SNMP) Applications</i>	All MIBs are supported except for the Proxy MIB.
RFC 3414, <i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i>	—
RFC 3415, <i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i>	—
RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i>	<b>NOTE:</b> RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.
RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>	—
RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i>	<b>NOTE:</b> RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.
RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	—
RFC 3826, <i>The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model</i>	—



**Table 619: Standard MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

RFC	Additional Information
RFC 4188, <i>Definitions of Managed Objects for Bridges</i>	<p>The QFX3500 and QFX3600 switches support 802.1D STP (1998) and the following subtrees and objects only:</p> <ul style="list-style-type: none"> <li>• dot1dTp subtree—dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus objects from the dot1dTpFdbTable table.</li> <li>• dot1dBase subtree—dot1dBasePort and dot1dBasePortIfIndex objects from the dot1dBasePortTable table.</li> </ul> <p><b>NOTE:</b> On QFX3500 and QFX3600 switches, the dot1dTpFdbTable table is populated only with MAC addresses learned on the default VLAN. To see the MAC addresses of all VLANs, specify the dot1qTpFdbTable table (RFC 4363b, <i>Q-Bridge VLAN MIB</i>) when you issue the <b>show snmp mib walk</b> command.</p>
RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>	Supports the ipAddrTable table only.
RFC 4318, <i>Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol</i>	Supports 802.1w and 802.1t extensions for RSTP.
RFC 4363b, <i>Q-Bridge VLAN MIB</i>	<p><b>NOTE:</b> On QFX3500 and QFX3600 switches, the dot1dTpFdbTable table (RFC 4188, <i>Definitions of Managed Objects for Bridges</i>) is populated only with MAC addresses learned on the default VLAN. To see the MAC addresses of all VLANs, specify the dot1qTpFdbTable table (in this MIB) when you issue the <b>show snmp mib walk</b> command.</p>
RFC 4444, <i>IS-IS MIB</i>	—
Internet Assigned Numbers Authority, <i>IANAiftype Textual Convention MIB</i> (referenced by RFC 2233)	See <a href="http://www.iana.org/assignments/ianaiftype-mib">http://www.iana.org/assignments/ianaiftype-mib</a> .
Internet draft draft-reeder-snmpv3-usm-3desede-00.txt, <i>Extension to the User-Based Security Model (USM) to Support Triple-DES EDE in 'Outside' CBC Mode</i>	—
Internet draft draft-ietf-idmr-igmp-mib-13.txt, <i>Internet Group Management Protocol (IGMP) MIB</i>	—
ESO Consortium MIB	<p><b>NOTE:</b> The ESO Consortium MIB has been replaced by RFC 3826. See <a href="http://www.snmp.com/eso/">http://www.snmp.com/eso/</a>.</p>

**Table 620: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis**

MIB	Description
Alarm MIB (mib-jnx-chassis-alarm)	<p>Provides support for alarms from the switch.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chassis-alarm.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chassis-alarm.txt</a>.</p> <p>For more information, see <i>Alarm MIB</i>.</p>
Analyzer MIB (mib-jnx-analyzer)	<p>Contains analyzer and remote analyzer data related to port mirroring.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-analyzer.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-analyzer.txt</a>.</p> <p>For more information, see <i>Analyzer MIB</i>.</p>
Chassis MIB (mib-jnx-chassis)	<p>Provides support for environmental monitoring (power supply state, board voltages, fans, temperatures, and airflow) and inventory support for the chassis, Flexible PIC Concentrators (FPCs), and PICs.</p> <p><b>NOTE:</b> The jnxLEDTable table has been deprecated.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chassis.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chassis.txt</a>.</p> <p>For more information, see <i>Chassis MIBs</i>.</p>
Chassis Definitions for Router Model MIB (mib-jnx-chas-defines)	<p>Contains the object identifiers (OIDs) that are used by the Chassis MIB to identify routing and switching platforms and chassis components. The Chassis MIB provides information that changes often, whereas the Chassis Definitions for Router Model MIB provides information that changes less often.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chas-defines.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-chas-defines.txt</a>.</p> <p>For more information, see <i>Chassis MIBs</i>.</p>
Class-of-Service MIB (mib-jnx-cos)	<p>Provides support for monitoring interface output queue statistics per interface and per forwarding class.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cos.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cos.txt</a>.</p> <p>For more information, see <i>Class-of-Service MIB</i>.</p>

**Table 620: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

MIB	Description
Configuration Management MIB (mib-jnx-cfgmgmt)	<p>Provides notification for configuration changes and rescue configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made.</p> <p>A history of the last 32 configuration changes is kept in jnxCmChgEventTable.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cfgmgmt.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cfgmgmt.txt</a>.</p> <p>For more information, see <i>Configuration Management MIB</i>.</p>
Ethernet MAC MIB (mib-jnx-mac)	<p>Monitors media access control (MAC) statistics on Gigabit Ethernet intelligent queuing (IQ) interfaces. It collects MAC statistics; for example, inoctets, inframes, outoctets, and outframes on each source MAC address and virtual LAN (VLAN) ID for each Ethernet port.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mac.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mac.txt</a>.</p> <p>For more information, see <i>Ethernet MAC MIB</i>.</p>
Event MIB (mib-jnx-event)	<p>Defines a generic trap that can be generated using an operations script or event policy. This MIB provides the ability to specify a system log string and raise a trap if that system log string is found.</p> <p>In Junos OS release 13.2X51-D10 or later, if you configured an event policy to raise a trap when a new SNMP trap target is added, the SNMPD_TRAP_TARGET_ADD_NOTICE trap is generated with information about the new target.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-event.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-event.txt</a>.</p> <p>For more information, see <i>Event MIB</i>.</p>
Firewall MIB (mib-jnx-firewall)	<p>Provides support for monitoring firewall filter counters.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-firewall.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-firewall.txt</a>.</p> <p>For more information, see <i>Firewall MIB</i>.</p>

**Table 620: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

MIB	Description
Host Resources MIB (mib-jnx-hostresources)	<p>Extends the hrStorageTable object, providing a measure of the usage of each file system on the switch as a percentage. Previously, the objects in the hrStorageTable measured the usage in allocation units—hrStorageUsed and hrStorageAllocationUnits—only. Using the percentage measurement, you can more easily monitor and apply thresholds on usage.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-hostresources.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-hostresources.txt</a>.</p> <p>For more information, see <i>Host Resources MIB</i>.</p>
Interface MIB (Extensions) (mib-jnx-if-extensions)	<p>Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information in the ifJnxTable and ifChassisTable tables.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-if-extensions.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-if-extensions.txt</a>.</p> <p>For more information, see <i>Interface MIB</i>.</p>
MPLS MIB (mib-jnx-mpls)	<p>Provides MPLS information and defines MPLS notifications.</p> <p><b>NOTE:</b> This MIB is not supported on the QFX5100 switch.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mpls.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mpls.txt</a>.</p> <p>For more information, see <i>MPLS MIB</i>.</p>
MPLS LDP MIB (mib-jnx-mpls-ldp)	<p>Contains object definitions as described in RFC 3815, <i>Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)</i>.</p> <p><b>NOTE:</b> This MIB is not supported on the QFX5100 switch.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mpls-ldp.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-mpls-ldp.txt</a>.</p> <p>For more information, see <i>MPLS LDP MIB</i>.</p>
Ping MIB (mib-jnx-ping)	<p>Extends the standard Ping MIB control table (RFC 2925). Items in this MIB are created when entries are created in pingCtlTable of the Ping MIB. Each item is indexed exactly as it is in the Ping MIB.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-ping.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-ping.txt</a>.</p> <p>For more information, see <i>PING MIB</i>.</p>

**Table 620: Juniper Networks Enterprise-Specific MIBs Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

MIB	Description
RMON Events and Alarms MIB (mib-jnx-rmon)	<p>Supports Junos OS extensions to the standard Remote Monitoring (RMON) Events and Alarms MIB (RFC 2819). The extension augments the alarmTable object with additional information about each alarm. Two additional traps are also defined to indicate when problems are encountered with an alarm.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-rmon.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-rmon.txt</a>.</p> <p>For more information, see <i>RMON Events and Alarms MIB</i>.</p>
Structure of Management Information MIB (mib-jnx-smi)	<p>Explains how the Juniper Networks enterprise-specific MIBs are structured.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-smi.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-smi.txt</a>.</p> <p>For more information, see <i>Structure of Management Information MIB</i>.</p>
System Log MIB (mib-jnx-syslog)	<p>Enables notification of an SNMP trap-based application when an important system log message occurs.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-syslog.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-syslog.txt</a>.</p> <p>For more information, see <i>System Log MIB</i>.</p>
Utility MIB (mib-jnx-util)	<p>Provides you with SNMP MIB container objects of the following types: 32-bit counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-util.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-util.txt</a>.</p> <p>For more information, see “Utility MIB” on page 6522 and “Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage” on page 6810.</p>
VLAN MIB (mib-jnx-vlan)	<p>Contains information about prestandard IEEE 802.10 VLANs and their association with LAN emulation clients.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-vlan.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-vlan.txt</a>.</p> <p>For more information, see <i>VLAN MIB</i>.</p>

### MIBs Supported on QFabric Systems

The QFabric systems support both standard MIBs and Juniper Networks enterprise-specific MIBs. For more information, see:

- [Table 621 on page 6540](#) for standard MIBs.
- [Table 622 on page 6543](#) for Juniper Networks enterprise-specific MIBs.

**Table 621: Standard MIBs Supported on QFabric Systems**

RFC	Additional Information
RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i>	—
RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>	—
RFC 1212, <i>Concise MIB Definitions</i>	—
RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i>	<p>The following areas are supported:</p> <ul style="list-style-type: none"> <li>• MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> <li>• Statistics counters</li> <li>• IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>)</li> <li>• ipAddrTable</li> <li>• SNMP management</li> <li>• Interface management</li> </ul> </li> <li>• SNMPv1 <b>Get</b>, <b>GetNext</b> requests, and version 2 <b>GetBulk</b> request</li> <li>• Junos OS-specific secured access list</li> <li>• Master configuration keywords</li> <li>• Reconfigurations upon SIGHUP</li> </ul>
RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>	Support is limited to MIB II SNMP version 1 traps and version 2 notifications.
RFC 1286, <i>Definitions of Managed Objects for Bridges</i>	—
RFC 1901, <i>Introduction to Community-based SNMPv2</i>	—
RFC 1905, <i>Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	—
RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>	<p><b>NOTE:</b> On the QFabric system, for the SNMP mibwalk request to work, you must configure the IP address of at least one interface besides the management Ethernet interfaces (me0 and me1) in the Director group.</p>

Table 621: Standard MIBs Supported on QFabric Systems (*continued*)

RFC	Additional Information
RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>	—
RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>	—
RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i>	<p><b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.</p> <p><b>NOTE:</b> The QFabric system supports the following objects only: ifNumber, ifTable, and ifxTable.</p>
RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access)	<b>NOTE:</b> RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.
RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access)	<b>NOTE:</b> RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.
RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	<b>NOTE:</b> RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.
RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>	—
RFC 2579, <i>Textual Conventions for SMIv2</i>	—
RFC 2580, <i>Conformance Statements for SMIv2</i>	—
RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>	<p>The QFabric system supports the following tables only:</p> <ul style="list-style-type: none"> <li>• dot3StatsTable—There is one row with statistics for each Ethernet-like interface in the QFabric system. The dot3StatsIndex is an interface index that is unique across the system.</li> <li>• dot3ControlTable—There is one row in this table for each Ethernet-like interface in the QFabric system that implements the MAC control sublayer. OIDs supported are dot3ControlFunctionsSupported and dot3ControlInUnknownOpcode.</li> <li>• dot3PauseTable—There is one row in this table for each Ethernet-like interface in the QFabric system that supports the MAC control PAUSE function. OIDs supported are dot3PauseAdminMode, dot3PauseOperMode, dot3InPauseFrames, and dot3OutPauseFrames.</li> </ul> <p><b>NOTE:</b> Scalar variables are not supported on the QFabric system.</p>

Table 621: Standard MIBs Supported on QFabric Systems (*continued*)

RFC	Additional Information
RFC 2863, <i>The Interfaces Group MIB</i>	<p><b>NOTE:</b> RFC 2233 has been replaced by RFC 2863. However, Junos OS supports both RFC 2233 and RFC 2863.</p> <p><b>NOTE:</b> The QFabric system supports the following objects only: ifNumber, ifTable, and ifxTable.</p>
RFC 2933, <i>Internet Group Management Protocol (IGMP) MIB</i>	—
RFC 3410, <i>Introduction and Applicability Statements for Internet Standard Management Framework</i>	—
RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i>	<b>NOTE:</b> RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571.
RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i>	<b>NOTE:</b> RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572.
RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i>	<b>NOTE:</b> RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.
RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>	—
RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i>	<b>NOTE:</b> RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.
RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	—
RFC 4188, <i>Definitions of Managed Objects for Bridges</i>	<p>The QFabric system support is limited to the following objects:</p> <ul style="list-style-type: none"> <li>Under the dot1dBase OID, the dot1dBasePortTable table supports only the first two columns in the table: dot1dBasePort and dot1dBasePortIfIndex.</li> <li>The system does not implement the optional traps supporting dot1dNotifications (dot1dBridge 0).</li> <li>Under the dot1dStp OID, supports only the dot1dStpPortTable table. Does not support the scalar variables under dot1dStp.</li> <li>The system does not support scalar variables under dot1dTp, but under that, the dot1dTpFdbTable table is supported (dot1dBridge 4).</li> <li>For OIDs with tables support only, scalar values that are returned by the SNMP agent may not be meaningful and are therefore not recommended for use.</li> </ul>



Table 621: Standard MIBs Supported on QFabric Systems (*continued*)

RFC	Additional Information
RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>	<p>Supports the ipAddrTable table only.</p> <p>On the QFabric system, supported objects in the ipAddrTable table include: ipAdEntAddr, ipAdEntIfIndex, ipAdEntNetMask, ipAdEntBcastAddr, and ipAdEntReasmMaxSize.</p> <p><b>NOTE:</b> On the QFabric system, for the SNMP mibwalk request to work, you must configure the IP address of at least one interface besides the management Ethernet interfaces (me0 and me1) in the Director group.</p>
RFC 4363b, <i>Q-Bridge VLAN MIB</i>	<p>The QFabric system supports the following tables only:</p> <ul style="list-style-type: none"> <li>• dot1qTpFdbTable</li> <li>• dot1qVlanStaticTable</li> <li>• dot1qPortVlanTable</li> <li>• dot1qFdbTable</li> </ul>

Table 622: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems

MIB	Description
Analyzer MIB (mib-jnx-analyzer)	<p>Contains analyzer and remote analyzer data related to port mirroring.</p> <p>The QFabric system supports:</p> <ul style="list-style-type: none"> <li>• Analyzer table—jnxAnalyzerName, jnxMirroringRatio, jnxLossPriority.</li> <li>• Analyzer input table—jnxAnalyzerInputValue, jnxAnalyzerInputOption, jnxAnalyzerInputType.</li> <li>• Analyzer output table—jnxAnalyzerOutputValue, jnxAnalyzerOutputType.</li> </ul> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-analyzer.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-analyzer.txt</a>.</p> <p>For more information, see <i>Analyzer MIB</i>.</p>
Chassis MIB (mib-jnx-chassis)	<p><b>NOTE:</b> The Chassis MIB has been deprecated for the QFabric system. We recommend that you use the Fabric Chassis MIB (mib-jnx-fabric-chassis) for information about the QFabric system.</p>

**Table 622: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems** (*continued*)

MIB	Description
Class-of-Service MIB (mib-jnx-cos)	<p>Provides support for monitoring interface output queue statistics per interface and per forwarding class.</p> <p>The QFabric system supports the following tables and objects:</p> <ul style="list-style-type: none"> <li>• Jnxcosifstatflagtable—jnxCosIfstatFlags and jnxCosIfIndex.</li> <li>• Jnxcosqstattable—jnxCosQstatTxedPkts, jnxCosQstatTxedPktRate, jnxCosQstatTxedBytes, and jnxCosQstatTxedByteRate.</li> <li>• Jnxcosfcidtable—jnxCosFcIdToFcName.</li> <li>• Jnxcosfctable—jnxCosFcQueueNr.</li> </ul> <p>The QFabric system does not support any traps for this MIB.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cos.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cos.txt</a>.</p> <p>For more information, see <i>Class-of-Service MIB</i>.</p>
Configuration Management MIB (mib-jnx-cfgmgmt)	<p>Provides notification for configuration changes and rescue configuration changes in the form of SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made.</p> <p>A history of the last 32 configuration changes is kept in jnxCmChgEventTable.</p> <p><b>NOTE:</b> On the QFabric system, these conditions apply:</p> <ul style="list-style-type: none"> <li>• All scalar variables under the jnxCmCfgChg table are supported.</li> <li>• Supported scalar OIDs are jnxCmCfgChgLatestIndex, jnxCmCfgChgLatestTime, jnxCmCfgChgLatestDate, jnxCmCfgChgLatestSource, jnxCmCfgChgLatestUser, and jnxCmCfgChgMaxEventEntries.</li> <li>• Scalar variables under the jnxCmRescueChg table are not supported.</li> </ul> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cfgmgmt.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-cfgmgmt.txt</a>.</p> <p>For more information, see <i>Configuration Management MIB</i>.</p>
Fabric Chassis MIB (mib-jnx-fabric-chassis)	<p>Provides hardware information about the QFabric system and its component devices. This MIB is based on the Juniper Networks enterprise-specific Chassis MIB but adds another level of indexing that provides information for QFabric system component devices.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-fabric-chassis.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-fabric-chassis.txt</a>.</p> <p>For more information, see “Fabric Chassis MIB” on page 6518.</p>

**Table 622: Juniper Networks Enterprise-Specific MIBs Supported on QFabric Systems (*continued*)**

MIB	Description
Host Resources MIB (mib-jnx-hostresources)	<p>Extends the hrStorageTable object, providing a measure of the usage of each file system on the switch as a percentage. Previously, the objects in the hrStorageTable measured the usage in allocation units—hrStorageUsed and hrStorageAllocationUnits—only. Using the percentage measurement, you can more easily monitor and apply thresholds on usage.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-hostresources.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-hostresources.txt</a>.</p> <p>For more information, see <i>Host Resources MIB</i>.</p>
Interface MIB (Extensions) (mib-jnx-if-extensions)	<p>Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information in the ifJnxTable and ifChassisTable tables.</p> <p><b>NOTE:</b> On the QFabric system, scalar variables are not supported.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-if-extensions.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-if-extensions.txt</a>.</p> <p>For more information, see <i>Interface MIB</i>.</p>
Power Supply Unit MIB (mib-jnx-power-supply-unit)	<p>Provides support for environmental monitoring of the power supply unit for the Interconnect device of the QFabric system.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-power-supply-unit.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-power-supply-unit.txt</a>.</p> <p>For more information, see <i>Power Supply Unit MIB</i>.</p> <p><b>NOTE:</b> On the QFabric system, scalar variables for the jnxPsuObjects 1 object ID in the jnxPsuScalars table are not supported.</p>
QFabric MIB (jnx-qf-smi)	<p>Explains how the Juniper Networks enterprise-specific QFabric MIBs are structured. Defines the MIB objects that are reported by the QFabric system and the contents of the traps that can be issued by the QFabric system.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-qf-smi.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-qf-smi.txt</a>.</p>
Utility MIB (mib-jnx-util)	<p>Provides you with SNMP MIB container objects of the following types: 32-bit counters, 64-bit counters, signed integers, unsigned integers, and octet strings. You can use these objects to store data that can be retrieved using other SNMP operations.</p> <p>For a downloadable version of this MIB, see <a href="http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-util.txt">http://www.juniper.net/techpubs/en_US/junos13.2/topics/reference/mibs/mib-jnx-util.txt</a>.</p> <p>For more information, see “Utility MIB” on page 6522 and “Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage” on page 6810.</p>

- Related Documentation
- [SNMP MIBs and Traps Reference](#)
  - [Understanding the Implementation of SNMP on page 6513](#)
  - [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
  - [SNMP Traps Support on page 6546](#)

SNMP Traps Support

The QFX Series standalone switches, QFX Series Virtual Chassis, and QFabric systems support standard SNMP traps and Juniper Networks enterprise-specific traps.

For more information, see:

- [SNMP Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis on page 6546](#)
- [SNMP Traps Supported on QFabric Systems on page 6554](#)

SNMP Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

QFX Series standalone switches and QFX Series Virtual Chassis support SNMPv1 and v2 traps. For more information, see:

- [SNMPv1 Traps on page 6546](#)
- [SNMPv2 Traps on page 6550](#)

SNMPv1 Traps

QFX Series standalone switches and QFX Series Virtual Chassis support both standard SNMPv1 traps and Juniper Networks enterprise-specific SNMPv1 traps. See:

- [Table 623 on page 6546](#) for standard SNMPv1 traps.
- [Table 624 on page 6549](#) for enterprise-specific SNMPv1 traps.

The traps are organized first by trap category and then by trap name. The system logging severity levels are listed for those traps that have them. Traps that do not have corresponding system logging severity levels are marked with an en dash (–).

Table 623: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag
Link Notifications						
RFC 1215, <i>Conventions for Defining Traps for</i>	linkDown	1.3.6.1.4.1.2636	2	0	Warning	SNMP_TRAP_LINK_DOWN

**Table 623: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag
<i>Use with the SNMP</i>	linkUp	1.3.6.1.4.1.2636	3	0	Info	SNMP_TRAP_LINK_UP
<b>Remote Operations Notifications</b>						
<i>RFC 2925, Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	pingProbeFailed	1.3.6.1.2.1.80.0	6	1	Info	SNMP_TRAP_PING_PROBE_FAILED
	pingTestFailed	1.3.6.1.2.1.80.0	6	2	Info	SNMP_TRAP_PING_TEST_FAILED
	pingTestCompleted	1.3.6.1.2.1.80.0	6	3	Info	SNMP_TRAP_PING_TEST_COMPLETED
	traceRoutePathChange	1.3.6.1.2.1.81.0	6	1	Info	SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE
	traceRouteTestFailed	1.3.6.1.2.1.81.0	6	2	Info	SNMP_TRAP_TRACE_ROUTE_TEST_FAILED
	traceRouteTestCompleted	1.3.6.1.2.1.81.0	6	3	Info	SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED
<b>RMON Alarms</b>						
<i>RFC 2819a, RMON MIB</i>	fallingAlarm	1.3.6.1.2.1.16	6	2	—	—
	risingAlarm	1.3.6.1.2.1.16	6	1	—	—
<b>Routing Notifications</b>						
<i>BGP 4 MIB</i>	bgpEstablished	1.3.6.1.2.1.15.7	6	1	—	—
	bgpBackwardTransition	1.3.6.1.2.1.15.7	6	2	—	—

**Table 623: Standard SNMP Version 1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag
<i>OSPF TRAP MIB</i>	ospfVirtIfStateChange	1.3.6.1.2.1.14.16.2	6	1	–	–
	ospfNbrStateChange	1.3.6.1.2.1.14.16.2	6	2	–	–
	ospfVirtNbrStateChange	1.3.6.1.2.1.14.16.2	6	3	–	–
	ospfIfConfigError	1.3.6.1.2.1.14.16.2	6	4	–	–
	ospfVirtIfConfigError	1.3.6.1.2.1.14.16.2	6	5	–	–
	ospfIfAuthFailure	1.3.6.1.2.1.14.16.2	6	6	–	–
	ospfVirtIfAuthFailure	1.3.6.1.2.1.14.16.2	6	7	–	–
	ospfIfRxBadPacket	1.3.6.1.2.1.14.16.2	6	8	–	–
	ospfVirtIfRxBadPacket	1.3.6.1.2.1.14.16.2	6	9	–	–
	ospfTxRetransmit	1.3.6.1.2.1.14.16.2	6	10	–	–
	ospfVirtIfTxRetransmit	1.3.6.1.2.1.14.16.2	6	11	–	–
	ospfMaxAgeLsa	1.3.6.1.2.1.14.16.2	6	13	–	–
	ospfIfStateChange	1.3.6.1.2.1.14.16.2	6	16	–	–
<b>Startup Notifications</b>						
RFC 1215, <i>Conventions for Defining Traps for Use with the SNMP</i>	authenticationFailure	1.3.6.1.4.1.2636	4	0	Notice	SNMPD_TRAP_GEN_FAILURE
	coldStart	1.3.6.1.4.1.2636	0	0	Critical	SNMPD_TRAP_COLD_START
	warmStart	1.3.6.1.4.1.2636	1	0	Error	SNMPD_TRAP_WARM_START
<b>VRRP Notifications</b>						
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	vrrpTrapNewMaster	1.3.6.1.2.1.68	6	1	Warning	VRRPD_NEW_MASTER_TRAP
	vrrpTrapAuthFailure	1.3.6.1.2.1.68	6	2	Warning	VRRPD_AUTH_FAILURE_TRAP

**Table 624: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis**

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag
<b>Chassis Notifications (Alarm Conditions)</b>						
<i>Chassis MIB</i> (jnx-chassis. mib)	jnxPowerSupplyFailure	1.3.6.1.4.1.2636.4.1	6	1	Warning	CHASSISD_ SNMP_ TRAP
	jnxFanFailure	1.3.6.1.4.1.2636.4.1	6	2	Critical	CHASSISD_ SNMP_ TRAP
	jnxOverTemperature	1.3.6.1.4.1.2636.4.1	6	3	Alert	CHASSISD_ SNMP_ TRAP
	jnxFruRemoval	1.3.6.1.4.1.2636.4.1	6	5	Notice	CHASSISD_ SNMP_ TRAP
	jnxFruInsertion	1.3.6.1.4.1.2636.4.1	6	6	Notice	CHASSISD_ SNMP_ TRAP
	jnxFruPowerOff	1.3.6.1.4.1.2636.4.1	6	7	Notice	CHASSISD_ SNMP_ TRAP
	jnxFruPowerOn	1.3.6.1.4.1.2636.4.1	6	8	Notice	CHASSISD_ SNMP_ TRAP
	jnxFruFailed	1.3.6.1.4.1.2636.4.1	6	9	Warning	CHASSISD_ SNMP_ TRAP
	jnxFruOffline	1.3.6.1.4.1.2636.4.1	6	10	Notice	CHASSISD_ SNMP_ TRAP
	jnxFruOnline	1.3.6.1.4.1.2636.4.1	6	11	Notice	CHASSISD_ SNMP_ TRAP
	jnxFruCheck	1.3.6.1.4.1.2636.4.1	6	12	Warning	CHASSISD_ SNMP_ TRAP
	jnxPowerSupplyOk	1.3.6.1.4.1.2636.4.2	6	1	Critical	CHASSISD_ SNMP_ TRAP
	jnxFanOK	1.3.6.1.4.1.2636.4.2	6	2	Critical	CHASSISD_ SNMP_ TRAP
	jnxTemperatureOK	1.3.6.1.4.1.2636.4.2	6	3	Alert	CHASSISD_ SNMP_ TRAP
<b>Configuration Notifications</b>						

**Table 624: Enterprise-Specific SNMPv1 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag
<i>Configuration Management MIB</i> (jnx- configmgmt. mib)	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5	6	1	–	–
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5	6	2	–	–
<b>Remote Operations</b>						
<i>Ping MIB</i> (jnx-ping.mib)	jnxPingRttThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	1	–	–
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	2	–	–
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	3	–	–
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	4	–	–
	jnxPingEgressStdDev ThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	5	–	–
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	6	–	–
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	7	–	–
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	8	–	–
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	9	–	–
<b>RMON Alarms</b>						
<i>RMON MIB</i> (jnx-rmon. mib)	jnxRmonAlarmGetFailure	1.3.6.1.4.1.2636.4.3	6	1	–	–
	jnxRmonGetOk	1.3.6.1.4.1.2636.4.3	6	2	–	–

**SNMPv2 Traps**

- [Table 625 on page 6551](#) lists the standard SNMP traps
- [Table 626 on page 6553](#) lists the Juniper Networks enterprise-specific traps



Table 625: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
<b>Link Notifications</b>				
RFC 2863, <i>The Interfaces Group MIB</i>	linkDown	1.3.6.1.6.3.1.1.5.3	Warning	SNMP_TRAP_LINK_DOWN
	linkUp	1.3.6.1.6.3.1.1.5.4	Info	SNMP_TRAP_LINK_UP
<b>Remote Operations Notifications</b>				
RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	pingProbeFailed	1.3.6.1.2.1.80.0.1	Info	SNMP_TRAP_PING_PROBE_FAILED
	pingTestFailed	1.3.6.1.2.1.80.0.2	Info	SNMP_TRAP_PING_TEST_FAILED
	pingTestCompleted	1.3.6.1.2.1.80.0.3	Info	SNMP_TRAP_PING_TEST_COMPLETED
	traceRoutePathChange	1.3.6.1.2.1.81.0.1	Info	SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE
	traceRouteTestFailed	1.3.6.1.2.1.81.0.2	Info	SNMP_TRAP_TRACE_ROUTE_TEST_FAILED
	traceRouteTestCompleted	1.3.6.1.2.1.81.0.3	Info	SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED
<b>RMON Alarms</b>				
RFC 2819a, <i>RMON MIB</i>	fallingAlarm	1.3.6.1.2.1.16.0.1	–	–
	risingAlarm	1.3.6.1.2.1.16.0.2	–	–
<b>Routing Notifications</b>				
<i>BGP 4 MIB</i>	bgpEstablished	1.3.6.1.2.1.15.7.1	–	–
	bgpBackwardTransition	1.3.6.1.2.1.15.7.2	–	–

Table 625: Standard SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
<i>OSPF Trap MIB</i>	ospfVirtIfStateChange	1.3.6.1.2.1.14.16.2.1	–	–
	ospfNbrStateChange	1.3.6.1.2.1.14.16.2.2	–	–
	ospfVirtNbrStateChange	1.3.6.1.2.1.14.16.2.3	–	–
	ospfIfConfigError	1.3.6.1.2.1.14.16.2.4	–	–
	ospfVirtIfConfigError	1.3.6.1.2.1.14.16.2.5	–	–
	ospfIfAuthFailure	1.3.6.1.2.1.14.16.2.6	–	–
	ospfVirtIfAuthFailure	1.3.6.1.2.1.14.16.2.7	–	–
	ospfIfRxBadPacket	1.3.6.1.2.1.14.16.2.8	–	–
	ospfVirtIfRxBadPacket	1.3.6.1.2.1.14.16.2.9	–	–
	ospfTxRetransmit	1.3.6.1.2.1.14.16.2.10	–	–
	ospfVirtIfTxRetransmit	1.3.6.1.2.1.14.16.2.11	–	–
	ospfMaxAgeLsa	1.3.6.1.2.1.14.16.2.13	–	–
	ospfIfStateChange	1.3.6.1.2.1.14.16.2.16	–	–
<b>Startup Notifications</b>				
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	coldStart	1.3.6.1.6.3.1.1.5.1	Critical	SNMPD_TRAP_COLD_START
	warmStart	1.3.6.1.6.3.1.1.5.2	Error	SNMPD_TRAP_WARM_START
	authenticationFailure	1.3.6.1.6.3.1.1.5.5	Notice	SNMPD_TRAP_GEN_FAILURE
<b>VRRP Notifications</b>				
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	vrrpTrapNewMaster	1.3.6.1.2.1.68.0.1	Warning	VRRPD_NEWMASTER_TRAP
	vrrpTrapAuthFailure	1.3.6.1.2.1.68.0.2	Warning	VRRPD_AUTH_FAILURE_TRAP

**Table 626: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis**

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<b>Chassis (Alarm Conditions) Notifications</b>				
<i>Chassis MIB</i> (mib-jnx-chassis)	jnxPowerSupplyFailure	1.3.6.1.4.1.2636.4.1.1	Alert	CHASSISD_SNMP_TRAP
	jnxFanFailure	1.3.6.1.4.1.2636.4.1.2	Critical	CHASSISD_SNMP_TRAP
	jnxOverTemperature	1.3.6.1.4.1.2636.4.1.3	Critical	CHASSISD_SNMP_TRAP
	jnxFruRemoval	1.3.6.1.4.1.2636.4.1.5	Notice	CHASSISD_SNMP_TRAP
	jnxFruInsertion	1.3.6.1.4.1.2636.4.1.6	Notice	CHASSISD_SNMP_TRAP
	jnxFruPowerOff	1.3.6.1.4.1.2636.4.1.7	Notice	CHASSISD_SNMP_TRAP
	jnxFruPowerOn	1.3.6.1.4.1.2636.4.1.8	Notice	CHASSISD_SNMP_TRAP
	jnxFruFailed	1.3.6.1.4.1.2636.4.1.9	Warning	CHASSISD_SNMP_TRAP
	jnxFruOffline	1.3.6.1.4.1.2636.4.1.10	Notice	CHASSISD_SNMP_TRAP
	jnxFruOnline	1.3.6.1.4.1.2636.4.1.11	Notice	CHASSISD_SNMP_TRAP
	jnxFruCheck	1.3.6.1.4.1.2636.4.1.12	Notice	CHASSISD_SNMP_TRAP
	jnxPowerSupplyOK	1.3.6.1.4.1.2636.4.2.1	Critical	CHASSISD_SNMP_TRAP
	jnxFanOK	1.3.6.1.4.1.2636.4.2.2	Critical	CHASSISD_SNMP_TRAP
	jnxTemperatureOK	1.3.6.1.4.1.2636.4.2.3	Alert	CHASSISD_SNMP_TRAP
<b>Configuration Notifications</b>				

**Table 626: Enterprise-Specific SNMPv2 Traps Supported on QFX Series Standalone Switches and QFX Series Virtual Chassis (*continued*)**

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<i>Configuration Management MIB</i> (mib-jnx-cfgmgmt)	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5.0.1	–	–
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5.0.2	–	–
<b>Remote Operations Notifications</b>				
<i>Ping MIB</i> (mib-jnx-ping)	jnxPingRttThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.1	–	–
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.2	–	–
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.3	–	–
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.4	–	–
	jnxPingEgressStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.5	–	–
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.6	–	–
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.7	–	–
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.8	–	–
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.9	–	–
<b>RMON Alarms</b>				
<i>RMON MIB</i> (mib-jnx-rmon)	jnxRmonAlarmGetFailure	1.3.6.1.4.1.2636.4.3.0.1	–	–
	jnxRmonGetOk	1.3.6.1.4.1.2636.4.3.0.2	–	–

**SNMP Traps Supported on QFabric Systems**

QFabric systems support standard SNMPv2 traps and Juniper Networks enterprise-specific SNMPv2 traps.



**NOTE:** QFabric systems do not support SNMPv1 traps.

For more information, see:

- [Table 627 on page 6555](#) for standard SNMPv2 traps
- [Table 628 on page 6556](#) for Juniper Networks enterprise-specific SNMPv2 traps

**Table 627: Standard SNMPv2 Traps Supported on QFabric Systems**

Defined in	Trap Name	SNMP Trap OID	System Logging Severity Level	Syslog Tag
<b>Link Notifications</b>				
RFC 2863, <i>The Interfaces Group MIB</i>	linkDown	1.3.6.1.6.3.1.1.5.3	Warning	SNMP_TRAP_LINK_DOWN
	linkUp	1.3.6.1.6.3.1.1.5.4	Info	SNMP_TRAP_LINK_UP
<b>Startup Notifications</b>				
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	coldStart	1.3.6.1.6.3.1.1.5.1	Critical	SNMPD_TRAP_COLD_START
	warmStart	1.3.6.1.6.3.1.1.5.2	Error	SNMPD_TRAP_WARM_START
	authenticationFailure	1.3.6.1.6.3.1.1.5.5	Notice	SNMPD_TRAP_GEN_FAILURE

Table 628: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<i>Fabric Chassis MIB</i> (mib-jnx-fabric-chassis)	<b>Fabric Chassis (Alarm Conditions) Notifications</b>			
	jnxFabricPowerSupplyFailure	1.3.6.1.4.1.2636.4.19.1	Warning	–
	jnxFabricFanFailure	1.3.6.1.4.1.2636.4.19.2	Critical	–
	jnxFabricOverTemperature	1.3.6.1.4.1.2636.4.19.3	Alert	–
	jnxFabricRedundancySwitchover	1.3.6.1.4.1.2636.4.19.4	Notice	–
	jnxFabricFruRemoval	1.3.6.1.4.1.2636.4.19.5	Notice	–
	jnxFabricFruInsertion	1.3.6.1.4.1.2636.4.19.6	Notice	–
	jnxFabricFruPowerOff	1.3.6.1.4.1.2636.4.19.7	Notice	–
	jnxFabricFruPowerOn	1.3.6.1.4.1.2636.4.19.8	Notice	–
	jnxFabricFruFailed	1.3.6.1.4.1.2636.4.19.9	Warning	–
	jnxFabricFruOffline	1.3.6.1.4.1.2636.4.19.10	Notice	–
	jnxFabricFruOnline	1.3.6.1.4.1.2636.4.19.11	Notice	–
	jnxFabricFruCheck	1.3.6.1.4.1.2636.4.19.12	Warning	–
	jnxFabricFEBSwitchover	1.3.6.1.4.1.2636.4.19.13	Warning	–
	jnxFabricHardDiskFailed	1.3.6.1.4.1.2636.4.19.14	Warning	–
	jnxFabricHardDiskMissing	1.3.6.1.4.1.2636.4.19.15	Warning	–
	jnxFabricBootFromBackup	1.3.6.1.4.1.2636.4.19.16	Warning	–
	<b>Fabric Chassis (Alarm Cleared Conditions) Notifications</b>			
	jnxFabricPowerSupplyOK	1.3.6.1.4.1.2636.4.20.1	Critical	–
	jnxFabricFanOK	1.3.6.1.4.1.2636.4.20.2	Critical	–
	jnxFabricTemperatureOK	1.3.6.1.4.1.2636.4.20.3	Alert	–
	jnxFabricFruOK	1.3.6.1.4.1.2636.4.20.4	–	–

Table 628: Enterprise-Specific SNMPv2 Traps Supported on QFabric Systems (*continued*)

Source MIB	Trap Name	SNMP Trap OID	System Logging Severity Level	System Log Tag
<i>QFabric MIB</i> (mib-jnx-qf-smi)	<b>QFabric MIB Notifications</b>			
	jnxQFabricDownloadIssued	1.3.6.1.4.1.2636.3.42.1.0.1	–	–
	jnxQFabricDownloadFailed	1.3.6.1.4.1.2636.3.42.1.0.2	–	–
	jnxQFabricDownloadSucceeded	1.3.6.1.4.1.2636.3.42.1.0.3	–	–
	jnxQFabricUpgradeIssued	1.3.6.1.4.1.2636.3.42.1.0.4	–	–
	jnxQFabricUpgradeFailed	1.3.6.1.4.1.2636.3.42.1.0.5	–	–
	jnxQFabricUpgradeSucceeded	1.3.6.1.4.1.2636.3.42.1.0.6	–	–
<b>Configuration Notifications</b>				
<i>Configuration Management MIB</i> (mib-jnx-cfgmgmt)	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5.0.1	–	–
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5.0.2	–	–
<b>Remote Operations Notifications</b>				
<i>Ping MIB</i> (mib-jnx-ping)	jnxPingRttThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.1	–	–
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.2	–	–
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.3	–	–
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.4	–	–
	jnxPingEgressStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.5	–	–
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.6	–	–
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.7	–	–
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.8	–	–
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.9	–	–

- Related Documentation**
- [SNMP MIBs and Traps Reference](#)
  - [Understanding the Implementation of SNMP on page 6513](#)
  - [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
  - [SNMP MIBs Support on page 6530](#)

## MIB Objects for the QFX Series

This topic lists the Juniper Networks enterprise-specific SNMP Chassis MIB definition objects for the QFX Series:

- [QFX Series Standalone Switches on page 6558](#)
- [QFabric Systems on page 6558](#)
- [QFabric System QFX3100 Director Device on page 6559](#)
- [QFabric System QFX3008-I Interconnect Device on page 6559](#)
- [QFabric System QFX3600-I Interconnect Device on page 6559](#)
- [QFabric System Node Devices on page 6560](#)

### QFX Series Standalone Switches

jnxProductLineQFXSwitch	OBJECT IDENTIFIER ::= { jnxProductLine 82 }
jnxProductNameQFXSwitch	OBJECT IDENTIFIER ::= { jnxProductName 82 }
jnxProductModelQFXSwitch	OBJECT IDENTIFIER ::= { jnxProductModel 82 }
jnxProductVariationQFXSwitch	OBJECT IDENTIFIER ::= { jnxProductVariation 82 }
jnxProductQFX3500s	OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 1 }
jnxProductQFX360016QS	OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 2 }
jnxProductQFX350048T4QS	OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 3 }
jnxProductQFX510024Q	OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 4 }
jnxProductQFX510048S6Q	OBJECT IDENTIFIER ::= { jnxProductVariationQFXSwitch 5 }
jnxChassisQFXSwitch	OBJECT IDENTIFIER ::= { jnxChassis 82 }
jnxSlotQFXSwitch	OBJECT IDENTIFIER ::= { jnxSlot 82 }
jnxQFXSwitchSlotFPC	OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 1 }
jnxQFXSwitchSlotHM	OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 2 }
jnxQFXSwitchSlotPower	OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 3 }
jnxQFXSwitchSlotFan	OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 4 }
jnxQFXSwitchSlotFPB	OBJECT IDENTIFIER ::= { jnxSlotQFXSwitch 5 }
jnxMediaCardSpaceQFXSwitch	OBJECT IDENTIFIER ::= { jnxMediaCardSpace 82 }
jnxQFXSwitchMediaCardSpacePIC	OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXSwitch 1 }

### QFabric Systems

jnxProductLineQFX3000	OBJECT IDENTIFIER ::= { jnxProductLine 84 }
jnxProductNameQFX3000	OBJECT IDENTIFIER ::= { jnxProductName 84 }
jnxProductModelQFX3000	OBJECT IDENTIFIER ::= { jnxProductModel 84 }
jnxProductVariationQFX3000	OBJECT IDENTIFIER ::= { jnxProductVariation 84 }
jnxProductQFX3000-G	OBJECT IDENTIFIER ::= { jnxProductVariationQFX3000 1 }
jnxProductQFX3000-M	OBJECT IDENTIFIER ::= { jnxProductVariationQFX3000 2 }
jnxChassisQFX3000	OBJECT IDENTIFIER ::= { jnxChassis 84 }



### QFabric System QFX3100 Director Device

```
jnxProductLineQFX3100 OBJECT IDENTIFIER ::= { jnxProductLine      100 }
jnxProductNameQFX3100 OBJECT IDENTIFIER ::= { jnxProductName      100 }
jnxProductModelQFX3100 OBJECT IDENTIFIER ::= { jnxProductModel    100 }
jnxProductVariationQFX3100 OBJECT IDENTIFIER ::= { jnxProductVariation 100 }
jnxChassisQFX3100      OBJECT IDENTIFIER ::= { jnxChassis         100 }

jnxSlotQFX3100          OBJECT IDENTIFIER ::= { jnxSlot           100 }
jnxQFX3100SlotCPU       OBJECT IDENTIFIER ::= { jnxSlotQFX3100    1 }
jnxQFX3100SlotMemory    OBJECT IDENTIFIER ::= { jnxSlotQFX3100    2 }
jnxQFX3100SlotPower     OBJECT IDENTIFIER ::= { jnxSlotQFX3100    3 }
jnxQFX3100SlotFan       OBJECT IDENTIFIER ::= { jnxSlotQFX3100    4 }
jnxQFX3100SlotHardDisk  OBJECT IDENTIFIER ::= { jnxSlotQFX3100    5 }
jnxQFX3100SlotNIC       OBJECT IDENTIFIER ::= { jnxSlotQFX3100    6 }
```

### QFabric System QFX3008-I Interconnect Device

```
jnxProductLineQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductLine      60 }
jnxProductNameQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductName      60 }
jnxProductModelQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductModel    60 }
jnxProductVariationQFXInterconnect OBJECT IDENTIFIER ::= { jnxProductVariation 60 }
jnxProductQFX3008          OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 1 }
jnxProductQFXC083008       OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 2 }
jnxProductQFX3008I         OBJECT IDENTIFIER ::= { jnxProductVariationQFXInterconnect 3 }

jnxChassisQFXInterconnect  OBJECT IDENTIFIER ::= { jnxChassis         60 }

jnxSlotQFXInterconnect     OBJECT IDENTIFIER ::= { jnxSlot           60 }
jnxQFXInterconnectSlotFPC  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect    1 }
jnxQFXInterconnectSlotHBM OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect    2 }
jnxQFXInterconnectSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect    3 }
jnxQFXInterconnectSlotFan  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect    4 }
jnxQFXInterconnectSlotCBD  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect    5 }
jnxQFXInterconnectSlotFPB  OBJECT IDENTIFIER ::= { jnxSlotQFXInterconnect    6 }

jnxMediaCardSpaceQFXInterconnect OBJECT IDENTIFIER ::= { jnxMediaCardSpace 60 }
jnxQFXInterconnectMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXInterconnect 1 }

jnxMidplaneQFXInterconnect OBJECT IDENTIFIER ::= { jnxBackplane       60 }
```

### QFabric System QFX3600-I Interconnect Device

```
jnxProductLineQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductLine      91 }
jnxProductNameQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductName      91 }
jnxProductModelQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductModel    91 }
jnxProductVariationQFXMInterconnect OBJECT IDENTIFIER ::= { jnxProductVariation 91 }
jnxProductQFX3600I         OBJECT IDENTIFIER ::= { jnxProductVariationQFXMInterconnect 1 }

jnxChassisQFXMInterconnect  OBJECT IDENTIFIER ::= { jnxChassis         91 }

jnxSlotQFXMInterconnect     OBJECT IDENTIFIER ::= { jnxSlot           91 }
jnxQFXMInterconnectSlotFPC  OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect    1 }
jnxQFXMInterconnectSlotHBM OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect    2 }
jnxQFXMInterconnectSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect    3 }
jnxQFXMInterconnectSlotFan  OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect    4 }
jnxQFXMInterconnectSlotFPB  OBJECT IDENTIFIER ::= { jnxSlotQFXMInterconnect    5 }
```

```
jnxMediaCardSpaceQFXMInterconnect OBJECT IDENTIFIER ::= { jnxMediaCardSpace 91 }
jnxQFXMInterconnectMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXMInterconnect 1 }
```

### QFabric System Node Devices

```
jnxProductLineQFXNode OBJECT IDENTIFIER ::= { jnxProductLine 61 }
jnxProductNameQFXNode OBJECT IDENTIFIER ::= { jnxProductName 61 }
jnxProductModelQFXNode OBJECT IDENTIFIER ::= { jnxProductModel 61 }
jnxProductVariationQFXNode OBJECT IDENTIFIER ::= { jnxProductVariation 61 }
jnxProductQFX3500 OBJECT IDENTIFIER ::= { jnxProductVariationQFXNode 1 }
jnxProductQFX360016Q OBJECT IDENTIFIER ::= { jnxProductVariationQFXNode 3 }

jnxChassisQFXNode OBJECT IDENTIFIER ::= { jnxChassis 61 }

jnxSlotQFXNode OBJECT IDENTIFIER ::= { jnxSlot 61 }
jnxQFXNodeSlotFPC OBJECT IDENTIFIER ::= { jnxSlotQFXNode 1 }
jnxQFXNodeSlotHM OBJECT IDENTIFIER ::= { jnxSlotQFXNode 2 }
jnxQFXNodeSlotPower OBJECT IDENTIFIER ::= { jnxSlotQFXNode 3 }
jnxQFXNodeSlotFan OBJECT IDENTIFIER ::= { jnxSlotQFXNode 4 }
jnxQFXNodeSlotFPB OBJECT IDENTIFIER ::= { jnxSlotQFXNode 5 }

jnxMediaCardSpaceQFXNode OBJECT IDENTIFIER ::= { jnxMediaCardSpace 61 }
jnxQFXNodeMediaCardSpacePIC OBJECT IDENTIFIER ::= { jnxMediaCardSpaceQFXNode 1 }
```

- Related Documentation**
- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
  - [Fabric Chassis MIB on page 6518](#)

## System Logging

- [Overview of Junos OS System Log Messages on page 6560](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Understanding the Implementation of System Log Messages on the QFabric System on page 6562](#)

### Overview of Junos OS System Log Messages

The Junos OS generates system log messages (also called *syslog messages*) to record events that occur on the switch, including the following:

- Routine operations, such as a user login into the configuration database.
- Failure and error conditions, such as failure to access a configuration file.
- Emergency or critical conditions, such as power-down of the switch due to excessive temperature.

Each system log message identifies the Junos OS process that generated the message and briefly describes the operation or error that occurred. For detailed information about specific system log messages, see the *Junos OS System Log Messages Reference*.

- Related Documentation**
- [Junos OS System Log Configuration Statements on page 6616](#)
  - [Junos OS Minimum System Logging Configuration on page 6616](#)

## Overview of Single-Chassis System Logging Configuration

The Junos OS system logging utility on the QFX Series is similar to the UNIX **syslogd** utility. This topic describes how to configure system logging for a single-chassis system that runs the Junos OS.

Each system log message belongs to a *facility*, which groups together related messages. Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects router functions. You always specify the facility and severity of the messages to include in the log. For more information, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6631](#).

You direct messages to one or more destinations by including the appropriate statement at the **[edit system syslog]** hierarchy level:

- To a named file in a local file system, by including the **file** statement. See [“Directing System Log Messages to a Log File” on page 6618](#).
- To the terminal session of one or more specific users (or all users) when they are logged in to the switch, by including the **user** statement. See [“Directing System Log Messages to a User Terminal” on page 6620](#).
- To the switch console, by including the **console** statement. See [“Directing System Log Messages to the Console” on page 6620](#).
- To a remote machine that is running the **syslogd** utility, by including the **host** statement. See [“Directing System Log Messages to a Remote Machine” on page 6619](#).

By default, messages are logged in a standard format, which is based on a UNIX system log format; for detailed information about message formatting, see the *Junos OS System Log Messages Reference*. You can alter the content and format of logged messages in the following ways:

- You can log messages to a file in structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from the message. For more information, see [“Logging Messages in Structured-Data Format” on page 6624](#).
- A message’s facility and severity level are together referred to as its *priority*. By default, the standard Junos OS format for messages does not include priority information (structured-data format includes a priority code by default). To include priority information in standard-format messages directed to a file or a remote destination, include the **explicit-priority** statement. For more information, see [“Including Priority Information in System Log Messages” on page 6622](#).
- By default, the standard Junos OS format for messages specifies the month, date, hour, minute, and second when the message was logged. You can modify the timestamp on standard-format system log messages to include the year, the millisecond, or both. (Structured-data format specifies the year and millisecond by default.) For more information, see [“Including the Year or Millisecond in Timestamps” on page 6623](#).

- When directing messages to a remote machine, you can specify the IP address that is reported in messages as their source. You can also configure features that make it easier to separate messages generated by Junos OS or messages generated on particular switches. For more information, see [“Directing System Log Messages to a Remote Machine” on page 6619](#).
- The predefined facilities group together related messages, but you can also use regular expressions to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination. For more information, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6637](#).



**NOTE:** During a commit check, warnings about the `traceoptions` configuration (for example, mismatch in trace file sizes or number of trace files) are not displayed on the console. However, these warnings are logged in the system log messages when the new configuration is committed.

#### Related Documentation

- [Examples: Configuring System Logging on page 6565](#)
- [Specifying the Facility and Severity of Messages to Include in the Log on page 6631](#)
- [Junos OS System Logging Facilities and Message Severity Levels on page 6633](#)
- [Directing System Log Messages to a Log File on page 6618](#)
- [Directing System Log Messages to a Remote Machine on page 6619](#)
- [Directing System Log Messages to a User Terminal on page 6620](#)
- [Directing System Log Messages to the Console on page 6620](#)

## Understanding the Implementation of System Log Messages on the QFabric System

This topic provides an overview of system log (syslog) messages as implemented on the QFabric system.

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

You configure system log messages by using the **host** and **file** statements at the **[edit system syslog]** hierarchy level. Use the **show log filename** operational mode command to view the messages.



**NOTE:** On the QFabric system, a syslog file named `messages` with a size of 100 MB is configured by default. If you do not configure a filename, you can use the default filename `messages` with the `show log filename` command.

All messages with a severity level of notice or higher are logged. Messages with a facility level of `interactive-commands` on Node devices are not logged.

The QFabric system supports the following system log message features:

- The `file filename` and `host hostname` statements at the `[edit system syslog]` hierarchy level are supported. Other statements at that hierarchy level are not supported.
- You can specify the maximum amount of data that is displayed when you issue the `show log filename` command by configuring the `file filename archive maximum-file-size` statement.
- You can specify that one or more system log message servers receive messages, which are sent to each server that is configured.
- If you configured an alias for a device or interface, the alias is displayed in the message for the device or interface.
- The level of detail that is included in a message depends on the facility and severity levels that are configured. Messages include the highest level of detail available for the configured facility and severity levels.
- The unit of time is measured and displayed in seconds, and not milliseconds. If you attempt to configure the `time-format` option in milliseconds, the log output displays `000`.

Starting in Junos OS Release 13.1, the QFabric system supports these additional syslog features:

- You can filter the output of the `show log filename` operational mode command by device type and device ID or device alias when you specify the `device-type (device-id | device-alias)` optional parameters. Device types include `director-device`, `infrastructure-device`, `interconnect-device`, and `node-device`.
- You can specify the syslog structured data output format when you configure the `structured-data` statement at the `[edit system syslog file filename]` and `[edit system syslog host hostname]` hierarchy levels.



**NOTE:** Information displayed in the structured data output for system logs originating from the Director software may not be complete.

- You can filter the types of logs that the Director group collects from a component device when you configure the `filter all facility severity` or `filter all match "regular-expression"` statements at the `[edit system syslog]` hierarchy level.

Unsupported syslog features include:

- File access to syslog messages
- Monitoring of syslog messages

**Related  
Documentation**

- [Example: Configuring System Log Messages on page 6568](#)
- [syslog \(QFabric System\) on page 6796](#)

## CHAPTER 77

# Configuration

- [Configuration Examples on page 6565](#)
- [Configuration Tasks for Network Management on page 6583](#)
- [Configuration Tasks for Automation on page 6587](#)
- [Configuration Tasks for Network Analytics on page 6590](#)
- [Configuration Tasks for sFlow Technology on page 6596](#)
- [Configuration Tasks for SNMP on page 6597](#)
- [Configuration Tasks for System Log Messages on page 6615](#)
- [Configuration Statements for Network Management on page 6639](#)
- [Configuration Statements for Automation on page 6647](#)
- [Configuration Statements for Network Analytics on page 6666](#)
- [Configuration Statements for sFlow Technology on page 6681](#)
- [Configuration Statements for SNMP on page 6689](#)
- [Configuration Statements for System Log Messages on page 6778](#)

### Configuration Examples

---

- [Examples: Configuring System Logging on page 6565](#)
- [Examples: Assigning an Alternative Facility on page 6567](#)
- [Example: Configuring System Log Messages on page 6568](#)
- [Example: Monitoring Network Traffic Using sFlow Technology on page 6571](#)
- [Example: Configuring SNMP on page 6575](#)
- [Example: Configuring Network Analytics on page 6577](#)

### Examples: Configuring System Logging

The system log provides an excellent way of tracking all management activity on the switch by recording events such as user authentication, access authorization, and command execution. Logged command executions include commands entered by users at the CLI prompt or by client applications such as the Junos XML protocol or NETCONF XML client. Because system log files contain information about commands executed on the switch and the user who executed the commands, checking system log files for failed authentication events can help identify attempts to hack in to the switch. You can also

analyze network activity by correlating executed commands with events and changes that occurred on the network at a particular time.

System log files are stored locally on the switch in the default `/var/log` directory.

The following example shows how to configure system log messages to record all commands entered by users and all authentication or authorization attempts. Logged commands include those entered by users at the CLI prompt and by client applications. Authentication and authorization attempts include events that are saved in the file named `cli-commands` and those that are sent to the terminal of a user who is logged in.

```
[edit system]
syslog {
  file cli-commands {
    interactive-commands info;
    authorization info;
  }
  user * {
    interactive-commands info;
    authorization info;
  }
}
```

The following example shows how to log all alarms state changes to the file `/var/log/alarms`:

```
[edit system]
syslog {
  file alarms {
    kernel warning;
  }
}
```

The following example shows how to configure the handling of messages of various types, as described in the comments. Information is logged to two files, to the terminal of user alex, to a remote machine, and to the console:

```
[edit system]
syslog {
  /* write all security-related messages to file /var/log/security */
  file security {
    authorization info;
    interactive-commands info;
  }
  /* write messages about potential problems to file /var/log/messages: */
  /* messages from "authorization" facility at level "notice"
  and above, */
  /* messages from all other facilities at level "warning" and above */
  file messages {
    authorization notice;
    any warning;
  }
  /* write all messages at level "critical" and above to terminal of user
  "alex" if */
  /* that user is logged in */
  user alex {
```



```

        any critical;
    }
    /* write all messages from the &ldquo;daemon&rdquo; facility at level &ldquo;info&rdquo;
       and above, and */
    /* messages from all other facilities at level &ldquo;warning&rdquo; and above, to the
       */
    /* machine monitor.mycompany.com */
    host monitor.mycompany.com {
        daemon info;
        any warning;
    }
    /* write all messages at level &ldquo;error&rdquo; and above to the system console */
    console {
        any error;
    }
}

```

The following example shows how to configure the handling of messages generated when users issue Junos OS CLI commands, by specifying the interactive-commands facility at the info, notice, and warning severity levels:

```

[edit system]
file user-actions {
    interactive-commands info;
}
user philip {
    interactive-commands notice;
}
console {
    interactive-commands warning;
}
}

```

The following list describes the security levels used in the example:

- **info**—Logs a message when users issue any command at the CLI operational or configuration mode prompt. The example writes the messages to the file `/var/log/user-actions`.
- **notice**—Logs a message when users issue the configuration mode command **commit**. The example writes the messages to the terminal of user philip.
- **warning**—Logs a message when users issue a command that restarts a software process. The example writes the messages to the console.

#### Related Documentation

- [Overview of Single-Chassis System Logging Configuration on page 6561](#)

### Examples: Assigning an Alternative Facility

This topic contains examples of configuring system log messages to use an alternative facility for logging.

The following example shows how to log all messages generated on the switch at the **error** level or higher to the **local0** facility on the remote host called **monitor.mycompany.com**:

```
[edit system syslog]
host monitor.mycompany.com {
  any error;
  facility-override local0;
}
```

The following example contains two sets of statements that show how to configure switches located in California and in New York to send messages to a single remote host called **central-logger.mycompany.com**. The messages from California are assigned to alternative facility **local0** and the messages from New York are assigned to alternative facility **local2**.

- The following statements configure the California switch to aggregate messages in the **local0** facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local0;
}
```

- The following statements configure the New York switch to aggregate messages in the **local2** facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local2;
}
```

On the remote host named **central-logger** you can subsequently configure the system logging utility to write messages from the **local0** facility to one file (for example, **california-config**) and the messages from the **local2** facility to another file (for example, **new-york-config**).

**Related  
Documentation**

- [Alternate Facilities for System Log Messages Directed to a Remote Destination on page 6635](#)

## Example: Configuring System Log Messages

The QFabric system monitors events that occur on its component devices and distributes system log messages about those events to all external system log message servers (hosts) that are configured. Component devices may include Node devices, Interconnect devices, Director devices, and the Virtual Chassis. Messages are stored for viewing only in the QFabric system database. To view the messages, issue the **show log** command.

This example describes how to configure system log messages on the QFabric system.

- [Requirements on page 6569](#)
- [Overview on page 6569](#)
- [Configuration on page 6569](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- QFabric system
- External servers that can be configured as system log message hosts

### Overview

Component devices that generate system log message events may include Node devices, Interconnect devices, Director devices, and the control plane switches. The following configuration example includes these components in the QFabric system:

- Director software running on the Director group
- Control plane switches
- Interconnect device
- Multiple Node devices

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set system syslog host 10.1.1.12 any error
set system syslog file qflogs
set system syslog file qflogs structured-data brief
set system syslog file qflogs archive size 1g
```

#### Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure system messages from the QFabric Director device:

1. Specify a host, any facility, and the **error** severity level.

```
[edit system syslog]
user@switch# set host 10.1.1.12 any error
```



**NOTE:** You can configure more than one system log message server (host). The QFabric system sends the messages to each server configured.

2. (Optional) Specify a filename to capture log messages.



**NOTE:** On the QFabric system, a syslog file named `messages` is configured implicitly with facility and severity levels of any any and a file size of 100 MBs. Therefore, you cannot specify the filename `messages` in your configuration, and automatic command completion does not work for that filename.

```
[edit system syslog]
user@switch# set file qflogs structured-data brief
user@switch# set file qflogs
```

3. (Optional) Configure the maximum size of your system log message archive file. This example specifies an archive size of 1 GB.

```
[edit system syslog]
user@switch# set file qflogs archive size 1g
```

**Results** From configuration mode, confirm your configuration by entering the **show system** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show system
syslog {
  file qflogs {
  }
  host 10.1.1.12 {
    any error;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

**Related Documentation**

- [Understanding the Implementation of System Log Messages on the QFabric System on page 6562](#)
- [syslog \(QFabric System\) on page 6796](#)
- [show log on page 948](#)

## Example: Monitoring Network Traffic Using sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

This example describes how to configure and use sFlow monitoring on a QFX3500 switch in standalone mode.

- [Requirements on page 6571](#)
- [Overview on page 6571](#)
- [Configuration on page 6572](#)
- [Verification on page 6573](#)

### Requirements

---

This example uses the following hardware and software components:

- Junos OS Release 11.3 or later
- One QFX3500 switch


### Overview

---

An sFlow monitoring system consists of an sFlow agent embedded in the device and a centralized collector on the network. The two main activities of the sFlow agent are random sampling and statistics gathering. The sFlow agent combines interface counters and flow samples and sends them to the IP address and UDP destination port of the sFlow collector in UDP datagrams.

[Figure 226 on page 6572](#) depicts the basic elements of an sFlow system.

Figure 226: sFlow Technology Monitoring System



### Configuration

#### CLI Quick Configuration

To quickly configure sFlow technology, copy the following commands and paste them into the terminal window of the switch:

```
[edit protocols sflow]
set collector 10.204.32.46 udp-port 5600
set interfaces xe-0/0/1.0
set polling-interval 20
set sample-rate 1000
```

#### Step-by-Step Procedure

To configure sFlow features using the CLI:

1. Configure the IP address and UDP port of at least one collector:

```
[edit protocols sflow]
user@switch# set collector 10.204.32.46 udp-port 5600
```

The default UDP port assigned is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces xe-0/0/1.0
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a LAG interface (for example, ae0), but you can enable sFlow technology on the member interfaces of the LAG (for example, xe-0/0/1).

3. Specify how often (in seconds) the sFlow agent polls all interfaces at the global level:

```
[edit protocols sflow]
user@switch# set polling-interval 20
```



**NOTE:** Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets must be sampled at the global level. The following example sets a sample rate of 1 in 1000 packets:

```
[edit protocols sflow]
user@switch# set sample-rate 1000
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show protocols
sflow {
  collector 10.204.32.46 {
    udp-port 5600;
  }
  interfaces xe-0/0/1.0 {
    polling-interval 20;
    sample-rate 1000;
  }
}
```

### Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That sFlow Technology Has Been Configured Properly on page 6573](#)
- [Verifying That sFlow Technology Is Enabled on an Interface on page 6574](#)
- [Verifying the sFlow Collector Configuration on page 6574](#)

#### *Verifying That sFlow Technology Has Been Configured Properly*

**Purpose** Verify that sFlow technology has been configured properly.

**Action** Enter the **show sflow** operational mode command:

```
user@switch> show sflow
```

```
sFlow           : Enabled
Sample limit    : 300 packets/second
Polling interval : 20 second
Sample rate     : 1:1000
Agent ID       : 10.1.1.2
```



**NOTE:** The sample limit cannot be configured and is set to 300 packets per second.

**Meaning** The output shows that sFlow technology is enabled and specifies the values for the sampling limit, polling interval, and sampling rate.

#### *Verifying That sFlow Technology Is Enabled on an Interface*

**Purpose** Verify that sFlow technology is enabled on interfaces and display the sampling parameters.

**Action** Enter the **show sflow interface** operational mode command:

```
user@switch> show sflow interface
Interface      Status      Sample   Polling
                rate      interval
xe-0/0/1.0     Enabled     1000     20
```

**Meaning** The output indicates that sFlow technology is enabled on the **Node1:xe-0/0/1.0** interface on the Node device with a sampling rate of 1000 and a polling interval of 20 seconds.

#### *Verifying the sFlow Collector Configuration*

**Purpose** Verify the sFlow collector configuration.

**Action** Enter the **show sflow collector** operational mode command:

```
user@switch> show sflow collector
Collector      Udp-port   No. of samples
address
10.204.32.46   5600       7516
```

**Meaning** The output displays the IP address of the collector, the UDP port, and the number of samples collected.

**Related Documentation**

- [Configuring sFlow Technology on page 6596](#)
- [Overview of sFlow Technology](#)



## Example: Configuring SNMP

By default, SNMP is disabled on devices running Junos OS. This example describes the steps for configuring SNMP on the QFabric system.

- [Requirements on page 6575](#)
- [Overview on page 6575](#)
- [Configuration on page 6575](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 12.2
- Network management system (NMS) (running the SNMP manager)
- QFabric system (running the SNMP agent) with multiple Node devices

### Overview

Because SNMP is disabled by default on devices running Junos OS, you must enable SNMP on your device by including configuration statements at the **[edit snmp]** hierarchy level. At a minimum, you must configure the **community public** statement. The community defined as public grants read-only access to MIB data to any client.

If no **clients** statement is configured, all clients are allowed. We recommend that you always include the **restrict** option to limit SNMP client access to the switch.

The network topology in this example includes an NMS, a QFabric system with four Node devices, and external SNMP servers that are configured for receiving traps.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set snmp name "snmp qfabric" description "qfabric0 switch"
set snmp location "Lab 4 Row 11" contact "qfabric-admin@qfabric0"
set snmp community public authorization read-only
set snmp client-list list0 192.168.0.0/24
set snmp community public client-list-name list0
set snmp community public clients 192.170.0.0/24 restrict
set snmp trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure SNMP on the QFabric system:



**NOTE:** If the name, description, location, contact, or community name contains spaces, enclose the text in quotation marks (" ").

1. Configure the SNMP system name:

```
[edit snmp]
user@switch# set name "snmp qfabric"
```

2. Specify a description.

```
[edit snmp]
user@switch# set description "qfabric0 system"
```

This string is placed into the MIB II sysDescription object.

3. Specify the physical location of the QFabric system.

```
[edit snmp]
user@switch# set location "Lab 4 Row 11"
```

This string is placed into the MIB II sysLocation object.

4. Specify an administrative contact for the SNMP system.

```
[edit snmp]
user@switch# set contact "qfabric-admin@qfabric0"
```

This name is placed into the MIB II sysContact object.

5. Specify a unique SNMP community name and the read-only authorization level.



**NOTE:** The read-write option is not supported on the QFabric system.

```
[edit snmp]
user@switch# set community public authorization read-only
```

6. Create a client list with a set of IP addresses that can use the SNMP community.

```
[edit snmp]
user@switch# set client-list list0 192.168.0.0/24
user@switch# set community public client-list-name list0
```

7. Specify IP addresses of clients that are restricted from using the community.

```
[edit snmp]
user@switch# set community public clients 192.170.0.0/24 restrict
```

8. Configure a trap group, destination port, and a target to receive the SNMP traps in the trap group.

```
[edit snmp]
user@switch# set trap-group "qf-traps" destination-port 155 targets 192.168.0.100
```



**NOTE:** You do not need to include the `destination-port` statement if you use the default port 162.

The trap group `qf-traps` is configured to send traps to 192.168.0.100.

**Results** From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@switch# show
snmp {
  name "snmp qfabric";
  description "qfabric0 system";
  location "Lab 4 Row 11";
  contact "qfabric-admin@qfabric0";
  client-list list0 {
    192.168.0.0/24;
  }
  community public {
    authorization read-only;
    clients {
      197.170.0.0/24 restrict;
    }
  }
  trap-group qf-traps {
    destination-port 155;
    targets {
      192.168.0.100;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)
  - [snmp on page 1454](#)

## Example: Configuring Network Analytics

This example shows how to configure network analytics which includes queue and traffic monitoring on a QFX3500 standalone switch.



**NOTE:** The configuration shown in this example is supported only on Junos OS Release 13.2X50-D15 and 13.2X51-D10.

- [Requirements on page 6578](#)
- [Overview on page 6578](#)
- [Configuration on page 6578](#)
- [Verification on page 6581](#)

---

## Requirements

This example uses the following hardware and software components:

- A QFX3500 standalone switch
- A external streaming server to collect data
- Junos OS Release 13.2X50-D15 software
- TCP server software (for remote streaming servers)

Before you configure network analytics, be sure you have:

- Junos OS Release 13.2X50-D15 or later software installed and running on the QFX3500 switch
- (Optional for streaming servers) TCP server software set up for processing records separated by a newline character (\n) on the remote streaming server
- All other devices running

---

## Overview

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed. You can enable network analytics by configuring queue and traffic statistics monitoring.

### **Topology**

In this example, the QFX3500 switch is connected to an external server used for streaming statistics data.

---

## Configuration

To configure network analytics, perform these tasks:

- [Configuring Queue and Traffic Statistics Monitoring on page 6579](#)
- [Configuring Local Statistics Files on page 6579](#)
- [Configuring Streaming Servers on page 6580](#)
- [Results on page 6580](#)

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
[edit]
set services analytics interfaces all queue-statistics
set services analytics interfaces all latency-threshold high 900 low 300
set services analytics interfaces xe-0/0/1 traffic-statistics
set services analytics queue-statistics file qstats1.qs files 3 size 10
set services analytics queue-statistics interval 10
set services analytics traffic-statistics file tstats1.ts files 3 size 10
set services analytics traffic-statistics interval 2
set services analytics streaming-servers address 10.94.198.11 port 50001 stream-format
json stream-type queue-statistics
set services analytics streaming-servers address 10.94.198.11 port 50005 stream-format
csv stream-type traffic-statistics
```

### *Configuring Queue and Traffic Statistics Monitoring*

**Step-by-Step Procedure** To configure queue and traffic monitoring on physical interfaces:



**NOTE:** You can configure queue and traffic monitoring on physical network interfaces only; logical interfaces and Virtual Chassis physical (VCP) interfaces are not supported.



**NOTE:** Disabling of the queue or traffic monitoring supersedes the configuration (enabling) of this feature. You disable monitoring by issuing the `no-queue-statistics` or `no-traffic-statistics` at the **[edit services analytics interfaces]** hierarchy level.

1. Configure all interfaces for queue monitoring and set the latency thresholds (in microseconds):

```
[edit]
set services analytics interfaces all queue-statistics
set services analytics interfaces all latency-threshold high 900 low 300
```

2. Configure one interface for traffic monitoring:

```
[edit]
set services analytics interfaces xe-0/0/1 traffic-statistics
```

### *Configuring Local Statistics Files*

**Step-by-Step Procedure** To configure local statistics files:

1. Configure the number of queue statistics files, and each file size in MB:

```
[edit]
```

```
set services analytics queue-statistics file qstats1.qs files 3 size 10m
```

2. Configure the queue statistics collection interval in milliseconds

```
[edit]
set services analytics queue-statistics interval 10
```

3. Configure the number of traffic statistics files, and each file size in MB:

```
[edit]
set services analytics traffic-statistics file tstats1.ts files 3 size 10m
```

4. Configure the traffic statistics collection interval in seconds:

```
[edit]
set services analytics traffic-statistics interval 2
```

### *Configuring Streaming Servers*

#### **Step-by-Step Procedure**

To configure streaming servers for receiving monitoring data:



**NOTE:** In addition to configuring streaming servers, you must also set up the TCP client software to process records that are separated by the newline character (\n) on the remote server.

1. Configure a server IP address and port for queue statistics monitoring:

```
[edit]
set services analytics streaming-servers address 10.94.198.11 port 50001
stream-format json stream-type queue-statistics
```

2. Configure a server IP address and port for traffic statistics monitoring:

```
[edit]
set services analytics streaming-servers address 10.94.198.11 port 50005
stream-format csv stream-type traffic-statistics
```

### *Results*

Display the results of the configuration:

```
[edit services analytics]
user@switch> show configuration
queue-statistics {
  file qstats1.qs size 10m files 3;
  interval 10;
}
traffic-statistics {
  file tstats1.ts size 10m files 3;
  interval 2;
}
interfaces {
  xe-0/0/1 {
    traffic-statistics;
  }
  all {
```

```

        queue-statistics;
        latency-threshold high 900 low 300;
    }
}

```

### Verification

Confirm that the configuration is correct and works as expected by performing these tasks:

- [Verifying the Network Analytics Configuration on page 6581](#)
- [Verifying the Network Analytics Status on page 6581](#)
- [Verifying Streaming Servers Configuration on page 6582](#)
- [Verifying Queue Statistics on page 6582](#)
- [Verifying Traffic Statistics on page 6582](#)

#### *Verifying the Network Analytics Configuration*

**Purpose** Verify the configuration for network analytics.

**Action** From operational mode, enter the **show analytics configuration** command to display the traffic and queue monitoring configuration.

```

user@host> show analytics configuration
Global configurations:
  Traffic statistics: Auto, Poll interval: 2 seconds
  Queue statistics: Enabled, Poll interval: 10 milliseconds
  Depth threshold high: 0 bytes, low: 0 bytes
  Latency threshold high: 900 microseconds, low: 300 microseconds

```

Interface	Traffic	Queue	Depth-threshold		Latency-threshold	
	Statistics	Statistics	High	Low	High	Low
			(bytes)		(microseconds)	
xe-0/0/1	Enabled	Auto	0	0	900	300

**Meaning** The output displays information about traffic and queue monitoring on the switch.

#### *Verifying the Network Analytics Status*

**Purpose** Verify the network analytics operational status of the switch.

**Action** From operational mode, enter the **show analytics status** command to display the traffic and queue monitoring status.

```
user@host> show analytics status
Global configurations:
  Traffic statistics: Auto, Poll interval: 2 seconds
  Queue statistics: Auto, Poll interval: 10 milliseconds
  Depth threshold high: 1228800 bytes, low: 1024 bytes
  Latency threshold high: 900 microseconds, low: 300 microseconds
```

Interface	Traffic	Queue	Depth-threshold		Latency-threshold	
	Statistics	Statistics	High	Low	High	Low
			(bytes)		(microseconds)	
xe-0/0/1	Enabled	Auto	1228800	1024	900	300
xe-0/0/7	Auto	Auto	1228800	1024	900	300
xe-0/0/8	Auto	Auto	1228800	1024	900	300

### *Verifying Streaming Servers Configuration*

**Purpose** Verify the configuration for streaming data to remote servers is working.

**Action** From operational mode, enter the **show analytics streaming-servers** command to display the streaming servers configuration.

```
user@host> show analytics streaming-servers
```

Address	Port	Stream-Format	Stream-Type	State	Sent
10.94.198.11	50001	json	QS	Established	1100
10.94.198.11	50005	csv	TS/QS	In Progress	0

**Meaning** The output displays information about the remote streaming server.

### *Verifying Queue Statistics*

**Purpose** Verify that queue statistics collection is working.

**Action** From operational mode, enter the **show analytics queue-statistics** command to display the queue statistics.

```
user@host> show analytics queue-statistics
```

Time	Interface	Queue-length (bytes)	Latency (us)
Apr 6 0:17:18.224	xe-0/0/1	1043952	835
Apr 6 0:17:18.234	xe-0/0/1	1053520	842
Apr 6 0:17:18.244	xe-0/0/1	1055184	844

**Meaning** The output displays queue-statistics information as expected.

### *Verifying Traffic Statistics*

**Purpose** Verify that traffic statistics collection is working.



**Action** From operational mode, enter the **show analytics traffic-statistics** command to display the traffic statistics.

```

user@host> show analytics traffic-statistics
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/1
Traffic Statistics:
Total octets:          4797548752936          408886273632
Total packet:          5658257464            3190613435
Octets per second:      0                    0
Packet per second:      0                    0
Octets dropped:         0                    252901000
Packet dropped:         0                    252901
Utilization:           0.0%                  0.0%
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/7
Traffic Statistics:
Total octets:          4790866253100          477139024
Total packet:          5624473639            477944
Octets per second:      0                    0
Packet per second:      0                    0
Octets dropped:         0                    166582000
Packet dropped:         0                    166582
Utilization:           0.0%                  0.0%
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/8
Traffic Statistics:
Total octets:          4789797668456          764910024
Total packet:          5623280870            765715
Octets per second:      0                    0
Packet per second:      0                    0
Octets dropped:         0                    156099000
Packet dropped:         0                    156099
Utilization:           0.0%                  0.0%

```

**Meaning** The output displays traffic-statistics information as expected.

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [analytics on page 6667](#)
  - [show analytics status on page 6841](#)
  - [show analytics streaming-servers on page 6845](#)

## Configuration Tasks for Network Management

- [Configuring Console and Auxiliary Port Properties on page 6583](#)
- [Configuring SSH Service for Remote Access to the Router or Switch on page 6584](#)
- [Configuring Telnet Service for Remote Access to a Switch on page 6586](#)

### Configuring Console and Auxiliary Port Properties

The console port and auxiliary port on a switch provide out-of-band remote access to the switch. You can configure the console and auxiliary ports so that an external data terminal may be connected to the switch. The console port is enabled by default, and its speed is 9600 baud. The auxiliary port is disabled by default.

By default, terminal connections to the console and auxiliary ports are secure. When you configure the console and auxiliary ports as insecure, root logins are not allowed to establish terminal connections, and superusers and anyone with a user identifier (UID) of 0 are not allowed to establish terminal connections in multiuser mode.

To configure the console and auxiliary port properties on the switch:

1. To specify that the console port session should terminate if the connection to the data carrier is lost:

```
[edit system ports]
user@switch# set console log-out-on-disconnect
```

2. To specify the auxiliary port terminal type:

```
[edit system ports]
user@switch# set auxiliary type (ansi | small-xterm | vt100 | xterm)
```

For example, to specify the auxiliary port terminal type of **xterm** with a display of 80 columns by 65 rows:

```
[edit system ports]
user@switch# set auxiliary type xterm
```

3. To check the configuration:

```
[edit system ports]
user@switch# show
console log-out-on-disconnect;
auxiliary type xterm;
```

- Related Documentation**
- [auxiliary on page 257](#)
  - [console \(Physical Port\) on page 266](#)
  - [ports on page 298](#)

## Configuring SSH Service for Remote Access to the Router or Switch

To configure the router or switch to accept SSH as an access service, include the **ssh** statement at the **[edit system services]** hierarchy level:

```
[edit system services]
ssh {
  ciphers [ cipher-1 cipher-2 cipher-3 ...]
  client-alive-count-max number;
  client-alive-interval seconds;
  connection-limit limit;
  hostkey-algorithm <algorithm | no-algorithm>;
  key-exchange algorithm;
  macs algorithm;
  max-sessions-per-connection number;
  no-passwords;
  no-tcp-forwarding;
  protocol-version [v1 v2];
  rate-limit limit;
  root-login <allow | deny | deny-password>;
```

```
}
```

By default, the router or switch supports a limited number of simultaneous SSH sessions and connection attempts per minute. Use the following statements to change the defaults:

- **connection-limit *limit***—Maximum number of simultaneous connections per protocol (IPv4 and IPv6). The range is a value from 1 through 250. The default is 75. When you configure a connection limit, the limit is applicable to the number of SSH sessions per protocol (IPv4 and IPv6). For example, a connection limit of 10 allows 10 IPv6 SSH sessions and 10 IPv4 SSH sessions.
- **max-sessions-per-connection *number***—Include this statement to specify the maximum number of SSH sessions allowed per single SSH connection. This allows you to limit the number of cloned sessions tunneled within a single SSH connection. The default value is 10.
- **rate-limit *limit***—Maximum number of connection attempts accepted per minute (a value from 1 through 250). The default is 150. When you configure a rate limit, the limit is applicable to the number of connection attempts per protocol (IPv4 and IPv6). For example, a rate limit of 10 allows 10 IPv6 SSH session connection attempts per minute and 10 IPv4 SSH session connection attempts per minute.

By default, a user can create an SSH tunnel over a CLI session to a router running Junos OS via SSH. This type of tunnel could be used to forward TCP traffic, bypassing any firewall filters or ACLs, allowing access to resources beyond the router. Use the **no-tcp-forwarding** option to prevent a user from creating an SSH tunnel to a router via SSH.

For information about other configuration settings, see the following topics:

- [Configuring the Root Login Through SSH on page 6585](#)
- [Configuring the SSH Protocol Version on page 6586](#)
- [Configuring the Client Alive Mechanism on page 6586](#)

### Configuring the Root Login Through SSH

By default, users are allowed to log in to the router or switch as **root** through SSH. To control user access through SSH, include the **root-login** statement at the **[edit systems services ssh]** hierarchy level:

```
[edit system services ssh]
  root-login (allow | deny | deny-password);
```

**allow**—Allows users to log in to the router or switch as root through SSH. The default is **allow**.

**deny**—Disables users from logging in to the router or switch as root through SSH.

**deny-password**—Allows users to log in to the router or switch as root through SSH when the authentication method (for example, RSA) does not require a password.

### Configuring the SSH Protocol Version

---

By default, both version 1 and version 2 of the SSH protocol are enabled. To configure the router or switch to use only version 1 of the SSH protocol, include the **protocol-version** statement and specify **v1** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 ];
```

To configure the router or switch to use only version 2 of the SSH protocol, include the **protocol-version** statement and specify **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v2 ];
```

To explicitly configure the router or switch to use version 1 and 2 of the SSH protocol, include the **protocol-version** statement and specify **v1** and **v2** at the **[edit system services ssh]** hierarchy level:

```
[edit system services ssh]
protocol-version [ v1 v2 ];
```

For J Series Services Routers, the export license software supports SSH version 1 only.

### Configuring the Client Alive Mechanism

---

The client alive mechanism is valuable when the client or server depends on knowing when a connection has become inactive. It differs from the standard keepalive mechanism because the client alive messages are sent through the encrypted channel. The client alive mechanism is not enabled at default. To enable it, configure the **client-alive-count-max** and the **client-alive-interval**. This option applies to SSH protocol version 2 only.

In the following example, unresponsive SSH clients will be disconnected after approximately 100 seconds (20 x 5).

```
[edit system services ssh]
client-alive-count-max 5;
client-alive-interval 20;
```

### Configuring Telnet Service for Remote Access to a Switch

Telnet provides unencrypted access to network devices. Configuring Telnet service for a switch enables in-band remote access to the switch.

By default, the switch supports a limited number of simultaneous Telnet sessions and connection attempts per minute. Optionally, you can change the default Telnet settings by configuring the connection limit and rate limit at the **[edit system services telnet]** hierarchy level.

The connection limit is the maximum number of simultaneous connections per protocol (IPv4). The range is from 1 through 250. The default is 75.

The rate limit is the maximum number of connection attempts accepted per minute per protocol. The range is from 1 through 250. The default is 150.

To configure Telnet service:

1. To specify the connection limit:

```
[edit system services]
user@switch# set telnet connection-limit connection-limit
```

2. To specify the rate limit:

```
[edit system services]
user@switch# set telnet rate-limit rate-limit
```

3. Check that the Telnet connection limit and rate limit show the values you specified:

```
[edit system services]
user@switch# show
telnet {
  connection-limit 50;
  rate-limit 100;
}
```

#### Related Documentation

- [Understanding Telnet on the QFabric System on page 6467](#)
- [Limiting the Number of User Login Attempts for SSH and Telnet Sessions on page 1369](#)
- [Example: Limiting the Number of Login Attempts for SSH and Telnet Sessions on page 1391](#)

## Configuration Tasks for Automation

- [Invoking the Python Interpreter on page 6587](#)
- [Controlling the Execution of Commit Scripts on page 6588](#)

### Invoking the Python Interpreter

The Python interpreter is available by default with the QFX5100 switch automation enhancements. You can invoke Python by entering the **python** command at the shell script.

To invoke the Python interpreter:

1. Start the shell interface:

```
user@switch> start shell
```

2. Enter the **python** command without any parameters:

```
% python
```



**NOTE:** The Python interpreter is designated with the prompt >>> at the beginning of a line or ... to indicate the continuation of a line.

- Related Documentation**
- [Overview of Python with QFX5100 Switch Automation Enhancements on page 6471](#)
  - [Overview of QFX5100 Switch Automation Enhancements on page 6470](#)
  - [Installing Junos OS Software with QFX5100 Switch Automation Enhancements](#)
  - [QFX5100 Switch with Automation Enhancements Frequently Asked Questions on page 6897](#)

## Controlling the Execution of Commit Scripts

This document describes the tasks that affect the way commit scripts are executed. In the QFabric system, commit scripts are stored in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory that is shared among Director devices in a Director group.

To determine which commit scripts are currently enabled on the QFabric system, use the `show` command to display the files included at the `[edit system scripts commit]` hierarchy level. To ensure that the enabled files are on the device, list the contents of the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit` directory using the `file list` operational mode command.

See the following tasks:

- [Enabling Commit Scripts to Execute on page 6588](#)
- [Removing Commit Scripts from the Configuration on page 6589](#)
- [Deactivating Commit Scripts on page 6590](#)
- [Activating Inactive Commit Scripts on page 6590](#)

---

### Enabling Commit Scripts to Execute

The commit operation requires that all scripts be included in configuration at the `[edit system scripts commit file]` hierarchy level for all QFabric Director devices.

If you need to temporarily remove a script from a commit operation but do not want to remove it from the configuration permanently, you may configure the `optional` statement at the `[edit system scripts commit file filename]` hierarchy level to enable the commit operation to succeed even if a script is missing from the commit script directory.



**CAUTION:** When you include the `optional` statement at the `[edit system scripts commit file filename]` hierarchy level, no error message is generated during the commit operation if the file does not exist. As a result, you might not be aware that a script has not been executed as expected.

---

The filename of a commit script written in SLAX must include the **.slax** extension for the script to be executed.

To enable a commit script to execute during a commit operation:

1. Ensure that the commit script is located in the correct directory:  
**/pbdata/mgd\_shared/partition-ip/var/db/scripts/commit** directory on the Director device.

```
[edit system scripts commit]
user@switch# set file filename <optional>
```

3. Commit the configuration.

```
[edit system scripts commit]
user@switch# top
[edit]
user@switch# commit
```

### Removing Commit Scripts from the Configuration

You can prevent commit scripts from executing during a commit operation by removing the scripts from the commit directory in the configuration.



**NOTE:** You can also deactivate a script using the **deactivate** statement instead of removing it from the configuration. Deactivated scripts may be reactivated later.

To prevent a commit script from executing during a commit operation:

1. Delete the commit script file from the commit directory in the configuration.

```
[edit system scripts commit]
user@switch# delete file filename
```

2. Commit the configuration.

```
[edit system scripts commit]
user@switch# top
[edit]
user@switch# commit
```

3. Remove the commit script from the **/pbdata/mgd\_shared/** directory on the Director device.



**BEST PRACTICE:** Although removing the commit script is not necessary, we recommend deleting unused files from the system.

### Deactivating Commit Scripts

---

Deactivating a commit script results in its being marked as inactive in the configuration. The script is not executed during the commit operation, but you can reactivate the script by using the **activate** statement.

To deactivate the commit script:

1. Deactivate the script.

```
[edit]
user@switch deactivate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@switch# commit
```

3. Verify that the commit script is deactivated.

```
[edit]
user@switch# show system scripts commit
inactive: file mycommit.slax
```

### Activating Inactive Commit Scripts

---

Deactivating a commit script results in its being marked as inactive in the configuration and is therefore not executed during the commit operation.

To activate an inactive commit script:

1. Activate the script.

```
[edit]
user@switch# activate system scripts commit file filename
```

2. Commit your changes.

```
[edit]
user@switch# commit
```

## Configuration Tasks for Network Analytics

---

- [Configuring Queue Monitoring on page 6591](#)
- [Configuring Traffic Monitoring on page 6593](#)
- [Configuring a Local File for Network Analytics Data on page 6594](#)
- [Configuring a Remote Collector for Streaming Analytics Data on page 6595](#)



## Configuring Queue Monitoring

Network analytics queue monitoring provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. You can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You enable queue monitoring by first defining a resource profile template, and then applying the profile to the system (for a global configuration) or to individual interfaces.



**NOTE:** You can configure queue monitoring on physical network interfaces only; logical interfaces and Virtual Chassis physical (VCP) interfaces are not supported.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure queue monitoring on a QFX Series standalone switch:

1. Configure the queue monitoring polling interval (in milliseconds) globally (for the system):

```
[edit]
set services analytics resource system polling-interval queue-monitoring interval
```

2. Configure a resource profile for the system, and enable queue monitoring:

```
[edit]
set services analytics resource-profiles profile-name queue-monitoring
```

3. Configure high and low values of the depth-threshold (in bytes) for queue monitoring in the system profile:

```
[edit]
set services analytics resource-profiles profile-name depth-threshold high number low number
```

For both high and low values, the range is from 1 to 1,250,000,000 bytes, and the default value is 0 bytes.



**NOTE:** You can configure either the depth-threshold or latency threshold for the system, but not both.

4. Apply the resource profile template to the system for a global configuration:

```
[edit]
set services analytics resource system resource-profile profile-name
```

5. Configure an interface-specific resource profile and enable queue monitoring for the interface:

```
[edit]  
set services analytics resource-profiles profile-name queue-monitoring
```

6. Configure the latency-threshold (high and low values) for queue monitoring in the interface-specific profile:

```
[edit]  
set services analytics resource-profiles profile-name latency-threshold high number  
low number
```

For both high and low values, the range is from 1 to 100,000,000 nanoseconds, and the default value is 1,000,000 nanoseconds.



**NOTE:** You can configure either the depth-threshold or latency threshold for interfaces, but not both.

---

7. Apply the resource profile template for interfaces to one or more interfaces:

```
[edit]  
set services analytics resource interfaces interface-name resource-profile profile-name
```



**NOTE:** If a conflict arises between the system and interface configurations, the interface-specific configuration supersedes the global (system) configuration.

---

#### Related Documentation

- [Network Analytics Overview on page 6490](#)
- [Example: Configuring Enhanced Network Analytics Features](#)
- [analytics on page 6667](#)

## Configuring Traffic Monitoring

Network analytics queue monitoring provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. You can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You enable traffic monitoring by first defining a resource profile template, and then applying the profile to the system (for a global configuration) or to individual interfaces.



**NOTE:** You can configure traffic monitoring on physical network interfaces only; logical interfaces and Virtual Chassis physical (VCP) interfaces are not supported.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure traffic monitoring on a QFX Series standalone switch:

1. Configure the traffic monitoring polling interval (in seconds) for the system:

[edit]

**set services analytics resource system polling-interval traffic-monitoring *interval***

2. Configure a resource profile for the system, and enable traffic monitoring in the profile:

[edit]

**set services analytics resource-profiles *profile-name* traffic-monitoring**

3. Apply the resource profile to the system for a global configuration:

[edit]

**set services analytics resource system resource-profile *profile-name***

4. Configure a resource profile for interfaces, and enable traffic monitoring in the profile:

[edit]

**set services analytics resource-profiles *profile-name* traffic-monitoring**



**NOTE:** If a conflict arises between the system and interface configurations, the interface-specific configuration supersedes the global (system) configuration.

5. Apply the resource profile template to one or more interfaces:

[edit]

**set services analytics resource interfaces *interface-name* resource-profile *profile-name***

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [Example: Configuring Enhanced Network Analytics Features](#)
  - [analytics on page 6667](#)

## Configuring a Local File for Network Analytics Data

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

To save the queue and traffic statistics data in a local file, you must configure a filename to store it.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure a local file for storing queue and traffic monitoring statistics:

1. Configure a filename:

[edit]

**set services analytics collector local file *filename***

There is no default filename. If you do not configure a filename, network analytics statistics are not saved locally.

2. Configure the number of files (from 2 to 1000 files):

[edit]

**set services analytics collector local file *filename* files *number***

3. Configure the file size (from 10 to 4095 MB) in the format of xm:

[edit]

**set services analytics collector local file an size *size***

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [Example: Configuring Enhanced Network Analytics Features](#)
  - [analytics on page 6667](#)

## Configuring a Remote Collector for Streaming Analytics Data

The network analytics feature provides visibility into the performance and behavior of the data center infrastructure. This feature collects data from the switch, analyzes the data using sophisticated algorithms, and captures the results in reports. Network administrators can use the reports to help troubleshoot problems, make decisions, and adjust resources as needed.

You can configure an export profile to define the stream format and type of data, and one or more remote servers (collectors) to receive streaming network analytics data.



**NOTE:** This procedure requires Junos OS Release 13.2X51-D15 or later to be installed on your device.

To configure a collector for receiving streamed analytics data:

1. Create an export profile and specify the stream format:

```
[edit]
set services analytics export-profiles profile-name stream-format format
```

2. Configure the export profile to include interface information:

```
[edit]
set services analytics export-profiles profile-name interface information
```

3. Configure the export profile to include interface queue statistics:

```
[edit]
set services analytics export-profiles profile-name interface statistics queue
```

4. Configure the export profile to include interface traffic statistics:

```
[edit]
set services analytics export-profiles profile-name interface statistics traffic
```

5. Configure the export profile to include interface status link information:

```
[edit]
set services analytics export-profiles profile-name interface status link
```

6. Configure the export profile to include system information:

```
[edit]
set services analytics export-profiles profile-name system information
```

7. Configure the export profile to include system queue status:

```
[edit]
set services analytics export-profiles profile-name system status queue
```

8. Configure the export profile to include system traffic status:

```
[edit]
set services analytics export-profiles profile-name system status traffic
```

9. Configure the transport protocol for the collector addresses and apply the export profile:

```
[edit]
set services analytics collector address ip-address port port transport protocol
export-profile profile-name
set services analytics collector address ip-address port port transport protocol
export-profile profile-name
```

---



**NOTE:** If you configure the `tcp` or `udp` option for the JSON, CSV, and TSV formats, you must also set up the TCP or UDP client software on the remote collector to process records that are separated by the newline character (`\n`) on the remote server.

If you configure the `tcp` or `udp` option for the GPB format, you must also set up the TCP or UDP build streaming server using the `analytics.proto` file.

---

**Related  
Documentation**

- [Network Analytics Overview on page 6490](#)
- [Example: Configuring Enhanced Network Analytics Features](#)
- [analytics on page 6667](#)

---

## Configuration Tasks for sFlow Technology

---

- [Configuring sFlow Technology on page 6596](#)

### Configuring sFlow Technology

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually; you cannot globally enable sFlow monitoring on all interfaces with a single configuration statement. Junos OS fully supports the sFlow technology standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

On the QFabric system, the sFlow monitoring global configuration that is defined on the Director device is distributed to Node groups that have sFlow sampling configured on the interfaces.

To configure sFlow features using the CLI:

1. Configure the IP address and UDP port of at least one collector:

```
[edit protocols sflow]
user@host# set collector ip-address udp-port port-number
```

The default UDP port assigned is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@host# set interfaces interface-name
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a LAG interface (for example ae0), but you can enable sFlow technology on the member interfaces of the LAG (for example, xe-0/0/1).

3. Specify how often (in seconds) the sFlow agent polls all interfaces at the global level:

```
[edit protocols sflow]
user@host# set polling-interval seconds
```



**NOTE:** Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets are sampled at the global level. For example, configuring a *number* of 1000 sets a sample rate of 1 in 1000 packets.

```
[edit protocols sflow]
user@host# set sample-rate number
```

5. (Optional) You can also configure the polling interval and sample rate at the interface level:

```
[edit protocols sflow]
user@host# set interfaces interface-name polling-interval seconds sample-rate number
```



**NOTE:** The interface-level configuration overrides the global configuration for the specified interface.

#### Related Documentation

- [Example: Monitoring Network Traffic Using sFlow Technology on page 6571](#)
- [Overview of sFlow Technology](#)

## Configuration Tasks for SNMP

- [Configuring SNMP on page 6598](#)
- [Configuring the SNMP Community String on page 6601](#)
- [Configuring SNMP Trap Groups on page 6602](#)
- [Adding a Group of Clients to an SNMP Community on page 6603](#)
- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 6604](#)
- [Configuring MIB Views on page 6605](#)
- [Configuring RMON Alarms and Events on page 6606](#)
- [Configuring Health Monitoring on page 6609](#)
- [Creating SNMPv3 Users on page 6609](#)
- [Configuring Access Privileges for a Group on page 6611](#)

- [Assigning a Security Name to a Group on page 6612](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 6613](#)
- [Configuring SNMP Informs on page 6614](#)

## Configuring SNMP

SNMP is implemented in the Junos OS Software running on the QFX Series products. By default, SNMP is not enabled. To enable SNMP, you must include the SNMP configuration statements at the **[edit]** hierarchy level.

To configure the minimum requirements for SNMP, include the following statements at the **[edit]** hierarchy level of the configuration:

```
[edit]
snmp {
  community public;
}
```

To configure complete SNMP features, include the following statements at the **[edit]** hierarchy level of the configuration:

```
snmp {
  client-list client-list-name {
    ip-addresses;
  }
  community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
      address restrict;
    }
    logical-system logical-system-name {
      routing-instance routing-instance-name {
        clients {
          addresses;
        }
      }
    }
    routing-instance routing-instance-name {
      clients {
        addresses;
      }
    }
    view view-name;
  }
  contact contact;
  description description;
  filter-duplicates;
  filter-interfaces;
  health-monitor {
    falling-threshold integer;
    interval seconds;
    rising-threshold integer;
  }
  interface [ interface-names ];
```



```

location location;
name name;
nonvolatile {
    commit-delay seconds;
}
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        falling-threshold-interval seconds;
        interval seconds;
        request-type;
        rising-event-index index;
        rising-threshold integer;
        sample-type (absolute-value | delta-value);
        startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
        syslog-subtag syslog-subtag;
        variable oid-variable;
    }
    event index {
        community community-name;
        description description;
        type type;
    }
    history history-index {
        bucket-size number;
        interface interface-name;
        interval seconds;
        owner owner-name;
    }
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
        regular-expression>;
    flag flag;
}
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance routing-instance-name;
    targets {
        address;
    }
    version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    source-address address;
}
v3 {
    notify name {
        tag tag-name;
        type trap;
    }
}

```

```
}
notify-filter profile-name {
    oid object-identifier (include | exclude);
}
snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
}
target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance routing-instance-name;
    tag-list tag-list;
    target-parameters target-parameters-name;
    timeout seconds;
}
target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
        message-processing-model (v1 | v2c | V3);
        security-level (authentication | none | privacy);
        security-model (usm | v1 | v2c);
        security-name security-name;
    }
}
usm {
    local-engine {
        user username {
            authentication-sha {
                authentication-password authentication-password;
            }
            authentication-md5 {
                authentication-password authentication-password;
            }
            authentication-none;
            privacy-aes128 {
                privacy-password privacy-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-none;
        }
    }
    remote-engine engine-id {
        user username {
            authentication-sha {
                authentication-password authentication-password;
            }
        }
    }
}
```

```

authentication-md5 {
    authentication-password authentication-password;
}
authentication-none;
privacy-aes128 {
    privacy-password privacy-password;
}
privacy-des {
    privacy-password privacy-password;
}
privacy-3des {
    privacy-password privacy-password;
}
privacy-none {
    privacy-password privacy-password;
}
}
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
}
view view-name {
    oid object-identifier (include | exclude);
}
}

```

- Related Documentation**
- [Understanding the Implementation of SNMP on page 6513](#)
  - [snmp on page 1454](#)

## Configuring the SNMP Community String

The SNMP community string defines the relationship between an SNMP server system and the client systems. This string acts like a password to control the clients' access to

the server. To configure a community string in a Junos OS configuration, include the **community** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
community name {
  authorization authorization;
  clients {
    default restrict;
    address restrict;
  }
  view view-name;
}
```

If the community name contains spaces, enclose it in quotation marks (" ").

The default authorization level for a community is **read-only**. To allow **Set** requests within a community, you need to define that community as **authorization read-write**. For **Set** requests, you also need to include the specific MIB objects that are accessible with read-write privileges using the **view** statement. The default view includes all supported MIB objects that are accessible with read-only privileges; no MIB objects are accessible with read-write privileges. For more information about the **view** statement, see [“Configuring MIB Views” on page 6605](#).

The **clients** statement lists the IP addresses of the clients (community members) that are allowed to use this community. If no **clients** statement is present, all clients are allowed. For **address**, you must specify an IPv4 address, not a hostname. Include the **default restrict** option to deny access to all SNMP clients for which access is not explicitly granted. We recommend that you always include the **default restrict** option to limit SNMP client access to the local switch.



**NOTE:** Community names must be unique within each SNMP system.

---

Related Documentation

- [Configuring SNMP on page 1356](#)

## Configuring SNMP Trap Groups

Before any SNMP traps can be sent, you must configure a trap group, the categories of traps the group can receive, and the targets (systems) that will receive the traps. To create and name an SNMP trap group, include the **trap-group** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  targets {
    address;
  }
  version (all | v1 | v2);
}
```

```
}
```

The trap group name can be any string and is embedded in the community name field of the trap. To configure your own trap group port, include the **destination-port** statement. The default destination port is port 162.

For each trap group that you define, you must include the **target** statement to define at least one system as the recipient of the SNMP traps in the trap group. Specify the IPv4 address of each recipient and not its hostname.

Specify the types of traps the trap group can receive in the **categories** statement.

A trap group can receive the following categories of traps:

- **authentication**—Authentication failures
- **chassis**—Chassis or environment notifications
- **configuration**—Configuration notifications
- **link**—Link-related notifications such as up-down transitions
- **remote-operations**—Remote operation notifications
- **startup**—System warm and cold starts

The **version** statement allows you to specify the SNMP version of the traps sent to targets of the trap group. If you specify **v1** only, SNMPv1 traps are sent. If you specify **v2** only, SNMPv2 traps are sent. If you specify **all**, both an SNMPv1 and an SNMPv2 trap are sent for every trap condition. For more information about the **version** statement, see [version](#).

#### Related Documentation

- *Standard SNMP Version 1 Traps*
- *Standard SNMP Version 2 Traps*
- *Juniper Networks Enterprise-Specific SNMP Version 1 Traps*
- *Juniper Networks Enterprise-Specific SNMP Version 2 Traps*

## Adding a Group of Clients to an SNMP Community

Junos OS enables you to add one or more groups of clients to an SNMP community. You can include the **client-list-name name** statement at the **[edit snmp community community-name]** hierarchy level to add all the members of the client list or prefix list to an SNMP community.

To define a list of clients, include the **client-list** statement followed by the IP addresses of the clients at the **[edit snmp]** hierarchy level:

```
[edit snmp]
  client-list client-list-name {
    ip-addresses;
  }
```

You can configure a prefix list at the **[edit policy options]** hierarchy level. Support for prefix lists in the SNMP community configuration enables you to use a single list to configure the SNMP and routing policies. For more information about the **prefix-list**

statement, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

To add a client list or prefix list to an SNMP community, include the **client-list-name** statement at the **[edit snmp community *community-name*]** hierarchy level:

```
[edit snmp community community-name]  
client-list-name client-list-name;
```



**NOTE:** The client list and prefix list must not have the same name.

The following example shows how to define a client list:

```
[edit]  
snmp {  
  client-list clentlist1 {  
    10.1.1.1/32;  
    10.2.2.2/32;  
  }  
}
```

The following example shows how to add a client list to an SNMP community:

```
[edit]  
snmp {  
  community community1 {  
    authorization read-only;  
    client-list-name clientlist1;  
  }  
}
```

The following example shows how to add a prefix list to an SNMP community:

```
[edit]  
policy-options {  
  prefix-list prefixlist {  
    10.3.3.3/32;  
    10.5.5.5/32;  
  }  
}  
snmp {  
  community community2 {  
    client-list-name prefixlist;  
  }  
}
```

- Related Documentation**
- [client-list on page 1408](#)
  - [client-list-name on page 1409](#)

## Configuring the Interfaces on Which SNMP Requests Can Be Accepted

By default, all router or switch interfaces have SNMP access privileges. To limit the access through certain interfaces only, include the **interface** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
interface [ interface-names ];
```

Specify the names of any logical or physical interfaces that should have SNMP access privileges. Any SNMP requests entering the router or switch from interfaces not listed are discarded.

#### Related Documentation

- *Configuring SNMP on a Device Running Junos OS*
- *Configuration Statements at the [edit snmp] Hierarchy Level*
- *Example: Configuring Secured Access List Checking*
- [Configuring SNMP on page 1356](#)

## Configuring MIB Views

SNMPv3 defines the concept of MIB views in RFC 3415, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*. MIB views provide an agent better control over who can access specific branches and objects within its MIB tree. A view consists of a name and a collection of SNMP object identifiers, which are either explicitly included or excluded. Once defined, a view is then assigned to an SNMPv3 group or SNMPv1/v2c community (or multiple communities), automatically masking which parts of the agent's MIB tree members of the group or community can (or cannot) access.

By default, an SNMP community grants read access and denies write access to all supported MIB objects (even communities configured as **authorization read-write**). To restrict or grant read or write access to a set of MIB objects, you must configure a MIB view and associate the view with a community.

To create a MIB view and assign it to a community:

1. Configure a MIB view.

Although most network management systems use SNMPv3, Junos OS allows the use of MIB views with both SNMPv1 and SNMPv2c communities.

The **view** statement defines a MIB view and identifies a group of MIB objects. Each MIB object of a view has a common object identifier (OID) prefix. Each object identifier represents a subtree of the MIB object hierarchy. The subtree can be represented either by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as **interfaces**). A configuration statement uses a view to specify a group of MIB objects on which to define access. You can also use a wildcard character asterisk (\*) to include OIDs that match a particular pattern in the SNMP view. To enable a view, you must associate the view with a community.

To remove an OID completely, use the **delete view all oid oid-number** command but omit the **include** parameter.

```
[edit groups global snmp]
user@host# set view view-name oid object-identifier (include | exclude)
```

The following example creates a MIB view called ping-mib-view. The **oid** statement does not require a dot at the beginning of the object identifier. The **snmp view** statement

includes the branch under the object identifier .1.3.6.1.2.1.80. This includes the entire DISMAN-PINGMIB subtree (as defined in RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*), which effectively permits access to any object under that branch.

```
[edit groups global snmp]
user@host# set view ping-mib-view oid 1.3.6.1.2.1.80 include
```

The following example adds a second branch in the same MIB view.

```
[edit groups global snmp]
user@host# set view ping-mib-view oid jnxPingMIB include
```

2. Assign a MIB view to a community that you want to control.

```
[edit groups global snmp community community-name]
user@host# set view view-name
```

This example creates a new community ping-mib which has read-write access to create entries within the DISMAN-PING-MIB.

```
[edit groups global snmp community ping-mib]
user@host# set authorization read-write
```

This example associate the MIB view created earlier with the new community.

```
[edit groups global snmp community ping-mib]
user@host# set view ping-mib-view
```

3. At the top level of the configuration, apply the configuration group.

If you use a configuration group, you must apply it for it to take effect.

```
[edit]
user@host# set apply-groups global
```

4. Commit the configuration.

```
user@host# commit
```

5. To verify, make sure that any member of the ping-mib community has read/write access to the branches that you specified under ping-mib-view.

#### Related Documentation

- *PING MIB*
- *Configuring SNMP on a Device Running Junos OS*
- *Configuration Statements at the [edit snmp] Hierarchy Level*
- *Example: Ping Proxy MIB*
- *SNMP MIBs and Traps Reference*

## Configuring RMON Alarms and Events

The Junos OS supports the *Remote Network Monitoring* (RMON) MIB (RFC 2819), which allows a management device to monitor the values of MIB objects, or variables, against configured thresholds. When the value of a variable crosses a threshold, an alarm and



its corresponding event are generated. The event can be logged and can generate an SNMP trap.

To configure RMON alarms and events using the CLI, perform these tasks:

1. [Configuring SNMP on page 6607](#)
2. [Configuring an Event on page 6607](#)
3. [Configuring an Alarm on page 6608](#)

## Configuring SNMP

To configure SNMP:

1. Grant read-only access to all SNMP clients:

```
[edit snmp]
user@switch# set community community-name authorization authorization
```

For example:

```
[edit snmp]
user@switch# set community public authorization read-only
```

2. Grant read-write access to the RMON and jnx-rmon MIBs:

```
[edit snmp]
user@switch# set view view-name oid object-identifier include
user@switch# set view view-name oid object-identifier include
user@switch# set community community-name authorization authorization view view-name
```

For example:

```
[edit snmp]
user@switch# set view rmon-mib-view oid .1.3.6.1.2.1.16 include
user@switch# set view rmon-mib-view oid .1.3.6.1.4.1.2636.13 include
user@switch# set community private authorization read-write view rmon-mib-view
```

OIDs 1.3.6.1.2.1.16 and 1.3.6.1.4.1.2636.13 correspond to the RMON and jnxRmon MIBs.

3. Configure an SNMP trap group:

```
[edit snmp]
user@switch# set trap-group group-name categories category
user@switch# set trap-group group-name targets address
```

For example:

```
[edit snmp]
user@switch# set trap-group rmon-trap-group categories rmon-alarm
user@switch# set trap-group rmon-trap-group targets 192.168.5.5
```

The trap group **rmon-trap-group** is configured to send RMON traps to 192.168.5.5.

## Configuring an Event

To configure an event:

1. Configure an event index, community name, and type:

```
[edit snmp rmon]
user@switch# set event index community community-name type type
```

For example:

```
[edit snmp rmon]
```

```
user@switch# set event 1 community rmon-trap-group type log-and-trap
```

The event community corresponds to the SNMP trap group and is not the same as an SNMP community. This event generates an SNMP trap and adds an entry to the **logTable** in the RMON MIB.

2. Configure a description for the event:

```
[edit snmp rmon]
user@switch# set event index description description
```

For example:

```
[edit snmp rmon]
user@switch# set event 1 description "rmon event"
```

---

## Configuring an Alarm

To configure an alarm:

1. Configure an alarm index, the variable to monitor, the rising and falling thresholds, and the corresponding rising and falling events:

```
[edit snmp rmon]
user@switch# set alarm index variable oid-variable falling-threshold integer rising-threshold
integer rising-event-index index falling-event-index index
```

For example:

```
[edit snmp rmon]
user@switch# set alarm 5 variable .1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0 falling-threshold 75
rising-threshold 90 rising-event-index 1 falling-event-index 1
```

The variable `.1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0` corresponds to the **jnxRmon** MIB object **jnxOperatingCPU**, which represents the CPU utilization of the Routing Engine. The falling and rising threshold integers are 75 and 90. The rising and falling events both generate the same event (event index 1).

2. Configure the sample interval and type and the alarm type:

```
[edit snmp rmon]
user@switch# set alarm index interval seconds sample-type (absolute-value | delta-value)
startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm)
```

For example:

```
[edit snmp rmon]
user@switch# set alarm 5 interval 30 sample-type absolute-value
startup-alarm rising-or-falling-alarm
```

The absolute value of the monitored variable is sampled every 30 seconds. The initial alarm can occur because of rising above the rising threshold or falling below the falling threshold.

### Related Documentation

- [Configuring SNMP on page 1356](#)
- [Juniper Networks Enterprise-Specific MIBs](#)
- [Monitoring RMON MIB Tables on page 6803](#)
- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6527](#)
- [Understanding RMON on page 6525](#)

## Configuring Health Monitoring

This topic describes how to configure the health monitor feature for QFX Series devices.

The health monitor feature extends the SNMP RMON alarm infrastructure to provide predefined monitoring for a selected set of object instances (such as file system usage, CPU usage, and memory usage) and dynamic object instances (such as Junos OS processes).

To configure health monitoring:

1. Configure the health monitor:

```
[edit snmp]
user@switch# set health-monitor
```

2. Configure the falling threshold:

```
[edit snmp]
user@switch# set health-monitor falling-threshold percentage
```

For example:

```
user@switch# set health-monitor falling-threshold 85
```

3. Configure the rising threshold:

```
[edit snmp]
user@switch# set health-monitor rising-threshold percentage
```

For example:

```
user@switch# set health-monitor rising-threshold 75
```

4. Configure the interval:

```
[edit snmp]
user@switch# set health-monitor interval seconds
```

For example:

```
user@switch# set health-monitor interval 600
```

### Related Documentation

- [Understanding Health Monitoring on page 6529](#)
- [falling-threshold on page 1415](#)
- [interval \(Health Monitor\) on page 1420](#)
- [rising-threshold \(Health Monitor\) on page 1451](#)

## Creating SNMPv3 Users

For each SNMPv3 user, you can specify the username, authentication type, authentication password, privacy type, and privacy password. After a user enters a password, a key based on the engine ID and password is generated and is written to the configuration file. After the generation of the key, the password is deleted from this configuration file.



**NOTE:** You can configure only one encryption type for each SNMPv3 user.

To create users, include the **user** statement at the **[edit snmp v3 usm local-engine]** hierarchy level:

```
[edit snmp v3 usm local-engine]
user username;
```

**username** is the name that identifies the SNMPv3 user.

To configure user authentication and encryption, include the following statements at the **[edit snmp v3 usm local-engine user username]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]
authentication-md5 {
  authentication-password authentication-password;
}
authentication-sha {
  authentication-password authentication-password;
}
authentication-none;
privacy-aes128 {
  privacy-password privacy-password;
}
privacy-des {
  privacy-password privacy-password;
}
privacy-3des {
  privacy-password privacy-password;
}
privacy-none;
```

**Related  
Documentation**

- [Complete SNMPv3 Configuration Statements](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524](#)
- [Example: Creating SNMPv3 Users](#)
- [Example: SNMPv3 Configuration](#)

## Configuring Access Privileges for a Group

In SNMPv3, you can configure a group that sets the same access privileges for one or more users. Configuring a group includes defining the security model and security level, and associating one or more MIB view permissions for the group.



**NOTE:** You must associate at least one MIB view with the group. You can associate multiple MIB views (read, notify, write) to authorize different permissions based on the view. The view name cannot exceed 32 characters.

To configure access privileges for a group:

1. To configure the group:

```
[edit snmp v3 vacm access]
user@switch# edit group group-name
```

2. To configure the context prefix of the SNMP instance for the group:

```
[edit snmp v3 vacm access group group-name]
user@switch# edit (default-context-prefix | context-prefix context-prefix)
```

For example, to configure the default context prefix:

```
[edit snmp v3 vacm access group group-name]
user@switch# edit default-context-prefix
```

3. To configure the security model:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
user@switch# edit security-model (any | usm | v1 | v2c)
```

For example, to configure the SNMPv3 user-based security model (USM):

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
user@switch# edit security-model usm
```

4. To configure the security level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c)]
user@switch# edit security-level (authentication | none | privacy)
```

For example, to configure a security level requiring user authentication and encryption:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c)]
user@switch# edit security-level privacy
```



**NOTE:** Access privileges are granted to all packets with a security level equal to or greater than that configured. If you are configuring the SNMPv1 or v2c security model, use *none* as your security level. If you are configuring the SNMPv3 security model (USM), use the *authentication*, *none*, or *privacy* security level.

5. (Optional) To associate a read-only MIB view with an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit read-view view-name
```

6. (Optional) To associate a MIB view with an SNMP notification permission for an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit notify-view view-name
```

7. (Optional) To associate a MIB view with write permission for an SNMP group:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication |
none | privacy)]
user@switch# edit write-view view-name
```

#### Related Documentation

- [SNMPv3 Overview on page 6523](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524](#)

## Assigning a Security Name to a Group

In SNMPv3, each username is associated with a security name. The security name, together with the SNMP engine ID, is included in SNMP messages to ensure messaging security.

Before you assign a security name to a group, first create the security name. For an SNMPv3 client, the security name is the username configured at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level. For SNMPv1 or v2c clients, the security name is the community string configured at the **[edit snmp v3 snmp-community *community-index*]** hierarchy level.

Assigning a security name to a group includes configuring a security model for the group, assigning the security name to the group, and configuring the group.

To assign an SNMP security name to a group:

1. To configure a security model for the group:

```
[edit snmp v3 vacm security-to-group]
user@switch# edit security-model (usm | v1 | v2c)
```

For example, to configure the SNMPv3 user-based security model (USM):

```
[edit snmp v3 vacm security-to-group]
user@switch# edit security-model usm
```

2. To associate the security name with a group:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]
user@switch# edit security-name security-name
```

3. To configure a group of SNMPv3 security names with the same security policy:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name
security-name]
user@switch# edit group group-name
```

#### Related Documentation

- [Creating SNMPv3 Users on page 6609](#)
- [group \(Associating a Security Name\) on page 6715](#)
- [security-model \(Group\) on page 6747](#)
- [security-name \(Community String\) on page 6749](#)
- [security-name \(Security Group\) on page 6750](#)

## Configuring SNMPv3 Traps on a Device Running Junos OS

In SNMPv3, you create traps and informs by configuring the **notify**, **target-address**, and **target-parameters** parameters. Traps are unconfirmed notifications, whereas informs are confirmed notifications. This section describes how to configure SNMP traps. For information about configuring SNMP informs, see [“Configuring SNMP Informs” on page 6614](#).

The target address defines a management application’s address and parameters to be used in sending notifications. Target parameters define the message processing and security parameters that are used in sending notifications to a particular management target. SNMPv3 also lets you define SNMPv1 and SNMPv2c traps.



**NOTE:** When you configure SNMP traps, make sure your configured access privileges allow the traps to be sent. Access privileges are configured at the `[edit snmp v3 vacm access]` and `[edit snmp v3 vacm security-to-group]` hierarchy levels.

To configure SNMP traps, include the following statements at the `[edit snmp v3]` hierarchy level:

```
[edit snmp v3]
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter name {
    oid object-identifier (include | exclude);
  }
```

```
target-address target-address-name {  
    address address;  
    address-mask address-mask;  
    logical-system (SNMP) logical-system;  
    port port-number;  
    routing-instance instance;  
    tag-list tag-list;  
    target-parameters target-parameters-name;  
}  
target-parameters target-parameters-name {  
    notify-filter profile-name;  
    parameters {  
        message-processing-model (v1 | v2c | v3);  
        security-level (authentication | none | privacy);  
        security-model (usm | v1 | v2c);  
        security-name security-name;  
    }  
}
```

**Related  
Documentation**

- *Configuring the SNMPv3 Trap Notification*
- *Configuring the Trap Notification Filter*
- *Configuring the Trap Target Address*
- *Defining and Configuring the Trap Target Parameters*
- [Configuring SNMP Informs on page 6614](#)
- *Configuring the Remote Engine and Remote User*
- *Configuring the Inform Notification Type and Target Address*
- *Complete SNMPv3 Configuration Statements*
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524](#)

## Configuring SNMP Informs

Junos OS supports two types of notifications: traps and informs. With traps, the receiver does not send any acknowledgment when it receives a trap. Therefore, the sender cannot determine if the trap was received. A trap may be lost because a problem occurred during transmission. To increase reliability, an inform is similar to a trap except that the inform is stored and retransmitted at regular intervals until one of these conditions occurs:


- The receiver (target) of the inform returns an acknowledgment to the SNMP agent.
- A specified number of unsuccessful retransmissions have been attempted and the agent discards the inform message.

If the sender never receives a response, the inform can be sent again. Thus, informs are more likely to reach their intended destination than traps are. Informs use the same communications channel as traps (same socket and port) but have different protocol data unit (PDU) types.



Informs are more reliable than traps, but they consume more network, router, and switch resources (see [Figure 227 on page 6615](#)). Unlike a trap, an inform is held in memory until a response is received or the timeout is reached. Also, traps are sent only once, whereas an inform may be retried several times. Use informs when it is important that the SNMP manager receive all notifications. However, if you are more concerned about network traffic, or router and switch memory, use traps.

**Figure 227: Inform Request and Response**



For information about configuring SNMP traps, see “[Configuring SNMPv3 Traps on a Device Running Junos OS](#)” on page 6613.

**Related Documentation**

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 6613](#)
- [Configuring the Remote Engine and Remote User](#)
- [Configuring the Inform Notification Type and Target Address](#)
- [Complete SNMPv3 Configuration Statements](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524](#)

## Configuration Tasks for System Log Messages

- [Junos OS Minimum System Logging Configuration on page 6616](#)
- [Junos OS System Log Configuration Statements on page 6616](#)
- [Adding a Text String to System Log Messages Directed to a Remote Destination on page 6617](#)
- [Directing System Log Messages to a Log File on page 6618](#)
- [Directing System Log Messages to a Remote Machine on page 6619](#)
- [Directing System Log Messages to a User Terminal on page 6620](#)
- [Directing System Log Messages to the Console on page 6620](#)
- [Disabling the System Logging of a Facility on page 6620](#)
- [Displaying a Log File from a Single-Chassis System on page 6621](#)
- [Including Priority Information in System Log Messages on page 6622](#)
- [Including the Year or Millisecond in Timestamps on page 6623](#)
- [Logging Messages in Structured-Data Format on page 6624](#)
- [Interpreting Messages Generated in Structured-Data Format on page 6625](#)

- [Interpreting Messages Generated in Standard Format on page 6628](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 6629](#)
- [Specifying the Facility and Severity of Messages to Include in the Log on page 6631](#)
- [Junos OS System Logging Facilities and Message Severity Levels on page 6633](#)
- [Default Facilities for System Log Messages Directed to a Remote Destination on page 6634](#)
- [Alternate Facilities for System Log Messages Directed to a Remote Destination on page 6635](#)
- [Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination on page 6636](#)
- [Using Regular Expressions to Refine the Set of Logged Messages on page 6637](#)

## Junos OS Minimum System Logging Configuration

To record or view system log messages, you must include the **syslog** statement at the **[edit system]** hierarchy level. Specify at least one destination for the messages, as described in [Table 629 on page 6616](#). For more information about the configuration statements, see *Single-Chassis System Logging Configuration Overview*.

**Table 629: Minimum Configuration Statements for System Logging**

Destination	Minimum Configuration Statements
File	<pre>[edit system syslog] file filename {   facility severity; }</pre>
Terminal session of one, several, or all users	<pre>[edit system syslog] user (username   *) {   facility severity; }</pre>
Router or switch console	<pre>[edit system syslog] console {   facility severity; }</pre>
Remote machine or the other Routing Engine on the router or switch	<pre>[edit system syslog] host (hostname   other-routing-engine) {   facility severity; }</pre>

### Related Documentation

- [Junos OS System Log Overview](#)
- [Overview of Junos OS System Log Messages on page 6560](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)

## Junos OS System Log Configuration Statements

To configure the switch to log system messages, include the **syslog** statement at the **[edit system]** hierarchy level:

```

[edit system]
syslog {
  archive <files number> <size size> <world-readable | no-world-readable>;
  console {
    facility severity;
  }
  file filename {
    facility severity;
    archive <archive-sites (ftp-url <password password>)> <files number> <size size>
      <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
      no-world-readable>;
    explicit-priority;
    match "regular-expression";
    structured-data {
      brief;
    }
  }
  host hostname {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string
    match "regular-expression";
  }
  source-address source-address;
  time-format (year | millisecond | year millisecond);
  user (username | *) {
    facility severity;
    match "regular-expression";
  }
}

```

**Related Documentation**

- [Overview of Junos OS System Log Messages on page 6560](#)

## Adding a Text String to System Log Messages Directed to a Remote Destination

To add a text string to every system log message directed to a remote machine or to the other Routing Engine, include the **log-prefix** statement at the **[edit system syslog host]** hierarchy level:

```

[edit system syslog host (hostname | other-routing-engine)]
facility severity;
log-prefix string;

```

The string can contain any alphanumeric or special character except the equal sign ( = ) and the colon ( : ). It also cannot include the space character; do not enclose the string in quotation marks ( " ") in an attempt to include spaces in it.

The Junos OS system logging utility automatically appends a colon and a space to the specified string when the system log messages are written to the log. The string is inserted after the identifier for the Routing Engine that generated the message.

The following example shows how to add the string M120 to all messages to indicate that the router is an M120 router, and direct the messages to the remote machine hardware-logger.mycompany.com:

```
[edit system syslog]
host hardware-logger.mycompany.com {
  any info;
  log-prefix M120;
}
```

When these configuration statements are included on an M120 router called origin1, a message in the system log on hardware-logger.mycompany.com looks like the following:

```
Mar 9 17:33:23 origin1 M120: mgd[477]: UI_CMDLINE_READ_LINE: user 'root', command 'run
show version'
```

**Related  
Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Specifying Log File Size, Number, and Archiving Properties on page 6629](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)

## Directing System Log Messages to a Log File

To direct system log messages to a file in the `/var/log` directory of the local Routing Engine, include the `file` statement at the `[edit system syslog]` hierarchy level:

```
[edit system syslog]
file filename {
  facility severity;
  archive <archive-sites (ftp-url <password password>) > <files number> <size size>
    <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |
    no-world-readable>;
  explicit-priority;
  match "regular-expression";
  structured-data {
    brief;
  }
}
```

For the list of facilities and severity levels, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6631](#).

To prevent log files from growing too large, the Junos OS system logging utility by default writes messages to a sequence of files of a defined size. By including the `archive` statement, you can configure the number of files, their maximum size, and who can read them, either for all log files or for a certain log file. For more information, see [“Specifying Log File Size, Number, and Archiving Properties” on page 6629](#).

For information about the following statements, see the indicated sections:

- `explicit-priority`—See [“Including Priority Information in System Log Messages” on page 6622](#)
- `match`—See [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6637](#)

- **structured-data**—See *Logging Messages in Structured-Data Format*

#### Related Documentation

- *Single-Chassis System Logging Configuration Overview*
- [Overview of Junos OS System Log Messages on page 6560](#)
- [Logging Messages in Structured-Data Format on page 6624](#)
- *Examples: Configuring System Logging*
- [Examples: Configuring System Logging on page 6565](#)

## Directing System Log Messages to a Remote Machine

To direct system log messages to a remote machine, include the **host** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
host (hostname | other-routing-engine) {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
}
source-address source-address;
```

To direct system log messages to a remote machine, include the **host hostname** statement to specify the remote machine's IP version 4 (IPv4) address or fully qualified hostname. The remote machine must be running the standard **syslogd** utility. We do not recommend directing messages to another Juniper Networks switch. In each system log message directed to the remote machine, the hostname of the local Routing Engine appears after the timestamp to indicate that it is the source for the message.

For the list of logging facilities and severity levels to configure under the **host** statement, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6631](#).

To record facility and severity level information in each message, include the **explicit-priority** statement. For more information, see [“Including Priority Information in System Log Messages” on page 6622](#).

For information about the **match** statement, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6637](#).

When directing messages to remote machines, you can include the **source-address** statement to specify the IP address of the switch that is reported in the messages as their source. In each **host** statement, you can also include the **facility-override** statement to assign an alternative facility and the **log-prefix** statement to add a string to each message.

#### Related Documentation

- [Overview of Single-Chassis System Logging Configuration on page 6561](#)

## Directing System Log Messages to a User Terminal

To direct system log messages to the terminal session of one or more specific users (or all users) when they are logged in to the local Routing Engine, include the **user** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
user (username | *) {
    facility severity;
    match "regular-expression";
}
```

Specify one or more Junos OS usernames, separating multiple values with spaces, or use the asterisk (\*) to indicate all users who are logged in to the local Routing Engine.

For the list of logging facilities and severity levels, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6631](#). For information about the **match** statement, see [“Using Regular Expressions to Refine the Set of Logged Messages” on page 6637](#).

### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Examples: Configuring System Logging](#)
- [Examples: Configuring System Logging on page 6565](#)

## Directing System Log Messages to the Console

To direct system log messages to the console of the local Routing Engine, include the **console** statement at the **[edit system syslog]** hierarchy level:

```
[edit system syslog]
console {
    facility severity;
}
```

For the list of logging facilities and severity levels, see [“Specifying the Facility and Severity of Messages to Include in the Log” on page 6631](#).

### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Examples: Configuring System Logging](#)
- [Examples: Configuring System Logging on page 6565](#)

## Disabling the System Logging of a Facility

To disable the logging of messages that belong to a particular facility, include the **facility none** statement in the configuration. This statement is useful when, for example, you want to log messages that have the same severity level and belong to all but a few facilities. Instead of including a statement for each facility you want to log, you can include

the **any severity** statement and then a **facility none** statement for each facility that you do not want to log. For example, the following logs all messages at the **error** level or higher to the console, except for messages from the **daemon** and **kernel** facilities. Messages from those facilities are logged to the file **>/var/log/internals** instead:

```
[edit system syslog]
console {
  any error;
  daemon none;
  kernel none;
}
file internals {
  daemon info;
  kernel info;
}
```

- Related Documentation**
- [Single-Chassis System Logging Configuration Overview](#)
  - [Overview of Single-Chassis System Logging Configuration on page 6561](#)

## Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysApp1ElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysApp1ElmtRunMemory.5.8.2292)
...
Nov  4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov  4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
```

```
Nov  4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```
user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
SNMP trap: cold start
...
```

#### Related Documentation

- [Interpreting Messages Generated in Standard Format on page 6628](#)
- [Interpreting Messages Generated in Structured-Data Format on page 6625](#)

## Including Priority Information in System Log Messages

The facility and severity level of a message are together referred to as its *priority*. By default, messages logged in the standard Junos OS format do not include information about priority. To include priority information in standard-format messages directed to a file, include the **explicit-priority** statement at the **[edit system syslog file *filename*]** hierarchy level:

```
[edit system syslog file filename]
  facility severity;
  explicit-priority;
```



**NOTE:** Messages logged in structured-data format include priority information by default. If you include the **structured-data** statement at the **[edit system syslog file *filename*]** hierarchy level along with the **explicit-priority** statement, the **explicit-priority** statement is ignored and messages are logged in structured-data format.

For information about the **structured-data** statement, see *Logging Messages in Structured-Data Format*. For information about the contents of a structured-data message, see the *Junos OS System Log Messages Reference*.

To include priority information in messages directed to a remote machine or the other Routing Engine, include the **explicit-priority** statement at the **[edit system syslog host (*hostname* | other-routing-engine)]** hierarchy level:

```
[edit system syslog host (hostname | other-routing-engine)]
  facility severity;
  explicit-priority;
```





**NOTE:** The `other-routing-engine` option does not apply to the QFX Series.

The priority recorded in a message always indicates the original, local facility name. If the **facility-override** statement is included for messages directed to a remote destination, the Junos OS system logging utility still uses the alternative facility name for the messages themselves when directing them to the remote destination. For more information, see [“Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination” on page 6636](#).

When the **explicit-priority** statement is included, the Junos OS logging utility prepends codes for the facility name and severity level to the message tag name, if the message has one:

*FACILITY-severity[-TAG]*

(The tag is a unique identifier assigned to some Junos OS system log messages; for more information, see the *Junos OS System Log Messages Reference*.)

In the following example, the **CHASSISD\_PARSE\_COMPLETE** message belongs to the **daemon** facility and is assigned severity **info** (6):

```
Aug 21 12:36:30 router1 chassisd[522]: %DAEMON-6-CHASSISD_PARSE_COMPLETE:
Using new configuration
```

When the **explicit-priority** statement is not included, the priority does not appear in the message:

```
Aug 21 12:36:30 router1 chassisd[522]: CHASSISD_PARSE_COMPLETE: Using new
configuration
```

For more information about message formatting, see the *Junos OS System Log Messages Reference*.

#### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Examples: Configuring System Logging](#)

## Including the Year or Millisecond in Timestamps

By default, the timestamp recorded in a standard-format system log message specifies the month, date, hour, minute, and second when the message was logged, as in the following example:

```
Aug 21 15:36:30
```

To include the year, the millisecond, or both, in the timestamp, include the **time-format** statement at the `[edit system syslog]` hierarchy level:

```
[edit system syslog]
time-format (year | millisecond | year millisecond);
```

The modified timestamp is used in messages directed to each destination configured by a **file**, **console**, or **user** statement at the **[edit system syslog]** hierarchy level, but not to destinations configured by a **host** statement.

The following example illustrates the format for a timestamp that includes both the millisecond (401) and the year (2010):

```
Aug 21 15:36:30.401 2010
```



**NOTE:** By default, messages logged in structured-data format include the year and millisecond. If you include the **structured-data** statement at the **[edit system syslog file *filename*]** hierarchy level along with the **time-format** statement, the **time-format** statement is ignored and messages are logged in structured-data format.

For information about the **structured-data** statement, see [“Logging Messages in Structured-Data Format” on page 6624](#). For information about interpreting messages in a structured-data format, see [“Interpreting Messages Generated in Structured-Data Format” on page 6625](#).

---

## Logging Messages in Structured-Data Format

You can log messages to a file in structured-data format instead of the standard Junos OS format. The structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.

The structured-data format complies with Internet draft **draft-ietf-syslog-protocol-21.txt**. The draft establishes a standard message format regardless of the source or transport protocol for logged messages.

To output messages to a file in structured-data format, include the **structured-data** statement at the **[edit system syslog file *filename*]** hierarchy level:

```
[edit system syslog file filename]  
  facility severity;  
  structured-data {  
    brief;  
  }
```

The optional **brief** statement suppresses the English-language text that appears by default at the end of a message to describe the error or event. For information about the fields in a structured-data-format message, see [“Interpreting Messages Generated in Structured-Data Format” on page 6625](#).

The structured format is used for all messages logged to the file that are generated by a Junos OS process or software library.



**NOTE:** If you include either or both of the **explicit-priority** and **time-format** statements along with the **structured-data** statement, they are ignored. These statements apply to the standard Junos OS system log format, not to structured-data format.

## Interpreting Messages Generated in Structured-Data Format

By default, Junos OS processes and software libraries write messages to the system log file in structured-data format. For information about the **structured-data** statement, see *Logging Messages in Structured-Data Format*.

Structured-format makes it easier for automated applications to extract information from the message. In particular, the standardized format for reporting the value of variables (elements in the English-language message that vary depending on the circumstances that triggered the message) makes it easy for an application to extract those values.

The structured-data format for a message includes the following fields (which appear here on two lines only for legibility):

```
<priority code>version timestamp hostname process processID TAG [junos@2636.platform
variable-value-pairs] message-text
```

Table 630 on page 6625 describes the fields. If the system logging utility cannot determine the value in a particular field, a hyphen ( - ) appears instead.

**Table 630: Fields in Structured-Data Messages**

Field	Description	Examples
<b>&lt;priority code&gt;</b>	Number that indicates the facility and severity of a message. It is calculated by multiplying the facility number by 8 and then adding the numerical value of the severity. For a mapping of the numerical codes to facility and severity, see “Specifying the Facility and Severity of Messages to Include in the Log” on page 6631.	<165> for a message from the <b>pfe</b> facility (facility=20) with severity <b>notice</b> (severity=5).
<b>version</b>	Version of the Internet Engineering Task Force (IETF) system logging protocol specification.	1 for the initial version

Table 630: Fields in Structured-Data Messages (*continued*)

Field	Description	Examples
<i>timestamp</i>	Time when the message was generated, in one of two representations: <ul style="list-style-type: none"> <li><i>YYYY-MM-DDTHH:MM:SS.MSZ</i> is the year, month, day, hour, minute, second and millisecond in Universal Coordinated Time (UTC)</li> <li><i>YYYY-MM-DDTHH:MM:SS.MS+/-HH:MM</i> is the year, month, day, hour, minute, second and millisecond in local time; the hour and minute that follows the plus sign (+) or minus sign (-) is the offset of the local time zone from UTC</li> </ul>	2007-02-15T09:17:15.719Z is 9:17 AM UTC on 15 February 2007. 2007-02-15T01:17:15.719-08:00 is the same timestamp expressed as Pacific Standard Time in the United States.
<i>hostname</i>	Name of the host that originally generated the message.	switch1
<i>process</i>	Name of the Junos OS process that generated the message.	mgd
<i>processID</i>	UNIX process ID (PID) of the Junos process that generated the message.	3046
<i>TAG</i>	Junos OS system log message tag, which uniquely identifies the message.	UI_DBASE_LOGOUT_EVENT
<i>junos@2636.platform</i>	An identifier for the type of hardware platform that generated the message. The <i>junos@2636</i> prefix indicates that the platform runs the Junos OS. It is followed by a dot-separated numerical identifier for the platform type.	junos@2636.1.1.1.2.18
<i>variable-value-pairs</i>	A variable-value pair for each element in the <i>message-text</i> string that varies depending on the circumstances that triggered the message. Each pair appears in the format <i>variable</i> = " <i>value</i> ".	username="regress"
<i>message-text</i>	English-language description of the event or error (omitted if the brief statement is included at the [edit system syslog file <i>filename</i> structured-data] hierarchy level).	User 'regress' exiting configuration mode

By default, the structured-data version of a message includes English text at the end, as in the following example (which appears on multiple lines only for legibility):

```
<165>1 2007-02-15T09:17:15.719Z router1 mgd 3046 UI_DBASE_LOGOUT_EVENT
[junos@2636.1.1.1.2.18 username="regress"] User 'regress' exiting configuration mode
```

When the brief statement is included at the [edit system syslog file *filename* structured-data ] hierarchy level, the English text is omitted, as in this example:

```
<165>1 2007-02-15T09:17:15.719Z router1 mgd 3046 UI_DBASE_LOGOUT_EVENT
[junos@2636.1.1.1.2.18 username="regress"]
```

Table 631 on page 6627 maps the codes that appear in the **priority-code** field to facility and severity level.



**NOTE:** Not all of the facilities and severities listed in Table 631 on page 6627 can be included in statements at the [edit system syslog] hierarchy level (some are used by internal processes). For a list of the facilities and severity levels that can be included in the configuration, see “[Specifying the Facility and Severity of Messages to Include in the Log](#)” on page 6631.

Table 631: Facility and Severity Codes in the priority-code Field

Facility (number)	Severity emergency	alert	critical	error	warning	notice	info	debug
kernel (0)	1	1	2	3	4	5	6	7
user (1)	8	9	10	11	12	13	14	15
mail (2)	16	17	18	19	20	21	22	23
daemon (3)	24	25	26	27	28	29	30	31
authorization (4)	32	33	34	35	36	37	38	39
syslog (5)	40	41	42	43	44	45	46	47
printer (6)	48	49	50	51	52	53	54	55
news (7)	56	57	58	59	60	61	62	63
uucp (8)	64	65	66	67	68	69	70	71
clock (9)	72	73	74	75	76	77	78	79
authorization-private (10)	80	81	82	83	84	85	86	87
ftp (11)	88	89	90	91	92	93	94	95
ntp (12)	96	97	98	99	100	101	102	103
security (13)	104	105	106	107	108	109	110	111
console (14)	112	113	114	115	116	117	118	119
local0 (16)	128	129	130	131	132	133	134	135
dfc (17)	136	137	138	139	140	141	142	143

Table 631: Facility and Severity Codes in the priority-code Field (*continued*)

Facility (number)	Severity emergency	alert	critical	error	warning	notice	info	debug
local2 (18)	144	145	146	147	148	149	150	151
firewall (19)	152	153	154	155	156	157	158	159
pfe (20)	160	161	162	163	164	165	166	167
conflict-log (21)	168	169	170	171	172	173	174	175
change-log (22)	176	177	178	179	180	181	182	183
interactive-commands (23)	184	185	186	187	188	189	190	191

## Interpreting Messages Generated in Standard Format

The syntax of a standard-format message generated by a Junos OS process or subroutine library depends on whether it includes priority information:

- When the **explicit-priority** statement is included at the `[edit system syslog file filename]` or `[edit system syslog host hostname]` hierarchy level, a system log message has the following syntax:

```
timestamp message-source: %facility-severity-TAG: message-text
```

- When directed to the console or to users, or when the **explicit-priority** statement is not included for files or remote hosts, a system log message has the following syntax:

```
timestamp message-source: TAG: message-text
```

Table 632 on page 6628 describes the message fields.

Table 632: Fields in Standard-Format Messages

Field	Description
<i>timestamp</i>	Time at which the message was logged.
<i>message-source</i>	Identifier of the process or component that generated the message and the routing platform on which the message was logged. This field includes two or more subfields: hostname, process and process ID (PID). If the process does not report its PID, the PID is not displayed. The message source subfields are displayed in the following format:  <i>hostname process[process-ID]</i>
<i>facility</i>	Code that specifies the facility to which the system log message belongs. For a mapping of codes to facility names, see Table: <b>Facility Codes Reported in Priority Information</b> in “Including Priority Information in System Log Messages” on page 6622.

Table 632: Fields in Standard-Format Messages (*continued*)

Field	Description
<b>severity</b>	Numerical code that represents the severity level assigned to the system log message. For a mapping of codes to severity names, see Table: <b>Numerical Codes for Severity Levels Reported in Priority Information</b> in “Including Priority Information in System Log Messages” on page 6622.
<b>TAG</b>	Text string that uniquely identifies the message, in all uppercase letters and using the underscore (_) to separate words. The tag name begins with a prefix that indicates the generating software process or library. The entries in this reference are ordered alphabetically by this prefix.  Not all processes on a routing platform use tags, so this field does not always appear.
<b>message-text</b>	Text of the message.

## Specifying Log File Size, Number, and Archiving Properties

To prevent log files from growing too large, by default the Junos system logging utility writes messages to a sequence of files of a defined size. The files in the sequence are referred to as *archive* files to distinguish them from the *active* file to which messages are currently being written. The default maximum size depends on the platform type:

- 128 kilobytes (KB) for EX Series switches and J Series routers
- 1 megabyte (MB) for M Series, MX Series, and T Series routers
- 10 MB for TX Matrix or TX Matrix Plus routers
- 1 MB for the QFX Series

When an active log file called **logfile** reaches the maximum size, the logging utility closes the file, compresses it, and names the compressed archive file **logfile.0.gz**. The logging utility then opens and writes to a new active file called **logfile**. This process is also known as file rotation. When the new **logfile** reaches the configured maximum size, **logfile.0.gz** is renamed **logfile.1.gz**, and the new **logfile** is closed, compressed, and renamed **logfile.0.gz**. By default, the logging utility creates up to 10 archive files in this manner. When the maximum number of archive files is reached and when the size of the active file reaches the configured maximum size, the contents of the last archived file are overwritten by the current active file. The logging utility by default also limits the users who can read log files to the **root** user and users who have the Junos OS **maintenance** permission.

Junos OS provides a configuration statement **log-rotate-frequency** that configures the system log file rotation frequency by configuring the time interval for checking the log file size. The frequency can be set to a value of 1 minute through 59 minutes. The default frequency is 15 minutes.

To configure the log rotation frequency, include the **log-rotate-frequency** statement at the **[edit system syslog]** hierarchy level.

You can include the **archive** statement to change the maximum size of each file, how many archive files are created, and who can read log files.

To configure values that apply to all log files, include the **archive** statement at the **[edit system syslog]** hierarchy level:

```
archive <files number> <size size> <world-readable | no-world-readable>;
```

To configure values that apply to a specific log file, include the **archive** statement at the **[edit system syslog file *filename*]** hierarchy level:

```
archive <archive-sites (ftp-url <password password>)> <files number> <size size>  
<start-time "YYYY-MM-DD.hh:mm"> <transfer-interval minutes> <world-readable |  
no-world-readable>;
```

**archive-sites *site-name*** specifies a list of archive sites that you want to use for storing files. The ***site-name*** value is any valid FTP URL to a destination. If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the specified log filename. For information about how to specify valid FTP URLs, see [“Format for Specifying Filenames and URLs in Junos OS CLI Commands” on page 42](#).

**binary-data** Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems). To restore the default setting, include the **no-binary-data** statement.

**files *number*** specifies the number of files to create before the oldest file is overwritten. The value can be from 1 through 1000.

**size *size*** specifies the maximum size of each file. The value can be from 64 KB (64k) through 1 gigabyte (1g); to represent megabytes, use the letter **m** after the integer. There is no space between the digits and the **k**, **m**, or **g** units letter.

**start-time "YYYY-MM-DD.hh:mm"** defines the date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

**transfer-interval *interval*** defines the amount of time the current log file remains open (even if it has not reached the maximum possible size) and receives new statistics before it is closed and transferred to an archive site. This interval value can be from 5 through 2880 minutes.

**world-readable** enables all users to read log files. To restore the default permissions, include the **no-world-readable** statement.

#### Related Documentation

- [Single-Chassis System Logging Configuration Overview](#)
- [Examples: Configuring System Logging](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)



## Specifying the Facility and Severity of Messages to Include in the Log

Each system log message belongs to a facility, which groups together messages that either are generated by the same source (such as a software process) or concern a similar condition or activity (such as authentication attempts). Each message is also preassigned a *severity level*, which indicates how seriously the triggering event affects routing platform functions.

When you configure logging for a facility and destination, you specify a severity level for each facility. Messages from the facility that are rated at that level and higher are logged to the following destination:

```
[edit system syslog]
(console | file filename | host destination | user username) {
  facility severity ;
}
```

For more information about the destinations, see [“Directing System Log Messages to a User Terminal” on page 6620](#), and, [“Directing System Log Messages to the Console” on page 6620](#).

To log messages belonging to more than one facility to a particular destination, specify each facility and associated severity as a separate statement within the set of statements for the destination.

[Table 633 on page 6631](#) lists the Junos system logging facilities that you can specify in configuration statements at the `[edit system syslog]` hierarchy level.

**Table 633: Junos OS System Logging Facilities**

Facility	Type of Event or Error
<b>any</b>	All (messages from all facilities)
<b>authorization</b>	Authentication and authorization attempts
<b>change-log</b>	Changes to the Junos OS configuration
<b>conflict-log</b>	Specified configuration is invalid on the router type
<b>daemon</b>	Actions performed or errors encountered by system processes
<b>dfc</b>	Events related to dynamic flow capture
<b>firewall</b>	Packet filtering actions performed by a firewall filter
<b>ftp</b>	Actions performed or errors encountered by the FTP process
<b>interactive-commands</b>	Commands issued at the Junos OS command-line interface (CLI) prompt or by a client application such as a Junos XML protocol or NETCONF XML client

Table 633: Junos OS System Logging Facilities (*continued*)

Facility	Type of Event or Error
<b>kernel</b>	Actions performed or errors encountered by the Junos OS kernel
<b>pfe</b>	Actions performed or errors encountered by the Packet Forwarding Engine
<b>user</b>	Actions performed or errors encountered by user-space processes

Table 634 on page 6632 lists the severity levels that you can specify in configuration statements at the **[edit system syslog]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Unlike the other severity levels, the **none** level disables logging of a facility instead of indicating how seriously a triggering event affects routing functions. For more information, see “Disabling the System Logging of a Facility” on page 6620.

Table 634: System Log Message Severity Levels

Severity Level	Description
<b>any</b>	Includes all severity levels
<b>none</b>	Disables logging of the associated facility to a destination
<b>emergency</b>	System panic or other condition that causes the router to stop functioning
<b>alert</b>	Conditions that require immediate correction, such as a corrupted system database
<b>critical</b>	Critical conditions, such as hard errors
<b>error</b>	Error conditions that generally have less serious consequences than errors at the emergency, alert, and critical levels
<b>warning</b>	Conditions that warrant monitoring
<b>notice</b>	Conditions that are not errors but might warrant special handling
<b>info</b>	Events or nonerror conditions of interest

#### Related Documentation

- [Junos OS System Logging Facilities and Message Severity Levels on page 6633](#)
- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Examples: Configuring System Logging](#)

## Junos OS System Logging Facilities and Message Severity Levels

Table 633 on page 6631 lists the Junos system logging facilities that you can specify in configuration statements at the **[edit system syslog]** hierarchy level.

**Table 635: Junos OS System Logging Facilities**

Facility	Type of Event or Error
<b>any</b>	All (messages from all facilities)
<b>authorization</b>	Authentication and authorization attempts
<b>change-log</b>	Changes to the Junos OS configuration
<b>conflict-log</b>	Specified configuration is invalid on the router type
<b>daemon</b>	Actions performed or errors encountered by system processes
<b>dfc</b>	Events related to dynamic flow capture
<b>firewall</b>	Packet filtering actions performed by a firewall filter
<b>ftp</b>	Actions performed or errors encountered by the FTP process
<b>interactive-commands</b>	Commands issued at the Junos OS command-line interface (CLI) prompt or by a client application such as a Junos XML protocol or NETCONF XML client
<b>kernel</b>	Actions performed or errors encountered by the Junos OS kernel
<b>pfe</b>	Actions performed or errors encountered by the Packet Forwarding Engine
<b>user</b>	Actions performed or errors encountered by user-space processes

Table 634 on page 6632 lists the severity levels that you can specify in configuration statements at the **[edit system syslog]** hierarchy level. The levels from **emergency** through **info** are in order from highest severity (greatest effect on functioning) to lowest.

Unlike the other severity levels, the **none** level disables logging of a facility instead of indicating how seriously a triggering event affects routing functions. For more information, see “Disabling the System Logging of a Facility” on page 6620.

**Table 636: System Log Message Severity Levels**

Severity Level	Description
<b>any</b>	Includes all severity levels
<b>none</b>	Disables logging of the associated facility to a destination

**Table 636: System Log Message Severity Levels (*continued*)**

Severity Level	Description
<b>emergency</b>	System panic or other condition that causes the router to stop functioning
<b>alert</b>	Conditions that require immediate correction, such as a corrupted system database
<b>critical</b>	Critical conditions, such as hard errors
<b>error</b>	Error conditions that generally have less serious consequences than errors at the emergency, alert, and critical levels
<b>warning</b>	Conditions that warrant monitoring
<b>notice</b>	Conditions that are not errors but might warrant special handling
<b>info</b>	Events or nonerror conditions of interest

**Related Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Examples: Configuring System Logging](#)

**Default Facilities for System Log Messages Directed to a Remote Destination**

Table 637 on page 6634 lists the default alternative facility name next to the Junos OS-specific facility name for which it is used. For facilities that are not listed, the default alternative name is the same as the local facility name.

**Table 637: Default Facilities for Messages Directed to a Remote Destination**

Junos OS-specific Local Facility	Default Facility When Directed to Remote Destination
change-log	local6
conflict-log	local5
dfc	local1
firewall	local3
interactive-commands	local7
pfe	local4

- Related Documentation**
- [Single-Chassis System Logging Configuration Overview](#)
  - [Overview of Single-Chassis System Logging Configuration on page 6561](#)

## Alternate Facilities for System Log Messages Directed to a Remote Destination

Table 638 on page 6635 lists the facilities that you can specify in the **facility-override** statement.

**Table 638: Facilities for the facility-override Statement**

Facility	Description
<b>authorization</b>	Authentication and authorization attempts
<b>daemon</b>	Actions performed or errors encountered by system processes
<b>ftp</b>	Actions performed or errors encountered by the FTP process
<b>kernel</b>	Actions performed or errors encountered by the Junos OS kernel
<b>local0</b>	Local facility number 0
<b>local1</b>	Local facility number 1
<b>local2</b>	Local facility number 2
<b>local3</b>	Local facility number 3
<b>local4</b>	Local facility number 4
<b>local5</b>	Local facility number 5
<b>local6</b>	Local facility number 6
<b>local7</b>	Local facility number 7
<b>user</b>	Actions performed or errors encountered by user-space processes

We do not recommend including the **facility-override** statement at the **[edit system syslog host other-routing-engine]** hierarchy level. It is not necessary to use alternative facility names when directing messages to the other Routing Engine, because its Junos OS system logging utility can interpret the Junos OS-specific names.

- Related Documentation**
- [Examples: Assigning an Alternative Facility to System Log Messages Directed to a Remote Destination](#)
  - [Single-Chassis System Logging Configuration Overview](#)
  - [Overview of Single-Chassis System Logging Configuration on page 6561](#)

## Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination

Some facilities assigned to messages logged on the local router or switch have Junos OS-specific names (see [“Junos OS System Logging Facilities and Message Severity Levels” on page 6633](#)). In the recommended configuration, a remote machine designated at the `[edit system syslog host hostname]` hierarchy level is not a Juniper Networks router or switch, so its syslogd utility cannot interpret the Junos OS-specific names. To enable the standard syslogd utility to handle messages from these facilities when messages are directed to a remote machine, a standard **localX** facility name is used instead of the Junos OS-specific facility name.

[Table 637 on page 6634](#) lists the default alternative facility name next to the Junos OS-specific facility name it is used for.

The syslogd utility on a remote machine handles all messages that belong to a facility in the same way, regardless of the source of the message (the Juniper Networks router or switch or the remote machine itself). For example, the following statements in the configuration of the router called **local-router** direct messages from the **authorization** facility to the remote machine *monitor.mycompany.com*:

```
[edit system syslog]
host monitor.mycompany.com {
  authorization info;
}
```

The default alternative facility for the local **authorization** facility is also **authorization**. If the syslogd utility on **monitor** is configured to write messages belonging to the **authorization** facility to the file `/var/log/auth-attempts`, then the file contains the messages generated when users log in to **local-router** and the messages generated when users log in to **monitor**. Although the name of the source machine appears in each system log message, the mixing of messages from multiple machines can make it more difficult to analyze the contents of the **auth-attempts** file.

To make it easier to separate the messages from each source, you can assign an alternative facility to all messages generated on **local-router** when they are directed to **monitor**. You can then configure the syslogd utility on **monitor** to write messages with the alternative facility to a different file from messages generated on **monitor** itself.

To change the facility used for all messages directed to a remote machine, include the **facility-override** statement at the `[edit system syslog host hostname]` hierarchy level:

```
[edit system syslog host hostname]
  facility severity;
  facility-override facility;
```

In general, it makes sense to specify an alternative facility that is not already in use on the remote machine, such as one of the **localX** facilities. On the remote machine, you must also configure the syslogd utility to handle the messages in the desired manner.

[Table 638 on page 6635](#) lists the facilities that you can specify in the **facility-override** statement.

We do not recommend including the **facility-override** statement at the **[edit system syslog host other-routing-engine]** hierarchy level. It is not necessary to use alternative facility names when directing messages to the other Routing Engine, because its Junos OS system logging utility can interpret the Junos OS-specific names.

The following example shows how to log all messages generated on the local router at the error level or higher to the local0 facility on the remote machine called `monitor.mycompany.com`:

```
[edit system syslog]
host monitor.mycompany.com {
  any error;
  facility-override local0;
}
```

The following example shows how to configure routers located in California and routers located in New York to send messages to a single remote machine called `central-logger.mycompany.com`. The messages from California are assigned to alternative facility `local0` and the messages from New York are assigned to alternative facility `local2`.

- Configure California routers to aggregate messages in the `local0` facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local0;
}
```

- Configure New York routers to aggregate messages in the `local2` facility:

```
[edit system syslog]
host central-logger.mycompany.com {
  change-log info;
  facility-override local2;
}
```

On `central-logger`, you can then configure the system logging utility to write messages from the `local0` facility to the file **change-log** and the messages from the `local2` facility to the file **new-york-config**.

#### Related Documentation

- [Table 637 on page 6634](#)
- [Alternate Facilities for System Log Messages Directed to a Remote Destination on page 6635](#)
- *Examples: Assigning an Alternative Facility to System Log Messages Directed to a Remote Destination*
- [Examples: Assigning an Alternative Facility on page 6567](#)

## Using Regular Expressions to Refine the Set of Logged Messages

The predefined facilities group together related messages, but you can also use regular expression matching to specify more exactly which messages from a facility are logged to a file, a user terminal, or a remote destination.

To specify the text string that must (or must not) appear in a message for the message to be logged to a destination, include the **match** statement and specify the regular expression which the text string must match:

```
match "regular-expression";
```

You can include this statement at the following hierarchy levels:

- **[edit system syslog file *filename*]** (for a file)
- **[edit system syslog user (*username* | \*)]** (for a specific user session or for all user sessions on a terminal)
- **[edit system syslog host (*hostname* | other-routing-engine)]** (for a remote destination)

In specifying the regular expression, use the notation defined in POSIX Standard 1003.2 for extended (modern) UNIX regular expressions. Explaining regular expression syntax is beyond the scope of this document, but POSIX standards are available from the Institute of Electrical and Electronics Engineers (IEEE, <http://www.ieee.org>).

Table 639 on page 6638 specifies which character or characters are matched by some of the regular expression operators that you can use in the match statement. In the descriptions, the term *term* refers to either a single alphanumeric character or a set of characters enclosed in square brackets, parentheses, or braces.



**NOTE:** The match statement is not case-sensitive.

**Table 639: Regular Expression Operators for the match Statement**

Operator	Matches
. (period)	One instance of any character except the space.
* (asterisk)	Zero or more instances of the immediately preceding term.
+ (plus sign)	One or more instances of the immediately preceding term.
? (question mark)	Zero or one instance of the immediately preceding term.
(pipe)	One of the terms that appears on either side of the pipe operator.
! (exclamation point)	Any string except the one specified by the expression, when the exclamation point appears at the start of the expression. Use of the exclamation point is Junos OS-specific.
^ (caret)	Start of a line, when the caret appears outside square brackets.  One instance of any character that does not follow it within square brackets, when the caret is the first character inside square brackets.
\$ (dollar sign)	End of a line.



**Table 639: Regular Expression Operators for the match Statement (*continued*)**

Operator	Matches
[ ] (paired square brackets)	One instance of one of the enclosed alphanumeric characters. To indicate a range of characters, use a hyphen ( - ) to separate the beginning and ending characters of the range. For example, [a-z0-9] matches any letter or number.
( ) (paired parentheses)	One instance of the evaluated value of the enclosed term. Parentheses are used to indicate the order of evaluation in the regular expression.

**Using Regular Expressions** Filter messages that belong to the **interactive-commands** facility, directing those that include the string **configure** to the terminal of the root user:

```
[edit system syslog]
user root {
  interactive-commands any;
  match ".*configure.*";
}
```

Messages like the following appear on the **root** user's terminal when a user issues a **configure** command to enter configuration mode:

```
timestamp router-name mgd[PID]: UI_CMDLINE_READ_LINE: User 'user', command
'configure private'
```

Filter messages that belong to the **daemon** facility and have a severity of **error** or higher, directing them to the file **/var/log/process-errors**. Omit messages generated by the SNMP process (snmpd), instead directing them to the file **/var/log/snmpd-errors**:

```
[edit system syslog]
file process-errors {
  daemon error;
  match "!(.*snmpd.*)";
}
file snmpd-errors {
  daemon error;
  match ".*snmpd.*";
}
```

**Related Documentation**

- [Single-Chassis System Logging Configuration Overview](#)
- [Overview of Single-Chassis System Logging Configuration on page 6561](#)
- [Examples: Configuring System Logging](#)
- [Examples: Configuring System Logging on page 6565](#)

## Configuration Statements for Network Management

- [connection-limit on page 6641](#)
- [destination-override on page 6642](#)

- [no-remote-trace on page 6642](#)
- [protocol-version on page 6643](#)
- [rate-limit on page 6644](#)
- [ssh on page 6645](#)
- [telnet on page 6646](#)
- [tracing on page 6647](#)

## connection-limit

<b>Syntax</b>	<code>connection-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	<code>[edit system services finger],</code> <code>[edit system services ftp],</code> <code>[edit system services netconf ssh],</code> <code>[edit system services ssh],</code> <code>[edit system services telnet],</code> <code>[edit system services xnm-clear-text],</code> <code>[edit system services xnm-ssl]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the maximum number of connections sessions for each type of system services (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) per protocol (either IPv6 or IPv4).
<b>Options</b>	<p><b>limit</b>—(Optional) Maximum number of established connections per protocol (either IPv6 or IPv4).</p> <p><b>Range:</b> 1 through 250</p> <p><b>Default:</b> 75</p>



**NOTE:** The actual number of maximum connections depends on the availability of system resources, and might be fewer than the configured `connection-limit` value if the system resources are limited.

<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</a></li> <li>• <a href="#">Configuring DTCP-over-SSH Service for the Flow-Tap Application</a></li> <li>• <a href="#">Configuring Finger Service for Remote Access to the Router</a></li> <li>• <a href="#">Configuring FTP Service for Remote Access to the Router or Switch</a></li> <li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1361</a></li> <li>• <a href="#">Configuring Telnet Service for Remote Access to a Router or Switch</a></li> </ul>

## destination-override

---

<b>Syntax</b>	<code>destination-override {   syslog host <i>ip-address</i>; }</code>
<b>Hierarchy Level</b>	[edit system tracing]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Override the system-wide configuration of the switch at the <b>[edit system tracing]</b> hierarchy level. This statement has no effect if system tracing is not configured.
<b>Options</b>	<b>syslog</b> —System process log files to send to the remote tracing host. <ul style="list-style-type: none"><li>• <b>syslog</b>—System process log files to send to the remote tracing host.</li><li>• <b>host <i>ip-address</i></b>—IP address to which to send tracing information.</li></ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Tracing and Logging Operations on page 6468</a></li><li>• <a href="#">tracing on page 330</a></li></ul>

## no-remote-trace

---

<b>Syntax</b>	<code>no-remote-trace</code>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the switch to disable remote tracing after remote tracing has been enabled.
<b>Default</b>	Remote tracing is disabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">tracing on page 330</a></li></ul>

---

## protocol-version

---

<b>Syntax</b>	<code>protocol-version <i>version</i>;</code>
<b>Hierarchy Level</b>	[edit system services ssh]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the secure shell (SSH) protocol version.
<b>Default</b>	<b>v2</b> —SSH protocol version 2 is the default, introduced in Junos OS Release 11.4.
<b>Options</b>	<b><i>version</i></b> —SSH protocol version: <b>v1</b> , <b>v2</b> , or both.
<b>Required Privilege Level</b>	<b>admin</b> —To view this statement in the configuration. <b>admin-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the SSH Protocol Version on page 1363</a></li></ul>

## rate-limit

---

<b>Syntax</b>	<code>rate-limit <i>limit</i>;</code>
<b>Hierarchy Level</b>	[edit system services finger], [edit system services ftp], [edit system services netconf ssh], [edit system services ssh], [edit system services telnet], [edit system services xnm-clear-text], [edit system services xnm-ssl]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the maximum number of connections attempts per protocol (either IPv6 or IPv4) on an access service.
<b>Default</b>	150 connections
<b>Options</b>	<b>rate-limit <i>limit</i></b> —(Optional) Maximum number of connection attempts allowed per minute, per IP protocol (either IPv4 or IPv6). <b>Range:</b> 1 through 250 <b>Default:</b> 150
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring clear-text or SSL Service for Junos XML Protocol Client Applications</i></li></ul>

## ssh

<b>Syntax</b>	<pre>ssh {   ciphers [ <i>cipher-1 cipher-2 cipher-3 ...</i>];   client-alive-count-max <i>seconds</i>;   client-alive-interval <i>seconds</i>;   connection-limit <i>limit</i>;   hostkey-algorithm &lt;<i>algorithm</i> no-<i>algorithm</i>&gt;;   key-exchange &lt;<i>algorithm</i>&gt;;   macs &lt;<i>algorithm</i>&gt;;   max-sessions-per-connection &lt;<i>number</i>&gt;;   no-passwords;   no-tcp-forwarding;   protocol-version [<i>v1 v2</i>];   rate-limit <i>limit</i>;   root-login (<i>allow</i>   <i>deny</i>   <i>deny-password</i>); }</pre>
<b>Hierarchy Level</b>	[edit system services]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>client-alive-interval</b> and <b>client-alive-max-count</b> statements introduced in Junos OS Release 12.2.</p> <p><b>no-passwords</b> statement introduced in Junos OS Release 13.3.</p>
<b>Description</b>	<p>Allow SSH requests from remote systems to the local router or switch.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring SSH Service for Remote Access to the Router or Switch on page 1361</a></li> </ul>


## telnet

---

<b>Syntax</b>	<pre>telnet {     connection-limit limit;     rate-limit limit; }</pre>
<b>Hierarchy Level</b>	[edit system services]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Provide Telnet connections from remote systems to the local router or switch.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Telnet Service for Remote Access to a Router or Switch</i></li></ul>



## tracing

<b>Syntax</b>	tracing { destination-override syslog host <i>ip-address</i> ; }
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the switch to enable remote tracing to a specified host IP address.
<div>  <b>NOTE:</b> The tracing statement is not supported on the QFX3000 QFabric system.         </div>	
<p>The following processes are supported:</p> <ul style="list-style-type: none"> <li>• <b>chassisd</b>—Chassis-control process</li> <li>• <b>eventd</b>—Event-processing process</li> <li>• <b>cosd</b>—Class-of-service process</li> </ul> <p>If you enabled remote tracing but wish to disable it for specific processes on the switch, use the <b>no-remote-trace</b> statement at the <b>[edit system process-name traceoptions]</b> hierarchy level.</p>	
<b>Default</b>	Remote tracing is disabled by default.
<b>Options</b>	<b>destination-override syslog host <i>ip-address</i></b> —Overrides the global configuration for system tracing and has no effect if the <b>tracing</b> statement is not configured.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Tracing and Logging Operations on page 6468</a></li> <li>• <a href="#">destination-override on page 271</a></li> </ul>

## Configuration Statements for Automation

- [allow-transients on page 6648](#)
- [apply-macro on page 6649](#)
- [checksum on page 6650](#)
- [command on page 6651](#)
- [commit on page 6652](#)

- [description on page 6653](#)
- [direct-access on page 6653](#)
- [file \(Commit Scripts\) on page 6654](#)
- [file \(Op Scripts\) on page 6655](#)
- [no-allow-url on page 6656](#)
- [op on page 6657](#)
- [optional on page 6658](#)
- [refresh \(Commit Scripts\) on page 6659](#)
- [refresh \(Op Scripts\) on page 6660](#)
- [refresh-from \(Commit Scripts\) on page 6661](#)
- [refresh-from \(Op Scripts\) on page 6662](#)
- [scripts on page 6663](#)
- [source \(Commit Scripts\) on page 6665](#)
- [source \(Op Scripts\) on page 6666](#)

---

## allow-transients

---

<b>Syntax</b>	allow-transients;
<b>Hierarchy Level</b>	[edit system scripts commit]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS commit scripts, enable transient configuration changes to be committed.
<b>Default</b>	Transient changes are disabled by default. If you do not include the <b>allow-transients</b> statement, and an enabled script generates transient changes, the command-line interface (CLI) generates an error message and the commit operation fails.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Generating a Persistent or Transient Change</i></li><li>• <i>Creating a Macro to Read the Custom Syntax and Generate Related Configuration Statements</i></li></ul>

## apply-macro

---

<b>Syntax</b>	<pre>apply-macro <i>apply-macro-name</i> {     <i>parameter-name parameter-value</i>; }</pre>
<b>Hierarchy Level</b>	All hierarchy levels
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	<p>With commit script macros, use custom syntax in your configuration.</p> <p>Macros work by locating <b>apply-macro</b> statements that you include in the candidate configuration and using the values specified in the <b>apply-macro</b> statement as parameters to a set of instructions (the macro) defined in a commit script. The commit script alters your configuration from one that contains custom syntax into a full configuration containing standard Junos OS statements.</p> <p>In effect, your custom configuration syntax serves a dual purpose. The syntax allows you to simplify your configuration tasks, and it provides data (or <i>hooks</i>) that are used by commit script macros.</p> <p>You can include the <b>apply-macro</b> statement at any level of the configuration hierarchy. You can include multiple <b>apply-macro</b> statements at each level of the configuration hierarchy; however, each must have a unique name.</p>
<b>Options</b>	<p><b><i>apply-macro-name</i></b>—Name of the <b>apply-macro</b> statement.</p> <p><b><i>parameter-name</i></b>—One or more parameters. Parameters can be any text you want to include in your configuration.</p> <p><b><i>parameter-value</i></b>—A value that corresponds to the parameter name. Parameter values can be any text you want to include in your configuration.</p>
<b>Required Privilege Level</b>	configure—To enter configuration mode; other required privilege levels depend on where the statement is located in the configuration hierarchy.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Overview of Creating Custom Configuration Syntax with Macros</i></li> </ul>

## checksum

---

<b>Syntax</b>	<code>checksum (md5   sha-256   sha1) hash;</code>
<b>Hierarchy Level</b>	[edit event-options event-script file <i>filename</i> ], [edit system <a href="#">scripts commit file filename</a> ], [edit system <a href="#">scripts op file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS commit scripts and op scripts, specify the MD5, SHA-1, or SHA-256 checksum hash. When it executes a local event, commit, or op script, Junos OS verifies the authenticity of the script by using the configured checksum hash.
<b>Options</b>	<b>md5 hash</b> —MD5 checksum of this script.  <b>sha-256 hash</b> —SHA-256 checksum of this script.  <b>sha1 hash</b> —SHA-1 checksum of this script.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Checksum Hashes for a Commit Script</i></li><li>• <i>Configuring Checksum Hashes for an Event Script</i></li><li>• <i>Configuring Checksum Hashes for an Op Script</i></li><li>• <i>Executing an Op Script from a Remote Site</i></li><li>• <a href="#">file checksum md5 on page 363</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha-256 on page 365</a> command in the <i>System Basics and Services Command Reference</i></li><li>• <a href="#">file checksum sha1 on page 364</a> command in the <i>System Basics and Services Command Reference</i></li></ul>

---

## command

---

<b>Syntax</b>	<code>command <i>filename-alias</i>;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op file</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS op scripts, configure a filename alias for the script file. This allows you to run the script by referencing either the script filename or the filename alias.
<b>Options</b>	<i>filename-alias</i> —Alias for the script file.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Enabling an Op Script and Defining a Script Alias</i></li></ul>

## commit

---

**Syntax** `commit {  
 allow-transients;  
 dampen {  
 dampen-options {  
 cpu-factor cpu-factor;  
 line-interval line-interval;  
 time-interval time-interval;  
 }  
 }  
 direct-access;  
 file filename {  
 checksum (md5 | sha-256 | sha1) hash;  
 optional;  
 refresh;  
 refresh-from url;  
 source url;  
 }  
 max-datasize  
 refresh;  
 refresh-from url;  
 traceoptions {  
 file <filename> <files number> <size size> <world-readable | no-world-readable>;  
 flag flag;  
 no-remote-trace;  
 }  
}`

**Hierarchy Level** [edit system [scripts](#)]

**Release Information** Statement introduced in Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** For Junos OS commit scripts, configure the commit-time scripting mechanism.

**Options** The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- *Storing and Enabling Scripts*

## description

---

<b>Syntax</b>	<code>description <i>descriptive-text</i>;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op file filename</a> ] [edit system <a href="#">scripts op file filename</a> arguments <i>argument-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS op scripts, provide a help-text string that appears in the command-line interface (CLI).
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Help Text for Op Scripts</i></li> <li>• <i>Declaring Arguments in Op Scripts</i></li> <li>• <a href="#">file (Op Scripts) on page 6655</a></li> </ul>

## direct-access

---

<b>Syntax</b>	<code>direct-access;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify that commit scripts read input configurations directly from the database when inspecting these scripts for errors.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Executing Large Commit Scripts</i></li> </ul>

## file (Commit Scripts)

---

<b>Syntax</b>	<pre>file <i>filename</i> {     checksum (md5   sha-256   sha1) <i>hash</i>;     optional;     refresh;     refresh-from <i>url</i>;     source <i>url</i>; }</pre>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS commit scripts, enable a commit script that is located in the <code>/var/db/scripts/commit</code> directory.
<b>Options</b>	<p><i>filename</i>—Name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing a commit script.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>maintenance—To view this statement in the configuration.</p> <p>maintenance-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><i>Controlling Execution of Commit Scripts During Commit Operations</i></li></ul>



## file (Op Scripts)

<b>Syntax</b>	<pre> file <i>filename</i> {   arguments {     <i>argument-name</i> {       <b>description</b> <i>descriptive-text</i>;     }   }   <b>checksum</b> (md5   sha-256   sha1) <i>hash</i>;   <b>command</b> <i>filename-alias</i>;   dampen {     dampen-options {       cpu-factor <i>cpu-factor</i>;       line-interval <i>line-interval</i>;       time-interval <i>time-interval</i>;     }   }   <b>description</b> <i>descriptive-text</i>;   <b>refresh</b>;   <b>refresh-from</b> <i>url</i>;   <b>source</b> <i>url</i>; } </pre>
<b>Hierarchy Level</b>	[edit system <b>scripts op</b> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	For Junos OS op scripts, enable an op script that is located in the <code>/var/db/scripts/op</code> directory.
<b>Options</b>	<p><b><i>filename</i></b>—The name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing an op script.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p><b>maintenance</b>—To view this statement in the configuration.</p> <p><b>maintenance-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Enabling an Op Script and Defining a Script Alias</i></li> </ul>

## no-allow-url

---

<b>Syntax</b>	no-allow-url;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS op scripts, prohibit the remote execution of scripts. When you include this configuration statement, the <b>op url</b> operational mode command generates an error and does not permit you to execute the op script from a remote site.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">file (Op Scripts) on page 6655</a></li><li>• <i>Executing an Op Script from a Remote Site</i></li></ul>

## op

```
Syntax  op {
        file filename {
            arguments {
                argument-name {
                    description descriptive-text;
                }
            }
        }
        checksum (md5 | sha-256 | sha1) hash;
        command filename-alias;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        description descriptive-text;
        max-datasize
        refresh;
        refresh-from url;
        source url;
    }
    no-allow-url
    refresh;
    refresh-from url;
    traceoptions {
        file <filename> <files number> <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

**Hierarchy Level** [edit system [scripts](#)]

**Release Information** Statement introduced in Junos OS Release 7.6.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** For Junos OS op scripts, configure an operation scripting mechanism.

**Options** The statements are explained separately.


**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**

- *Storing and Enabling Scripts*

## optional

---

<b>Syntax</b>	optional;
<b>Hierarchy Level</b>	[edit system <b>scripts commit</b> file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS commit scripts, allow a commit operation to succeed even if the script specified in the <b>file</b> statement is missing from the <b>/var/db/scripts/commit</b> directory on the device.
<div> <b>NOTE:</b> On the QFabric system, commit scripts are stored in the <b>/pbdata/mgd_shared/partition-ip/var/db/scripts/commit/</b> directory on the Director device.</div>	
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><i>Controlling Execution of Commit Scripts During Commit Operations</i></li></ul>

## refresh (Commit Scripts)

<b>Syntax</b>	refresh;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ], [edit system <a href="#">scripts commit</a> file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at the source URL, as specified in the <b>source</b> statement at the same hierarchy level.



**NOTE:** Issuing the `set refresh` command does not add the `refresh` statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.


The `set refresh` command is unique in the Junos OS CLI in that it behaves like an operational mode command and yet it can be executed from within configuration mode. All other Junos OS CLI operational mode commands can only be executed from command mode. The functionality is provided in this manner as a convenience to users developing commit scripts.

On the QFabric system, commit scripts are stored in the `/pbdata/mgd_shared/partition-ip/var/db/scripts/commit/` directory on the Director device.


<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using a Master Source Location for a Script</i></li> <li>• <a href="#">refresh-from (Commit Scripts) on page 6661</a></li> <li>• <a href="#">source (Commit Scripts) on page 6665</a></li> </ul>

## refresh (Op Scripts)

---


<b>Syntax</b>	refresh;
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op</a> ], [edit system <a href="#">scripts op file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 on the QFX Series.
<b>Description</b>	<p>For Junos OS op scripts, overwrite the local copy of all enabled op scripts or a single enabled script located in the <code>/var/db/scripts/op</code> directory with the copy located at the source URL, specified in the <b>source</b> statement at the same hierarchy level.</p> <p>The update operation occurs as soon as you issue the <b>set refresh</b> configuration mode command. Issuing the <b>set refresh</b> command does not add the <b>refresh</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<hr/>	
<div> <b>NOTE:</b> On the QFabric system, op scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/op/</code> directory on the Director device.</div> <hr/>	
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using a Master Source Location for a Script</i></li><li>• <a href="#">refresh-from (Op Scripts) on page 6662</a></li><li>• <a href="#">source (Op Scripts) on page 6666</a></li></ul>

## refresh-from (Commit Scripts)

<b>Syntax</b>	<code>refresh-from url;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit</a> ], [edit system <a href="#">scripts commit</a> file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>For Junos OS commit scripts, overwrite the local copy of all enabled commit scripts or a single enabled script located in the <code>/var/db/scripts/commit</code> directory with the copy located at a URL other than the URL specified in the <b>source</b> statement.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from url</b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<div>  <b>NOTE:</b> This statement is not supported on the QFabric system.         </div>	
<b>Options</b>	<b>url</b> —The source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using an Alternate Source Location for a Script</i></li> <li>• <a href="#">refresh (Commit Scripts) on page 6659</a></li> <li>• <a href="#">source (Commit Scripts) on page 6665</a></li> </ul>

## refresh-from (Op Scripts)

---

<b>Syntax</b>	<code>refresh-from url;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op</a> ], [edit system <a href="#">scripts op file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 on the QFX Series.
<b>Description</b>	<p>For Junos OS op scripts, overwrite the local copy of all enabled op scripts or a single enabled script located in the <code>/var/db/scripts/op</code> directory with the copy located at a URL other than the URL specified in the <b>source</b> statement.</p> <p>The update operation occurs as soon as you issue the <b>set refresh-from url</b> configuration mode command. Issuing the <b>set refresh-from</b> command does not add the <b>refresh-from</b> statement to the configuration. Thus the command behaves like an operational mode command by executing an operation, instead of adding a statement to the configuration.</p>
<div> <b>NOTE:</b> This statement is not supported on the QFabric system.</div>	
<b>Options</b>	<b>url</b> —Source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Using an Alternate Source Location for a Script</i></li><li>• <a href="#">refresh (Op Scripts) on page 6660</a></li><li>• <a href="#">source (Op Scripts) on page 6666</a></li></ul>



## scripts

```

Syntax  scripts {
        commit {
            allow-transients;
        dampen {
            dampen-options {
                cpu-factor cpu-factor;
                line-interval line-interval;
                time-interval time-interval;
            }
        }
        direct-access;
        file filename {
            checksum (md5 | sha-256 | sha1) hash;
            optional;
            refresh;
            refresh-from url;
            source url;
        }
        max-datasize
        refresh;
        refresh-from url;
        traceoptions {
            file <filename> <files number> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
    load-scripts-from-flash;
    op {
        file filename {
            arguments {
                argument-name {
                    description descriptive-text;
                }
            }
            checksum (md5 | sha-256 | sha1) hash;
            command filename-alias;
            dampen {
                dampen-options {
                    cpu-factor cpu-factor;
                    line-interval line-interval;
                    time-interval time-interval;
                }
            }
            description descriptive-text;
            max-datasize
            refresh;
            refresh-from url;
            source url;
        }
        no-allow-url
        refresh;
    }

```

```
refresh-from url;  
traceoptions {  
  file <filename> <files number> <size size> <world-readable | no-world-readable>;  
  flag flag;  
  no-remote-trace;  
}  
}  
synchronize;  
}
```

**Hierarchy Level** [edit system]

**Release Information** Statement introduced in Junos OS Release 7.4.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** For Junos OS commit or op scripts, configure scripting mechanisms.



**NOTE:** The traceoptions statement is not supported on QFabric systems.

---


**Options** The statements are explained separately.

**Required Privilege Level** maintenance—To view this statement in the configuration.  
maintenance-control—To add this statement to the configuration.

**Related Documentation**


- *Storing and Enabling Scripts*

## source (Commit Scripts)

<b>Syntax</b>	<code>source url;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts commit file</a> <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS commit scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/commit</code> directory. When you include the <b>refresh</b> statement at the same hierarchy level and commit the configuration, the local copy is overwritten by the version stored at the specified URL.
<div>  <b>NOTE:</b> On the QFabric system, commit scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/op/</code> directory on the Director device. </div>	
<b>Options</b>	<i>url</i> —The source specified as an HTTP URL, FTP URL, or scp-style remote file specification.
<b>Required Privilege Level</b>	maintenance—To view this statement in the configuration. maintenance-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Using a Master Source Location for a Script</i></li> <li>• <i>Overview of Updating Scripts from a Remote Source</i></li> <li>• <a href="#">refresh (Commit Scripts) on page 6659</a></li> <li>• <a href="#">refresh-from (Commit Scripts) on page 6661</a></li> </ul>

## source (Op Scripts)

---

<b>Syntax</b>	<code>source url;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">scripts op file filename</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	For Junos OS op scripts, specify the location of the source file for an enabled script located in the <code>/var/db/scripts/op</code> directory. When you include the <b>refresh</b> statement at the same hierarchy level, the local copy is overwritten by the version stored at the specified URL.
<hr/>	
<div> <b>NOTE:</b> On the QFabric system, commit scripts are stored in the <code>/pbdata/mgd_shared/partition-ip/var/db/scripts/op/</code> directory on the Director device.</div> <hr/>	
<b>Options</b>	<b>url</b> —Master source file for an op script specified as an HTTP URL, FTP URL, or scp-style remote file specification.
<b>Required Privilege Level</b>	<b>maintenance</b> —To view this statement in the configuration. <b>maintenance-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Using a Master Source Location for a Script</a></li><li>• <a href="#">refresh (Op Scripts) on page 6660</a></li><li>• <a href="#">refresh-from (Op Scripts) on page 6662</a></li></ul>

## Configuration Statements for Network Analytics

---

- [analytics on page 6667](#)
- [depth-threshold on page 6671](#)
- [interfaces \(Analytics\) on page 6672](#)
- [latency-threshold on page 6674](#)
- [queue-statistics on page 6676](#)
- [streaming-servers on page 6678](#)
- [traceoptions \(Analytics\) on page 6679](#)
- [traffic-statistics on page 6680](#)

## analytics

**Syntax** *Junos OS Release 13.2X51-D15 and later:*

```
analytics {
  collector {
    local {
      file filename {
        size size;
        files number;
      }
    }
    address ip-address {
      port number {
        transport protocol {
          export-profile profile-name;
        }
      }
    }
  }
  export-profiles {
    profile-name {
      interface {
        information;
        statistics {
          queue;
          traffic;
        }
        status {
          link;
          queue;
          traffic;
        }
      }
    }
    stream-format format;
    system {
      information;
      status {
        queue;
        traffic;
      }
    }
  }
  resource {
    interfaces {
      interface-name {
        resource-profile name;
      }
    }
    system {
      polling-interval {
        queue-monitoring interval;
        traffic-monitoring interval;
      }
    }
  }
}
```

```
        resource-profile name;  
    }  
}  
resource-profiles {  
    profile-name {  
        depth-threshold {  
            high number;  
            low number;  
        }  
        latency-threshold {  
            high number;  
            low number;  
        }  
        no-queue-monitoring;  
        no-traffic-monitoring;  
        queue-monitoring;  
        traffic-monitoring;  
    }  
}  
traceoptions {  
    file filename {  
        files number;  
        size size;  
    }  
}
```

*Junos OS Release 13.2X50-D15 and 13.2X51-D10 only:*

```
analytics {
  interfaces {
    all {
      depth-threshold high number low number;
      latency-threshold high number low number;
      queue-statistics;
      no-queue-statistics;
      traffic-statistics;
      no-traffic-statistics;
    }
    interface-name {
      depth-threshold high number low number;
      latency-threshold high number low number;
      queue-statistics;
      no-queue-statistics;
      traffic-statistics;
      no-traffic-statistics;
    }
  }
  queue-statistics {
    file filename {
      files number-of-files;
      size size;
    }
    interval interval;
  }
  streaming-servers {
    address ip-address {
      port number {
        stream-format format;
        stream-type type
      }
    }
  }
  traceoptions {
    file filename {
      files number;
      size size;
    }
  }
  traffic-statistics {
    file filename {
      files number-of-files;
      size size;
    }
    interval interval;
  }
}
```

**Hierarchy Level** [edit services]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure the network analytics feature that includes monitoring for traffic and queue statistics. The network analytics processes running on the Packet Forwarding Engine and Routing Engine collect and analyze the data, and generate reports that may be saved in log files or sent as streaming data to remote servers.

The remaining statements are explained separately.


**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [show analytics traffic-statistics on page 6847](#)
- [show analytics collector on page 6833](#)
- [show analytics status on page 6841](#)
- [show analytics queue-statistics on page 6839](#)
- [show analytics configuration on page 6835](#)



## depth-threshold

<b>Syntax</b>	depth-threshold { high <i>number</i> ; low <i>number</i> ; }
<b>Hierarchy Level</b>	[edit services analytics interfaces] [edit services analytics resource-profiles]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement in the <b>[edit services analytics resource-profiles]</b> hierarchy level introduced in Junos OS Release 13.2X51-D15. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	If network analytics queue statistics monitoring is enabled, specify the high and low values (in bytes) of the queue depth (buffer) threshold. If you configure a depth threshold, you cannot configure the latency threshold. You can configure the depth threshold for one interface or all interfaces. Specify the high and low queue depth threshold numbers:
<div>  <b>NOTE:</b> The configuration for a specific interface supersedes the global configuration for all interfaces. </div>	
<b>Options</b>	<p><b>high <i>number</i></b>—Specify the maximum value for the depth threshold.</p> <p><b>Range:</b> 1 to 1,250,000,000 bytes</p> <p><b>Default:</b></p> <ul style="list-style-type: none"> <li>Junos OS Release 13.2X51-D10 or later—0 bytes</li> <li>Junos OS Release 13.2X50-D15—14,680,064 bytes (14 MB)</li> </ul> <p><b>low <i>number</i></b>—Specify the minimum value for the depth threshold.</p> <p><b>Range:</b> 1 to 1,250,000,000 bytes</p> <p><b>Default:</b></p> <ul style="list-style-type: none"> <li>Junos OS Release 13.2X51-D10 or later—0 bytes</li> <li>Junos OS Release 13.2X50-D15—1024 bytes (1 KB)</li> </ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Network Analytics Overview on page 6490</a></li> <li><a href="#">analytics on page 6667</a></li> <li><a href="#">latency-threshold on page 6674</a></li> <li><i>resource-profiles (Analytics)</i></li> </ul>

## interfaces (Analytics)

---

**Syntax**

```
interfaces {  
  all {  
    depth-threshold high number low number;  
    latency-threshold high number low number;  
    queue-statistics;  
    no-queue-statistics;  
    traffic-statistics;  
    no-traffic-statistics;  
  }  
  interface-name {  
    depth-threshold high number low number;  
    latency-threshold high number low number;  
    queue-statistics;  
    no-queue-statistics;  
    traffic-statistics;  
    no-traffic-statistics;  
  }  
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure physical interfaces for monitoring traffic and queue statistics by the network analytics processes running on the Packet Forwarding Engine and Routing Engine. You may specify one interface or all interfaces in your configuration.



**NOTE:** The configuration for a specific interface supersedes the global configuration for all interfaces. You can configure traffic and queue monitoring for physical interfaces only; logical interfaces and Virtual Chassis port (VCP) interfaces are not supported.



**NOTE:** Disabling the queue or traffic monitoring (using the `no-queue-statistics` or `no-traffic-statistics` configuration statements) supersedes the configuration (enabling) of the feature.

**Options** `all`—Configure all interfaces on the device for high-frequency monitoring.

`interface-name`—Name of the interface to configure for high-frequency monitoring.

`no-queue-statistics`—Disable the collection of queue statistics.



**NOTE:** The `no-queue-statistics` statement supersedes the `queue-statistics` statement.

**no-traffic-statistics**—Disable the collection of traffic statistics.



**NOTE:** The `no-traffic-statistics` statement supersedes the `traffic-statistics` statement.

**queue-statistics**—Enable the collection of queue statistics for a specific interface or all interfaces.

**traffic-statistics**—Enable the collection of traffic statistics for a specific interface or all interfaces.

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Network Analytics Overview on page 6490</a>
	• <a href="#">analytics on page 6667</a>

## latency-threshold

---

<b>Syntax</b>	latency-threshold { high <i>number</i> ; low <i>number</i> ; }
<b>Hierarchy Level</b>	[edit services analytics interfaces] [edit services analytics resource-profiles]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement in the <b>[edit services analytics resource-profiles]</b> hierarchy level introduced in Junos OS Release 13.2X51-D15. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	If network analytics queue statistics monitoring is enabled, specify the high and low values (in microseconds) of the latency threshold of the queue. If you configure a latency threshold, you cannot configure the depth threshold. You can configure the latency threshold for one interface or all interfaces. Specify the high and low latency threshold numbers:



**NOTE:** The configuration for a specific interface supersedes the global configuration for all interfaces.

---

<b>Options</b>	<b>high <i>number</i></b> —Specify the maximum value for the latency threshold. <b>Range:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—1 to 100,000,000 nanoseconds (0.001 to 100,000 microseconds)</li><li>Junos OS Release 13.2X51-D10 or earlier—1 to 100,000 microseconds</li></ul> <b>Default:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—1,000,000 nanoseconds (1000 microseconds or 1 millisecond)</li><li>Junos OS Release 13.2X51-D10—1000 microseconds</li><li>Junos OS Release 13.2X50-D15—900 microseconds</li></ul> <b>low <i>number</i></b> —Specify the minimum value for the latency threshold. <b>Range:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—1 to 100,000,000 nanoseconds</li><li>Junos OS Release 13.2X51-D10 or earlier—1 to 100,000 microseconds</li></ul> <b>Default:</b> <ul style="list-style-type: none"><li>Junos OS Release 13.2X51-D15 or later—100 nanoseconds (0.1 microseconds)</li></ul>
----------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

- Junos OS Release 13.2X51-D10—50 microseconds
- Junos OS Release 13.2X50-D15—300 microseconds

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)
- [depth-threshold on page 6671](#)

## queue-statistics

---

**Syntax**    `queue-statistics {  
              file filename {  
                  files number-of-files;  
                  size size;  
              }  
              interval interval;  
          }`

**Hierarchy Level**    [edit services analytics]

**Release Information**    Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description**    Enable the logging of queue statistics in a local file. This statement does not enable queue statistics monitoring.

To enable queue monitoring, you must specify the **queue-statistics** configuration statement at the [edit services analytics interfaces] hierarchy level.

**Default**    This feature is disabled by default.

**Options**    `interval interval`—Configure the polling interval in milliseconds.



**NOTE:** You can configure the polling interval for queue statistics globally for all interfaces only. Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.

---

**Range:**

- Junos OS Release 13.2X50-D15—8 to 1000 milliseconds (8 milliseconds to 1 second)
- Junos OS Release 13.2X51-D10 or later—10 to 1000 milliseconds (10 milliseconds to 1 second)



**NOTE:** In Junos OS Release 13.2X51-D10 or later, if you configured an interval of less than 10 milliseconds, the following warning messages appear during the commit process: Queue statistics polling interval can not be less than 10 milliseconds and Setting Queue statistics polling interval to 10 milliseconds. These messages do not stop the commit operation, but the interval is automatically set to 10 milliseconds.

---

**Default:**

- Junos OS Release 13.2X50-D15—8 milliseconds

- Junos OS Release 13.2X51-D10 or later—10 milliseconds

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Network Analytics Overview on page 6490</a>
	• <a href="#">analytics on page 6667</a>

## streaming-servers

---

**Syntax**

```
streaming-servers {  
  address ip-address {  
    port number {  
      stream-format format;  
      stream-type type  
    }  
  }  
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure remote servers to receive streaming output for the network analytics monitoring of traffic and queue statistics. The streaming function supports TCP connections only, and sends records separated by a newline character.



**NOTE:** Before you use the remote server to receive streaming data, you must set up the TCP server software to process records that are separated by the newline character (\n).

You can configure multiple servers and multiple ports on each server to receive the streaming data. You can configure different streaming data types and formats for different ports on a server, but you can configure only one streaming type and one format for each port on a server.

**Options** **address *ip-address***—IP address of the remote server receiving the streaming data.

**port *number***—Port number of the remote server receiving the streaming data.

**stream-format *format***—Format of the streaming data being sent to a server. Only one format can be sent to each port on a server.

**Values:**

- **csv**—Comma-separated Values (CSV). Records sent in this format contain a “q” for a queue statistics, or a “t” for a traffic statistics.
- **json**—JavaScript Object Notification (JSON). Records sent in this format contain “queue-statistics” or “traffic-statistics” in the “record type” field.
- **tsv**—Tab-separated Values (TSV). Records sent in this format contain a “q” for a queue statistics, or a “t” for a traffic statistics.

**stream-type *type***—Type of streaming data sent to a port. You can specify different types of streaming data to be sent to different ports on the same server.

**Values:**



- `queue-statistics`
- `traffic-statistics`

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [Understanding Network Analytics Streaming Data on page 6499](#)
- [analytics on page 6667](#)

## traceoptions (Analytics)

**Syntax**

```
traceoptions {
    file filename;
    files number-of-files;
    size size;
}
```

**Hierarchy Level** [edit services analytics]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Configure traceoptions for the network analytics daemon (analyticsd) running on the Routing Engine.

**Options** **file *filename***—Specify a filename for storing the traceoptions data. The file is stored in the `/var/log/` directory of your device.  
If you do not specify a filename, the data is not stored in a file.

**files *number-of-files***—Specify the number of files to store locally. After the number files with the maximum file size is reached, the system starts over and writes the data to the first file.

**Range:** 2 to 1,000 files.

**size *size***—Configure the file size in megabytes (MB).

**Syntax:** `xm` to specify MB.

**Range:** 10 to 4095 MB

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)

## traffic-statistics

---

<b>Syntax</b>	<pre>traffic-statistics {     file <i>filename</i> {         files <i>number-of-files</i>;         size <i>size</i>;     }     interval <i>interval</i>; }</pre>
<b>Hierarchy Level</b>	[edit services analytics]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	<p>Enable the logging of traffic statistics in a local file. This statement does not enable traffic statistics monitoring.</p> <p>To enable the monitoring of traffic statistics, configure the <b>traffic-statistics</b> configuration statement at the [edit services analytics interfaces] hierarchy level.</p>
<b>Default</b>	This feature is disabled by default.
<b>Options</b>	<p><b>file <i>filename</i></b>—Specify a filename for storing the traffic statistics in the JavaScript Object Notification (JSON) format. The file is stored in the <b>/var/log/</b> directory of your device. If you do not specify a filename, the data is not stored in a file.</p> <p><b>files <i>number-of-files</i></b>—Specify the number of files to store locally. After the number files with the maximum file size is reached, the system starts over and writes the data to the first file.</p> <p><b>Range:</b> 2 to 1,000 files.</p> <p><b>interval <i>interval</i></b>—Configure the polling interval in seconds.</p>



**NOTE:** You can configure the polling interval for traffic statistics globally for all interfaces only. Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.

---

### Range:

- Junos OS Release 13.2X51-D10 or later—2 to 300 seconds (2 seconds to 5 minutes)
- Junos OS Release 13.2X50-D15—1 to 300 seconds (1 second to 5 minutes)



**NOTE:** In Junos OS Release 13.2X51-D10 or later, if you configured an interval of less than 2 seconds, the following warning messages appear during the commit process:

Traffic statistics polling interval can not be less than 2 seconds, and

Setting Traffic statistics polling interval to 2 seconds.

These messages do not stop the commit operation, but the interval is automatically set to 2 seconds.

**Default:**

- Junos OS Release 13.2X50-D15—1 second
- Junos OS Release 13.2X51-D10 or later—2 seconds

**size size**—Configure the file size in megabytes (MB).

**Syntax:** *xm* to specify MB.

**Range:** 10 to 4095 MB

**Required Privilege Level** interface—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)

## Configuration Statements for sFlow Technology

- [agent-id on page 6682](#)
- [collector \(sFlow Technology\) on page 6682](#)
- [interfaces \(sFlow\) on page 6683](#)
- [polling-interval on page 6684](#)
- [sample-rate on page 6685](#)
- [sflow on page 6686](#)
- [source-ip on page 6687](#)
- [traceoptions \(sFlow Technology\) on page 6688](#)
- [udp-port on page 6689](#)

## agent-id

---

<b>Syntax</b>	<code>agent-id ip-address;</code>
<b>Hierarchy Level</b>	[edit protocols sflow]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the IP address of the sFlow agent. If you do not configure the sFlow agent ID, the IP address for the agent is dynamically created using the IP address of an interface configured on the QFX Series device.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6596</a></li><li>• <a href="#">sflow on page 6686</a></li></ul>

## collector (sFlow Technology)

---

<b>Syntax</b>	<code>collector ip-address {     udp-port port-number; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">protocols sflow</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>Configure a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the configured collector for analysis. You can configure up to four collectors on the device. You specify the IP address for each collector you configure.</p> <p>The remaining statement is explained separately.</p>
<b>Options</b>	<i>ip-address</i> —IP address of the collector.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6596</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li></ul>

---

## interfaces (sFlow)

---

<b>Syntax</b>	<code>interfaces <i>interface-name</i> {     polling-interval <i>seconds</i>;     sample-rate <i>number</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">protocols sflow</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>Configure sFlow network traffic monitoring on the specified interface on the device. You can configure sFlow parameters (polling interval, sample rate) with different values on different interfaces.</p> <p>The remaining statements are explained separately.</p>
<b>Options</b>	<i>interface-name</i> —Name of the interface on which to configure sFlow parameters.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6596</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li></ul>

## polling-interval

---

<b>Syntax</b>	<code>polling-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">protocols sflow</a> ], [edit <a href="#">protocols sflow interfaces</a> <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the rate (in seconds) at which successive samples of interface statistics (counters) are taken.
<b>Default</b>	If no polling interval is configured for a particular interface, the device uses the global polling interval configured at the <a href="#">[edit protocols sflow]</a> hierarchy level. If no global interval is configured, the device uses the default polling interval of 20 seconds.
<b>Options</b>	<b><i>seconds</i></b> —Number of seconds between successive samples of interface statistics. Specifying a value of <b>0</b> (zero) disables the polling. <b>Range:</b> 0 through 3600 seconds
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6596</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li></ul>

## sample-rate

---

<b>Syntax</b>	<code>sample-rate <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">protocols sflow</a>],</code> <code>[edit <a href="#">protocols sflow interfaces</a> <i>interface-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Specify the denominator ( <i>number</i> ) of the ratio that is the sample rate in sFlow traffic monitoring. For example, to configure a sample rate of 1 in 1000 packets, you specify a <i>number</i> of 1000.
<b>Default</b>	If no sample rate is configured for a particular interface, the device uses the global sample rate configured at the <code>[edit <a href="#">protocols sflow</a>]</code> hierarchy level. If no global rate is configured, the device uses the default sample rate of 1 in 2000 packets.
<b>Options</b>	<i>number</i> —Denominator of the ratio representing the sample rate (one packet out of <i>number</i> ). <b>Range:</b> 1 through 16,777,215
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring sFlow Technology on page 6596</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li> </ul>

## sflow

---

<b>Syntax</b>	<pre>sflow {   agent-id <i>ip-address</i>;   collector <i>ip-address</i> {     udp-port <i>port-number</i>;   }   interfaces <i>interface-name</i> {     polling-interval <i>number</i>;     sample-rate {       egress <i>number</i>;       ingress <i>number</i>;     }   }   polling-interval <i>number</i>;   sample-rate {     egress <i>number</i>;     ingress <i>number</i>;   }   source-ip <i>ip-address</i>;   traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable         no-world-readable&gt;;     flag <i>flag</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit protocols]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>Configure sFlow technology to monitor traffic continuously on specified interfaces simultaneously. sFlow data can be used to characterize network activity.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	The sFlow protocol is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6596</a></li><li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li></ul>



---

## source-ip

---

<b>Syntax</b>	source-ip <i>ip-address</i> ;
<b>Hierarchy Level</b>	[edit protocols sflow]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Configure the source IP address to be used for sFlow datagrams. If you do not configure a source IP address, it is dynamically created based on the IP address of an Ethernet interface configured on the QFX Series device.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring sFlow Technology on page 6596</a></li><li>• <a href="#">sflow on page 6686</a></li></ul>

## tracoptions (sFlow Technology)

---

<b>Syntax</b>	<pre>tracoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i>; }</pre>
<b>Hierarchy Level</b>	[edit protocols <a href="#">sflow</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Define tracing operations for sFlow technology.
<b>Default</b>	The <b>tracoptions</b> feature is disabled.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Output files are located in the <b>/var/log/</b> directory.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>. Incoming trace file data is logged in the now empty <b>trace-file</b>. When <b>trace-file</b> again reaches its maximum size, <b>trace-file.0</b> is renamed <b>trace-file.1</b> and <b>trace-file</b> is renamed <b>trace-file.0</b>. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify the maximum number of files, you must also specify the maximum file size using the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 1 trace file</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"><li>• <b>all</b>—Trace all sFlow monitoring events.</li><li>• <b>client-server</b>—Trace sFlow monitoring client-server events.</li><li>• <b>configuration</b>—Trace sFlow monitoring configuration events.</li><li>• <b>interface</b>—Trace sFlow monitoring interface events.</li><li>• <b>rtsock</b>—Trace routing socket code events.</li></ul> <p><b>no-stamp</b>—(Optional) Do not place timestamp information at the beginning of each line in the trace file.</p> <p><b>no-world-readable</b>—(Optional) Prevent any user from reading the trace file.</p> <p><b>replace</b>—(Optional) Replace an existing trace file if there is one.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named <b>trace-file</b> reaches its maximum size, it</p>

is renamed **trace-file.0**. Incoming trace file data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size of 4 GB

**Default:** 128 KB

**world-readable**—(Optional) Allow any user to read the trace file.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Overview of sFlow Technology](#)

## udp-port

**Syntax** `udp-port port-number;`

**Hierarchy Level** [edit protocols [sflow collector](#)]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure the UDP port for a remote collector for sFlow network traffic monitoring. The device sends sFlow UDP datagrams to the collector for analysis.

**Default** Port 6343

**Options** *port-number*—UDP port number for this collector.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring sFlow Technology on page 6596](#)
- [Example: Monitoring Network Traffic Using sFlow Technology on page 6571](#)

## Configuration Statements for SNMP

- [access \(SNMP\) on page 6693](#)
- [address \(SNMP\) on page 6693](#)
- [address-mask on page 6694](#)
- [agent-address on page 6694](#)
- [alarm \(SNMP RMON\) on page 6695](#)

- [authentication-md5 on page 6696](#)
- [authentication-none on page 6697](#)
- [authentication-password on page 6698](#)
- [authentication-sha on page 6699](#)
- [authorization on page 6700](#)
- [bucket-size on page 6701](#)
- [categories on page 6701](#)
- [client-list on page 6702](#)
- [client-list-name on page 6702](#)
- [clients on page 6703](#)
- [commit-delay on page 6703](#)
- [community \(SNMP\) on page 6704](#)
- [community \(RMON\) on page 6705](#)
- [community-name \(SNMP\) on page 6706](#)
- [contact on page 6707](#)
- [description \(SNMP\) on page 6707](#)
- [description \(RMON\) on page 6708](#)
- [destination-port \(SNMP\) on page 6708](#)
- [engine-id on page 6709](#)
- [event on page 6710](#)
- [falling-event-index \(RMON\) on page 6711](#)
- [falling-threshold \(Health Monitor\) on page 6712](#)
- [falling-threshold \(RMON\) on page 6713](#)
- [falling-threshold-interval on page 6714](#)
- [filter-duplicates on page 6714](#)
- [filter-interfaces on page 6715](#)
- [group \(Associating a Security Name\) on page 6715](#)
- [group \(Configuring Access Privileges\) on page 6716](#)
- [health-monitor on page 6717](#)
- [history on page 6718](#)
- [interface \(SNMP\) on page 6719](#)
- [interface \(RMON\) on page 6720](#)
- [interval \(Health Monitor\) on page 6720](#)
- [interval \(RMON\) on page 6721](#)
- [local-engine on page 6722](#)
- [location on page 6723](#)
- [message-processing-model on page 6723](#)

- [name](#) on page 6724
- [nonvolatile](#) on page 6724
- [notify](#) on page 6725
- [notify-filter \(Applying to the Management Target\)](#) on page 6726
- [notify-filter \(Configuring the Profile Name\)](#) on page 6726
- [notify-view](#) on page 6727
- [oid](#) on page 6727
- [oid \(SNMPv3\)](#) on page 6728
- [owner](#) on page 6729
- [parameters](#) on page 6729
- [port \(SNMP\)](#) on page 6730
- [privacy-3des](#) on page 6731
- [privacy-aes128](#) on page 6732
- [privacy-des](#) on page 6733
- [privacy-none](#) on page 6733
- [privacy-password](#) on page 6734
- [read-view](#) on page 6735
- [remote-engine](#) on page 6736
- [request-type](#) on page 6737
- [retry-count \(SNMPv3\)](#) on page 6738
- [rising-event-index](#) on page 6739
- [rising-threshold \(Health Monitor\)](#) on page 6740
- [rising-threshold \(RMON\)](#) on page 6741
- [rmon](#) on page 6742
- [sample-type](#) on page 6743
- [security-level \(Defining Access Privileges\)](#) on page 6744
- [security-level \(Generating SNMP Notifications\)](#) on page 6745
- [security-model \(Access Privileges\)](#) on page 6746
- [security-model \(Group\)](#) on page 6747
- [security-model \(SNMP Notifications\)](#) on page 6748
- [security-name \(Community String\)](#) on page 6749
- [security-name \(Security Group\)](#) on page 6750
- [security-name \(SNMP Notifications\)](#) on page 6751
- [security-to-group](#) on page 6752
- [snmp](#) on page 6753
- [snmp-community](#) on page 6757
- [source-address \(SNMP\)](#) on page 6757

- [startup-alarm on page 6758](#)
- [syslog-subtag on page 6759](#)
- [tag \(Configuring Notification Targets\) on page 6759](#)
- [tag \(Configuring the SNMP Community\) on page 6760](#)
- [tag-list on page 6760](#)
- [target-address on page 6761](#)
- [target-parameters on page 6762](#)
- [targets on page 6763](#)
- [timeout on page 6763](#)
- [traceoptions \(SNMP\) on page 6764](#)
- [trap-group on page 6766](#)
- [trap-options on page 6767](#)
- [type \(RMON Notification\) on page 6768](#)
- [type \(SNMPv3\) on page 6769](#)
- [user on page 6769](#)
- [usm on page 6770](#)
- [v3 on page 6772](#)
- [vacm on page 6774](#)
- [variable on page 6775](#)
- [version on page 6776](#)
- [view \(Configuring a MIB View\) on page 6777](#)
- [view \(Associating MIB View with a Community\) on page 6778](#)
- [write-view on page 6778](#)

## access (SNMP)

<b>Syntax</b>	<pre> access {   group group-name {     (default-context-prefix   context-prefix <i>context-prefix</i>) {       security-model (any   usm   v1   v2c) {         security-level (authentication   none   privacy) {           notify-view <i>view-name</i>;           read-view <i>view-name</i>;           write-view <i>view-name</i>;         }       }     }   } } </pre>
<b>Hierarchy Level</b>	[edit snmp v3 vacm]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Set SNMP access limits.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>

## address (SNMP)

<b>Syntax</b>	address <i>address</i> ;
<b>Hierarchy Level</b>	[edit snmp v3 target-address <i>target-address-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the SNMP target address for receiving traps or informs.
<b>Options</b>	<b>address</b> —IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 6516</a></li> <li>• <a href="#">Configuring SNMP on page 1356</a></li> <li>• <a href="#">Example: Configuring SNMP on page 6575</a></li> </ul>

## address-mask

---

<b>Syntax</b>	<code>address-mask address-mask;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 target-address target-address-name]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 on the QFX Series.
<b>Description</b>	Define and verify the source addresses for a group of target addresses for SNMP traps and informs.
<b>Options</b>	<b>address-mask</b> —Define a range of addresses.
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Address Mask</i></li></ul>

## agent-address

---

<b>Syntax</b>	<code>agent-address outgoing-interface;</code>
<b>Hierarchy Level</b>	<code>[edit snmp trap-options]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is <b>outgoing-interface</b> , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
<b>Options</b>	<b>outgoing-interface</b> —Value of the agent address of all SNMPv1 traps generated by this router or switch. The <b>outgoing-interface</b> option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. <b>Default:</b> Disabled (the agent address is not specified in SNMPv1 traps).
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Agent Address for SNMP Traps</i></li></ul>



## alarm (SNMP RMON)

<b>Syntax</b>	<pre>alarm <i>index</i> {     description <i>description</i>;     falling-event-index <i>index</i>;     falling-threshold <i>integer</i>;     falling-threshold-interval <i>seconds</i>;     interval <i>seconds</i>;     request-type (get-next-request   get-request   walk-request);     rising-event-index <i>index</i>;     rising-threshold <i>integer</i>;     sample-type (absolute-value   delta-value);     startup-alarm (falling-alarm   rising-alarm   rising-or-falling alarm);     syslog-subtag <i>syslog-subtag</i>;     variable <i>oid-variable</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp rmon]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure RMON alarm entries.
<b>Options</b>	<p><b><i>index</i></b>—Identifies this alarm entry as an integer.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring an RMON Alarm Entry and Its Attributes</a></li> <li>• <a href="#">event (SNMP)</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> </ul>

## authentication-md5

---

<b>Syntax</b>	<code>authentication-md5 {     authentication-password authentication-password; }</code>
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> ], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure MD5 as the authentication type for the SNMPv3 user.



**NOTE:** You can only configure one authentication type for each SNMPv3 user.

---

The remaining statement is explained separately.

<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring MD5 Authentication</i></li></ul>

## authentication-none

---

<b>Syntax</b>	authentication-none;
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> ], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure that there should be no authentication for the SNMPv3 user.



**NOTE:** You can configure only one authentication type for each SNMPv3 user.

<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring No Authentication</i></li> </ul>

## authentication-password

---

<b>Syntax</b>	<code>authentication-password <i>authentication-password</i>;</code>
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> authentication-md5], [edit snmp v3 usm local-engine user <i>username</i> authentication-sha], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-md5], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-sha]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the password for user authentication.
<b>Options</b>	<p><b><i>authentication-password</i></b>—Password that a user enters. The password is then converted into a key that is used for authentication.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"><li>• The password must be at least eight characters long.</li><li>• The password can include lowercase letters, uppercase letters, numbers, and the following special characters: <code>.,/\&lt;&gt;;:'[]{}~!@#\$%^*_+=-`</code></li></ul> <p>In addition, the following special characters are also supported, but you must enclose them within quotation marks ("" ) if you enter them on the CLI; if you use a Network Management System to enter the password, the quotation marks are not required:</p> <p><code>  &amp; ( ) ?</code></p> <p>Control characters—entered by simultaneously pressing the Ctrl key and additional keys—are not supported.</p>
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring MD5 Authentication</i></li><li>• <i>Configuring SHA Authentication</i></li></ul>

## authentication-sha

<b>Syntax</b>	authentication-sha { authentication-password authentication-password; }
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> ], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the secure hash algorithm (SHA) as the authentication type for the SNMPv3 user.




**NOTE:** You can configure only one authentication type for each SNMPv3 user.

The remaining statement is explained separately.

<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring SHA Authentication</li> </ul>

## authorization

---

<b>Syntax</b>	<code>authorization <i>authorization</i>;</code>
<b>Hierarchy Level</b>	[edit snmp community <i>community-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the access authorization for SNMP <b>Get</b> , <b>GetBulk</b> , <b>GetNext</b> , and <b>Set</b> requests.
<b>Options</b>	<p><i>authorization</i>—Access authorization level:</p> <ul style="list-style-type: none"><li>• <b>read-only</b>—Enable <b>Get</b>, <b>GetNext</b>, and <b>GetBulk</b> requests.</li><li>• <b>read-write</b>—Enable all requests, including <b>Set</b> requests. You must configure a view to enable <b>Set</b> requests.</li></ul>
	<div> <b>NOTE:</b> The read-write option is not supported on the QFX3000 QFabric system.</div>
	<b>Default:</b> read-only
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the SNMP Community String on page 6601</a></li></ul>

## bucket-size

---

<b>Syntax</b>	<code>bucket-size <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon history <i>index</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.
<b>Default</b>	50
<b>Options</b>	<i>number</i> —Number of discrete samples of Ethernet statistics requested.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## categories

---

<b>Syntax</b>	<code>categories {     <i>category</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit snmp trap-group <i>group-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the types of traps that are sent to the targets of the named trap group.
<b>Default</b>	If you omit the <b>categories</b> statement, all trap types are included in trap notifications.
<b>Options</b>	<i>category</i> —Name of a trap type: <b>authentication</b> , <b>chassis</b> , <b>configuration</b> , <b>link</b> , <b>remote-operations</b> , <b>rmon-alarm</b> , or <b>startup</b> .
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li> </ul>

## client-list

---

<b>Syntax</b>	<code>client-list <i>client-list-name</i> {     <i>ip-addresses</i>; }</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define a list of SNMP clients.
<b>Options</b>	<i>client-list-name</i> —Name of the client list.  <i>ip-addresses</i> —IP addresses of the SNMP clients to be added to the client list,
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Adding a Group of Clients to an SNMP Community on page 6603</a></li></ul>

## client-list-name

---

<b>Syntax</b>	<code>client-list-name <i>client-list-name</i>;</code>
<b>Hierarchy Level</b>	[edit snmp community <i>community-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Add a client list or prefix list to an SNMP community.
<b>Options</b>	<i>client-list-name</i> —Name of the client list or prefix list.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Adding a Group of Clients to an SNMP Community on page 6603</a></li></ul>



## clients

<b>Syntax</b>	clients { <i>address</i> <restrict>; }
<b>Hierarchy Level</b>	[edit snmp community <i>community-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
<b>Default</b>	If you omit the <b>clients</b> statement, all SNMP clients using this community string are authorized to access the switch.
<b>Options</b>	<p><b>address</b>—Address of an SNMP client that is authorized to access this switch. You must specify an address, not a hostname. To specify more than one client, include multiple <b>address</b> options.</p> <p><b>restrict</b>—(Optional) Do not allow the specified SNMP client to access the switch.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring SNMP Communities</i></li> </ul>

## commit-delay

<b>Syntax</b>	commit-delay <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit snmp nonvolatile]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the timer for the SNMP <b>Set</b> reply and start of the commit.
<b>Options</b>	<p><b>seconds</b>—Delay between an affirmative SNMP <b>Set</b> reply and start of the commit operation.</p> <p><b>Default:</b> 5 seconds</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Commit Delay Timer</i></li> </ul>

## community (SNMP)

---

**Syntax**    `community community-name {  
                  authorization authorization;  
                  client-list-name client-list-name;  
                  clients {  
                      address restrict;  
                  }  
                  view view-name;  
                  }`

**Hierarchy Level**    [edit snmp]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.



**NOTE:** The **authorization read-write** option is not supported on the QFX3000 QFabric system.

---

The SNMP client application specifies an SNMP community name in **Get**, **GetBulk**, **GetNext**, and **Set** SNMP requests.

**Default**    If you omit the **community** statement, all SNMP requests are denied.

**Options**    **community-name**—Community string. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
                                  snmp-control—To add this statement to the configuration.

**Related Documentation**    • [Configuring the SNMP Community String on page 6601](#)

## community (RMON)

---

<b>Syntax</b>	<code>community <i>community-name</i>;</code>
<b>Hierarchy Level</b>	[edit snmp rmon event <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure the SNMP trap group that is used when generating a trap (if the eventType object is configured to send traps). If that trap group has the rmon-alarm trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of eventCommunity). If nothing is configured, traps are sent to each group that has the rmon-alarm category configured.</p> <p>The event community is not the same as an SNMP community.</p>
<b>Options</b>	<b><i>community-name</i></b> —Name of the trap group that is used when generating a trap if the event is configured to send traps.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## community-name (SNMP)

---

<b>Syntax</b>	<code>community-name <i>community-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 snmp-community <i>community-index</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11. for the QFX Series.
<b>Description</b>	Define an SNMP community to authorize SNMPv1 or SNMPv2c clients in an SNMPv3 system. When you configure a community in SNMPv3, you can also specify a security name. The access privileges associated with the security name determine which MIB objects are available and which operations (read, write, or notify) are allowed on those objects.
<b>Options</b>	<b><i>community-name</i></b> —Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose the name in quotation marks (" ").



**NOTE:** Community names must be unique. You cannot configure the same community name at the `[edit snmp community]` and `[edit snmp v3 snmp-community community-index]` hierarchy levels.

The community name at the `[edit snmp v3 snmp-community community-index]` hierarchy level is encrypted and not displayed in the command-line interface (CLI).

---

<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the SNMPv3 Community</i></li></ul>

## contact

<b>Syntax</b>	<code>contact <i>contact</i>;</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the value of the MIB II <b>sysContact</b> object, which is the contact person for the managed system.
<b>Options</b>	<b>contact</b> —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the System Contact on a Device Running Junos OS</i></li> </ul>

## description (SNMP)

<b>Syntax</b>	<code>description <i>description</i>;</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the value of the MIB II sysDescription object, which is the description of the system being managed.
<b>Default</b>	<p>By default, the sysDescription object includes the following information:  Juniper Networks, Inc. <i>platform</i>, <i>build</i>, Build date: <i>date</i> UTC Copyright (c) <i>date-range</i>  Juniper Networks, Inc.</p> <p>For example:</p> <pre>sysDescr.0 = Juniper Networks, Inc. m7i internet router, kernel JUNOS 13.2-20130530_ib_13_3_psd.1, Build date: 2013-05-30 22:48:07 UTC Copyright (c) 1996-2013 Juniper Networks, Inc.</pre>
<b>Options</b>	<b>description</b> —System description. If the name includes spaces, enclose it in quotation marks (" ").
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the System Description on a Device Running Junos OS</i></li> </ul>

## description (RMON)

---


<b>Syntax</b>	<code>description</code> <i>description</i> ;
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ], [edit snmp rmon event <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Text description of alarm or event.
<b>Options</b>	<b><i>description</i></b> —Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>

## destination-port (SNMP)

---

<b>Syntax</b>	<code>destination-port</code> <i>port-number</i> ;
<b>Hierarchy Level</b>	[edit snmp trap-group]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Assign a trap port number other than the default.
<b>Default</b>	If you omit this statement, the default port is 162.
<b>Options</b>	<b><i>port-number</i></b> —SNMP trap port number.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li></ul>

## engine-id

<b>Syntax</b>	engine-id { (local <i>engine-id-suffix</i>   use-default-ip-address   use-mac-address); }
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Define a unique identifier for an SNMPv3 engine by configuring the suffix of the engine ID. The engine ID is used for identification only and not for addressing. There are two parts of an engine ID: the prefix and the suffix. The prefix is formatted according to the specifications defined in RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> and cannot be configured. The suffix is configured here.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> SNMPv3 authentication and encryption keys are generated based on the associated user passwords and the engine ID. If you configure or change the engine ID, you must commit the user passwords and new engine ID before you configure SNMPv3 users, or the authentication will fail.</p> <p>By default, the engine ID suffix is configured with the MAC address of the management interface (the <i>use-mac-address</i> option) on the QFX Series. You can override this configuration by using the local <i>engine-id-suffix</i> or <i>use-default-ip-address</i> option.</p> </div>
<b>Default</b>	use-mac-address
<b>Options</b>	<p><i>local engine-id-suffix</i>—The engine ID suffix is set based on the data entered.</p> <p><i>use-default-ip-address</i>—The engine ID suffix is generated from the default IP address.</p> <p><i>use-mac-address</i>—The engine ID suffix is generated from the MAC address of the management interface on the switch.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SNMPv3 Overview on page 6523</a></li> <li>• <a href="#">Configuring SNMP on page 1356</a></li> <li>• <a href="#">Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524</a></li> </ul>

## event

---

<b>Syntax</b>	<pre>event <i>index</i> {     <b>community</b> <i>community-name</i>;     <b>description</b> <i>description</i>;     <b>type</b> (RMON Notification) <i>type</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp rmon]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure RMON event entries.
<b>Options</b>	<p><i>index</i>—Identifier for a specific event entry.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>



---

## falling-event-index (RMON)

---

<b>Syntax</b>	<code>falling-event-index <i>index</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon alarm <i>index</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the index number of the event entry that is used when a falling threshold is crossed. You specify the falling-event index when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.
<b>Options</b>	<b><i>index</i></b> —Index of the event entry that is used when a falling threshold is crossed. <b>Range:</b> 0 through 65,535 <b>Default:</b> 0
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>

## falling-threshold (Health Monitor)

---

<b>Syntax</b>	<code>falling-threshold <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit snmp health-monitor]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
<b>Options</b>	<b><i>percentage</i></b> —Lower threshold for the alarm entry. <b>Range:</b> 1 through 100 <b>Default:</b> 70 percent of the maximum possible value
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">rising-threshold on page 1451</a></li><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li></ul>

## falling-threshold (RMON)

---

<b>Syntax</b>	<code>falling-threshold <i>integer</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon alarm <i>index</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the lower threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
<b>Options</b>	<p><b><i>integer</i></b>—Lower threshold for the alarm entry.</p> <p><b>Range:</b> -2,147,483,648 through 2,147,483,647</p> <p><b>Default:</b> 20 percent less than the <b>rising-threshold</b> value</p>
<b>Required Privilege Level</b>	<p><b>snmp</b>—To view this statement in the configuration.</p> <p><b>snmp-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## falling-threshold-interval

---

<b>Syntax</b>	<code>falling-threshold-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon alarm <i>index</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the interval between samples after the rising threshold is exceeded and the value of the sample starts to drop. If the value of the sample drops and exceeds the falling threshold, the regular sampling interval is used.
<b>Options</b>	<b><i>interval</i></b> —Time between samples, in seconds. <b>Range:</b> 1 through 2,147,483,647 seconds <b>Default:</b> 60 seconds
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>

## filter-duplicates

---

<b>Syntax</b>	<code>filter-duplicates;</code>
<b>Hierarchy Level</b>	<code>[edit snmp]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Filter duplicate <b>Get</b> , <b>GetNext</b> , or <b>GetBulk</b> SNMP requests.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 6516</a></li><li>• <a href="#">Example: Configuring SNMP on page 6575</a></li></ul>

## filter-interfaces

<b>Syntax</b>	<code>filter-interfaces {     all-internal-interfaces;     interfaces <i>interface</i> }</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Filter out information related to specific interfaces from the output of SNMP <b>Get</b> and <b>GetNext</b> requests performed on interface-related MIBs.
<b>Options</b>	<p><b>all-internal-interfaces</b>—Filter out information from SNMP <b>Get</b> and <b>GetNext</b> requests for all internal interfaces.</p> <p><b>interfaces</b>—Filter out information from SNMP <b>Get</b> and <b>GetNext</b> requests for the specified interface.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Filtering Interface Information Out of SNMP Get and GetNext Output</i></li> </ul>

## group (Associating a Security Name)

<b>Syntax</b>	<code>group <i>group-name</i>;</code>
<b>Hierarchy Level</b>	[edit snmp v3 vacm security-to-group security-model (usm   v1   v2c) <i>security-name security-name</i> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Associate a security name with a group composed of users with the same access privileges. The security name is used during authentication of SNMP messages, and is mapped to a username.
<b>Options</b>	<b>group-name</b> —Collection of SNMP security names that share the same SNMPv3 access privileges.
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Group</i></li> </ul>

## group (Configuring Access Privileges)

---

**Syntax**    `group group-name {  
                  (default-context-prefix | context-prefix context-prefix){  
                    security-model (any | usm | v1 | v2c) {  
                      security-level (authentication | none | privacy) {  
                        notify-view view-name;  
                        read-view view-name;  
                        write-view view-name;  
                      }  
                    }  
                  }  
                  }  
                  }`

**Hierarchy Level**    [edit snmp v3 vacm access]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The **default-context-prefix** statement, when included, adds all the contexts configured on the device to the group, whereas the **context-prefix context-prefix** statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series.) When the context prefix is specified as default (for example, **context-prefix default**), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as **logical system/routing instance**. For example, to specify routing instance ri1 in logical system ls1, include **context-prefix ls1/ri1**.

The remaining statements under this hierarchy are explained separately.

**Options**    *group-name*—SNMPv3 group name created for the SNMPv3 group.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring the Group*

---

## health-monitor

---

<b>Syntax</b>	health-monitor { falling-threshold <i>percentage</i> ; interval <i>seconds</i> ; rising-threshold <i>percentage</i> ; }
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure health monitoring.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li><li>• <a href="#">Understanding Health Monitoring on page 6529</a></li></ul>

## history

---

<b>Syntax</b>	<pre>history <i>history-index</i> {     <i>bucket-size</i> <i>number</i>;     <i>interface</i> <i>interface-name</i>;     <i>interval</i> <i>seconds</i>;     <i>owner</i> <i>owner-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp rmon]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent on the network to monitor all the traffic flowing among devices on all connected LAN segments. The RMON history feature collects statistics in accordance with user-configurable parameters.</p> <p>The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. If you use the <b>history</b> statement, you must also configure the <b>interface</b> <i>interface-name</i> statement.</p>
<b>Options</b>	<p><b>history-index</b>—Provide a number for this history entry.</p> <p><b>Range:</b> 1 through 65535</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>



---

## interface (SNMP)

---

<b>Syntax</b>	<code>interface [ <i>interface-names</i> ];</code>
<b>Hierarchy Level</b>	<code>[edit snmp]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the interfaces on which SNMP requests can be accepted.
<b>Default</b>	If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.
<b>Options</b>	<i>interface-names</i> —Names of one or more logical interfaces.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 6604</a></li></ul>

## interface (RMON)

---

<b>Syntax</b>	<code>interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon history <i>history-index</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Specify the interface to be monitored in the specified RMON history entry.</p> <p>Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created.</p>
<b>Options</b>	<i>interface-name</i> —Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>

## interval (Health Monitor)

---

<b>Syntax</b>	<code>interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp health-monitor]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the interval between sampling of the object being monitored by the health monitor.
<b>Options</b>	<i>seconds</i> —Time between samples, in seconds. <b>Range:</b> 1 through 2147483647 seconds <b>Default:</b> 300 seconds
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li></ul>

---

## interval (RMON)

---

<b>Syntax</b>	<code>interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ], [edit snmp rmon history <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the interval over which data is to be sampled for the specified alarm or interface.
<b>Default</b>	60 sec for alarm sampling.  1800 sec for history sampling.
<b>Options</b>	<i>seconds</i> —Interval at which data is to be sampled for the specified alarm or interface.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>

## local-engine

---

**Syntax**    local-engine {  
              user *username* {  
                  authentication-md5 {  
                    authentication-password *authentication-password*;  
                  }  
                  authentication-none;  
                  authentication-sha {  
                    authentication-password *authentication-password*;  
                  }  
                  privacy-aes128 {  
                    privacy-password *privacy-password*;  
                  }  
                  privacy-des {  
                    privacy-password *privacy-password*;  
                  }  
                  privacy-3des {  
                    privacy-password *privacy-password*;  
                  }  
                  privacy-none {  
                    privacy-password *privacy-password*;  
                  }  
              }  
          }

**Hierarchy Level**    [edit snmp v3 [usm](#)]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure local engine information for the user-based security model (USM).  
  
                      The remaining statements are explained separately.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
                                  snmp-control—To add this statement to the configuration.

**Related Documentation**    • [Creating SNMPv3 Users on page 6609](#)

## location

---

<b>Syntax</b>	<code>location <i>location</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the value of the MIB II <b>sysLocation</b> object, which is the physical location of the managed system.
<b>Options</b>	<b>location</b> —Location of the local system. You must enclose the name within quotation marks (" ").
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the System Location for a Device Running Junos OS</i></li> </ul>

## message-processing-model

---

<b>Syntax</b>	<code>message-processing-model (v1   v2c   v3);</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 target-parameters <i>target-parameter-name</i> parameters]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the message processing model to be used when generating SNMP notifications.
<b>Options</b>	<b>v1</b> —SNMPv1 message process model.  <b>v2c</b> —SNMPv2c message process model.  <b>v3</b> —SNMPv3 message process model.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the Message Processing Model</i></li> </ul>

## name

---

<b>Syntax</b>	<code>name <i>name</i>;</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the system name from the command-line interface.
<b>Options</b>	<i>name</i> —System name override.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the System Name</i></li></ul>

## nonvolatile

---

<b>Syntax</b>	<code>nonvolatile {     <code>commit-delay</code> <i>seconds</i>; }</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure options for SNMP <b>Set</b> requests.  The statement is explained separately.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Commit Delay Timer</i></li><li>• <i>commit-delay</i></li></ul>

## notify

---

<b>Syntax</b>	<pre> notify <i>name</i> {     tag <i>tag-name</i>;     type (trap   inform); } </pre>
<b>Hierarchy Level</b>	[edit snmp v3]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>type inform</b> option added in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.
<b>Options</b>	<p><b><i>name</i></b>—Name assigned to the notification.</p> <p><b><i>tag-name</i></b>—Notifications are sent to all targets configured with this tag.</p> <p><b><i>type</i></b>—Notification type is <b>trap</b> or <b>inform</b>. Traps are unconfirmed notifications. Informs are confirmed notifications.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Inform Notification Type and Target Address</i></li> <li>• <i>Configuring the SNMPv3 Trap Notification</i></li> </ul>

## notify-filter (Applying to the Management Target)

---

<b>Syntax</b>	<code>notify-filter <i>profile-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 <b>target-parameters</b> <i>target-parameters-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the notify filter applied to a specific set of SNMPv3 target parameters. Target parameters are the message processing and security parameters for notifications sent to a target SNMP manager.
<b>Options</b>	<b><i>profile-name</i></b> —Name of the notify filter to apply to notifications.
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Applying the Trap Notification Filter</i></li></ul>

## notify-filter (Configuring the Profile Name)

---

<b>Syntax</b>	<code>notify-filter <i>profile-name</i> {     oid <i>oid</i> (include   exclude); }</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.
<b>Options</b>	<b><i>profile-name</i></b> —Name assigned to the notify filter.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Trap Notification Filter</i></li><li>• <i>oid (SNMP)</i></li></ul>



## notify-view

---

<b>Syntax</b>	<code>notify-view <i>view-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
<b>Options</b>	<b><i>view-name</i></b> —Name of the view to which the SNMP user group has access.
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views on page 6605</a></li> <li>• <a href="#">Configuring the Notify View</a></li> </ul>

## oid

---

<b>Syntax</b>	<code>oid <i>object-identifier</i> (exclude  include);</code>
<b>Hierarchy Level</b>	<code>[edit snmp view <i>view-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify an object identifier (OID) used to represent a subtree of MIB objects.
<b>Options</b>	<b>exclude</b> —Exclude the subtree of MIB objects represented by the specified OID.  <b>include</b> —Include the subtree of MIB objects represented by the specified OID.  <b><i>object-identifier</i></b> —OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.
<b>Required Privilege Level</b>	<b>snmp</b> —To view this statement in the configuration. <b>snmp-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views on page 6605</a></li> </ul>

## oid (SNMPv3)

---

<b>Syntax</b>	oid <i>oid</i> (include   exclude);
<b>Hierarchy Level</b>	[edit snmp v3 notify-filter <i>profile-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.
<b>Options</b>	<p><b>exclude</b>—Exclude the subtree of MIB objects represented by the specified OID.</p> <p><b>include</b>—Include the subtree of MIB objects represented by the specified OID.</p> <p><b>oid</b>—Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">SNMPv3 Overview on page 6523</a></li><li>• <a href="#">Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524</a></li><li>• <a href="#">Configuring SNMP on page 1356</a></li><li>• <a href="#">Configuring the SNMPv3 Trap Notification</a></li></ul>

## owner

---

<b>Syntax</b>	<code>owner owner-name;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon history index]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the user or group responsible for this RMON history configuration.
<b>Options</b>	<p><b>owner-name</b>—User or group responsible for this configuration.</p> <p><b>Range:</b> 0 through 32 alphanumeric characters</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## parameters

---

<b>Syntax</b>	<pre>parameters {   message-processing-model (v1   v2c   v3);   security-level (none   authentication   privacy);   security-model (usm   v1   v2c);   security-name security-name; }</pre>
<b>Hierarchy Level</b>	<code>[edit snmp v3 target-parameters target-parameters-name]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure a set of target parameters for message processing and security.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Defining and Configuring the Trap Target Parameters</a></li> </ul>

## port (SNMP)

---

<b>Syntax</b>	<code>port <i>port-number</i>;</code>
<b>Hierarchy Level</b>	[edit snmp v3 target-address <i>target-address-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a UDP port number for an SNMP target.
<b>Default</b>	If you omit this statement, the default port is 162.
<b>Options</b>	<i>port-number</i> —Port number for the SNMP target.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Port</i></li></ul>

## privacy-3des

---

<b>Syntax</b>	<pre>privacy-3des {   <b>privacy-password</b> <i>privacy-password</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure the triple Data Encryption Standard (3DES) as the privacy type for the SNMPv3 user.</p>
<b>Options</b>	<p><b>privacy-password</b> <i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> <li>• The password must be at least eight characters long.</li> <li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li> </ul>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the SNMPv3 Encryption Type</i></li> </ul>

## privacy-aes128

---

<b>Syntax</b>	<pre>privacy-aes128 {     <b>privacy-password</b> <i>privacy-password</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> ], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the Advanced Encryption Standard encryption algorithm (CFB128-AES-128 Privacy Protocol) for the SNMPv3 user.
<b>Options</b>	<p><b>privacy-password</b> <i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"><li>• The password must be at least eight characters long.</li><li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li></ul>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the SNMPv3 Encryption Type</i></li></ul>

## privacy-des

<b>Syntax</b>	<code>privacy-des {     <b>privacy-password</b> <i>privacy-password</i>; }</code>
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> ], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the Data Encryption Standard (DES) as the privacy type for the SNMPv3 user.
<b>Options</b>	<p><b>privacy-password</b> <i>privacy-password</i>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"> <li>• The password must be at least eight characters long.</li> <li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li> </ul>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the SNMPv3 Encryption Type</i></li> </ul>

## privacy-none

<b>Syntax</b>	<code>privacy-none;</code>
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> ], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure that no encryption be used for the SNMPv3 user.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the SNMPv3 Encryption Type</i></li> </ul>

## privacy-password

---

<b>Syntax</b>	<code>privacy-password <i>privacy-password</i>;</code>
<b>Hierarchy Level</b>	[edit snmp v3 usm local-engine user <i>username</i> privacy-3des], [edit snmp v3 usm local-engine user <i>username</i> privacy-aes128], [edit snmp v3 usm local-engine user <i>username</i> privacy-des], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-3des], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-aes128], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-des]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a privacy password for the SNMPv3 user.
<b>Options</b>	<p><b><i>privacy-password</i></b>—Password that a user enters. The password is then converted into a key that is used for encryption.</p> <p>SNMPv3 has special requirements when you create plain-text passwords on a router or switch:</p> <ul style="list-style-type: none"><li>• The password must be at least eight characters long.</li><li>• The password can include alphabetic, numeric, and special characters, but it cannot include control characters.</li></ul>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the SNMPv3 Encryption Type</i></li></ul>



---

## read-view

---

<b>Syntax</b>	<code>read-view <i>view-name</i>;</code>
<b>Hierarchy Level</b>	[ <code>edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)</code> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
<b>Options</b>	<b><i>view-name</i></b> —The name of the view to which the SNMP user group has access.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Read View</i></li><li>• <a href="#">Configuring MIB Views on page 6605</a></li></ul>

## remote-engine

---

**Syntax**    `remote-engine engine-id {  
                  user username {  
                    authentication-md5 {  
                      authentication-password authentication-password;  
                    }  
                    authentication-none;  
                    authentication-sha {  
                      authentication-password authentication-password;  
                    }  
                    privacy-aes128 {  
                      privacy-password privacy-password;  
                    }  
                    privacy-des {  
                      privacy-password privacy-password;  
                    }  
                    privacy-3des {  
                      privacy-password privacy-password;  
                    }  
                    privacy-none {  
                      privacy-password privacy-password;  
                    }  
                  }  
                }`

**Hierarchy Level**    [edit snmp v3 usm]

**Release Information**    Statement introduced in Junos OS Release 7.4.  
                              Statement introduced in Junos OS Release 9.0 for EX Series switches.  
                              Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Configure the remote engine information for the user-based security model (USM). To send inform messages to an SNMPv3 user on a remote device, you must configure the engine identifier for the SNMP agent on the remote device where the user resides.

**Options**    *engine-id*—Specify engine identifier in hexadecimal format. Used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host.  
  
              The remaining statements are explained separately.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
                                  snmp-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring the Remote Engine and Remote User*

## request-type

---

<b>Syntax</b>	request-type (get-next-request   get-request   walk-request);
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Extend monitoring to a specific SNMP object instance (get-request), to all object instances belonging to a MIB branch (walk-request), or to the next object instance after the instance specified in the configuration (get-next-request).
<b>Default</b>	walk-request
<b>Options</b>	<p><b>get-next-request</b>—Perform an SNMP get next request.</p> <p><b>get-request</b>—Perform an SNMP get request.</p> <p><b>walk-request</b>—Perform an SNMP walk request.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## retry-count (SNMPv3)

---

<b>Syntax</b>	<code>retry-count <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the retry count for SNMP informs.
<b>Options</b>	<b><i>number</i></b> —Maximum number of times the inform is transmitted if no acknowledgment is received. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded. <b>Default:</b> 3 times
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Informs on page 6614</a></li><li>• <i>timeout</i></li></ul>

## rising-event-index

---

<b>Syntax</b>	<code>rising-event-index <i>index</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">snmp rmon alarm index</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the index of the event entry that is used when a rising alarm threshold is exceeded. The rising-event index is specified when you configure an SNMP RMON alarm. If this value is zero, no event is triggered.
<b>Options</b>	<p><b><i>index</i></b>—Index of the event entry that is used when a rising threshold is exceeded.</p> <p><b>Range:</b> 0 through 65,535</p> <p><b>Default:</b> 0</p>
<b>Required Privilege Level</b>	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## rising-threshold (Health Monitor)

---

<b>Syntax</b>	<code>rising-threshold <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit snmp health-monitor]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.
<b>Options</b>	<p><b><i>percentage</i></b>—Upper threshold for the alarm entry.</p> <p><b>Range:</b> 1 through 100</p> <p><b>Default:</b> 80 percent of the maximum possible value</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Health Monitoring on page 6609</a></li><li>• <a href="#">falling-threshold on page 1415</a></li></ul>

## rising-threshold (RMON)

---

<b>Syntax</b>	<code>rising-threshold <i>integer</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp rmon alarm <i>index</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the upper threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
<b>Options</b>	<p><i>integer</i>—Upper threshold for the alarm entry.</p> <p><b>Range:</b> –2,147,483,648 through 2,147,483,647</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## rmon

```
Syntax  rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
            variable oid-variable;
        }
        event index {
            community community-name;
            description description;
            type (RMON Notification) type;
        }
        history history-index {
            bucket-size number;
            interface interface-name;
            interval seconds;
            owner owner-name;
        }
    }
```

**Hierarchy Level** [edit snmp]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Provide comprehensive network fault diagnosis, planning, and performance tuning information. RMON delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.

Junos OS supports the RMON statistics, history, alarm, and event groups.

The remaining statements are explained separately.

**Default** Disabled.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [RMON MIB Event, Alarm, Log, and History Control Tables on page 6527](#)
- [Monitoring RMON MIB Tables on page 6803](#)
- [Understanding RMON on page 6525](#)



- [Junos OS Network Management Configuration Guide](#)

## sample-type

---

<b>Syntax</b>	sample-type (absolute-value   delta-value);
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the method of sampling the selected variable (monitored object). When you configure an SNMP RMON alarm, you can specify the sample type.
<b>Options</b>	<p><b>absolute-value</b>—Actual value of the selected variable is used when comparing against the thresholds.</p> <p><b>delta-value</b>—Difference between samples of the selected variable is used when comparing against the thresholds.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## security-level (Defining Access Privileges)

---

<b>Syntax</b>	<code>security-level (authentication   none   privacy) {     <b>notify-view</b> <i>view-name</i>;     <b>read-view</b> <i>view-name</i>;     <b>write-view</b> <i>view-name</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>) security-model (any   usm   v1   v2c)]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define the security level used for access privileges.
<b>Default</b>	<code>none</code>
<b>Options</b>	<b>authentication</b> —Provide authentication but no encryption.  <b>none</b> —No authentication and no encryption.  <b>privacy</b> —Provide authentication and encryption.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Security Level</i></li></ul>

---

## security-level (Generating SNMP Notifications)

---

<b>Syntax</b>	security-level (authentication   none   privacy);
<b>Hierarchy Level</b>	[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the security level to use when generating SNMP notifications.
<b>Default</b>	none
<b>Options</b>	<b>authentication</b> —Provide authentication but no encryption.  <b>none</b> —No authentication and no encryption.  <b>privacy</b> —Provide authentication and encryption.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Security Level</i></li></ul>

## security-model (Access Privileges)

---

<b>Syntax</b>	<code>security-model (usm   v1   v2c);</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix   context-prefix <i>context-prefix</i>)]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.
<b>Options</b>	<code>usm</code> —SNMPv3 security model.  <code>v1</code> —SNMPv1 security model.  <code>v2c</code> —SNMPv2c security model.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Security Model</i></li></ul>

## security-model (Group)

<b>Syntax</b>	<pre>security-model (usm   v1   v2c) {   security-name security-name {     group group-name;   } }</pre>
<b>Hierarchy Level</b>	[edit snmp v3 vacm <a href="#">security-to-group</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Define a security model for an SNMPv3 group and associate the security name of a user with the group. All users in the group have the same access privileges.
<b>Options</b>	<p><b>usm</b>—SNMPv3 security model.</p> <p><b>v1</b>—SNMPv1 security model.</p> <p><b>v2c</b>—SNMPv2c security model.</p>
<b>Required Privilege Level</b>	<p><b>snmp</b>—To view this statement in the configuration.</p> <p><b>snmp-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring the Security Model</i></li> </ul>

## security-model (SNMP Notifications)

---

<b>Syntax</b>	security-model (usm   v1   v2c);
<b>Hierarchy Level</b>	[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.
<b>Options</b>	<b>usm</b> —SNMPv3 security model.  <b>v1</b> —SNMPv1 security model.  <b>v2c</b> —SNMPv2c security model.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Security Model</i></li></ul>

## security-name (Community String)

<b>Syntax</b>	<code>security-name <i>security-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 <i>snmp-community</i> <i>community-index</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.
<b>Options</b>	<i>security-name</i> —Name that is used for messaging security and user access control.



**NOTE:** The security name must match the configured security name at the `[edit snmp v3 target-parameters target-parameters-name parameters]` hierarchy level when you configure traps or informs.

<b>Required Privilege</b>	snmp—To view this statement in the configuration.
<b>Level</b>	snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Security Names</i></li> </ul>


## security-name (Security Group)

---

<b>Syntax</b>	<code>security-name security-name {     group group-name; }</code>
<b>Hierarchy Level</b>	[edit snmp v3 vacm security-to-group <b>security-model</b> (usm   v1   v2c)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Associate the security name of a user (for SNMPv3 clients) or a community string (for SNMPv1 and SNMPv2c clients) with a configured security group.
<b>Options</b>	<b>security-name</b> —SNMPv3 secure username configured at the [edit snmp v3 usm local-engine user <b>username</b> ] hierarchy level that is used for messaging security. For SNMPv1 and SNMPv2c, the security name is the community string configured at the [edit snmp v3 snmp-community <b>community-index</b> ] hierarchy level.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Assigning Security Names to Groups</i></li><li>• <a href="#">Assigning a Security Name to a Group on page 6612</a></li></ul>



## security-name (SNMP Notifications)

<b>Syntax</b>	<code>security-name <i>security-name</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the security name used when generating SNMP notifications.
<b>Options</b>	<b><i>security-name</i></b> —If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification.
<div>  <p><b>NOTE:</b> The access privileges for the group associated with this security name must allow this notification to be sent.</p> <p>If you are using the v1 or v2 security models, the security name at the <code>[edit snmp v3 vacm security-to-group]</code> hierarchy level must match the security name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.</p> </div>	
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Security Name</i></li> </ul>

## security-to-group

---

<b>Syntax</b>	<pre>security-to-group {   security-model (usm   v1   v2c) {     group group-name;     security-name security-name;   } }</pre>
<b>Hierarchy Level</b>	[edit snmp v3 vacm]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Assigning Security Model and Security Name to a Group</i></li></ul>

## snmp

```
Syntax  snmp {
    client-list client-list-name {
        ip-addresses;
    }
    community community-name {
        authorization authorization;
        client-list-name client-list-name;
        clients {
            address restrict;
        }
        logical-system logical-system-name {
            routing-instance routing-instance-name {
                clients {
                    addresses;
                }
            }
        }
        routing-instance routing-instance-name {
            clients {
                addresses;
            }
        }
        view view-name;
    }
    contact contact;
    description description;
    filter-duplicates;
    filter-interfaces;
    health-monitor {
        falling-threshold integer;
        interval seconds;
        rising-threshold integer;
    }
    interface [ interface-names ];
    location location;
    name name;
    nonvolatile {
        commit-delay seconds;
    }
    rmon {
        alarm index {
            description description;
            falling-event-index index;
            falling-threshold integer;
            falling-threshold-interval seconds;
            interval seconds;
            request-type;
            rising-event-index index;
            rising-threshold integer;
            sample-type (absolute-value | delta-value);
            startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
            syslog-subtag syslog-subtag;
        }
    }
}
```

```
    variable oid-variable;
  }
  event index {
    community community-name;
    description description;
    type type;
  }
  history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regular-expression>;
  flag flag;
}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance routing-instance-name;
  targets {
    address;
  }
  version (all | v1 | v2);
}
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
  snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance routing-instance-name;
    tag-list tag-list;
    target-parameters target-parameters-name;
```

```

    timeout seconds;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
  usm {
    local-engine {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none;
      }
    }
    remote-engine engine-id {
      user username {
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-3des {
          privacy-password privacy-password;
        }
        privacy-none {
          privacy-password privacy-password;
        }
      }
    }
  }
}

```

```
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
view view-name {
  oid object-identifier (include | exclude);
}
}
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure SNMP.

The remaining statements are explained separately.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding the Implementation of SNMP on page 6513](#)
- [Configuring SNMP on page 1356](#)

## snmp-community

<b>Syntax</b>	snmp-community <i>community-index</i> { <i>community-name</i> <i>community-name</i> ; <i>security-name</i> <i>security-name</i> ; tag <i>tag-name</i> ; }
<b>Hierarchy Level</b>	[edit snmp v3]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the SNMP community which authorizes SNMPv1 or SNMPv2c clients in an SNMPv3 system.
<b>Options</b>	<i>community-index</i> —(Optional) String that identifies an SNMP community.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the SNMPv3 Community</i></li> </ul>

## source-address (SNMP)

<b>Syntax</b>	source-address <i>address</i> ;
<b>Hierarchy Level</b>	[edit snmp trap-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the source address of every SNMP trap packet sent by this switch to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.
<b>Options</b>	<p><i>address</i>—Source address of SNMP traps. You can configure the source address of trap packets two ways: <b>lo0</b> or a valid IPv4 address configured on one of the interfaces. The value <b>lo0</b> indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface <b>lo0</b>.</p> <p><b>Default:</b> Disabled. (The source address is the address of the outgoing interface.)</p>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Source Address for SNMP Traps</i></li> </ul>

## startup-alarm

---

<b>Syntax</b>	startup-alarm (falling-alarm   rising-alarm   rising-or-falling-alarm);
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set an initial alarm that is sent after the configured SNMP RMON alarm becomes active.
<b>Default</b>	rising-or-falling-alarm
<b>Options</b>	<p><b>falling-alarm</b>—Generated if the first sample after the alarm becomes active is equal to or greater than the falling threshold.</p> <p><b>rising-alarm</b>—Generated if the first sample after the alarm becomes active is equal to or greater than the rising threshold.</p> <p><b>rising-or-falling-alarm</b>—Generated if the first sample after the alarm entry becomes active is equal to or greater than either the rising threshold or the falling threshold.</p>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>



## syslog-subtag

---

<b>Syntax</b>	<code>syslog-subtag <i>syslog-subtag</i>;</code>
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Add the <b>syslog-subtag</b> tag to the system log message. The tag should not exceed 80 uppercase characters.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## tag (Configuring Notification Targets)

---

<b>Syntax</b>	<code>tag <i>tag-name</i>;</code>
<b>Hierarchy Level</b>	[edit snmp v3 notify <i>name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a set of target addresses to receive SNMP traps or informs (for IPv4 packets only).
<b>Options</b>	<b>tag-name</b> —Define the target addresses to which an SNMP notification is sent. Target addresses containing the same tag in their tag list are sent the same notification. The <b>tag-name</b> is not included in the notification.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">SNMPv3 Overview on page 6523</a></li> <li>• <a href="#">Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524</a></li> <li>• <a href="#">Configuring SNMP on page 1356</a></li> <li>• <a href="#">Configuring the SNMPv3 Trap Notification</a></li> </ul>

## tag (Configuring the SNMP Community)

---

<b>Syntax</b>	<code>tag tag-name;</code>
<b>Hierarchy Level</b>	[edit snmp v3 snmp-community <i>community-index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure a set of SNMP managers that are authorized to use a community string.
<b>Options</b>	<i>tag-name</i> —Identify the set of addresses for the SNMP managers authorized to use the community string.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">SNMPv3 Overview on page 6523</a></li><li>• <a href="#">Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524</a></li><li>• <a href="#">Configuring SNMP on page 1356</a></li><li>• <a href="#">Configuring the SNMPv3 Trap Notification</a></li></ul>

## tag-list

---

<b>Syntax</b>	<code>tag-list tag-list;</code>
<b>Hierarchy Level</b>	[edit snmp v3 target-address <i>target-address-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure an SNMP tag list used to select target addresses.
<b>Options</b>	<i>tag-list</i> —Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Trap Target Address</a></li></ul>

## target-address

---

<b>Syntax</b>	<pre>target-address <i>target-address-name</i> {   address <i>address</i>;   address-mask <i>address-mask</i>;   port <i>port-number</i>;   retry-count <i>number</i>;   tag-list <i>tag-list</i>;   target-parameters <i>target-parameters-name</i>;   timeout <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit snmp v3]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the address of an SNMP management application and the parameters to be used in sending notifications.
<b>Options</b>	<p><b><i>target-address-name</i></b>—String that identifies the target address.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding the Implementation of SNMP on page 6513</a></li> <li>• <a href="#">SNMP MIBs Support on page 6530</a></li> <li>• <a href="#">SNMP Traps Support on page 6546</a></li> <li>• <a href="#">snmp on page 1454</a></li> <li>• <a href="#">Configuring SNMP on page 1356</a></li> <li>• <a href="#">Monitoring SNMP on page 1475</a></li> <li>• <a href="#">Example: Configuring SNMP on page 6575</a></li> </ul>

## target-parameters

---

**Syntax** At the **[edit snmp v3]** hierarchy level:

```
target-parameters target-parameters-name {  
  profile-name;  
  parameters {  
    message-processing-model (v1 | v2c | V3);  
    security-level (authentication | none | privacy);  
    security-model (usm | v1 | v2c);  
    security-name security-name;  
  }  
}
```

At the **[edit snmp v3 target-address *target-address-name*]** hierarchy level:

```
target-parameters target-parameters-name;
```

**Hierarchy Level** [edit snmp v3]  
[edit snmp v3 target-address *target-address-name*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the **[edit snmp v3]** hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the **[edit snmp v3 target-parameters *target-parameters-name*]** hierarchy level to the target address configuration at the **[edit snmp v3]** hierarchy level.

**Required Privilege Level** snmp—To view this statement in the configuration.  
snmp-control—To add this statement to the configuration.

**Related Documentation**

- *Defining and Configuring the Trap Target Parameters*
- *Applying Target Parameters*

## targets

---

<b>Syntax</b>	<code>targets {     <i>address</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit snmp trap-group <i>group-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure one or more systems to receive SNMP traps.
<b>Options</b>	<b><i>address</i></b> —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li> </ul>

## timeout

---

<b>Syntax</b>	<code>timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the timeout period (in seconds) for SNMP informs.
<b>Default</b>	15 seconds
<b>Options</b>	<b><i>seconds</i></b> —Number of seconds to wait for an inform acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding the Implementation of SNMP on page 6513</a></li> <li>• <a href="#">Configuring SNMP Informs on page 6614</a></li> <li>• <a href="#">retry-count (SNMPv3) on page 6738</a></li> </ul>

## traceoptions (SNMP)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;match <i>regular-expression</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable           no-world-readable&gt;;     flag <i>flag</i>;     no-remote-trace; }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Track the activities of SNMP agents on the switch and record the information in log files.



**NOTE:** The **traceoptions** statement is not supported on the QFabric system.

The output of the tracing operations is placed into log files in the **/var/log** directory. Each log file is named after the SNMP agent that generates it. The following logs are created in the **/var/log** directory when the **traceoptions** statement is used:

- chassisd
- craftd
- ilmid
- mib2d
- rmopd
- serviced
- snmpd

**Options** **file *filename***—By default, the name of the log file that records trace output is the name of the process being traced (for example, mib2d or snmpd). Use this option to specify another name.

**files *number***—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, snmpd) reaches its maximum size, it is archived by being renamed to snmpd.0. The previous snmpd.1 is renamed to snmpd.2, and so on. The oldest archived file is deleted.

**Range:** 2 through 1000 files

**Default:** 10 files

**flag *flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Log all SNMP events.

- **configuration**—Log reading of configuration at the **[edit snmp]** hierarchy level.
- **database**—Log events involving storage and retrieval in the events database.
- **events**—Log important events.
- **general**—Log general events.
- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **policy**—Log policy processing.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **server**—Log communication with processes that are generating events.
- **subagent**—Log subagent restarts.
- **timer-events**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

**match *regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

**size *size***—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

**Range:** 10 KB through 1 GB

**Default:** 1000 KB

**world-readable | no-world-readable**—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

**Required Privilege Level**    snmp—To view this statement in the configuration.  
                                      snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Tracing and Logging Operations on page 6468](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 6807](#)

## trap-group

---

<b>Syntax</b>	<pre>trap-group group-name {     categories {         category;     }     destination-port port-number;     targets {         address;     } }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for QFX Series switches.
<b>Description</b>	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
<b>Options</b>	<p><b>group-name</b>—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li></ul>



---

## trap-options

---

<b>Syntax</b>	<pre>trap-options {     agent-address outgoing-interface;     source-address address; }</pre>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	Disabled
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring SNMP Trap Options</i></li></ul>

## type (RMON Notification)

---

<b>Syntax</b>	<code>type type;</code>
<b>Hierarchy Level</b>	[edit snmp rmon event <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the type of notification generated when a rising or falling threshold is crossed.
<b>Default</b>	<code>log-and-trap</code>
<b>Options</b>	<p><b>type</b>—Type of notification. It can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>log</b>—Add an entry to the <b>logTable</b> object.</li><li>• <b>log-and-trap</b>—Send an SNMP trap and add a log entry.</li><li>• <b>none</b>—No notifications are sent.</li><li>• <b>snmptrap</b>—Send an SNMP trap.</li></ul>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li></ul>

## type (SNMPv3)

---

<b>Syntax</b>	<code>type (inform   trap);</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 notify <i>name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>inform</b> option added in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the type of SNMP notification.
<b>Options</b>	<b>inform</b> —Defines the type of notification as an inform. SNMP informs are confirmed notifications.  <b>trap</b> —Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring SNMP Informs on page 6614</a></li> <li>• <a href="#">Configuring the SNMPv3 Trap Notification</a></li> </ul>

## user

---

<b>Syntax</b>	<code>user <i>username</i>;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 usm local-engine],</code> <code>[edit snmp v3 usm remote-engine <i>engine-id</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a user associated with an SNMPv3 group on a local or remote SNMP engine.
<b>Options</b>	<b><i>username</i></b> —SNMPv3 user-based security model (USM) username.
<b>Required Privilege Level</b>	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating SNMPv3 Users on page 6609</a></li> </ul>

## usm

---

```
Syntax  usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
        remote-engine engine-id {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
    }
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure user-based security model (USM) information.

The remaining statements are explained separately.

**Required Privilege** snmp—To view this statement in the configuration.  
**Level** snmp-control—To add this statement to the configuration.

**Related** • [Creating SNMPv3 Users on page 6609](#)  
**Documentation** • *Configuring the Remote Engine and Remote User*

## v3

---

```
Syntax  v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        port port-number;
        retry-count number;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | v3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
}

usm {
    local-engine {
        user username {
            authentication-md5 {
                authentication-password authentication-password;
            }
            authentication-sha {
                authentication-password authentication-password;
            }
            authentication-none;
            privacy-aes128 {
                privacy-password privacy-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-none;
        }
    }
}
```

```

}
remote-engine engine-id {
  user username {
    authentication-md5 {
      authentication-password authentication-password;
    }
    authentication-sha {
      authentication-password authentication-password;
    }
    authentication-none;
    privacy-aes128 {
      privacy-password privacy-password;
    }
    privacy-des {
      privacy-password privacy-password;
    }
    privacy-3des {
      privacy-password privacy-password;
    }
    privacy-none {
      privacy-password privacy-password;
    }
  }
}
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix) {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
  security-to-group {
    security-model (usm | v1 | v2c) {
      security-name security-name {
        group group-name;
      }
    }
  }
}
}

```

**Hierarchy Level** [edit snmp]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure SNMPv3.

The remaining statements are explained separately.

**Required Privilege** snmp—To view this statement in the configuration.  
**Level** snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524](#)

---

## vacm

**Syntax**

```
vacm {  
  access {  
    group group-name {  
      (default-context-prefix | context-prefix context-prefix){  
        security-model (any | usm | v1 | v2c) {  
          security-level (authentication | none | privacy) {  
            notify-view view-name;  
            read-view view-name;  
            write-view view-name;  
          }  
        }  
      }  
    }  
  }  
  security-to-group {  
    security-model (usm | v1 | v2c);  
    security-name security-name {  
      group group-name;  
    }  
  }  
}
```

**Hierarchy Level** [edit snmp v3]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure view-based access control model (VACM) information, including access privileges such as security model and security level for a group of users.

The remaining statements are explained separately.

**Required Privilege** snmp—To view this statement in the configuration.  
**Level** snmp-control—To add this statement to the configuration.

**Related Documentation**

- [Defining Access Privileges for an SNMP Group](#)



## variable

---


<b>Syntax</b>	<code>variable <i>oid-variable</i>;</code>
<b>Hierarchy Level</b>	[edit snmp rmon alarm <i>index</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Set the object identifier (OID) of the MIB object (also called variable) to be monitored when you configure an SNMP RMON alarm. If the value of the monitored variable exceeds the configured rising threshold or falling threshold, an alarm is triggered and a corresponding event may be generated.
<b>Options</b>	<i>oid-variable</i> —OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1) or the name of the MIB object—for example, <code>ifInOctets.1</code> .
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> </ul>

## version

---

<b>Syntax</b>	version (all   v1   v2);
<b>Hierarchy Level</b>	[edit snmp trap-group <i>group-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the version number of SNMP traps.
<b>Default</b>	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.
<b>Options</b>	all—Send an SNMPv1 and SNMPv2 trap for every trap condition.  v1—Send SNMPv1 traps only.  v2—Send SNMPv2 traps only.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring SNMP Trap Groups on page 6602</a></li></ul>

## view (Configuring a MIB View)

<b>Syntax</b>	<code>view <i>view-name</i> {     oid <i>object-identifier</i> (include   exclude); }</code>
<b>Hierarchy Level</b>	[edit snmp]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The <b>view</b> statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the <b>view</b> statement at the <b>[edit snmp community <i>community-name</i>]</b> hierarchy level.
<div>  <b>NOTE:</b> To remove an OID completely, use the <code>delete view all oid oid-number</code> command but omit the <code>include</code> parameter. </div>	
<b>Options</b>	<p><b><i>view-name</i></b>—Name of the view.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring MIB Views on page 6605</a></li> <li>• <a href="#">Associating MIB Views with an SNMP User Group</a></li> <li>• <a href="#">community on page 1411</a></li> </ul>

## view (Associating MIB View with a Community)

---

<b>Syntax</b>	<code>view view-name;</code>
<b>Hierarchy Level</b>	<code>[edit snmp community community-name]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Associate a view with a community. A view represents a group of MIB objects.
<b>Options</b>	<b>view-name</b> —Name of the view. You must use a view name already configured in the <b>view</b> statement at the <b>[edit snmp]</b> hierarchy level.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring SNMP Communities</i></li></ul>

## write-view

---

<b>Syntax</b>	<code>write-view view-name;</code>
<b>Hierarchy Level</b>	<code>[edit snmp v3 vacm access group group-name (default-context-prefix   context-prefix context-prefix) security-model (any   usm   v1   v2c) security-level (authentication   none   privacy)]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series switches.
<b>Description</b>	Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
<b>Options</b>	<b>view-name</b> —Name of the view for which the SNMP user group has write permission.
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring MIB Views on page 6605</a></li><li>• <i>Configuring the Write View</i></li></ul>

## Configuration Statements for System Log Messages

---

- [archive \(All System Log Files\) on page 6780](#)
- [archive \(Individual System Log File\) on page 6782](#)
- [archive \(QFabric System\) on page 6783](#)

- [console \(System Logging\) on page 6784](#)
- [explicit-priority on page 6785](#)
- [facility-override on page 6785](#)
- [file \(QFabric System\) on page 6786](#)
- [file \(System Logging\) on page 6787](#)
- [files on page 6788](#)
- [host \(System\) on page 6789](#)
- [log-prefix \(System\) on page 6791](#)
- [match on page 6791](#)
- [size \(System\) on page 6792](#)
- [structured-data on page 6793](#)
- [syslog \(System\) on page 6794](#)
- [syslog \(QFabric System\) on page 6796](#)
- [time-format on page 6797](#)
- [user \(System Logging\) on page 6798](#)

## archive (All System Log Files)

---

<b>Syntax</b>	<code>archive &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;start-time <i>time</i>&gt; &lt;transfer-interval <i>interval</i>&gt; &lt;binary-data   no-binary-data&gt;; &lt;world-readable   no-world-readable&gt; ;</code>
<b>Hierarchy Level</b>	[edit system <a href="#">syslog</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure archiving properties for all system log files.
<b>Options</b>	<p><b>files <i>number</i></b>—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <b>logfile</b>, it closes the file, compresses it, and renames it <b>logfile.0.gz</b> (the amount of data is determined by the <b>size</b> statement at this hierarchy level). The utility then opens and writes to a new file called <b>logfile</b>. When the new file reaches the maximum size, the <b>logfile.0.gz</b> file is renamed to <b>logfile.1.gz</b>, and the new file is closed, compressed, and renamed <b>logfile.0.gz</b>. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).</p> <p><b>Range:</b> 1 through 1000</p> <p><b>Default:</b> 10 files</p> <p><b>size <i>size</i></b>—Maximum amount of data that the Junos OS logging utility writes to a log file <b>logfile</b> before archiving it (closing it, compressing it, and changing its name to <b>logfile.0.gz</b>). The utility then opens and writes to a new file called <b>logfile</b>.</p> <p><b>Syntax:</b> <i>x k</i> to specify the number of kilobytes, <i>x m</i> for the number of megabytes, or <i>x g</i> for the number of gigabytes</p> <p><b>Range:</b> 64 KB through 1 GB</p> <p><b>Default:</b></p> <ul style="list-style-type: none"><li>• 128 KB for EX Series switches and J Series routers</li><li>• 1 MB for M Series, MX Series, and T Series routers, and the QFX3500 switch</li><li>• 10 MB for TX Matrix and TX Matrix Plus routers</li></ul> <p><b>binary-data   no-binary-data</b>—Mark file as containing binary data. This allows proper archiving of binary files, such as WTMP files (login records for UNIX based systems)..</p> <p><b>Default:</b> no-binary-data</p> <p><b>world-readable   no-world-readable</b>—Grant all users permission to read archived log files, or restrict the permission only to the <b>root</b> user and users who have the Junos OS <b>maintenance</b> permission.</p> <p><b>Default:</b> no-world-readable</p>

**Required Privilege** system—To view this statement in the configuration.  
**Level** system-control—To add this statement to the configuration.

**Related Documentation** • [Specifying Log File Size, Number, and Archiving Properties on page 6629](#)

## archive (Individual System Log File)

---

Syntax	archive <archive-sites ( <i>ftp-url</i> <password <i>password</i> >)> <files <i>number</i> > <size <i>size</i> > <start-time "YYYY-MM-DD.hh:mm"> <transfer-interval <i>minutes</i> > <world-readable   no-world-readable>;
Hierarchy Level	[edit system <b>syslog file</b> <i>filename</i> ]
Release Information	Statement introduced before Junos OS Release 7.4. <b>start-time</b> and <b>transfer-interval</b> statements introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure archiving properties for a specific system log file.
Options	<p><b>archive-sites</b> <i>site-name</i>—FTP URL representing the destination for the archived log file (for information about how to specify valid FTP URLs, see <a href="#">“Format for Specifying Filenames and URLs in Junos OS CLI Commands” on page 42</a>). If more than one site name is configured, a list of archive sites for the system log files is created. When a file is archived, the router attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with the filename specified at the <b>[edit system syslog]</b> hierarchy level.</p> <p><b>files</b> <i>number</i>—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i>, it closes the file, compresses it, and renames it <i>logfile.0.gz</i> (the amount of data is determined by the <b>size</b> statement at this hierarchy level). The utility then opens and writes to a new file called <i>logfile</i>. When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i>, and the new file is closed, compressed, and renamed <i>logfile.0.gz</i>. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).</p> <p><b>Range:</b> 1 through 1000</p> <p><b>Default:</b> 10 files</p> <p><b>password</b> <i>password</i>—Password for authenticating with the site specified by the <b>archive-sites</b> statement.</p> <p><b>size</b> <i>size</i>—Maximum amount of data that the Junos OS logging utility writes to a log file <i>logfile</i> before archiving it (closing it, compressing it, and changing its name to <i>logfile.0.gz</i>). The utility then opens and writes to a new file called <i>logfile</i>.</p> <p><b>Syntax:</b> <b>xk</b> to specify the number of kilobytes, <b>xm</b> for the number of megabytes, or <b>xg</b> for the number of gigabytes</p> <p><b>Range:</b> 64 KB through 1 GB</p> <p><b>Default:</b> 128 KB for J Series routers; 1 MB for M Series, MX Series, and T Series routers, and the QFX3500 switch; 10 MB for TX Matrix and TX Matrix Plus routers</p>



**start-time "YYYY-MM-DD.hh:mm"**—Date and time in the local time zone for a one-time transfer of the active log file to the first reachable site in the list of sites specified by the **archive-sites** statement.

**transfer-interval *interval***—Interval at which to transfer the log file to an archive site.

**Range:** 5 through 2880 minutes

**world-readable | no-world-readable**—Grant all users permission to read archived log files, or restrict the permission only to the **root** user and users who have the Junos OS **maintenance** permission.

**Default:** no-world-readable

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Specifying Log File Size, Number, and Archiving Properties on page 6629](#)

## archive (QFabric System)

**Syntax** archive {  
size *size*;  
}

**Hierarchy Level** [edit system [syslog](#) file *filename*]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure the archiving properties for the system message log file.

**Options** **size *size***—Maximum amount of system log message data that the QFabric system stores in the log file.

**Syntax:** *xk* to specify the number of kilobytes, *xm* for the number of megabytes, or *xg* for the number of gigabytes

**Range:** 65 KB through 1 GB

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [syslog on page 6796](#)

## console (System Logging)

---

<b>Syntax</b>	<code>console {     <i>facility severity</i>; }</code>
<b>Hierarchy Level</b>	[edit system <a href="#">syslog</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the logging of system messages to the system console.
<b>Options</b>	<p><b><i>facility</i></b>—Class of messages to log. To specify multiple classes, include multiple <b><i>facility severity</i></b> statements. For a list of the facilities, see <a href="#">“Junos OS System Logging Facilities and Message Severity Levels”</a> on page 6633.</p> <p><b><i>severity</i></b>—Severity of the messages that belong to the facility specified by the paired <b><i>facility</i></b> name. Messages with severities of the specified level and higher are logged. For a list of the severities, see <a href="#">“Junos OS System Logging Facilities and Message Severity Levels”</a> on page 6633.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Directing System Log Messages to the Console</a> on page 6620</li><li>• <a href="#">Junos OS System Log Messages Reference</a></li></ul>

## explicit-priority

<b>Syntax</b>	explicit-priority;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ], [edit logical-systems <i>logical-system-name</i> system syslog host], [edit system syslog file <i>filename</i> ], [edit system syslog host]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination.  When the <b>structured-data</b> statement is also included at the [edit system syslog file <i>filename</i> ] hierarchy level, this statement is ignored for the file.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Including Priority Information in System Log Messages on page 6622</a></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">structured-data on page 6793</a></li> </ul>

## facility-override

<b>Syntax</b>	facility-override <i>facility</i> ;
<b>Hierarchy Level</b>	[edit system syslog host]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Substitute an alternate facility for the default facilities used when messages are directed to a remote destination.
<b>Options</b>	<i>facility</i> —Alternate facility to substitute for the default facilities. For a list of the possible facilities, see <a href="#">Table 638 on page 6635</a> .
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Changing the Alternative Facility Name for System Log Messages Directed to a Remote Destination on page 6636</a></li> <li>• <i>Junos OS System Log Messages Reference</i></li> </ul>

## file (QFabric System)

---

Syntax	<pre>file <i>filename</i> {   archive {     <i>size</i> <i>maximum-file-size</i>;   }   <i>explicit-priority</i>;   <i>facility</i> <i>severity</i>;   <i>match</i> "<i>regular-expression</i>";   <i>structured-data</i> {     <i>brief</i>;   } }</pre>
Hierarchy Level	[edit system <a href="#">syslog</a> ]
Release Information	Statement introduced in Junos OS Release 11.3 for the QFX Series.
Description	Configure the logging of system messages to a file.
Options	<p><i>facility</i>—Class of messages to log. To specify multiple classes, include multiple <i>facility severity</i> statements.</p> <p><i>filename</i>—Filename that you specify with the <b>show log</b> command.</p> <p><b>Default:</b> Filename <b>messages</b></p> <p><i>severity</i>—Severity of the messages that belong to the facility specified by the paired <i>facility</i> name. Messages with severities at the specified level and higher are logged.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">syslog on page 6796</a></li></ul>

## file (System Logging)

<b>Syntax</b>	<pre> file <i>filename</i> {     <i>facility severity</i>;     archive {         files <i>number</i>;         size <i>size</i>;         (no-world-readable   world-readable);     }     explicit-priority;     match "<i>regular-expression</i>";     structured-data {         brief;     } } </pre>
<b>Hierarchy Level</b>	[edit system <a href="#">syslog</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the logging of system messages to a file.
<b>Options</b>	<p><b><i>facility</i></b>—Class of messages to log. To specify multiple classes, include multiple <b><i>facility severity</i></b> statements. For a list of the facilities, see <a href="#">“Junos OS System Logging Facilities and Message Severity Levels” on page 6633</a>.</p> <p><b><i>file filename</i></b>—File in the <b><i>severity</i></b> directory in which to log messages from the specified facility. To log messages to more than one file, include more than one <b><i>file</i></b> statement.</p> <p><b><i>severity</i></b>—Severity of the messages that belong to the facility specified by the paired <b><i>facility</i></b> name. Messages with severities of the specified level and higher are logged. For a list of the severities, see <a href="#">“Junos OS System Logging Facilities and Message Severity Levels” on page 6633</a>.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Directing System Log Messages to a Log File on page 6618</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> </ul>

## files

---

<b>Syntax</b>	<code>files <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit system syslog archive], [edit system syslog file <i>filename</i> <a href="#">archive</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series switches.
<b>Description</b>	Configure the maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file <i>logfile</i> , it closes the file, compresses it, and renames it to <i>logfile.0.gz</i> (for information about the maximum file size, see <a href="#">size</a> ). The utility then opens and writes to a new file called <i>logfile</i> . When the new file reaches the maximum size, the <i>logfile.0.gz</i> file is renamed to <i>logfile.1.gz</i> , and the new file is closed, compressed, and renamed <i>logfile.0.gz</i> . By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).
<b>Options</b>	<i>number</i> —Maximum number of archived files. <b>Range:</b> 1 through 1000 <b>Default:</b> 10 files
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Specifying Log File Size, Number, and Archiving Properties on page 6629</a></li><li>• <a href="#">Junos OS System Log Messages Reference</a></li><li>• <a href="#">size on page 6792</a></li></ul>

## host (System)

<b>Syntax</b>	<pre> host (hostname   other-routing-engine) {     facility severity;     exclude-hostname     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     source-address source-address;     structured-data {         brief;     } } </pre>
<b>QFX Series</b>	<pre> host (hostname {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>
<b>TX Matrix Router and EX Series Switches</b>	<pre> host (hostname   other-routing-engine   scc-master) {     facility severity;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>
<b>TX Matrix Plus Router</b>	<pre> host (hostname   other-routing-engine   sfc0-master) {     facility severity;     allow-duplicates;     explicit-priority;     facility-override facility;     log-prefix string;     match "regular-expression";     port;     source-address source-address; } </pre>
<b>Hierarchy Level</b>	<pre> [edit logical-systems logical-system-name system syslog], [edit system syslog] </pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Configure the logging of system messages to a remote destination.

**Options** *facility*—Class of messages to log. To specify multiple classes, include multiple *facility severity* statements. For a list of the facilities, see [“Junos OS System Logging Facilities and Message Severity Levels” on page 6633](#).

*hostname*—IPv4 address, IPv6 address, or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a **host** statement for each one.

**other-routing-engine**—Direct messages to the other Routing Engine on a router or switch with two Routing Engines installed and operational.



**NOTE:** The **other-routing-engine** option is not applicable to the QFX Series.

---

**port**—Port number of the remote syslog server that can be modified.

**scc-master**—(TX Matrix routers only) On a T640 router that is part of a routing matrix, direct messages to the TX Matrix router.

*severity*—Severity of the messages that belong to the facility specified by the paired *facility* name. Messages with severities of the specified level and higher are logged. For a list of the severities, see [“Junos OS System Logging Facilities and Message Severity Levels” on page 6633](#).

**sfc0-master**—(TX Matrix Plus routers only) On a T1600 or T4000 router that is part of a routing matrix, direct messages to the TX Matrix Plus router.

The remaining statements are explained separately.

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <i>Directing System Log Messages to a Remote Machine or the Other Routing Engine</i></li><li>• <i>Directing Messages to a Remote Destination from the Routing Matrix Based on the TX Matrix Router</i></li><li>• <i>Directing Messages to a Remote Destination from the Routing Matrix Based on a TX Matrix Plus Router</i></li><li>• <i>Junos OS System Log Messages Reference</i></li></ul> |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|



## log-prefix (System)

<b>Syntax</b>	<code>log-prefix <i>string</i>;</code>
<b>Hierarchy Level</b>	[edit system syslog host]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Include a text string in each message directed to a remote destination.
<b>Options</b>	<i>string</i> —Text string to include in each message.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Adding a Text String to System Log Messages Directed to a Remote Destination on page 6617</a></li> <li>• <i>Junos OS System Log Messages Reference</i></li> </ul>

## match

<b>Syntax</b>	<code>match "regular-expression";</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ], [edit logical-systems <i>logical-system-name</i> system syslog user ( <i>username</i>   *)], [edit system syslog file <i>filename</i> ], [edit system syslog host <i>hostname</i>   other-routing-engine  scc-master)], [edit system syslog user ( <i>username</i>   *)]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Using Regular Expressions to Refine the Set of Logged Messages on page 6637</a></li> </ul>


## size (System)

---

<b>Syntax</b>	<code>size size;</code>
<b>Hierarchy Level</b>	[edit system syslog archive], [edit system syslog file <i>filename</i> archive]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the maximum amount of data that the Junos OS logging utility writes to a log file <b>logfile</b> before archiving it (closing it, compressing it, and changing its name to <b>logfile.0.gz</b> ). The utility then opens and writes to a new file called <b>logfile</b> . For information about the number of archive files that the utility creates in this way, see <a href="#">files</a> .
<b>Options</b>	<b>size</b> —Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). <b>Syntax:</b> <b>xk</b> to specify the number of kilobytes, <b>xm</b> for the number of megabytes, or <b>xg</b> for the number of gigabytes <b>Range:</b> 64 KB through 1 GB <b>Default:</b> 1 MB for MX Series routers and the QFX Series
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Specifying Log File Size, Number, and Archiving Properties on page 6629</a></li><li>• <a href="#">Junos OS System Log Messages Reference</a></li><li>• <a href="#">files on page 6788</a></li></ul>

## structured-data

---

<b>Syntax</b>	structured-data { brief; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> system syslog file <i>filename</i> ], [edit system syslog file <i>filename</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, <i>The syslog Protocol</i> ( <a href="http://tools.ietf.org/html/draft-ietf-syslog-protocol-23">http://tools.ietf.org/html/draft-ietf-syslog-protocol-23</a> ).
<div>  <p><b>NOTE:</b> When this statement is included, other statements that specify the format for messages written to the file are ignored (the <code>explicit-priority</code> statement at the [edit system syslog file <i>filename</i>] hierarchy level and the <code>time-format</code> statement at the [edit system syslog] hierarchy level).</p> </div>	
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Logging Messages in Structured-Data Format</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">explicit-priority on page 6785</a></li> <li>• <a href="#">time-format on page 6797</a></li> </ul>

## syslog (System)

---

```
Syntax  syslog {
        allow-duplicates;
        archive {
            (binary-data | no-binary-data);
            files number;
            size maximum-file-size;
            start-time "YYYY-MM-DD.hh:mm";
            transfer-interval minutes;
            (world-readable | no-world-readable);
        }
        console {
            facility severity;
        }
        file filename {
            facility severity;
            explicit-priority;
            match "regular-expression";
            archive {
                (binary-data | no-binary-data);
                files number;
                size maximum-file-size;
                start-time "YYYY-MM-DD.hh:mm";
                transfer-interval minutes;
                (world-readable | no-world-readable);
            }
            structured-data {
                brief;
            }
        }
        host (hostname | other-routing-engine | scc-master) {
            facility severity;
            explicit-priority;
            facility-override facility;
            log-prefix string;
            match "regular-expression";
            source-address source-address;
            structured-data {
                brief;
            }
            port port number;
        }
        log-rotate-frequency frequency;
        server server name;
        source-address source-address;
        time-format (millisecond | year | year millisecond);
        user (username | *) {
            facility severity;
            match "regular-expression";
        }
    }
```

Hierarchy Level    [edit system]


<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	<p>Configure the types of system log messages to send to files, to a remote destination, to user terminals, or to the system console.</p> <p>The remaining statements are explained separately.</p>
<b>Options</b>	<p><b>archive</b>—Define parameters for archiving log messages.</p> <p><b>console</b>—Send log messages of a specified class and severity to the console.</p> <p><b>file</b>—Send log messages to a named file.</p> <p><b>host</b> —Remote location to be notified of specific log messages.</p> <p><b>log-rotate-frequency</b>—Configure the interval for checking logfile size and archiving messages.</p> <p><b>server</b>—Name of the system log server in the inet.0 routing instance.</p> <p><b>source-address</b>—Include a specified address as the source address for log messages.</p> <p><b>time-format</b>—Additional information to include in the system log time stamp.</p> <p><b>user</b>—Notify a specific user of the log event.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Junos OS System Log Overview</i></li> <li>• <i>Junos OS System Log Messages Reference</i></li> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6561</a></li> </ul>

## syslog (QFabric System)

---

<b>Syntax</b>	<pre>syslog {   file <i>filename</i> {     archive {       size <i>maximum-file-size</i>;     }     explicit-priority;     <i>facility severity</i>;     match "<i>regular-expression</i>";     structured-data;   }   filter all {     <i>facility severity</i>;     match "<i>regular-expression</i>";   }   host <i>hostname</i> {     explicit-priority;     <i>facility severity</i>;     facility-override <i>facility</i>;     log-prefix <i>string</i>;     match "<i>regular-expression</i>";     structured-data;   } }</pre>
<b>Hierarchy Level</b>	[edit system]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	<p>Configure system log messages for the QFabric system.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding the Implementation of System Log Messages on the QFabric System on page 6562</a></li><li>• <a href="#">Directing System Log Messages to a Remote Machine on page 6619</a></li></ul>

## time-format

<b>Syntax</b>	<code>time-format (year   millisecond   year millisecond);</code>
<b>Hierarchy Level</b>	<code>[edit system syslog]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a <b>file</b>, <b>console</b>, or <b>user</b> statement at the <code>[edit system syslog]</code> hierarchy level. As of Junos OS Release 11.4, the additional time information is also sent to destinations configured by a <b>host</b> statement.</p> <p>By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, <b>Aug 21 12:36:30</b>. However, the timestamp for traceoption messages is specified in milliseconds by default, and is independent of the <code>[edit system syslog time-format]</code> statement.</p>
	<p> <b>NOTE:</b> When the <b>structured-data</b> statement is included at the <code>[edit system syslog file <i>filename</i>]</code> hierarchy level, this statement is ignored for the file.</p>
<b>Options</b>	<p><b>millisecond</b>—Include the millisecond in the timestamp.</p> <p><b>year</b>—Include the year in the timestamp.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Including the Year or Millisecond in Timestamps on page 181</a></li> <li>• <a href="#">Junos OS System Log Messages Reference</a></li> <li>• <a href="#">structured-data on page 6793</a></li> </ul>

## user (System Logging)

---

<b>Syntax</b>	<pre>user (username   *) {     facility severity;     match "regular-expression"; }</pre>
<b>Hierarchy Level</b>	[edit system syslog]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Configure the logging of system messages to user terminals.
<b>Options</b>	<p><b>*</b> (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.</p> <p><b>facility</b>—Class of messages to log. To specify multiple classes, include multiple <b>facility severity</b> statements. For a list of the facilities, see <a href="#">“Junos OS System Logging Facilities and Message Severity Levels” on page 6633</a>.</p> <p><b>severity</b>—Severity of the messages that belong to the facility specified by the paired <b>facility</b> name. Messages with severities the specified level and higher are logged. For a list of the severities, see <a href="#">“Junos OS System Logging Facilities and Message Severity Levels” on page 6633</a>.</p> <p><b>username</b>—Junos OS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one <b>user</b> statement.</p> <p>The remaining statement is explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Directing System Log Messages to a User Terminal on page 6620</a></li><li>• <a href="#">Junos OS System Logging Facilities and Message Severity Levels on page 6633</a></li><li>• <a href="#">Junos OS System Log Messages Reference</a></li></ul>



## CHAPTER 78

# Administration

- [Monitoring Tasks on page 6799](#)
- [Commands for General Monitoring on page 6814](#)
- [Commands for Network Analytics on page 6828](#)
- [Commands for sFlow Technology on page 6849](#)
- [Commands for SNMP on page 6855](#)
- [Commands for Syslog on page 6888](#)

### Monitoring Tasks

---

- [Displaying a Log File from a Single-Chassis System on page 6799](#)
- [Monitoring Traffic Through the Router or Switch on page 6800](#)
- [Monitoring RMON MIB Tables on page 6803](#)
- [Monitoring SNMP on page 6804](#)
- [Monitoring System Log Messages on page 6805](#)
- [Pinging Hosts on page 6806](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 6807](#)
- [Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage on page 6810](#)
- [Displaying Commit Script Output on page 6812](#)

### Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
```

```
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysApp1ElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysApp1ElmtRunMemory.5.8.2292)
...
Nov  4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov  4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```
user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
SNMP trap: cold start
...
```

#### Related Documentation

- [Interpreting Messages Generated in Standard Format on page 6628](#)
- [Interpreting Messages Generated in Structured-Data Format on page 6625](#)

## Monitoring Traffic Through the Router or Switch

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. [Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 6800](#)
2. [Displaying Real-Time Statistics About an Interface on the Router or Switch on page 6801](#)

### Displaying Real-Time Statistics About All Interfaces on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through all interfaces on the router or switch.

**Action** To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

### Sample Output

```
user@host> monitor interface traffic
host name          Seconds:15          Time: 12:31:09
Interface  Link  Input packets  (pps)  Output packets  (pps)
so-1/0/0   Down    0             (0)      0             (0)
so-1/1/0   Down    0             (0)      0             (0)
so-1/1/1   Down    0             (0)      0             (0)
so-1/1/2   Down    0             (0)      0             (0)
so-1/1/3   Down    0             (0)      0             (0)
t3-1/2/0   Down    0             (0)      0             (0)
t3-1/2/1   Down    0             (0)      0             (0)
t3-1/2/2   Down    0             (0)      0             (0)
t3-1/2/3   Down    0             (0)      0             (0)
so-2/0/0   Up      211035        (1)      36778         (0)
so-2/0/1   Up      192753        (1)      36782         (0)
so-2/0/2   Up      211020        (1)      36779         (0)
so-2/0/3   Up      211029        (1)      36776         (0)
so-2/1/0   Up      189378        (1)      36349         (0)
so-2/1/1   Down    0             (0)      18747         (0)
so-2/1/2   Down    0             (0)      16078         (0)
so-2/1/3   Up      0             (0)      80338         (0)
at-2/3/0   Up      0             (0)      0             (0)
at-2/3/1   Down    0             (0)      0             (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

**Meaning** The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the C key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

### Displaying Real-Time Statistics About an Interface on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through an interface on the router or switch.

**Action** To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface interface-name
```

### Sample Output

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C3 Traffic statistics:
  Input bytes:          5856541 (88 bps)
  Output bytes:         6271468 (96 bps)
  Input packets:        157629 (0 pps)
  Output packets:       157024 (0 pps)
Encapsulation statistics:
```

```

Input keepalives:          42353
Output keepalives:        42320
LCP state: Opened
Error statistics:
Input errors:              0
Input drops:               0
Input framing errors:      0
Input runs:                0
Input giants:              0
Policed discards:          0
L3 incompletes:            0
L2 channel errors:         0
L2 mismatch timeouts:      0
Carrier transitions:       1
Output errors:             0
Output drops:              0
Aged packets:              0
Active alarms : None
Active defects: None
SONET error counts/seconds:
LOS count                  1
LOF count                  1
SEF count                  1
ES-S                       77
SES-S                      77
SONET statistics:
BIP-B1                     0
BIP-B2                     0
REI-L                      0
BIP-B3                     0
REI-P                      0
Received SONET overhead: F1      : 0x00 J0      : 0xZ

```

**Meaning** The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in [Table 640 on page 6802](#).

**Table 640: Output Control Keys for the monitor interface Command**

Action	Key
Display information about the next interface. The <b>monitor interface</b> command scrolls through the physical or logical interfaces in the same order that they are displayed by the <b>show interfaces terse</b> command.	<b>N</b>
Display information about a different interface. The command prompts you for the name of a specific interface.	<b>I</b>
Freeze the display, halting the display of updated statistics.	<b>F</b>
Thaw the display, resuming the display of updated statistics.	<b>T</b>

**Table 640: Output Control Keys for the monitor interface Command (*continued*)**

Action	Key
Clear (zero) the current delta counters since <b>monitor interface</b> was started. It does not clear the accumulative counter.	C
Stop the <b>monitor interface</b> command.	Q

See the [CLI Explorer](#) for details on using match conditions with the **monitor traffic** command.

## Monitoring RMON MIB Tables

**Purpose** Monitor remote monitoring (RMON) alarm, event, and log tables.

**Action** To display the RMON tables:

```
user@switch> show snmp rmon
Alarm
Index  Variable description                               Value State
      5 monitor
      jnxOperatingCPU.9.1.0.0                        5 falling threshold

Event
Index  Type                               Last Event
      1 log and trap                     2010-07-10 11:34:17 PDT
Event Index: 1
      Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
      Time: 2010-07-10 11:34:07 PDT
      Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
      Time: 2010-07-10 11:34:17 PDT
```

**Meaning** The display shows that an alarm has been defined to monitor jnxRmon MIB object jnxOperatingCPU, which represents the CPU utilization of the Routing Engine. The alarm is configured to generate an event that sends an SNMP trap and adds an entry to the logTable in the RMON MIB. The log table shows that two occurrences of the event have been generated—one for rising above a threshold of 90 percent, and one for falling below a threshold of 75 percent.

**Related Documentation**

- [Configuring RMON Alarms and Events on page 6606](#)
- [show snmp rmon on page 6877](#)
- [show snmp rmon history on page 6881](#)
- [clear snmp statistics on page 6857](#)
- [clear snmp history on page 6856](#)

## Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.
- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
Alarm
Index  Variable description                                     Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                             58 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                             0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                               0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0                            35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon   0 active
      Chassis daemon                                       50 active
      Firewall daemon                                       0 active
      Interface daemon                                      5 active
      SNMP daemon   11 active
      MIB2 daemon   42 active
      ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system

sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx

Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0   = 24444184
sysContact.0  = J Smith
sysName.0     = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics
```

SNMP statistics:

Input:

```
Packets: 0, Bad versions: 0, Bad community names: 0,
Bad community uses: 0, ASN parse errors: 0,
Too bigs: 0, No such names: 0, Bad values: 0,
Read onlys: 0, General errors: 0,
Total request varbinds: 0, Total set varbinds: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0,
Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
Throttle drops: 0, Duplicate request drops: 0
```

Output:

```
Packets: 0, Too bigs: 0, No such names: 0,
Bad values: 0, General errors: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0
```

- Related Documentation**
- [health-monitor on page 1416](#)
  - [show snmp mib on page 6874](#)
  - [show snmp statistics on page 1503](#)

## Monitoring System Log Messages

**Purpose** Display system log messages about the QFX Series. By looking through a system log file for any entries pertaining to the interface that you are interested in, you can further investigate a problem with an interface on the switch.

**Action** To view system log messages:

```
user@switch1> show log messages
```

## Sample Output

```
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Management process): new instance detected (variable:
sysApp1ElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Command-line interface): new instance detected (variable:
sysApp1ElmtRunMemory.5.8.2292)
...
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
```

```
'exit '  
Nov 4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting  
configuration mode  
Nov 4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command  
'show log messages'
```

**Meaning** The sample output shows the following entries in the **messages** file:

- A new log file was created when the previous file reached the maximum size of 128 kilobytes (KB).
- The fan speed for Fan 1, 2, and 3 is set at 65 percent.
- Health monitoring activity is detected.
- CLI commands were entered by the user jsmith.

**Related  
Documentation**

- [Overview of Junos OS System Log Messages on page 6560](#)
- [Understanding the Implementation of System Log Messages on the QFabric System on page 6562](#)
- [Example: Configuring System Log Messages on page 6568](#)
- [clear log on page 350](#)
- [show log on page 948](#)
- [syslog on page 313](#)

## Pinging Hosts

**Purpose** Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The switch sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to a specified host and receives ICMP echo responses.

**Action** To use the **ping** command to send four requests (ping count) to host3:  
**ping host count number**

## Sample Output

```
ping host3 count 4  
user@switch> ping host3 count 4  
PING host3.site.net (176.26.232.111): 56 data bytes  
64 bytes from 176.26.232.111: icmp_seq=0 ttl=122 time=0.661 ms  
64 bytes from 176.26.232.111: icmp_seq=1 ttl=122 time=0.619 ms  
64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms  
64 bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms  
  
--- host3.site.net ping statistics ---  
4 packets transmitted, 4 packets received, 0% packet loss  
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

**Meaning** • The **ping** results show the following information:



- Size of the ping response packet (in bytes).
- IP address of the host from which the response was sent.
- Sequence number of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
- Time-to-live (ttl) hop-count value of the ping response packet.
- Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also called round-trip time.
- Number of ping requests (probes) sent to the host.
- Number of ping responses received from the host.
- Packet loss percentage.
- Round-trip time statistics: minimum, average, maximum, and standard deviation of the round-trip time.

**Related  
Documentation**

- [Troubleshooting Overview on page 6895](#)
- [Understanding Troubleshooting Resources on page 6893](#)

## Tracing SNMP Activity on a Device Running Junos OS

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
  - chassisd
  - craftd
  - ilmid
  - mib2d
  - rmopd
  - serviced
  - snmpd
- When a trace file named **filename** reaches its maximum size, it is renamed **filename.0**, then **filename.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)

- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
traceoptions {
  file <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  memory-trace;
  no-remote-trace;
  no-default-memory-trace;
}
```

These statements are described in the following sections:

- [Configuring the Number and Size of SNMP Log Files on page 6808](#)
- [Configuring Access to the Log File on page 6808](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 6809](#)
- [Configuring the Trace Operations on page 6809](#)

### Configuring the Number and Size of SNMP Log Files

---

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed **filename.0**, then **filename.1**, and so on, until there are three trace files. Then the oldest trace file (**filename.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (**filename**) reaches 2 MB, **filename** is renamed **filename.0**, and a new file called **filename** is created. When the new **filename** reaches 2 MB, **filename.0** is renamed **filename.1** and **filename** is renamed **filename.0**. This process repeats until there are 20 trace files. Then the oldest file (**filename.19**) is overwritten by the newest file (**filename.0**).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

### Configuring Access to the Log File

---

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file no-world-readable;
```

### Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file filename]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

### Configuring the Trace Operations

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
  all;
  configuration;
  database;
  events;
  general;
  interface-stats;
  nonvolatile-sets;
  pdu;
  policy;
  protocol-timeouts;
  routing-socket;
  server;
  subagent;
  timer;
  varbind-error;
}
```

Table 641 on page 6809 describes the meaning of the SNMP tracing flags.

**Table 641: SNMP Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Log all operations.	Off
<b>configuration</b>	Log reading of the configuration at the <b>[edit snmp]</b> hierarchy level.	Off
<b>database</b>	Log events involving storage and retrieval in the events database.	Off
<b>events</b>	Log important events.	Off

Table 641: SNMP Tracing Flags (*continued*)

Flag	Description	Default Setting
<b>general</b>	Log general events.	Off
<b>interface-stats</b>	Log physical and logical interface statistics.	Off
<b>nonvolatile-set</b>	Log nonvolatile SNMP set request handling.	Off
<b>pdu</b>	Log SNMP request and response packets.	Off
<b>policy</b>	Log policy processing.	Off
<b>protocol-timeouts</b>	Log SNMP response timeouts.	Off
<b>routing-socket</b>	Log routing socket calls.	Off
<b>server</b>	Log communication with processes that are generating events.	Off
<b>subagent</b>	Log subagent restarts.	Off
<b>timer</b>	Log internal timer events.	Off
<b>varbind-error</b>	Log variable binding errors.	Off

To display the end of the log for an agent, issue the **show log agentd | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where **agent** is the name of an SNMP agent.

#### Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level](#)
- [Example: Tracing SNMP Activity](#)
- [Configuring SNMP on page 1356](#)

## Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage

Even though the Junos OS has built-in performance metrics and monitoring options, you might need to have customized performance metrics. To make it easier for you to monitor such customized data through a standard monitoring system, the Junos OS provides you with an enterprise-specific Utility MIB that can store such data and thus extend SNMP support for managing and monitoring the data of your choice.

The enterprise-specific Utility MIB provides you with container objects of the following types: **32-bit counters**, **64-bit counters**, **signed integers**, **unsigned integers**, and **octet strings**.

You can use these container MIB objects to store the data that are otherwise not supported for SNMP operations. You can populate data for these objects either by using CLI commands or with the help of Op scripts and an RPC API that can invoke the CLI commands.

The following CLI commands enable you to set and clear Utility MIB object values:

- **request snmp utility-mib set** instance *name* object-type <counter | counter 64 | integer | string | unsigned integer> object-value *value*
- **request snmp utility-mib clear** instance *name* object-type <counter | counter 64 | integer | string | unsigned integer>

The *instance name* option of the **request snmp utility-mib <set | clear>** command specifies the name of the data instance and is the main identifier of the data. The **object-type <counter | counter 64 | integer | string | unsigned integer>** option enables you specify the object type, and the **object-value *value*** option enables you to set the value of the object.

To automate the process of populating Utility MIB data, you can use a combination of an event policy and event script. The following examples show the configuration for an event policy to run **show system buffers** every hour and to store the **show system buffers** data in Utility MIB objects by running an event script (**check-mbufs.slax**).

#### Event Policy Configuration

To configure an event policy that runs the **show system buffers** command every hour and invokes **check-mbufs.slax** to store the **show system buffers** data into Utility MIB objects, include the following statements at the **[edit]** hierarchy level:

```
event-options {
  generate-event {
    1-HOUR time-interval 3600;
  }
  policy MBUFS {
    events 1-HOUR;
    then {
      event-script check-mbufs.slax; # script stored at /var/db/scripts/event/
    }
  }
  event-script {
    file check-mbufs.slax;
  }
}
```

#### check-mbufs.slax Script

The following example shows the **check-mbufs.slax** script that is stored under **/var/db/scripts/event/**:

```
----- script START -----
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";
ns ext = "http://xmlsoft.org/XSLT/namespace";

match / {
  <op-script-results>{
```

```

var $cmd = <command> "show system buffers";
var $out = jcs:invoke($cmd);

var $lines = jcs:break_lines($out);
for-each ($lines) {
    if (contains(., "current/peak/max")) {
        var $pattern = "([0-9]+)/([0-9]+)/([0-9]+) mbufs";
        var $split = jcs:regex($pattern, .);
        var $result = $split[2];

        var $rpc = <request-snmp-utility-mib-set> {
            <object-type> "integer";
            <instance> "current-mbufs";
            <object-value> $result;
        }
        var $res = jcs:invoke($rpc);
    }
}
}
----- script END -----

```

You can run the following command to check the data stored in the Utility MIB as a result of the event policy and script shown in the preceding examples:

```

user@host> show snmp mib walk jnxUtilData ascii jnxUtilIntegerValue."current-mbufs"
= 0 jnxUtilIntegerTime."current-mbufs" = 07 da 05 0c 03 14 2c 00 2d 07 00
regress@caramels>

```



**NOTE:** The `show snmp mib walk` command is not available on the QFabric system, but you can use external SNMP client applications to perform this operation.

#### Related Documentation

- [Understanding the Implementation of SNMP on the QFabric System on page 6516](#)

## Displaying Commit Script Output

[Table 642 on page 6812](#) summarizes the Junos OS command-line interface (CLI) commands you can use to monitor and troubleshoot commit scripts. For more information about the `cscrip.log` file, see *Tracing Commit Script Processing*.



**NOTE:** Tracing commit script processing, including the `cscrip.log` file, is not supported on the QFX3000-G QFabric system.

**Table 642: Commit Script Configuration and Operational Mode Commands**

Task	Command
<b>Configuration Mode Commands</b>	
Display errors and warnings generated by commit scripts.	<code>commit</code> or <code>commit check</code>

**Table 642: Commit Script Configuration and Operational Mode Commands (*continued*)**

Task	Command
Display detailed information.	<b>commit   display detail</b>
Display the underlying Extensible Markup Language (XML) data.	<b>commit   display xml</b>
Display the postinheritance contents of the configuration database. This view includes transient changes, but does not include changes made in configuration groups.	<b>show   display commit-scripts</b>
Display the postinheritance contents of the configuration database. This view excludes transient changes.	<b>show   display commit-scripts no-transients</b>
Display the postinheritance configuration in XML format.  Viewing the configuration in XML format can be helpful when you are writing XML Path Language (XPath) expressions and configuration element tags.	<b>show   display commit-scripts view</b>
Display the postinheritance configuration in XML format, but exclude transient changes.	<b>show   display commit-scripts view   display commit-scripts no-transients</b>
Display all configuration groups data, including script-generated changes to the groups.	<b>show groups   display commit-scripts</b>
Display a particular configuration group, including script-generated changes to the group.	<b>show groups <i>group-name</i>   display commit-scripts</b>
<b>Operational Mode Commands</b>	
Display logging data associated with all commit script processing.	<b>show log cscript.log</b>
Display processing for only the most recent commit operation.	<b>show log cscript.log   last</b>
Display processing for script errors.	<b>show log cscript.log   match error</b>
Display processing for a particular script.	<b>show log cscript.log   match <i>filename</i></b>

**Related Documentation**

- *Tracing Commit Script Processing*

## Commands for General Monitoring

---

- [monitor traffic](#)
- [ping](#)



## monitor traffic

**Syntax** monitor traffic  
 <brief | detail | extensive>  
 <absolute-sequence>  
 <count *count*>  
 <interface *interface-name*>  
 <layer2-headers>  
 <matching *matching*>  
 <no-domain-names>  
 <no-promiscuous>  
 <no-resolve>  
 <no-timestamp>  
 <print-ascii>  
 <print-hex>  
 <resolve-timeout>  
 <size *size*>

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Display packet headers or packets received and sent from the Routing Engine.



### NOTE:

- Using the **monitor-traffic** command can degrade router or switch performance.
- Delays from DNS resolution can be eliminated by using the **no-resolve** option.



**NOTE:** This command is not supported on the QFabric system.

**Options** **none**—(Optional) Display packet headers transmitted through **fxp0**. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

**brief | detail | extensive**—(Optional) Display the specified level of output.

**absolute-sequence**—(Optional) Display absolute TCP sequence numbers.

**count *count***—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The monitor traffic command quits automatically after displaying the number of packets specified.

**interface *interface-name***—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

**layer2-headers**—(Optional) Display the link-level header on each line.

**matching *matching***—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

**no-domain-names**—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only **team** for the hostname **team.company.net**.

**no-promiscuous**—(Optional) Do not put the interface into promiscuous mode.

**no-resolve**—(Optional) Suppress reverse lookup of the IP addresses.

**no-timestamp**—(Optional) Suppress timestamps on displayed packets.

**print-ascii**—(Optional) Display each packet in ASCII format.

**print-hex**—(Optional) Display each packet, except the link-level header, in hexadecimal format.

**resolve-timeout *timeout***—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

**size *size***—(Optional) Read but do not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

**Additional Information** In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

**monitor traffic matching "*expression*"**

Replace ***expression*** with one or more of the match conditions listed in [Table 643 on page 6817](#).

Table 643: Match Conditions for the monitor traffic Command

Match Type	Condition	Description
Entity	<b>host</b> [ <i>address</i>   <i>hostname</i> ]	Matches packets that contain the specified address or hostname.  The protocol match conditions <b>arp</b> , <b>ip</b> , or <b>rarp</b> , or any of the directional match conditions can be prepended to the <b>host</b> match condition.
	<b>net</b> <i>address</i>	Matches packets with source or destination addresses containing the specified network address.
	<b>net</b> <i>address mask mask</i>	Matches packets containing the specified network address and subnet mask.
	<b>port</b> ( <i>port-number</i>   <i>port-name</i> )	Matches packets containing the specified source or destination TCP or UDP port number or port name.  In place of the numeric port address, you can specify a text synonym, such as <b>bgp</b> (179), <b>dhcp</b> (67), or <b>domain</b> (53) (the port numbers are also listed).
Directional	<b>dst</b>	Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.
	<b>src</b>	Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.
	<b>src and dst</b>	Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.
	<b>src or dst</b>	Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.
Packet Length	<b>less</b> <i>value</i>	Matches packets shorter than or equal to the specified value, in bytes.
	<b>greater</b> <i>value</i>	Matches packets longer than or equal to the specified value, in bytes.

Table 643: Match Conditions for the monitor traffic Command (*continued*)

Match Type	Condition	Description
Protocol	<b>amt</b>	Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.
	<b>arp</b>	Matches all ARP packets.
	<b>ether</b>	Matches all Ethernet packets.
	<b>ether (broadcast   multicast)</b>	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with <b>src</b> and <b>dst</b> .
	<b>ether protocol (address   (arp   ip   rarp))</b>	Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The <b>ether protocol</b> arguments <b>arp</b> , <b>ip</b> , and <b>rarp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ether protocol</b> match condition.
	<b>icmp</b>	Matches all ICMP packets.
	<b>ip</b>	Matches all IP packets.
	<b>ip (broadcast   multicast)</b>	Matches broadcast or multicast IP packets.
	<b>ip protocol (address   (icmp   igmp   tcp   udp))</b>	Matches packets with the specified address or protocol type. The <b>ip protocol</b> arguments <b>icmp</b> , <b>tcp</b> , and <b>udp</b> are also independent match conditions, so they must be preceded by a backslash (\) when used in the <b>ip protocol</b> match condition.
	<b>isis</b>	Matches all IS-IS routing messages.
	<b>rarp</b>	Matches all RARP packets.
	<b>tcp</b>	Matches all TCP datagrams.
	<b>udp</b>	Matches all UDP datagrams.

To combine expressions, use the logical operators listed in [Table 644 on page 6818](#).

Table 644: Logical Operators for the monitor traffic Command

Logical Operator (Highest to Lowest Precedence)	Description
<b>!</b>	Logical NOT. If the first condition does not match, the next condition is evaluated.

Table 644: Logical Operators for the monitor traffic Command (*continued*)

Logical Operator (Highest to Lowest Precedence)	Description
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.
	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
()	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0"arithmetic_expression relational_operator arithmetic_expression
```

The packet data accessor uses the following syntax:

```
protocol [byte-offset <size>]
```

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in [Table 645 on page 6820](#).



**NOTE:** Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The **monitor traffic** command cannot apply match conditions to inbound traffic.
- The **monitor traffic interface** command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (`| match`) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see [Table 643 on page 6817](#).
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the **monitor traffic interface ae[x]** command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

**Table 645: Arithmetic and Relational Operators for the monitor traffic Command**

Arithmetic or Relational Operator	Description
<b>Arithmetic Operator</b>	
+	Addition operator.
-	Subtraction operator.
/	Division operator.
&	Bitwise AND.
*	Bitwise exclusive OR.
	Bitwise inclusive OR.
<b>Relational Operator (Highest to Lowest Precedence)</b>	
<=	If the first expression is less than or equal to the second, the packet matches.
>=	If the first expression is greater than or equal to the second, the packet matches.
<	If the first expression is less than the second, the packet matches.
>	If the first expression is greater than the second, the packet matches.
=	If the compared expressions are equal, the packet matches.
!=	If the compared expressions are unequal, the packet matches.

**Required Privilege Level** trace  
maintenance

**List of Sample Output** [monitor traffic count on page 6821](#)  
[monitor traffic detail count on page 6821](#)  
[monitor traffic extensive \(Absolute Sequence\) on page 6821](#)  
[monitor traffic extensive \(Relative Sequence\) on page 6821](#)  
[monitor traffic extensive count on page 6821](#)  
[monitor traffic interface on page 6822](#)  
[monitor traffic matching on page 6822](#)  
[monitor traffic \(TX Matrix Plus Router\) on page 6822](#)  
[monitor traffic \(QFX3500 Switch\) on page 6823](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

### monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0
win 17680 (DF) [tos 0x10] (ttl 6)
```

### monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0] (ttl )
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 207.17.136.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951
4935628>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10] (ttl 64, id 5394)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

### monitor traffic extensive count

```
user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
```

```
reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)
```

### monitor traffic interface

```
user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650 In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813 In host30.lab.home.net.syslog > host40.home0
18:17:28.817846 In host30.lab.home.net.syslog > host40.home0
...
```

### monitor traffic matching

```
user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511 In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864 In arp who-has 192.168.1.17 tell 192.168.1.9
...
```

### monitor traffic (TX Matrix Plus Router)

```
user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes
04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog > sv-log-01.englab.juniper.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog >
sv-log-02.englab.juniper.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.englab.juniper.net.65235 >
```



```

summit-em0.englab.juniper.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.englab.juniper.net.telnet > aj-em0.englab.juniper.net.65235: P
1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: . ack 241 win 33304 <nop,nop,timestamp
42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell11.juniper.net.46182 > summit-em0.englab.juniper.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
...

```

### monitor traffic (QFX3500 Switch)

```

user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.

```

```
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```


## ping

**List of Syntax**   [Syntax on page 6825](#)  
                               [Syntax \(QFX Series\) on page 6825](#)

**Syntax**   `ping host`  
                   `<bypass-routing>`  
                   `<count requests>`  
                   `<detail>`  
                   `<do-not-fragment>`  
                   `<inet | inet6>`  
                   `<interface source-interface>`  
                   `<interval seconds>`  
                   `<logical-system logical-system-name>`  
                   `<loose-source value>`  
                   `<mac-address mac-address>`  
                   `<no-resolve>`  
                   `<pattern string>`  
                   `<rapid>`  
                   `<record-route>`  
                   `<routing-instance routing-instance-name>`  
                   `<size bytes>`  
                   `<source source-address>`  
                   `<strict >`  
                   `<strict-source value.>`  
                   `<tos type-of-service>`  
                   `<ttl value>`  
                   `<verbose>`  
                   `<vpls instance-name>`  
                   `<wait seconds>`

**Syntax (QFX Series)**   `ping host`  
                               `<bypass-routing>`  
                               `<count requests>`  
                               `<detail>`  
                               `<do-not-fragment>`  
                               `<inet>`  
                               `<interface source-interface>`  
                               `<interval seconds>`  
                               `<logical-system logical-system-name>`  
                               `<loose-source value>`  
                               `<mac-address mac-address>`  
                               `<no-resolve>`  
                               `<pattern string>`  
                               `<rapid>`  
                               `<record-route>`  
                               `<routing-instance routing-instance-name>`  
                               `<size bytes>`  
                               `<source source-address>`  
                               `<strict>`  
                               `< strict-source value>`  
                               `<tos type-of-service>`  
                               `<ttl value>`  
                               `<verbose>`

<wait *seconds*>

<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Check host reachability and network connectivity. The <b>ping</b> command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.
<b>Options</b>	<p><b>host</b>—IP address or hostname of the remote system to ping.</p> <p><b>bypass-routing</b>—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.</p> <p><b>count requests</b>—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p><b>detail</b>—(Optional) Include in the output the interface on which the ping reply was received.</p> <p><b>do-not-fragment</b>—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div><p><b>NOTE:</b> In Junos OS Release 11.1 and later, when issuing the <b>ping</b> command for an IPv6 route with the <b>do-not-fragment</b> option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p></div> <p><b>inet</b>—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p><b>inet6</b>—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p><b>interface source-interface</b>—(Optional) Interface to use to send the ping requests.</p> <p><b>interval seconds</b>—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p><b>logical-system logical-system-name</b>—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the <b>set cli logical-system logical-system-name</b> command and then run the <b>ping</b> command. To return to the main router or switch, enter the <b>clear cli logical-system</b> command.</p>

**loose-source *value***—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

**mac-address *mac-address***—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

**no-resolve**—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

**pattern *string***—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

**rapid**—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

**record-route**—(Optional) Record and report the packet's path (IPv4).

**routing-instance *routing-instance-name***—(Optional) Name of the routing instance for the ping attempt.

**size *bytes***—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

**strict**—(Optional) Use the strict source route option (IPv4).

**strict-source *value***—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

**tos *type-of-service***—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

If the device configuration includes the **dscp-code-point *value*** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

**ttl *value***—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

**verbose**—(Optional) Display detailed output.

**vpls *instance-name***—(Optional) Ping the instance to which this VPLS belongs.

**wait *seconds***—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level	network
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</a></li></ul>
List of Sample Output	<a href="#">ping hostname on page 6828</a> <a href="#">ping hostname rapid on page 6828</a> <a href="#">ping hostname size count on page 6828</a>
Output Fields	When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

### ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

### ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

---

## Commands for Network Analytics

- [monitor start \(Analytics\)](#)
- [show analytics collector](#)

- `show analytics configuration`
- `show analytics queue-statistics`
- `show analytics status`
- `show analytics streaming-servers`
- `show analytics traffic-statistics`

## monitor start (Analytics)

**Syntax** `monitor start filename`

**Release Information** Command introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Start the display of the queue statistics or traffic statistics file if you had enabled queue or traffic monitoring on your device. The output is displayed in the JavaScript Object Notation (JSON) format.



**NOTE:** This topic describes the local file output in Junos OS Release 13.2X50-D15 and 13.2X51-D10 only. For information about 13.2X51-D15 and later, see [“Understanding Enhanced Analytics Local File Output” on page 6506](#)

**Options** *filename*—Name of the queue statistics or traffic statistics file.

**Required Privilege Level** trace

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)

**List of Sample Output**

- [monitor start Using the Queue Statistics File \(Junos OS Release 13.2X51-D10\) on page 6831](#)
- [monitor start Using the Queue Statistics File \(Junos OS Release 13.2X50-D15\) on page 6832](#)
- [monitor start Using the Traffic Statistics File \(Junos OS Release 13.2X51-D10\) on page 6832](#)
- [monitor start Using the Traffic Statistics File \(Junos OS Release 13.2X50-D15\) on page 6832](#)

**Output Fields** [Table 646 on page 6830](#) describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

**Table 646: monitor start Command Output Fields**

Field	Description
hostname (used in Junos OS Release 13.2X50-D15 only)	Name of the network analytics host device.
record type	Type of statistics. May be queue statistics or traffic statistics.
time	Time at which the statistics were captured.
router-id	ID of the network analytics host device.



Table 646: monitor start Command Output Fields (*continued*)

Field	Description
latency	For queue statistics only. Traffic queue latency in milliseconds.
port	Name of the physical port configured for network analytics.
queue depth	For queue statistics only. Depth of the traffic queue in bytes.
rxpkt	For traffic statistics monitoring only. Total packets received.
rxpps	For traffic statistics monitoring only. Total packets received per second.
rxbyte	For traffic statistics monitoring only. Total bytes received.
rxbps	For traffic statistics monitoring only. Total bytes received per second.
rxdrop	For traffic statistics monitoring only. Total incoming packets dropped.
rxerr	For traffic statistics monitoring only. Total packets with errors.
rxutil (in Junos OS Release 13.2X50-D15 only)	For traffic statistics monitoring only. Total percent of traffic utilization for incoming traffic.
txpkt	For traffic statistics monitoring only. Total packets transmitted.
txpps	For traffic statistics monitoring only. Total packets transmitted per second.
txbyte	For traffic statistics monitoring only. Total bytes transmitted.
txbps	For traffic statistics monitoring only. Total bytes transmitted per second.
txdrop	For traffic statistics monitoring only. Total transmitted bytes dropped.
txerr	For traffic statistics monitoring only. Total transmitted packets with errors (dropped).
txutil (in Junos OS Release 13.2X50-D15 only)	For traffic statistics monitoring only. Total percent of traffic utilization for outgoing traffic.

## Sample Output

### monitor start Using the Queue Statistics File (Junos OS Release 13.2X51-D10)

```

user@host> monitor start analytics.qs
{"record-type":"queue-stats","time":"2013 Nov 3 4:40:42.840",
"router-id":"qfx5100-switch","port":"xe-0/0/18","latency":0,"queue-depth":208}

{"record-type":"queue-stats","time":"2013 Nov 3 4:40:44.887",
"router-id":"qfx5100-switch","port":"xe-0/0/18","latency": 1110,"queue-depth":
1387568}

```

### monitor start Using the Queue Statistics File (Junos OS Release 13.2X50-D15)

```
user@host> monitor start analytics.qs
{"hostname":"sw-la-pb-03","latency":566,"port":"xe-0/0/9","queue depth":708656,
"record type":"queue-stats","time":"Apr 11 20:18:40.329"}
```

## Sample Output

### monitor start Using the Traffic Statistics File (Junos OS Release 13.2X51-D10)

```
user@host> monitor start analytics.ts
{"record-type":"traffic-stats","time":"2013 Nov 3 4:39:53.910",
"router-id":"qfx5100-switch","port":"xe-0/0/18","rxpkt":23193749091,"rxpps":8299889,

"rxbyte":2968799876957,"rxbps":824002992,"rxdrop":0,"rxerr":0,"txpkt":1029323986,
"txpps":82671,"txbyte":131753470470,"txbps":85598256,"txdrop":0,"txerr":0}
```

### monitor start Using the Traffic Statistics File (Junos OS Release 13.2X50-D15)

```
user@host> monitor start analytics.ts
{"hostname":"sw-la-pb-03","port":"xe-0/0/9","record type":"traffic-statistics",
"time":"Apr 11 20:13:48.545", "rxpkt":601024640, "rxpps": 840315,
"rxbyte":76931153920,
"rxbps":863997032, "rxdrop":0, "rxerr":0, "rxutil":8.32,"txpkt":336551380309,
"txpps":405395,"txbyte":23369872265951,"txbps":3240000976,"txdrop":1010566660824,
"txerr":69920099883860,"txutil":32.76}
```

## show analytics collector

**Syntax** show analytics collector

**Release Information** Command introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Show the list of network analytics remote collectors and related information. Remote collectors can be configured to receive streaming output for queue statistics and traffic statistics from the network analytics process (Analyticsd) running on the Routing Engine.



**NOTE:** The `show analytics collector` command is not available in Junos OS Releases prior to 13.2X51-D15.

**Required Privilege Level** interface-control

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)
- *address (Analytics Collector)*

**List of Sample Output** [show analytics collector on page 6834](#)

**Output Fields** [Table 647 on page 6833](#) describes the output fields for the `show analytics collector` command.

**Table 647: show analytics collector Command Output Fields**

Field	Description
Address	IP Address of the collector that is configured for receiving the streaming data.
Port	Port number of the collector receiving the streaming data.
Transport	Transport protocol: <ul style="list-style-type: none"> <li>• tcp—Transmission Control Protocol</li> <li>• udp—User Datagram Protocol</li> </ul> <p><b>NOTE:</b> The connection state of a port configured with the <code>udp</code> transport protocol is always displayed as <code>n/a</code>.</p>
Stream format	Format of the data that is sent to the server: <ul style="list-style-type: none"> <li>• csv—Comma-separated values</li> <li>• gpb—Google Protocol Buffer</li> <li>• json—JavaScript Object Notation</li> <li>• tsv—Tab-separated values</li> </ul>

Table 647: show analytics collector Command Output Fields (*continued*)

Field	Description
State	Connection state of the streaming server.
Sent	Number of bytes sent to the streaming server.

## Sample Output

### show analytics collector

```
user@host> show analytics collector
Address      Port    Transport Stream format State      Sent
10.94.184.25 50013   udp      gpb        n/a       8710
10.94.184.25 50040   tcp      gpb        Not initialized 0
10.94.184.25 50050   tcp      gpb        Established 405
10.94.184.62 50010   tcp      csv        Established 18
10.94.184.62 50020   udp      json       n/a       17
```

## show analytics configuration

<b>Syntax</b>	show analytics configuration
<b>Release Information</b>	Command introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	Show the network analytics configuration details for the global and interface configurations.
<b>Required Privilege Level</b>	interface-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 6490</a></li> <li>• <a href="#">analytics on page 6667</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show analytics configuration (Junos OS Release 13.2X51-D15 and Later) on page 6838</a> <a href="#">show analytics configuration (Junos OS Release 13.2X51-D10 and Earlier) on page 6838</a>
<b>Output Fields</b>	describes the output fields for the <b>show analytics configuration</b> command in Junos OS Release 13.2X51-D15 and later.

**Table 648: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D15 and Later)**

Field	Descriptions
<b>Global Configurations</b>	
Traffic monitoring status	Settings are enabled or disabled. If traffic statistics monitoring is not enabled, this field is not shown.
Traffic monitoring polling interval	Interval for traffic statistics polling in seconds.  <b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.
Queue monitoring status	Settings are enabled or disabled. If queue statistics monitoring is not enabled, this field is not shown.
Queue monitoring polling interval	Interval for queue statistics polling in milliseconds.  <b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.
Queue depth high threshold	Upper limit of the depth threshold configuration in number of bytes.  If the queue depth threshold is not configured, this field is not shown.
Queue depth low threshold	Lower limit of the depth threshold configuration in number of bytes.  If the queue depth threshold is not configured, this field is not shown.

**Table 648: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D15 and Later) (continued)**

Field	Descriptions
Queue latency high threshold	Upper limit of the latency threshold configuration in nanoseconds.  If the queue latency threshold is not configured, this field is not shown.
Queue latency low threshold	Lower limit of the latency threshold configuration in microseconds.  If the queue latency threshold is not configured, this field is not shown.
<b>Interface Configurations</b>	
Interface	Name of interface that is configured for network analytics. The interface configuration overrides the global network analytics configuration.
Traffic Statistics	Settings are Enabled or Disabled for the interface.
Queue Statistics	Settings are Enabled or Disabled for the interface.
Queue depth threshold High	Upper limit of the depth threshold configuration in number of bytes.  If the queue depth threshold is not configured, <b>n/a</b> is displayed.
Queue depth threshold Low	Lower limit of the depth threshold configuration in number of bytes.  If the queue depth threshold is not configured, <b>n/a</b> is displayed.
Latency threshold High	Upper limit of the latency threshold configuration in nanoseconds.  If the latency threshold is not configured, <b>n/a</b> is displayed.
Latency threshold Low	Lower limit of the latency threshold configuration in nanoseconds.  If the latency threshold is not configured, <b>n/a</b> is displayed.

[Table 649 on page 6836](#) describes the output fields for the **show analytics configuration** command in Junos OS Release 13.2X51-D10 and 13.2X50-D15.

**Table 649: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier)**

Field	Descriptions
<b>Global Configurations</b>	
Traffic statistics	Settings are Auto, Enabled, or Disabled.  If <b>Auto</b> is displayed, traffic statistics monitoring is not enabled.

**Table 649: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier) (continued)**

Field	Descriptions
Poll interval (traffic statistics)	<p>Interval for traffic statistics polling in seconds.</p> <p>If the output displays a setting of 0 seconds, the polling interval was not configured, and the default interval applies.</p> <p><b>NOTE:</b> The default interval is 1 second in Junos OS Release 13.2X50-D15 and later, except for EX4300 switches, on which the default interval is 5 seconds, and 2 seconds in Junos OS Release 13.2X51-D10.</p> <p><b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.</p>
Queue statistics	<p>Settings are Auto, Enabled, or Disabled.</p> <p>If <b>Auto</b> is displayed, queue statistics monitoring is not enabled.</p>
Poll interval (queue statistics)	<p>Interval for queue statistics polling in milliseconds.</p> <p><b>NOTE:</b> The default interval is 8 milliseconds in Junos OS Release 13.2X50-D15 and later, and 10 milliseconds in Junos OS Release 13.2X51-D10 or later.</p> <p><b>NOTE:</b> Due to limitations and variations in hardware capability in different devices, there might be a difference in value between the actual interval and configured interval.</p>
Depth threshold high	<p>Upper limit of the depth threshold configuration in number of bytes.</p> <p>If <b>0</b> is displayed, depth threshold is not enabled.</p>
Depth threshold low	<p>Lower limit of the depth threshold configuration in number of bytes.</p> <p>If <b>0</b> is displayed, depth threshold is not enabled.</p>
Latency threshold high	<p>Upper limit of the latency threshold configuration in microseconds.</p> <p>If <b>0</b> is displayed, latency threshold is not enabled.</p>
Latency threshold low	<p>Lower limit of the latency threshold configuration in microseconds.</p> <p>If <b>0</b> is displayed, latency threshold is not enabled.</p>
<b>Interface Configurations</b>	
Interface	Name of interface that is configured for network analytics. The interface configuration overrides the global network analytics configuration.
Traffic Statistics	Settings are Enabled or Disabled for the interface.
Queue Statistics	Settings are Enabled or Disabled for the interface.
Depth-threshold High	<p>Upper limit of the depth threshold configuration in number of bytes.</p> <p>If <b>0</b> is displayed, depth threshold is not enabled.</p>

Table 649: show analytics configuration Command Output Fields (Junos OS Release 13.2X51-D10 and earlier) (continued)

Field	Descriptions
Depth-threshold Low	Lower limit of the depth threshold configuration in number of bytes.  If 0 is displayed, depth threshold is not enabled.
Latency-threshold High	Upper limit of the latency threshold configuration in microseconds.  If 0 is displayed, latency threshold is not enabled.
Latency-threshold Low	Lower limit of the latency threshold configuration in microseconds.  If 0 is displayed, latency threshold is not enabled.

## Sample Output

### show analytics configuration (Junos OS Release 13.2X51-D15 and Later)

```
user@host> show analytics configuration
Traffic monitoring status is enabled
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring polling interval : 1000 milliseconds
Queue depth high threshold : 99999 bytes
Queue depth low threshold : 99 bytes
```

Interface	Traffic Statistics	Queue Statistics	Queue depth threshold		Latency threshold	
			High (bytes)	Low	High (nanoseconds)	Low
xe-0/0/16	enabled	enabled	n/a	n/a	2300	20
xe-0/0/18	enabled	enabled	n/a	n/a	2300	20
xe-0/0/19	enabled	enabled	n/a	n/a	2300	20

### show analytics configuration (Junos OS Release 13.2X51-D10 and Earlier)

```
user@host> show analytics configuration
Global configurations:
  Traffic statistics: Enabled, Poll interval: 2 seconds
  Queue statistics: Auto, Poll interval: 10 milliseconds
  Depth threshold high: 0 bytes, low: 0 bytes
  Latency threshold high: 0 microseconds, low: 0 microseconds
```

Interface	Traffic Statistics	Queue Statistics	Depth-threshold		Latency-threshold	
			High (bytes)	Low	High (microseconds)	Low
xe-0/0/0	Auto	Auto	204800	10	0	0



## show analytics queue-statistics

<b>Syntax</b>	<code>show analytics queue-statistics</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	Show the queue statistics (queue length and latency) that are collected for all interfaces that are enabled for network analytics on a device. Optionally, if you wish to see the queue statistics for one interface only, you may specify the interface.
<b>Options</b>	<code>interface <i>interface-name</i></code> —(Optional) Display the queue statistics for the specified interface only.
<b>Required Privilege Level</b>	interface-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 6490</a></li> <li>• <a href="#">analytics on page 6667</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show analytics queue-statistics (Junos OS Release 13.2X51-D15 and Later) on page 6839</a> <a href="#">show analytics queue-statistics (Junos OS Release 13.2X51-D10) on page 6840</a> <a href="#">show analytics queue-statistics (Junos OS Release 13.2X50-D15) on page 6840</a>
<b>Output Fields</b>	<a href="#">Table 610 on page 6507</a> describes the output fields for the <code>show analytics queue-statistics</code> command.

Table 650: show analytics queue-statistics Command Output Fields

Field	Description
Time	Date and time at which the queue statistics are collected.
Interface	Name of the interface at which the queue statistics are collected.
Queue-length or queue-depth (bytes)	Queue depth (length) in number of bytes.
Latency	Queue depth in nanoseconds (Junos OS Release 13.2X51-D15 and later) or microseconds (Junos OS Release 13.2X51-D10 and earlier).

## Sample Output

### show analytics queue-statistics (Junos OS Release 13.2X51-D15 and Later)

```

user@host> show analytics queue-statistics
CLI issued at 2014-01-07 17:20:29.978561
Time                Interface      Queue-depth      Latency
                    (bytes)         (nanoseconds)
00:00:00.870058 ago  xe-0/0/19      1369680          1095744

```

00:00:01.875049 ago	xe-0/0/19	1381952	1105561
00:00:02.875053 ago	xe-0/0/19	1387776	1110220
00:00:03.876047 ago	xe-0/0/19	1387568	1110054
00:00:04.873045 ago	xe-0/0/19	1388192	1110553
00:00:05.871044 ago	xe-0/0/19	1385904	1108723
00:00:06.873354 ago	xe-0/0/19	1371552	1097241

#### show analytics queue-statistics (Junos OS Release 13.2X51-D10)

```
user@host> show analytics queue-statistics
```

Time	Interface	Queue-length (bytes)	Latency (us)
2013 Nov 3 3:52:26.272	xe-0/0/9	208	0
2013 Nov 3 3:52:26.292	xe-0/0/9	208	0
2013 Nov 3 3:52:26.372	xe-0/0/9	208	0
2013 Nov 3 3:52:26.392	xe-0/0/9	208	0
2013 Nov 3 3:52:26.432	xe-0/0/9	208	0
2013 Nov 3 3:52:26.492	xe-0/0/9	208	0
2013 Nov 3 3:52:26.572	xe-0/0/9	208	0
2013 Nov 3 4:30:24.584	xe-0/0/9	1387152	1109
2013 Nov 3 4:30:24.604	xe-0/0/9	1372384	1097
2013 Nov 3 4:30:24.624	xe-0/0/9	1384864	1107


## Sample Output

#### show analytics queue-statistics (Junos OS Release 13.2X50-D15)

```
user@host> show analytics queue-statistics
```

Time	Interface	Queue-length (bytes)	Latency (us)
Apr 6 0:17:18.224	xe-0/0/9	1043952	835
Apr 6 0:17:18.234	xe-0/0/9	1053520	842
Apr 6 0:17:18.244	xe-0/0/9	1055184	844

## show analytics status

<b>Syntax</b>	show analytics status <global>
<b>Release Information</b>	Command introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	Show the status of the network analytics components that are configured on a device.
<b>Options</b>	<b>none</b> —Show the global and interface status for network analytics.  <b>global</b> —Show the global status only for network analytics.
<div>  <p><b>NOTE:</b> The <b>global</b> option is not available in Junos OS Releases prior to 13.2X51-D15.</p> </div>	
<b>Required Privilege Level</b>	interface-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 6490</a></li> <li>• <a href="#">analytics on page 6667</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show analytics status (Junos OS Release 13.2X51-D15 or Later) on page 6843</a> <a href="#">show analytics status global (Junos OS Release 13.2X51-D15 or Later) on page 6843</a> <a href="#">show analytics status (Junos OS Release 13.2X50-D15 and 13.2X51-D10) on page 6843</a>
<b>Output Fields</b>	<a href="#">Table 651 on page 6841</a> describes the output fields for the <b>show analytics status</b> command.

**Table 651: show analytics status Command Output Fields**

Field	Descriptions
<b>Global Configurations</b>	
Traffic statistics or Traffic monitoring status	<p>Settings are Auto, Enabled, or Disabled.</p> <p>If <b>Auto</b> is displayed, traffic statistics monitoring is not enabled.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>
Poll interval or Traffic monitoring polling interval	<p>Interval for traffic statistics polling in seconds.</p> <p><b>NOTE:</b> Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.</p>

Table 651: show analytics status Command Output Fields (*continued*)

Field	Descriptions
Queue statistics or Queue monitoring status	<p>Can be Auto, Enabled, or Disabled.</p> <p>If <b>Auto</b> is displayed, queue statistics monitoring is not enabled.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>
Poll interval or Queue monitoring polling interval	<p>Interval for queue statistics polling in milliseconds.</p> <p><b>NOTE:</b> Due to limitations and variations in the hardware capability of different devices, you might see a difference in value between the actual interval and configured interval.</p>
Depth threshold high or Queue depth high threshold	<p>Upper limit of the depth threshold configuration in number of bytes.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p>
Depth threshold low or Queue depth low threshold	<p>Lower limit of the depth threshold configuration in number of bytes.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p>
Latency threshold high	<p>Upper limit of the latency threshold configuration in microseconds.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p>
Latency threshold low	<p>Lower limit of the latency threshold configuration in microseconds.</p> <p>In Junos OS Release 13.2X51-D15 or later, if this parameter is not configured, this field is not shown.</p> <p>In Junos OS Release 13.2X51-D10 or earlier, if this parameter is not configured, a value of <b>0</b> is displayed.</p>
<b>Interface Configurations</b>	
Interface	Name of an interface that is configured for network analytics. The interface configuration overrides the global network analytics configuration.
Traffic Statistics	<p>Settings are Enabled or Disabled for the interface.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>
Queue Statistics	<p>Settings are Enabled or Disabled for the interface.</p> <p><b>NOTE:</b> The Disabled setting always supersedes the Enabled setting.</p>

Table 651: show analytics status Command Output Fields (*continued*)

Field	Descriptions
Depth-threshold High or Queue depth threshold high	Upper limit of the depth threshold configuration in number of bytes.  If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release.
Depth-threshold Low or Queue depth threshold low	Lower limit of the depth threshold configuration in number of bytes.  If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release.
Latency-threshold High	Upper limit of the latency threshold configuration in nanoseconds or microseconds.  If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release.
Latency-threshold Low	Lower limit of the latency threshold configuration in nanoseconds or microseconds.  If this parameter is not configured, an output of <b>n/a</b> or <b>0</b> is displayed in this column, depending on the software release.

## Sample Output

### show analytics status (Junos OS Release 13.2X51-D15 or Later)

```

user@host> show analytics status
Traffic monitoring status is auto
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring status polling interval : 1000 milliseconds
Queue depth high threshold : 1000000000 bytes
Queue depth low threshold : 99 bytes

```

Interface	Traffic Statistics	Queue Statistics	Queue depth threshold		Latency threshold	
			High (bytes)	Low	High (nanoseconds)	Low
xe-0/0/16	enabled	enabled	1000000000	99	n/a	n/a
xe-0/0/18	disabled	enabled	1000000000	99	n/a	n/a
xe-0/0/19	enabled	enabled	1000000000	99	n/a	n/a

### show analytics status global (Junos OS Release 13.2X51-D15 or Later)

```

user@host> show analytics status global

Traffic monitoring status is auto
Traffic monitoring polling interval : 5 seconds
Queue monitoring status is enabled
Queue monitoring status polling interval : 1000 milliseconds
Queue depth high threshold : 1000000000 bytes
Queue depth low threshold : 99 bytes

```

### show analytics status (Junos OS Release 13.2X50-D15 and 13.2X51-D10)

```

user@host> show analytics status

```

## Global configurations:

Traffic statistics: Auto, Poll interval: 2 seconds

Queue statistics: Auto, Poll interval: 10 milliseconds

Depth threshold high: 0 bytes, low: 0 bytes

Latency threshold high: 1000 microseconds, low: 50 microseconds

Interface	Traffic Statistics	Queue Statistics	Depth-threshold		Latency-threshold	
			High	Low	High	Low
			(bytes)		(microseconds)	
xe-0/0/6	Enabled	Enabled	0	0	1000	50
xe-0/0/7	Enabled	Enabled	204800	10	0	0
xe-0/0/8	Enabled	Enabled	0	0	1000	50

## show analytics streaming-servers

**Syntax** show analytics streaming-servers

**Release Information** Command introduced in Junos OS Release 13.2 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.

**Description** Show the list of streaming servers that are configured for network analytics. Streaming servers receive streaming output for queue statistics and traffic statistics from the network analytics process (Analyticsd) running on the Routing Engine.



**NOTE:** The show analytics streaming-servers command is available in Junos OS Release 13.2X50-D15 and 13.2X51-D10 only.

**Required Privilege Level** interface-control

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)
- [show analytics collector on page 6833](#)

**List of Sample Output** [show analytics streaming-servers on page 6846](#)

**Output Fields** [Table 652 on page 6845](#) describes the output fields for the **show analytics streaming-servers** command.

**Table 652: show analytics streaming-servers Command Output Fields**

Field	Description
Address	IP Address of the streaming server that is configured for receiving the streaming data.
Port	Port number of the streaming server receiving the streaming data.
Stream-Format	Format of the data that is sent to the server. Values are: <ul style="list-style-type: none"> <li>• csv—Comma-separated values.</li> <li>• json—JavaScript Object Notification.</li> <li>• tsv—Tab-separated values.</li> </ul>
Stream-Type	Type of data that is sent to the a port on the streaming server: <ul style="list-style-type: none"> <li>• QS—Queue statistics.</li> <li>• TS—Traffic statistics.</li> </ul>
State	Connection state of the streaming server.
Sent	Number of bytes sent to the streaming server.

## Sample Output

### show analytics streaming-servers

```
user@host> show analytics streaming-servers
```

Address	Port	Stream-Format	Stream-Type	State	Sent
10.94.198.14	50001	json	QS	Established	0
10.94.198.14	50005	csv	TS	Established	1185
172.17.28.28	50005	tsv	TS/QS	In Progress	0



## show analytics traffic-statistics

<b>Syntax</b>	<code>show analytics traffic-statistics</code> <code>&lt;interface <i>interface-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D25 for EX Series switches.
<b>Description</b>	Show the traffic statistics that are collected for all interfaces that are enabled for network analytics on a device. Optionally, if you wish to see the traffic statistics for one interface only, you may specify the interface.
<b>Options</b>	<code>interface <i>interface-name</i></code> —(Optional) Display the traffic statistics for the specified interface only.
<b>Required Privilege Level</b>	interface-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Network Analytics Overview on page 6490</a></li> <li>• <a href="#">analytics on page 6667</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show analytics traffic-statistics (Junos OS Release 13.2X51-D15 or Later) on page 6848</a> <a href="#">show analytics traffic-statistics (Junos OS Release 13.2X51-D10) on page 6848</a> <a href="#">show analytics traffic-statistics (Junos OS Release 13.2X50-D15) on page 6848</a>
<b>Output Fields</b>	<a href="#">Table 653 on page 6847</a> describes the output fields for the <code>show analytics traffic-statistics</code> command.

**Table 653: show analytics traffic-statistics Command Output Fields**

Field	Description
Time	The date and time at which the traffic statistics are generated.
Physical interface	Name of the interface at which the traffic statistics are collected.
Total octets	Total number of octets that are received and transmitted.
Total packets	Total number of packets that are received and transmitted.
Octets per second	Number of octets received and transmitted per second.
Packet per second	Number of packets received and transmitted per second.
CRC/Align errors or Octets dropped	Number of cyclic redundancy check (CRC) errors or octets dropped. <ul style="list-style-type: none"> <li>• Junos OS Release 13.2X51-D15 or later—Number of cyclic redundancy check (CRC) errors.</li> <li>• Junos OS Release 13.2X51-D10 and earlier—Number of octets dropped.</li> </ul>
Packets dropped	Number of packets dropped.

## Sample Output

### show analytics traffic-statistics (Junos OS Release 13.2X51-D15 or Later)

```

user@host> show analytics traffic-statistics
CLI issued at 2014-01-07 17:22:28.952677
Time: 00:00:03.480244 ago, Physical interface: xe-0/0/19
Traffic Statistics:
Total octets:          3929946593792      393001011519232
Total packets:         30702707784       3070320402462
Unicast packet:        30702707784       3070320402462
Multicast packets:     0                 0
Broadcast packets:     0                 0
Octets per second:     86407016          59044064
Packets per second:    84787             8469688
CRC/Align errors:      0                 392986110751744
Packets dropped:       0                 3070203990248

```

### show analytics traffic-statistics (Junos OS Release 13.2X51-D10)

```

user@host> show analytics traffic-statistics
Time: 2013 Nov 3 4:36:55.542, Physical interface: xe-0/0/8
Traffic Statistics:
Total octets:          2777524779008      101855533467
Total packet:          21699412289       795746503
Octets per second:     904001272         0
Packet per second:     8399574          0
Octets dropped:         0                 0
Packet dropped:         0                 0
Time: 2013 Nov 3 4:36:57.559, Physical interface: xe-0/0/10
Traffic Statistics:
Total octets:          2777546444381      129840936198
Total packet:          21699581650       1014382311
Octets per second:     90400211          86403728
Packet per second:     8400382           84438
Octets dropped:         0                 0
Packet dropped:         0                 0

```

### show analytics traffic-statistics (Junos OS Release 13.2X50-D15)

```

user@host> show analytics traffic-statistics
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/8
Traffic Statistics:
Total octets:          4797548752936      408886273632
Total packet:          5658257464        3190613435
Octets per second:     0                 0
Packet per second:     0                 0
Octets dropped:         0                 252901000
Packet dropped:         0                 252901
Utilization:           0.0%              0.0%
Time: Apr 5 19:52:48.549, Physical interface: xe-0/0/10
Traffic Statistics:
Total octets:          4790866253100      477139024
Total packet:          5624473639        477944
Octets per second:     0                 0
Packet per second:     0                 0
Octets dropped:         0                 166582000
Packet dropped:         0                 166582
Utilization:           0.0%              0.0%

```

## Commands for sFlow Technology

---

- `clear sflow collector statistics`
- `show sflow`
- `show sflow collector`
- `show sflow interface`

## clear sflow collector statistics

---

**Syntax** clear sflow collector statistics

**Release Information** Command introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Clear the sample counters for all sFlow collectors.

**Required Privilege Level** view

**Related Documentation**

- [Example: Monitoring Network Traffic Using sFlow Technology on page 6571](#)
- [Configuring sFlow Technology on page 6596](#)
- [show sflow collector on page 6853](#)

**List of Sample Output** [clear sflow collector statistics on page 6850](#)

### Sample Output

#### clear sflow collector statistics

The following example shows two output examples for the **show sflow collector** command, one before and one after the **clear sflow collector statistics** command was issued.

```
user@host> show sflow collector
Collector      Udp-port      No. of samples
address
10.1.1.1       6343          3174
10.1.2.1       6343          3562
```

```
user@host> clear sflow collector statistics
```

```
user@host> show sflow collector
Collector      Udp-port      No. of samples
address
10.1.1.1       6343          0
10.1.2.1       6343          0
```

## show sflow

<b>Syntax</b>	show sflow <collector> <interface>
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display sFlow configuration information.
<b>Options</b>	<p><b>none</b>—Display all sFlow configuration information.</p> <p><b>collector</b>—(Optional) Display a list of configured sFlow collectors and their properties.</p> <p><b>interface</b>—(Optional) Display the interfaces on which sFlow technology is enabled and the sampling parameters.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show sflow interface on page 6854</a></li> <li>• <a href="#">show sflow collector on page 6853</a></li> <li>• <a href="#">clear sflow collector statistics on page 6850</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6596</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show sflow on page 6852</a>
<b>Output Fields</b>	<a href="#">Table 654 on page 6851</a> lists the output fields for the <b>show sflow</b> command. Output fields are listed in the approximate order in which they appear.

**Table 654: show sflow Output Fields**

Field Name	Field Description	Level of Output
sFlow	Status of the feature: <b>Enabled</b> or <b>Disabled</b> .	All levels
Sample limit	Number of packets sampled per second. This sample limit cannot be configured and is set to 300 packets per second.	All levels
Polling interval	Interval at which the sFlow agent polls the interface.	All levels
Sample rate egress	Rate at which egress packets are sampled.	All levels
Sample rate ingress	Rate at which ingress packets are sampled.	All levels
Agent ID	IP address assigned to the sFlow agent.	All levels
Source IP address	Source IP address for the sFlow packets.	All levels

## Sample Output

show sflow

```
user@host> show sflow
```

```
sFlow           : Enabled
Sample limit    : 300 packets/second
Polling interval : 20 second
Sample rate egress : 1:2048: Disabled
Sample rate ingress : 1:1000: Enabled
Agent ID        : 10.93.54.7
Source IP address : 10.93.54.7
```

## show sflow collector

<b>Syntax</b>	show sflow collector
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display a list of configured sFlow collectors and their properties.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear sflow collector statistics on page 6850</a></li> <li>• <a href="#">show sflow on page 6851</a></li> <li>• <a href="#">show sflow interface on page 6854</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6596</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show sflow collector on page 6853</a>
<b>Output Fields</b>	<a href="#">Table 655 on page 6853</a> lists the output fields for the <b>show sflow collector</b> command. Output fields are listed in the approximate order in which they appear.

**Table 655: show sflow collector Output Fields**

Field Name	Field Description	Level of Output
Collector address	IP address of the collector.	All levels
UDP-Port	UDP port number of the collector.	All levels
No. of samples	Number of samples collected.	All levels

## Sample Output

### show sflow collector

```

user@host> show sflow collector

Collector      Udp-port    No. of samples
address
10.204.32.46   6343        1000
100.204.32.76 3400        1000

```

## show sflow interface

<b>Syntax</b>	show sflow interface
<b>Release Information</b>	Command introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Display the interfaces on which sFlow is enabled and the sampling parameters for the interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show sflow on page 6851</a></li> <li>• <a href="#">show sflow collector on page 6853</a></li> <li>• <a href="#">Example: Monitoring Network Traffic Using sFlow Technology on page 6571</a></li> <li>• <a href="#">Configuring sFlow Technology on page 6596</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show sflow interface (QFX3500 Switch in Standalone Mode) on page 6854</a> <a href="#">show sflow interface (QFabric System) on page 6855</a>
<b>Output Fields</b>	Table 656 on page 6854 lists the output fields for the <b>show sflow interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 656: show sflow interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which sFlow technology is enabled.	All levels
Status Egress	Indicates whether an egress sample rate is enabled.	All levels
Status Ingress	Indicates whether an ingress sample rate is enabled.	All levels
Sample rate Egress	Rate at which egress packets are sampled.	All levels
Sample rate Ingress	Rate at which ingress packets are sampled.	All levels
Adapted sample rate Egress	Adapted rate at which egress packets are sampled.	All levels
Adapted sample rate Ingress	Adapted rate at which ingress packets are sampled.	All levels
Polling-interval	Interval at which the sFlow agent polls the interface.	All levels

## Sample Output

### show sflow interface (QFX3500 Switch in Standalone Mode)

```
user@host> show sflow interface
```



Interface	Status	Sample rate		Adapted sample rate		Polling-interval	
		Egress	Ingress	Egress	Ingress		
xe-0/0/0.0	Enabled	Disabled	1000	2048	1000	2048	20
xe-1/0/1.0	Enabled	Disabled	1000	2048	1000	2048	20

## Sample Output

### show sflow interface (QFabric System)

```

user@host> show sflow interface
Interface  Status      Sample rate    Adapted sample rate  Polling-interval
           Egress Ingress  Egress Ingress  Egress Ingress
node1:xe-0/0/0.0  Enabled Disabled 1000 2048 1000 2048 20
node2:xe-1/0/1.0  Enabled Disabled 1000 2048 1000 2048 20
node4:xe-1/0/0.0  Enabled Disabled 1000 2048 1000 2048 20

```

## Commands for SNMP

- clear snmp history
- clear snmp statistics
- request snmp spoof-trap
- request snmp utility-mib clear instance
- request snmp utility-mib set instance
- show snmp health-monitor
- show snmp inform-statistics
- show snmp mib
- show snmp rmon
- show snmp rmon history
- show snmp statistics
- show snmp v3

## clear snmp history

---

**Syntax**    clear snmp history (*index* | all)

**Release Information**    Command introduced in Junos OS Release 11.1 for the QFX Series.

**Description**    Delete the samples of Ethernet statistics collected for a history group.

**Options**    all—Clear all the entries in the history index.

*index*—Clear the contents of the specified entry in the history index.

**Required Privilege Level**    clear

**Related Documentation**    • [clear snmp statistics on page 6857](#)

## clear snmp statistics

<b>Syntax</b>	clear snmp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Clear Simple Network Management Protocol (SNMP) statistics.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show snmp statistics on page 1503</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear snmp statistics on page 6857</a>
<b>Output Fields</b>	See <a href="#">show snmp statistics</a> for an explanation of output fields.

## Sample Output

### clear snmp statistics

In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290
```

```
user@host> clear snmp statistics
```

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
```

```
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops 0  
Output:  
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

## request snmp spoof-trap

<b>Syntax</b>	<b>request snmp spoof-trap</b> <b>&lt;trap&gt; variable-bindings &lt;object&gt; &lt;instance&gt; &lt;value&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.
<b>Options</b>	<p><b>&lt;trap&gt;</b>—Name of the trap to spoof.</p> <p><b>variable-bindings &lt;object&gt; &lt;instance&gt; &lt;value&gt;</b>—(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, <b>ifIndex[14] = 14</b>). Enclose the list of variable bindings in quotation marks ( " ") and use a comma to separate each object name, instance, and value definition (for example, <b>variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"</b>). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.</p> <p><b>&lt;dummy name&gt;</b>—A dummy trap name to display the list of available traps.</p> <p><b>Question mark (?)</b>—Question mark? to display possible completions.</p>
<b>Required Privilege Level</b>	request
<b>List of Sample Output</b>	<a href="#">request snmp spoof-trap (with Variable Bindings) on page 6859</a> <a href="#">request snmp spoof-trap (Illegal Trap Name) on page 6859</a> <a href="#">request snmp spoof-trap (Question Mark ?) on page 6863</a>

## Sample Output

### request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmp spoof-trap linkUp variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"
Spoof trap request result: trap sent successfully
```

### request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found
```

```
Allowed Traps:
ads1AtucInitFailureTrap
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
```

ads1AturPerfLofsThreshTrap  
ads1AturPerfLossThreshTrap  
ads1AturPerfLprsThreshTrap  
ads1AturRateChangeTrap  
apsEventChannelMismatch  
apsEventFEPLF  
apsEventModeMismatch  
apsEventPSBF  
apsEventSwitchover  
authenticationFailure  
bfdSessDown  
bfdSessUp  
bgpBackwardTransition  
bgpEstablished  
coldStart  
dlswTrapCircuitDown  
dlswTrapCircuitUp  
dlswTrapTConnDown  
dlswTrapTConnPartnerReject  
dlswTrapTConnProtViolation  
dlswTrapTConnUp  
dsx1LineStatusChange  
dsx3LineStatusChange  
entConfigChange  
fallingAlarm  
frDLCIStatusChange  
ggsnTrapChanged  
ggsnTrapCleared  
ggsnTrapNew  
gmp1sTunnelDown  
ifMauJabberTrap  
ipv6IfStateChange  
isisAreaMismatch  
isisAttemptToExceedMaxSequence  
isisAuthenticationFailure  
isisAuthenticationTypeFailure  
isisCorruptedLSPDetected  
isisDatabaseOverload  
isisIDLenMismatch  
isisLSPTooLargeToPropagate  
isisManualAddressDrops  
isisMaxAreaAddressesMismatch  
isisOriginatingLSPBufferSizeMismatch  
isisOwnLSPPurge  
isisProtocolsSupportedMismatch  
isisRejectedAdjacency  
isisSequenceNumberSkip  
isisVersionSkew  
jnxAccessAuthServerDisabled  
jnxAccessAuthServerEnabled  
jnxAccessAuthServiceDown  
jnxAccessAuthServiceUp  
jnxBfdSessDetectionTimeHigh  
jnxBfdSessTxIntervalHigh  
jnxBgpM2BackwardTransition  
jnxBgpM2Established  
jnxCmCfgChange  
jnxCmRescueChange  
jnxCollFlowOverload  
jnxCollFlowOverloadCleared  
jnxCollFtpSwitchover

jnxCollMemoryAvailable  
jnxCollMemoryUnavailable  
jnxCollUnavailableDest  
jnxCollUnavailableDestCleared  
jnxCollUnsuccessfulTransfer  
jnxDfcHardMemThresholdExceeded  
jnxDfcHardMemUnderThreshold  
jnxDfcHardPpsThresholdExceeded  
jnxDfcHardPpsUnderThreshold  
jnxDfcSoftMemThresholdExceeded  
jnxDfcSoftMemUnderThreshold  
jnxDfcSoftPpsThresholdExceeded  
jnxDfcSoftPpsUnderThreshold  
jnxEventTrap  
jnxExampleStartup  
jnxFEBSwitchover  
jnxFanFailure  
jnxFanOK  
jnxFruCheck  
jnxFruFailed  
jnxFruInsertion  
jnxFruOK  
jnxFruOffline  
jnxFruOnline  
jnxFruPowerOff  
jnxFruPowerOn  
jnxFruRemoval  
jnxHardDiskFailed  
jnxHardDiskMissing  
jnxJsAvPatternUpdateTrap  
jnxJsChassisClusterSwitchover  
jnxJsFwAuthCapacityExceeded  
jnxJsFwAuthFailure  
jnxJsFwAuthServiceDown  
jnxJsFwAuthServiceUp  
jnxJsNatAddrPoolThresholdStatus  
jnxJsScreenAttack  
jnxJsScreenCfgChange  
jnxLdpLspDown  
jnxLdpLspUp  
jnxLdpSesDown  
jnxLdpSesUp  
jnxMIMstCistPortLoopProtectStateChangeTrap  
jnxMIMstCistPortRootProtectStateChangeTrap  
jnxMIMstErrTrap  
jnxMIMstGenTrap  
jnxMIMstInvalidBpduRxdTrap  
jnxMIMstMstiPortLoopProtectStateChangeTrap  
jnxMIMstMstiPortRootProtectStateChangeTrap  
jnxMIMstNewRootTrap  
jnxMIMstProtocolMigrationTrap  
jnxMIMstRegionConfigChangeTrap  
jnxMIMstTopologyChgTrap  
jnxMacChangedNotification  
jnxMplsLdpInitSesThresholdExceeded  
jnxMplsLdpPathVectorLimitMismatch  
jnxMplsLdpSessionDown  
jnxMplsLdpSessionUp  
jnxOspfV3IfConfigError  
jnxOspfV3IfRxBadPacket  
jnxOspfV3IfStateChange

jnxOspfV3LsdbApproachingOverflow  
jnxOspfV3LsdbOverflow  
jnxOspfV3NbrRestartHelperStatusChange  
jnxOspfV3NbrStateChange  
jnxOspfV3NssaTranslatorStatusChange  
jnxOspfV3RestartStatusChange  
jnxOspfV3VirtIfConfigError  
jnxOspfV3VirtIfRxBadPacket  
jnxOspfV3VirtIfStateChange  
jnxOspfV3VirtNbrRestartHelperStatusChange  
jnxOspfV3VirtNbrStateChange  
jnxOtnAlarmCleared  
jnxOtnAlarmSet  
jnxOverTemperature  
jnxPMonOverloadCleared  
jnxPMonOverloadSet  
jnxPingEgressJitterThresholdExceeded  
jnxPingEgressStdDevThresholdExceeded  
jnxPingEgressThresholdExceeded  
jnxPingIngressJitterThresholdExceeded  
jnxPingIngressStdDevThresholdExceeded  
jnxPingIngressThresholdExceeded  
jnxPingRttJitterThresholdExceeded  
jnxPingRttStdDevThresholdExceeded  
jnxPingRttThresholdExceeded  
jnxPortBpduErrorStatusChangeTrap  
jnxPortLoopProtectStateChangeTrap  
jnxPortRootProtectStateChangeTrap  
jnxPowerSupplyFailure  
jnxPowerSupplyOK  
jnxRedundancySwitchover  
jnxRmonAlarmGetFailure  
jnxRmonGetOk  
jnxSecAccessIfMacLimitExceeded  
jnxSecAccessSdsRateLimitCrossed  
jnxSonetAlarmCleared  
jnxSonetAlarmSet  
jnxSpSvcSetCpuExceeded  
jnxSpSvcSetCpuOk  
jnxSpSvcSetZoneEntered  
jnxSpSvcSetZoneExited  
jnxStormEventNotification  
jnxSyslogTrap  
jnxTemperatureOK  
jnxVccpPortDown  
jnxVccpPortUp  
jnxVpnIfDown  
jnxVpnIfUp  
jnxVpnPwDown  
jnxVpnPwUp  
jnxl2aldGlobalMacLimit  
jnxl2aldInterfaceMacLimit  
jnxl2aldRoutingInstMacLimit  
linkDown  
linkUp  
lldpRemTablesChange  
mfrMibTrapBundleLinkMismatch  
mplsLspChange  
mplsLspDown  
mplsLspInfoChange  
mplsLspInfoDown



```

mplsLspInfoPathDown
mplsLspInfoPathUp
mplsLspInfoUp
mplsLspPathDown
mplsLspPathUp
mplsLspUp
mplsNumVrfRouteMaxThreshExceeded
mplsNumVrfRouteMidThreshExceeded
mplsNumVrfSecIllglLb1ThrshExcd
mplsTunnelDown
mplsTunnelReoptimized
mplsTunnelRerouted
mplsTunnelUp
mplsVrfIfDown
mplsVrfIfUp
mplsXCDown
mplsXCUp
msdpBackwardTransition
msdpEstablished
newRoot
ospfIfAuthFailure
ospfIfConfigError
ospfIfRxBadPacket
ospfIfStateChange
ospfLsdbApproachingOverflow
ospfLsdbOverflow
ospfMaxAgeLsa
ospfNbrStateChange
ospfOriginateLsa
ospfTxRetransmit
ospfVirtIfAuthFailure
ospfVirtIfConfigError
ospfVirtIfRxBadPacket
ospfVirtIfStateChange
ospfVirtIfTxRetransmit
ospfVirtNbrStateChange
pethMainPowerUsageOffNotification
pethMainPowerUsageOnNotification
pethPsePortOnOffNotification
pingProbeFailed
pingTestCompleted
pingTestFailed
ptopoConfigChange
risingAlarm
rpMauJabberTrap
sd1cLSStatusChange
sd1cPortStatusChange
topologyChange
traceRoutePathChange
traceRouteTestCompleted
traceRouteTestFailed
vrrpTrapAuthFailure
vrrpTrapNewMaster
warmStart

```

#### request snmp spoof-trap (Question Mark ?)

```

user@host> request snmp spoof-trap ?
Possible completions:
<trap>                The name of the trap to spoof
ads1AtucInitFailureTrap

```

ads1AtucPerfESsThreshTrap  
ads1AtucPerfLofsThreshTrap  
ads1AtucPerfLolsThreshTrap  
ads1AtucPerfLossThreshTrap  
ads1AtucPerfLprsThreshTrap  
ads1AtucRateChangeTrap  
ads1AturPerfESsThreshTrap  
ads1AturPerfLofsThreshTrap  
ads1AturPerfLossThreshTrap  
ads1AturPerfLprsThreshTrap  
ads1AturRateChangeTrap  
apsEventChannelMismatch  
apsEventFEPLF  
apsEventModeMismatch  
apsEventPSBF  
apsEventSwitchover  
authenticationFailure  
bfdSessDown  
bfdSessUp  
bgpBackwardTransition  
bgpEstablished  
coldStart  
dlswTrapCircuitDown  
dlswTrapCircuitUp  
---(more 10%)---

---

## request snmp utility-mib clear instance

---

<b>Syntax</b>	request snmp utility-mib clear instance <i>name</i> object-type <i>type</i>
<b>Release Information</b>	Command introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	Clear the data stored in the specified container object in the SNMP Utility MIB.
<b>Options</b>	<p><b><i>name</i></b>—Name of the SNMP instance that is used to identify the data stored in the container object.</p> <p><b>object-type <i>type</i></b>—Type of container object in which the data is stored. The following container object types are supported:</p> <ul style="list-style-type: none"><li>• <b>counter</b>—Stores a 32-bit counter value.</li><li>• <b>counter64</b>—Stores a 64-bit counter value.</li><li>• <b>integer</b>—Stores a 32-bit signed integer value.</li><li>• <b>unsigned-integer</b>—Stores a 32-bit unsigned integer value.</li></ul>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Utility MIB on page 6522</a></li><li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 6516</a></li><li>• <a href="#">request snmp utility-mib set instance on page 6866</a></li></ul>

## request snmp utility-mib set instance

---

<b>Syntax</b>	<code>request snmp utility-mib set instance <i>name</i></code> <code>object-type <i>type</i></code> <code>object-value <i>value</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 12.2 for the QFX Series.
<b>Description</b>	Store data in the specified container object in the SNMP Utility MIB. The data may be retrieved by SNMP operations.
<b>Options</b>	<p><b><i>name</i></b>—Name of the SNMP instance that is used to identify the data stored in the container object.</p> <p><b><i>object-type type</i></b>—Type of container object in which to store data. The following container object types are supported:</p> <ul style="list-style-type: none"><li>• <b>counter</b>—Stores a 32-bit counter value.</li><li>• <b>counter64</b>—Stores a 64-bit counter value.</li><li>• <b>integer</b>—Stores a 32-bit signed integer value.</li><li>• <b>unsigned-integer</b>—Stores a 32-bit unsigned integer value.</li><li>• <b>string</b>—Stores an octet string value.</li></ul> <p><b><i>object-value value</i></b>—Data that is stored in the container object.</p>
<b>Required Privilege Level</b>	request
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Utility MIB on page 6522</a></li><li>• <a href="#">Understanding the Implementation of SNMP on the QFabric System on page 6516</a></li><li>• <a href="#">request snmp utility-mib clear instance on page 6865</a></li></ul>

## show snmp health-monitor

<b>Syntax</b>	show snmp health-monitor <alarms (brief   detail)   logs>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.
<b>Options</b>	<p><b>none</b>—Display information about all health monitor alarms and logs.</p> <p><b>alarms (brief   detail)</b>—(Optional) Display information about health monitor alarms. Optionally, specify brief or detailed information about the alarms.</p> <p><b>logs</b>—(Optional) Display information about health monitor logs.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Health Monitoring on page 6529</a></li> <li>• <a href="#">Configuring Health Monitoring on page 6609</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show snmp health-monitor on page 6869</a></p> <p><a href="#">show snmp health-monitor alarms detail on page 6869</a></p>
<b>Output Fields</b>	Table 657 on page 6867 describes the output fields for the <b>show snmp health-monitor</b> command. Output fields are listed in the approximate order in which they appear.

**Table 657: show snmp health-monitor Output Fields**

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
Variable description	Description of the health monitor object instance being monitored.	All levels
Variable name	Name of the health monitor object instance being monitored.	All levels
Value	Current value of the monitored variable in the most recent sample interval.	All levels

Table 657: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <ul style="list-style-type: none"> <li>Alarms: <ul style="list-style-type: none"> <li>active—Entry is fully configured and activated.</li> <li>falling threshold crossed—Value of the variable has crossed the lower threshold limit.</li> <li>rising threshold crossed—Value of the variable has crossed the upper threshold limit.</li> <li>under creation—Entry is being configured and is not yet activated.</li> <li>startup—Alarm is waiting for the first sample of the monitored variable.</li> <li>object not available—Monitored variable of that type is not available to the health monitor agent.</li> <li>instance not available—Monitored variable's instance is not available to the health monitor agent.</li> <li>object type invalid—Monitored variable is not a numeric value.</li> <li>object processing errored—An error occurred when the monitored variable was processed.</li> <li>unknown—State is not one of the above.</li> </ul> </li> </ul>	All levels
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value <i>absolute value</i> or <i>delta value</i> .	detail
Startup alarm	<p>Alarm that might be sent when this entry is first activated, depending on the following criteria:</p> <ul style="list-style-type: none"> <li>Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. <i>falling alarm</i></li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is either <i>falling alarm</i> or <i>rising or falling alarm</i>.</li> </ul> </li> <li>Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is <i>falling alarm</i>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is <i>rising alarm</i>.</li> <li>Value of the alarm is between the thresholds.</li> </ul> </li> </ul>	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has <b>monitor</b> prepended to it.	detail
Creator	Mechanism by which the entry was configured (Health Monitor).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value as a percentage of the maximum possible value.	detail

Table 657: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Falling threshold	Lower limit threshold value as a percentage of the maximum possible value.	detail
Rising event index	Index number of the event triggered when the rising threshold is crossed.	detail
Falling event index	Index number of the event triggered when the falling threshold is crossed. Details include the value of the falling event instance and the state of the falling event instance.	detail

## Sample Output

### show snmp health-monitor

```

user@switch> show snmp health-monitor

Alarm
Index  Variable description                                Value State
-----
32768  Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                          59 active

32769  Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                          0 active

32770  Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                            9 falling threshold

32772  Health Monitor: RE 0 memory utilization
      jnxOperatingBuffer.9.1.0.0                         23 active

32774  Health Monitor: Max Kernel Memory Used (%)
      jnxBoxKernelMemoryUsedPercent.0                    3 active
Event Index: 32768
Description: Health Monitor: RE 0 CPU utilization crossed falling threshold
70 (value: 5), (variable: jnxOperatingCPU.9.1.0.0)
Time: 2011-01-09 19:18:35 PST

```

### show snmp health-monitor alarms detail

```

user@switch> show snmp health-monitor alarms detail

Alarm Index 32768:
Variable name      jnxHrStoragePercentUsed.1
Variable OID       1.3.6.1.4.1.2636.3.31.1.1.1.1.1
Sample type        absolute value
Startup alarm      rising alarm
Owner              Health Monitor: root file system
                  utilization
Creator            Health Monitor
State              active
Sample interval    300 seconds
Rising threshold   80

```

Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 59  
Instance State: active

Alarm Index 32769:

Variable name jnxHrStoragePercentUsed.2  
Variable OID 1.3.6.1.4.1.2636.3.31.1.1.1.2  
Sample type absolute value  
Startup alarm rising alarm  
Owner Health Monitor: /config file system  
utilization  
Creator Health Monitor  
State active  
Sample interval 300 seconds  
Rising threshold 80  
Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 0  
Instance State: active

Alarm Index 32770:

Variable name jnxOperatingCPU.9.1.0.0  
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0  
Sample type absolute value  
Startup alarm rising alarm  
Owner Health Monitor: RE 0 CPU utilization  
Creator Health Monitor  
State active  
Sample interval 300 seconds  
Rising threshold 80  
Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 9  
Instance State: falling threshold

Alarm Index 32772:

Variable name jnxOperatingBuffer.9.1.0.0  
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0  
Sample type absolute value  
Startup alarm rising alarm  
Owner Health Monitor: RE 0 memory utilization  
Creator Health Monitor  
State active  
Sample interval 300 seconds  
Rising threshold 80  
Falling threshold 70  
Rising event index 32768  
Falling event index 32768  
Instance Value: 23  
Instance State: active

Alarm Index 32774:

Variable name jnxBoxKernelMemoryUsedPercent.0  
Variable OID 1.3.6.1.4.1.2636.3.1.16.0  
Sample type absolute value



Startup alarm	rising alarm
Owner	Health Monitor: Max Kernel Memory Used (%)
Creator	Health Monitor
State	active
Sample interval	300 seconds
Rising threshold	80
Falling threshold	70
Rising event index	32768
Falling event index	32768
Instance Value:	3
Instance State:	active

## show snmp inform-statistics

<b>Syntax</b>	show snmp inform-statistics
<b>Release Information</b>	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about Simple Network Management Protocol (SNMP) inform requests.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show snmp inform-statistics on page 6872</a>
<b>Output Fields</b>	<a href="#">Table 658 on page 6872</a> describes the output fields for the <b>show snmp inform-statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 658: show snmp inform-statistics Output Fields**

Field Name	Field Description
<b>Target Name</b>	Name of the device configured to receive and respond to SNMP informs.
<b>Address</b>	IP address of the target device.
<b>Sent</b>	Number of informs sent to the target device and acknowledged by the target device.
<b>Pending</b>	Number of informs held in memory pending a response from the target device.
<b>Discarded</b>	Number of informs discarded after the specified number of retransmissions to the target device were attempted.
<b>Timeouts</b>	Number of informs that did not receive an acknowledgement from the target device within the timeout specified.
<b>Probe Failures</b>	Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address).

## Sample Output

### show snmp inform-statistics

```

user@host> show snmp inform-statistics
Inform Request Statistics:
  Target Name: TA1_v3_md5_none Address: 172.17.20.184
    Sent: 176, Pending: 0
    Discarded: 0, Timeouts: 0, Probe Failures: 0
  Target Name: TA2_v3_sha_none Address: 192.168.110.59

```

Sent: 0, Pending: 4  
Discarded: 84, Timeouts: 0, Probe Failures: 258  
Target Name: TA5\_v2\_none Address: 172.17.20.184  
Sent: 0, Pending: 0  
Discarded: 2, Timeouts: 10, Probe Failures: 0

## show snmp mib

---

<b>Syntax</b>	<code>show snmp mib (get   get-next   walk) (ascii   decimal) <i>object-id</i></code>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>ascii</b> and <b>decimal</b> options introduced in Junos OS Release 9.6.</p> <p><b>ascii</b> and <b>decimal</b> options introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
<b>Description</b>	Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.
<b>Options</b>	<p><b>get</b>—Retrieve and display one or more SNMP object values.</p> <p><b>get-next</b>—Retrieve and display the next SNMP object values.</p> <p><b>walk</b>—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.</p> <p><b>ascii</b>—Display the SNMP object's string indices as an ASCII-key representation.</p> <p><b>decimal</b>—Display the SNMP object values in the decimal (default) format. The <b>decimal</b> option is the default option for this command. Therefore, issuing the <b>show snmp mib (get   get-next   walk) decimal object-id</b> and the <b>show snmp mib (get   get-next   walk) object-id</b> commands display the same output.</p> <p><b>object-id</b>—The object can be represented by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as <b>interfaces</b>). When entering multiple objects, enclose the objects in quotation marks.</p>
<b>Required Privilege Level</b>	snmp—To view this statement in the configuration.
<b>List of Sample Output</b>	<p><a href="#">show snmp mib get on page 6875</a></p> <p><a href="#">show snmp mib get (Multiple Objects) on page 6875</a></p> <p><a href="#">show snmp mib get (Layer 2 Policer) on page 6875</a></p> <p><a href="#">show snmp mib get-next on page 6875</a></p> <p><a href="#">show snmp mib get-next (Specify an OID) on page 6875</a></p> <p><a href="#">show snmp mib walk on page 6875</a></p> <p><a href="#">show snmp mib walk (QFX Series) on page 6875</a></p> <p><a href="#">show snmp mib walk decimal on page 6876</a></p> <p><a href="#">show snmp mib walk (ASCII) on page 6876</a></p> <p><a href="#">show snmp mib walk (Multiple Indices) on page 6876</a></p> <p><a href="#">show snmp mib walk decimal (Multiple Indices) on page 6876</a></p>
<b>Output Fields</b>	<p>Table 659 on page 6875 describes the output fields for the <b>show snmp mib</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 659: show snmp mib Output Fields

Field Name	Field Description
<i>name</i>	Object name and numeric instance value.
<i>object value</i>	Object value. The Junos OS translates OIDs into the corresponding object names.

## Sample Output

### show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

### show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get ?sysObjectID.0 sysUpTime.0?
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

### show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

### show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

### show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
Networks, Inc.
```

### show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004
Juniper Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

### show snmp mib walk (QFX Series)

```
user@switch> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS
11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 138980301
sysContact.0 = System Contact
```

```
sysName.0      = LabQFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

#### show snmp mib walk decimal

```
user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100
```

#### show snmp mib walk (ASCII)

```
show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100
```

#### show snmp mib walk (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

#### show snmp mib walk decimal (Multiple Indices)

```
show snmp mib walk decimal jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

## show snmp rmon

<b>Syntax</b>	show snmp rmon <alarms (brief   detail)> <events (brief   detail)> <logs>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms, events, and logs.
<b>Options</b>	<p><b>none</b>—Display information about all RMON alarms and events.</p> <p><b>brief   detail</b>—(Optional) Display brief or detailed information about RMON alarms or events.</p> <p><b>alarms</b>—(Optional) Display information about RMON alarms.</p> <p><b>events</b>—(Optional) Display information about RMON events.</p> <p><b>logs</b>—(Optional) Display information about RMON monitoring logs.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li> <li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li> <li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li> <li>• <a href="#">Understanding RMON on page 6525</a></li> <li>• <a href="#">clear snmp statistics on page 6857</a></li> <li>• <a href="#">clear snmp history on page 6856</a></li> <li>• <a href="#">show snmp rmon history on page 6881</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show snmp rmon on page 6879</a> <a href="#">show snmp rmon alarms detail on page 6880</a> <a href="#">show snmp rmon events detail on page 6880</a> <a href="#">show snmp rmon logs on page 6880</a>
<b>Output Fields</b>	Table 660 on page 6877 describes the output fields for the <b>show snmp rmon</b> command. Output fields are listed in the approximate order in which they appear.

**Table 660: show snmp rmon Output Fields**

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels

Table 660: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	<p>State of the alarm or event entry:</p> <p>Alarms:</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Entry is fully configured and activated.</li> <li>• <b>falling threshold crossed</b>—Value of the variable has crossed the lower threshold limit.</li> <li>• <b>rising threshold crossed</b>—Value of the variable has crossed the upper threshold limit.</li> <li>• <b>under creation</b>—Entry is being configured and is not yet activated.</li> <li>• <b>startup</b>—Alarm is waiting for the first sample of the monitored variable.</li> <li>• <b>object not available</b>—Monitored variable of that type is not available to the SNMP agent.</li> <li>• <b>instance not available</b>—Monitored variable's instance is not available to the SNMP agent.</li> <li>• <b>object type invalid</b>—Monitored variable is not a numeric value.</li> <li>• <b>object processing errored</b>—An error occurred when the monitored variable was processed.</li> <li>• <b>unknown</b>—State is not one of the above.</li> </ul> <p>Events:</p> <ul style="list-style-type: none"> <li>• <b>active</b>—Entry has been fully configured and activated.</li> <li>• <b>under creation</b>—Entry is being configured and is not yet activated.</li> <li>• <b>unknown</b>—State is not one of the above.</li> </ul>	All levels
<b>Variable name</b>	Name of the SNMP object instance being monitored.	All levels
<b>Event Index</b>	Event identifier.	All levels
<b>Type</b>	<p>Type of notification made when an event is triggered. It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>log</b>—A system log message is generated and an entry is made to the log table.</li> <li>• <b>snmptrap</b>—An SNMP trap is sent to the configured destination.</li> <li>• <b>log and trap</b>—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination.</li> <li>• <b>none</b>—Neither log nor trap will be sent.</li> </ul>	<b>detail</b>
<b>Last Event</b>	Date and time of the last event. It has the format <i>yyyy-mm-dd hh:mm:ss timezone</i> .	<b>brief</b>
<b>Community</b>	Trap group used for sending the SNMP trap.	<b>detail</b>
<b>Variable OID</b>	Object ID to which the variable name is resolved. The format is x.x.x.x.	<b>detail</b>
<b>Sample type</b>	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of <b>absolute value</b> or <b>delta value</b> .	<b>detail</b>



Table 660: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Startup alarm</b>	Alarm that might be sent when this entry is first activated, depending on the following criteria: <ul style="list-style-type: none"> <li>Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is either <b>rising alarm</b> or <b>rising or falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is either <b>falling alarm</b> or <b>rising or falling alarm</b>.</li> </ul> </li> <li>Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> <li>Value of the alarm is above or equal to the rising threshold and the startup type is <b>falling alarm</b>.</li> <li>Value of the alarm is below or equal to the falling threshold and the startup type is <b>rising alarm</b>.</li> <li>Value of the alarm is between the thresholds.</li> </ul> </li> </ul>	<b>detail</b>
<b>Owner</b>	Name of the entry configured by the user. If the entry was created through the CLI, the owner has <b>monitor</b> prepended to it.	<b>detail</b>
<b>Creator</b>	Mechanism by which the entry was configured ( <b>CLI</b> or <b>SNMP</b> ).	<b>detail</b>
<b>Sample interval</b>	Time period between samples (in seconds).	<b>detail</b>
<b>Rising threshold</b>	Upper limit threshold value configured by the user.	<b>detail</b>
<b>Falling threshold</b>	Lower limit threshold value configured by the user.	<b>detail</b>
<b>Rising event index</b>	Event triggered when the rising threshold is crossed.	<b>detail</b>
<b>Falling event index</b>	Event triggered when the falling threshold is crossed.	<b>detail</b>
<b>Current value</b>	Current value of the monitored variable in the most recent sample interval.	<b>detail</b>

## Sample Output

### show snmp rmon

```

user@host> show snmp rmon
Alarm
Index  Variable description                               Value State

      5  monitor
         jnxOperatingCPU.9.1.0.0                      5 falling threshold

Event
Index  Type                               Last Event
      1  log and trap                     2009-07-10 11:34:17 PDT
Event Index: 1
Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
Time: 2009-07-10 11:34:07 PDT

```

Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,  
(variable: jnxOperatingCPU.9.1.0.0, value: 5)  
Time: 2009-07-10 11:34:17 PDT

#### show snmp rmon alarms detail

```
user@host> show snmp rmon alarms detail
Alarm Index 5:
  Variable name           jnxOperatingCPU.9.1.0.0
  Variable OID            1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
  Sample type             absolute value
  Startup alarm           rising or falling alarm
  Owner                   monitor

  Creator                 CLI
  State                   active
  Sample interval         5 seconds
  Rising threshold        90
  Falling threshold       75
  Rising event index      1
  Falling event index     1
  Instance Value: 4
  Instance State: falling threshold
```

#### show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
  Description             rmon event
  Type                    log and trap
  Community               rmon-trap-group
  Last event              2009-07-10 11:34:17 PDT
  Creator                 CLI
  State                   active
```

#### show snmp rmon logs

```
user@host> show snmp rmon logs
Event Index: 1
  Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
  Time: 2009-07-10 11:34:07 PDT
  Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
  Time: 2009-07-10 11:34:17 PDT
```

---

## show snmp rmon history

---

<b>Syntax</b>	show snmp rmon history <history-index> sample-index <sample-index>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the contents of the RMON history group.
<b>Options</b>	<p><b>none</b>—Display all the entries in the RMON history group.</p> <p><b>history-index</b>—(Optional) Display the contents of the specified entry in the RMON history group.</p> <p><b>sample-index sample-index</b>—(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">RMON MIB Event, Alarm, Log, and History Control Tables on page 6527</a></li><li>• <a href="#">Monitoring RMON MIB Tables on page 6803</a></li><li>• <a href="#">Configuring RMON Alarms and Events on page 6606</a></li><li>• <a href="#">Understanding RMON on page 6525</a></li><li>• <a href="#">clear snmp statistics on page 6857</a></li><li>• <a href="#">clear snmp history on page 6856</a></li><li>• <a href="#">show snmp rmon on page 6877</a></li></ul>

## show snmp statistics

<b>Syntax</b>	show snmp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear snmp statistics on page 6857</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show snmp statistics on page 6885</a>
<b>Output Fields</b>	<a href="#">Table 91 on page 1503</a> describes the output fields for the <b>show snmp statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 661: show snmp statistics Output Fields

Field Name	Field Description
<b>Input</b>	<p>Information about received packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets(snmplnPkts)</b>—Total number of messages delivered to the SNMP entity from the transport service.</li> <li>• <b>Bad versions—(snmplnBadVersions)</b> Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version.</li> <li>• <b>Bad community names—(snmplnBadCommunityNames)</b> Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity.</li> <li>• <b>Bad community uses—(snmplnBadCommunityUses)</b> Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.</li> <li>• <b>ASN parse errors—(snmplnASNParseErrs)</b> Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.</li> <li>• <b>Too big—(snmplnTooBig)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmplnNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmplnBadValues)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>Read only—(snmplnReadOnly)</b> Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of <b>readOnly</b>. Only incorrect implementations of SNMP generate this error.</li> </ul>

Table 661: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> <li>• <b>General errors—(snmpInGenErrs)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Total requests varbinds—(snmpInTotalReqVars)</b> Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP <b>GetRequest</b> and <b>GetNext</b> PDUs.</li> <li>• <b>Total set varbinds—(snmpInSetVars)</b> Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP <b>SetRequest</b> PDUs.</li> <li>• <b>Get requests—(snmpInGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get nexts—(snmpInGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Set requests—(snmpInSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Get responses—(snmpInGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs that have been accepted and processed by the SNMP entity.</li> <li>• <b>Traps—(snmpInTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> <li>• <b>Silent drops—(snmpSilentDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.</li> <li>• <b>Proxy drops—(snmpProxyDrops)</b> Total number of <b>GetRequest</b>, <b>GetNextRequest</b>, <b>GetBulkRequest</b>, <b>SetRequests</b>, and <b>InformRequest</b> PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.</li> <li>• <b>Commit pending drops</b>—Number of SNMP packets for <b>Set</b> requests dropped because of a previous pending SNMP <b>Set</b> request on the committed configuration.</li> <li>• <b>Throttle drops</b>—Number of SNMP packets for any requests dropped reaching the throttle limit.</li> </ul>

Table 661: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> <li>• <b>Unknown security models—(snmpUnknownSecurityModels)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine.</li> <li>• <b>Invalid messages—(snmpInvalidMsgs)</b> Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message.</li> <li>• <b>Unknown pdu handlers—(snmpUnknownPDUHandlers)</b> Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type.</li> <li>• <b>Unavailable contexts—(snmpUnavailableContexts)</b> Number of requests received for a context that is known to the SNMP engine, but is currently unavailable.</li> <li>• <b>Unknown contexts—(snmpUnknownContexts)</b> Total number of requests received for a context that is unknown to the SNMP engine.</li> <li>• <b>Unsupported security levels—(usmStatsUnsupportedSecLevels)</b> Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable).</li> <li>• <b>Not in time windows—(usmStatsNotInTimeWindows)</b> Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window.</li> <li>• <b>Unknown user names—(usmStatsUnknownUserNames)</b> Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine.</li> <li>• <b>Unknown engine ids—(usmStatsUnknownEngineIDs)</b> Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine.</li> <li>• <b>Wrong digests—(usmStatsWrongDigests)</b> Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value.</li> <li>• <b>Decryption errors—(usmStatsDecryptionErrors)</b> Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.</li> </ul>

Table 661: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Output</b>	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> <li>• <b>Packets—(snmpOutPkts)</b> Total number of messages passed from the SNMP entity to the transport service.</li> <li>• <b>Too big—(snmpOutTooBig)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>tooBig</b>.</li> <li>• <b>No such names—(snmpOutNoSuchNames)</b> Total number of SNMP PDUs delivered to the SNMP entity with an error status field of <b>noSuchName</b>.</li> <li>• <b>Bad values—(snmpOutBadValues)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>badValue</b>.</li> <li>• <b>General errors—(snmpOutGenErrs)</b> Total number of SNMP PDUs generated by the SNMP entity with an error status field of <b>genErr</b>.</li> <li>• <b>Get requests—(snmpOutGetRequests)</b> Total number of SNMP <b>GetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get nexts—(snmpOutGetNexts)</b> Total number of SNMP <b>GetNext</b> PDUs generated by the SNMP entity.</li> <li>• <b>Set requests—(snmpOutSetRequests)</b> Total number of SNMP <b>SetRequest</b> PDUs generated by the SNMP entity.</li> <li>• <b>Get responses—(snmpOutGetResponses)</b> Total number of SNMP <b>GetResponse</b> PDUs generated by the SNMP entity.</li> <li>• <b>Traps—(snmpOutTraps)</b> Total number of SNMP traps generated by the SNMP entity.</li> </ul>

## Sample Output

### show snmp statistics

```

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0

```

## show snmp v3

---

<b>Syntax</b>	<code>show snmp v3</code> <code>&lt;access &lt;brief   detail&gt;   community   general   groups   notify &lt;filter&gt;   target &lt;address   parameters&gt;   users&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.
<b>Options</b>	<p><b>none</b>—Display all of the SNMPv3 operating configuration.</p> <p><b>access</b>—(Optional) Display SNMPv3 access information.</p> <p><b>brief   detail</b>—(Optional) Display brief or detailed information about SNMPv3 access information.</p> <p><b>community</b>—(Optional) Display SNMPv3 community information.</p> <p><b>general</b>—(Optional) Display SNMPv3 general information.</p> <p><b>groups</b>—(Optional) Display SNMPv3 security-to-group information.</p> <p><b>notify &lt;filter&gt;</b>—(Optional) Display SNMPv3 notify information and, optionally, notify filter information.</p> <p><b>target &lt;address   parameters&gt;</b>—(Optional) Display SNMPv3 target information and, optionally, either target address or target parameter information.</p> <p><b>users</b>—(Optional) Display SNMPv3 user information.</p>
<b>Additional Information</b>	To edit the default display of the <b>show snmp v3</b> command, specify options in the <b>show</b> statement at the <b>[edit snmp v3]</b> hierarchy level.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">SNMPv3 Overview on page 6523</a></li><li>• <a href="#">Minimum SNMPv3 Configuration on a Device Running Junos OS on page 6524</a></li><li>• <a href="#">Configuring Access Privileges for a Group on page 6611</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show snmp v3 on page 6887</a>
<b>Output Fields</b>	<a href="#">Table 662 on page 6887</a> describes the output fields for the <b>show snmp v3</b> command. Output fields are listed in the approximate order in which they appear.



Table 662: show snmp v3 Output Fields

Field Name	Field Description
Local engine	<p>Information about the local SNMP engine configuration:</p> <ul style="list-style-type: none"> <li>• <b>Local engine ID</b>—Unique Identifier of the local SNMPv3 engine.</li> <li>• <b>Engine boots</b>—Number of times the local SNMPv3 engine has rebooted or reinitialized since this engine ID was configured.</li> <li>• <b>Engine time</b>—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized.</li> <li>• <b>Max msg size</b>—Maximum message size the sender can accommodate.</li> </ul>
Engine ID (local engine)	<p>Information about the local SNMP engine ID and the associated users:</p> <ul style="list-style-type: none"> <li>• <b>User</b>—SNMPv3 username.</li> <li>• <b>Auth/Priv</b>—Authentication and encryption algorithm that is configured for the user.</li> <li>• <b>Storage</b>—Indicates whether a username is saved to the configuration file (nonvolatile) or not saved (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of the user as listed in the SNMPv3 user table. Only rows with an active status in the table are used by the SNMPv3 engine.</li> </ul>
Engine ID (remote engine)	<p>Information about a remote SNMP engine, associated users, user groups, and user access policies:</p> <ul style="list-style-type: none"> <li>• <b>User</b>—SNMPv3 username.</li> <li>• <b>Auth/Priv</b>—Authentication and encryption algorithm that is configured for the user.</li> <li>• <b>Storage</b>—Indicates whether a username is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of a new user that has been activated. Only users with an active status can use SNMPv3.</li> <li>• <b>Group name</b>—Name of a group of users for which the configured access privileges apply.</li> <li>• <b>Security model</b>—Security model (such as <b>usm</b>, <b>v1</b>, <b>v2c</b>, or <b>any</b>) that is configured for the group. The security model is used with the security name to ensure messaging security.</li> <li>• <b>Security name</b>—Security name that is associated with a user, and which is used with the security model to ensure messaging security.</li> <li>• <b>Storage type</b>—Indicates whether a username is saved to the configuration file (nonvolatile) or not saved (volatile). Applies only to users with active status.</li> <li>• <b>Status</b>—Status of a user in a group. Only users with an active status can use SNMPv3.</li> </ul>
Access control	<p>Information about access control:</p> <ul style="list-style-type: none"> <li>• <b>Group name</b>—Name of a group of users for which the configured access privileges apply.</li> <li>• <b>Context prefix</b>—SNMPv3 context for which the configured access privileges apply.</li> <li>• <b>Security model/level</b>—Security model and security level combination that is configured for user access privileges.</li> <li>• <b>Read view</b>—Identifies the MIB view used for SNMPv3 read operations.</li> <li>• <b>Write view</b>—Identifies the MIB view used for SNMPv3 write operations.</li> <li>• <b>Notify view</b>—Identifies the MIB view used for outbound SNMP notifications.</li> </ul>

## Sample Output

### show snmp v3

```
user@host> show snmp v3
```

```
Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots:      38
Engine time:       64583 seconds
Max msg size:      2048 bytes
```

Engine ID: local

User	Auth/Priv	Storage	Status
user1	md5/des	nonvolatile	active
user2	sha/none	nonvolatile	active
user3	none/none	nonvolatile	active

Engine ID: 81 00 0a 4c 04 64 64 64 64

User	Auth/Priv	Storage	Status
UNEW	md5/none	nonvolatile	active

Group name	Security model	Security name	Storage type	Status
g1	usm	user1	nonvolatile	active
g2	usm	user2	nonvolatile	active
g3	usm	user3	nonvolatile	active

Access control:

Group	Context prefix	Security model/level	Read view	Write view	Notify view
g1		usm/privacy	v1	v1	
g2		usm/authent	v1	v1	
g3		usm/none	v1	v1	

---

## Commands for Syslog

- [show log](#)

## show log

<b>List of Syntax</b>	<a href="#">Syntax on page 6889</a> <a href="#">Syntax (QFabric System) on page 6889</a> <a href="#">Syntax (TX Matrix Routers) on page 6889</a>
<b>Syntax</b>	<pre>show log &lt;filename   user &lt;username&gt;&gt;</pre>
<b>Syntax (QFabric System)</b>	<pre>show log filename &lt;device-type (device-id   device-alias)&gt;</pre>
<b>Syntax (TX Matrix Routers)</b>	<pre>show log &lt;all-lcc   lcc number   scc&gt; &lt;filename   user &lt;username&gt;&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <i>device-type (device-id   device-alias)</i> is introduced in Junos OS Release 13.1 for the QFX Series.</p>
<b>Description</b>	List log files, display log file contents, or display information about users who have logged in to the router or switch.
<b>Options</b>	<p><b>none</b>—List all log files.</p> <p><b>&lt;all-lcc   lcc number   scc&gt;</b>—(TX Matrix routers only) (Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace <i>number</i> with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).</p> <p><b>device-type</b>—(QFabric system only) (Optional) Display log messages for only one of the following device types:</p> <ul style="list-style-type: none"> <li>• <b>director-device</b>—Display logs for Director devices.</li> <li>• <b>infrastructure-device</b>—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).</li> <li>• <b>interconnect-device</b>—Display logs for Interconnect devices.</li> <li>• <b>node-device</b>—Display logs for Node devices.</li> </ul>



**NOTE:** If you specify the *device-type* optional parameter, you must also specify either the *device-id* or *device-alias* optional parameter.

**(device-id | device-alias)**—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

**filename**—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.



**NOTE:** The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of messages.

**user <username>**—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include *username*, display logging information about the specified user.

**Required Privilege Level** trace

**List of Sample Output** [show log on page 6890](#)  
[show log filename on page 6890](#)  
[show log filename \(QFabric System\) on page 6891](#)  
[show log user on page 6891](#)

## Sample Output

### show log

```
user@host> show log
total 57518
-rw-r--r-- 1 root bin      211663 Oct  1 19:44 dcd
-rw-r--r-- 1 root bin      999947 Oct  1 19:41 dcd.0
-rw-r--r-- 1 root bin      999994 Oct  1 17:48 dcd.1
-rw-r--r-- 1 root bin      238815 Oct  1 19:44 rpd
-rw-r--r-- 1 root bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r-- 1 root bin      1061095 Oct  1 12:13 rpd.1
-rw-r--r-- 1 root bin      1052026 Oct  1 06:08 rpd.2
-rw-r--r-- 1 root bin      1056309 Sep 30 18:21 rpd.3
-rw-r--r-- 1 root bin      1056371 Sep 30 14:36 rpd.4
-rw-r--r-- 1 root bin      1056301 Sep 30 10:50 rpd.5
-rw-r--r-- 1 root bin      1056350 Sep 30 07:04 rpd.6
-rw-r--r-- 1 root bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r-- 1 root bin        19656 Oct  1 19:37 wtmp
```

### show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT rcv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT rcv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT rcv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT rcv len 144 V9 seq 151 op chnge Type ifdev devindex 44
```

```

Oct  1 18:00:19 KRT rcv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT rcv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT rcv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

```

### show log filename (QFabric System)

```

user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_REO_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

### show log user

```

user@host> show log user
darius  mg2546                Thu Oct  1 19:37   still logged in
darius  mg2529                Thu Oct  1 19:08 - 19:36 (00:28)
darius  mg2518                Thu Oct  1 18:53 - 18:58 (00:04)
root    mg1575                Wed Sep 30 18:39 - 18:41 (00:02)
root    ttyp2      jun.site.per Wed Sep 30 18:39 - 18:41 (00:02)
alex    ttyp1      192.168.1.2   Wed Sep 30 01:03 - 01:22 (00:19)

```



## CHAPTER 79

# Troubleshooting

- [Troubleshooting Overview on page 6893](#)
- [Troubleshooting Procedures on page 6899](#)

### Troubleshooting Overview

---

- [Understanding Troubleshooting Resources on page 6893](#)
- [Troubleshooting Overview on page 6895](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions on page 6897](#)

### Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

[Table 663 on page 6893](#) provides a list of some of the troubleshooting resources.

**Table 663: Troubleshooting Resources on the QFX Series**

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	<a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 7192</a>
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	<a href="#">Chassis Status LEDs on a QFX3500 Device</a>
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	<a href="#">“Interface Alarm Messages” on page 7195</a>
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	<a href="#">“Understanding Alarms” on page 7191</a>

Table 663: Troubleshooting Resources on the QFX Series (*continued*)

Troubleshooting Resource	Description	Documentation
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6561</a></li> <li>• <a href="#">Junos OS System Log Configuration Statements on page 6616</a></li> </ul>
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the <b>traceroute monitor</b> command to locate points of failure in a network.	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 333</a></li> <li>• <a href="#">Monitoring System Properties on page 334</a></li> <li>• <a href="#">traceroute monitor</a></li> </ul>
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	<i>Junos OS Automation Library</i>
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	<i>Junos XML API Operational Developer Reference</i>
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b> , <b>set</b> , and <b>commit</b> to perform those operations.	<i>NETCONF XML Management Protocol Developer Guide</i>
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	<ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs Support on page 6530</a></li> <li>• <a href="#">SNMP Traps Support on page 6546</a></li> <li>• <a href="#">Using the Traceroute MIB for SNMP Remote Operations</a></li> </ul>
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	<a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a>
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	<i>Service Automation</i>



Table 663: Troubleshooting Resources on the QFX Series (*continued*)

Troubleshooting Resource	Description	Documentation
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	<i>Service Automation</i>
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	<a href="http://kb.juniper.net">http://kb.juniper.net</a>

## Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series product.

[Table 664 on page 6895](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 664: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	See <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 7192</a> .
	Fan tray LED is blinking amber.	See <i>Fan Tray LED on a QFX3500 Device</i> .
	Chassis status LED for the power is blinking amber.	See <i>Chassis Status LEDs on a QFX3500 Device</i> .
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. See <i>Chassis Status LEDs on a QFX3500 Device</i> .

Table 664: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	<p>Check whether the port is a valid Gigabit Ethernet port (6 through 41).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a Fibre Channel port.	<p>Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a 10-Gigabit Ethernet port.	<p>If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.</p> <p>If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a 40-Gbps QSFP+ interface.	<p>The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.
Initial device configuration	Cannot configure management Ethernet ports.	<p>Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.</p> <p><b>NOTE:</b> The management ports are on the front panel of the QFX3500 switch. They are labeled <b>C0</b> and <b>C1</b> on the front panel. In the CLI they are referred to as <b>me0</b> and <b>me1</b>.</p> <p>See “<a href="#">Configuring a QFX3500 Device as a Standalone Switch</a>” on page 175.</p>

Table 664: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Software upgrade and configuration	Failed software upgrade.	See <a href="#">“Recovering from a Failed Software Installation” on page 121.</a>
	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: <ul style="list-style-type: none"> <li>• <a href="#">Loading a Previous Configuration File on page 1252</a></li> <li>• <a href="#">Reverting to the Default Factory Configuration on page 188</a></li> <li>• <a href="#">Reverting to the Rescue Configuration on page 189</a></li> <li>• <a href="#">Performing a Recovery Installation on page 116</a></li> </ul>
	Root password is lost or forgotten.	Recover the root password. See <a href="#">“Recovering the Root Password” on page 1233.</a>
Network interfaces	An aggregated Ethernet interface is down.	See <a href="#">“Troubleshooting an Aggregated Ethernet Interface” on page 1234.</a>
	Interface on built-in network port is down.	See <a href="#">“Troubleshooting Network Interfaces” on page 1234.</a>
	Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See <a href="#">“Troubleshooting Ethernet Switching” on page 1895.</a>
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See <a href="#">“Troubleshooting Firewall Filter Configuration” on page 5411.</a>

## QFX5100 Switch with Automation Enhancements Frequently Asked Questions

This FAQ addresses questions regarding using QFX5100 switches with automation enhancements, which were introduced at Junos OS Release 13.2X51-D15.

This FAQ covers the following questions:

- [Who Should You Contact If You Have Problems with Loading, Installing or Updating Libraries? on page 6898](#)
- [Who Should You Contact If You Have Problems with Puppet for Junos OS? on page 6898](#)
- [Who Should You Contact If You Have Problems with Chef for Junos OS? on page 6898](#)
- [What Happens to the User Partition If You Downgrade a QFX5100 Switch That Is Running the jinstall-qfx-5-flex-x.tgz Software Bundle to a QFX Switch That Is Running a Different QFX5100 Software Bundle? on page 6898](#)
- [How Do You Recover Junos OS Binaries That You Have Deleted? on page 6898](#)

- [How Do You Recover from a System Crash?](#) on page 6898
- [How Can You Verify That a QFX5100 Switch Is Running a jinstall-qfx-5-flex-x.tgz Software Bundle?](#) on page 6898

---

### Who Should You Contact If You Have Problems with Loading, Installing or Updating Libraries?

---

Contact Customer Support at <http://www.juniper.net/support>.

---

### Who Should You Contact If You Have Problems with Puppet for Junos OS?

---

You can obtain support for Puppet for Junos OS through the J-Net Forum for Puppet at [http://forums.juniper.net/t5/Puppet-for-JunOS/bd-p/puppet\\_junos](http://forums.juniper.net/t5/Puppet-for-JunOS/bd-p/puppet_junos).

---

### Who Should You Contact If You Have Problems with Chef for Junos OS?

---

You can obtain support for Chef for Junos OS through the J-Net Forum for Chef at [http://forums.juniper.net/t5/Chef-for-JunOS/bd-p/chef\\_junos](http://forums.juniper.net/t5/Chef-for-JunOS/bd-p/chef_junos).

---

### What Happens to the User Partition If You Downgrade a QFX5100 Switch That Is Running the jinstall-qfx-5-flex-x.tgz Software Bundle to a QFX Switch That Is Running a Different QFX5100 Software Bundle?

---

In this case, the user partition remains intact.



**NOTE:** If you make changes to the user partition while performing a unified in-service software upgrade (unified ISSU), the changes might be lost.

---

---

### How Do You Recover Junos OS Binaries That You Have Deleted?

---

You must reinstall the software package.

---

### How Do You Recover from a System Crash?

---

You must reinstall the software package.

---

### How Can You Verify That a QFX5100 Switch Is Running a jinstall-qfx-5-flex-x.tgz Software Bundle?

---

You cannot use the **show version** command to verify that a QFX5100 switch is running the jinstall-qfx-5-flex-x.tgz software bundle. However, there are two other ways to verify this.

- Use the **show configuration** command to check that you are running a Layer 3 configuration. See *Installing Junos OS Software with QFX5100 Switch Automation Enhancements*.
- Go to the shell and confirm that you can invoke Python. See “[Invoking the Python Interpreter](#)” on page 6587.

#### Related Documentation

- [Overview of QFX5100 Switch Automation Enhancements on page 6470](#)
- [Installing Junos OS Software with QFX5100 Switch Automation Enhancements](#)

- [Invoking the Python Interpreter on page 6587](#)
- [Chef for Junos Getting Started Guide](#)
- [Puppet for Junos OS Documentation](#)

## Troubleshooting Procedures

---

- [Recovering from a Failed Software Installation on page 6899](#)
- [Loading a Previous Configuration File on page 6900](#)
- [Reverting to the Default Factory Configuration on page 6900](#)
- [Reverting to the Rescue Configuration on page 6901](#)
- [Recovering the Root Password on page 6901](#)
- [Troubleshooting a Deprecated Network Analytics Configuration on page 6903](#)

### Recovering from a Failed Software Installation

**Problem** **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution** If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [- --format] [- --external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:

- Network address of the server and the path on the server; for example, **ftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
- Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Loading a Previous Configuration File

You can use the **rollback <number>** command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

### Syntax

**rollback <number>**

### Options

- **none**— Return to the most recently saved configuration.
- **number**—Configuration to return to.
  - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
  - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

**Related Documentation**

- [Configuration File Terms on page 11](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

#### Related Documentation

- [Understanding Configuration Files on page 1242](#)
- [Loading a Previous Configuration File on page 1252](#)
- [Reverting to the Rescue Configuration on page 189](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.
 

```
[edit]
user@switch# load override filename
```
2. Commit your changes.
 

```
[edit]
user@switch# commit filename
```

#### Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1261](#)
- [Reverting to the Default Factory Configuration on page 188](#)
- [Configuration File Terms on page 11](#)

## Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



**NOTE:** The root password cannot be recovered on a QFabric system.



**NOTE:** You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:  
  
Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...
11. At the following prompt, enter **boot -s** to start up the system in single-user mode.  
  
ok **boot -s**
12. At the following prompt, enter **recovery** to start the root password recovery procedure.  
  
Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: **recovery**
13. Enter configuration mode in the CLI.
14. Set the root password. For example:



```
user@switch# set system root-authentication plain-text-password
```

15. At the following prompt, enter the new root password. For example:

```
New password: juniper1
Retype new password:
```

16. At the second prompt, reenter the new root password.

17. After you have finished configuring the password, commit the configuration.

```
root@host# commit
commit complete
```

18. Exit configuration mode in the CLI.

19. Exit operational mode in the CLI.

20. At the prompt, enter **y** to reboot the device.

```
Reboot the system? [y/n] y
```

**Related Documentation**

- [Configuring the Root Password on page 1354](#)

## Troubleshooting a Deprecated Network Analytics Configuration

**Problem** **Description:** After a software upgrade to Junos OS Release 13.2X51-D15 from an earlier release, the network analytics configuration is no longer valid and the feature is disabled.

**Symptoms:** The network analytics configuration used in Junos OS Release 13.2X51-D10 has been deprecated in Release 13.2X51-D15. Issuing the **show services analytics** command results in the following output:

```
root@qfx5100# show services analytics

queue-statistics { ## Warning: 'queue-statistics' is deprecated
    interval 1;
}
```

**Cause** Junos OS Release 13.2X51-D15 added enhancements to the network analytics feature, resulting in significant changes in the CLI. The updated **[edit services analytics]** hierarchy level contains some statements that have replaced those that were previously released. As a result, the earlier configuration does not work in the new release.

**Solution** Use the new CLI statements to reconfigure the network analytics feature.

**Related Documentation**

- [Network Analytics Overview on page 6490](#)
- [analytics on page 6667](#)



## PART 22

# Virtual Chassis

- [Overview on page 6907](#)
- [Configuration on page 6931](#)
- [Administration on page 6975](#)



## CHAPTER 80

# Overview

- [Virtual Chassis Overview on page 6907](#)

### Virtual Chassis Overview

---

- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)
- [Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric on page 6913](#)
- [Understanding How the Master in a Virtual Chassis Is Elected on page 6917](#)
- [Understanding Software Upgrades in a QFX Series Virtual Chassis on page 6918](#)
- [Understanding Global Management of a Virtual Chassis on page 6919](#)
- [Understanding Nonvolatile Storage in a Virtual Chassis on page 6921](#)
- [Understanding QFX Series Virtual Chassis Port Link Aggregation on page 6921](#)
- [Understanding Split and Merge in a Virtual Chassis on page 6922](#)
- [Understanding Automatic Software Update on Virtual Chassis Member Switches on page 6925](#)
- [Understanding MAC Address Assignment on a Virtual Chassis on page 6928](#)

### Understanding QFX Series Virtual Chassis

This topic discusses QFX Series Virtual Chassis. A QFX Series Virtual Chassis allows you to interconnect up to ten QFX3500, QFX3600, or QFX5100 switches into one logical device and manage the device as a single chassis. EX4300 switches can also be interconnected into a Virtual Chassis with QFX3500, QFX3600, and QFX5100 switches.

This topic does not discuss Virtual Chassis Fabric (VCF). For information on understanding VCF, see [“Virtual Chassis Fabric Overview” on page 7033](#).

This topic includes:

- [QFX Virtual Chassis Overview on page 6908](#)
- [QFX5100 Switches in a Virtual Chassis on page 6908](#)
- [QFX3500 and QFX3600 Switches in a Virtual Chassis on page 6909](#)
- [EX4300 Switches in a QFX Series Virtual Chassis on page 6909](#)

## QFX Virtual Chassis Overview

---

The QFX Series Virtual Chassis brings the Virtual Chassis flexible, scaling switch solution to QFX3500, QFX3600, and QFX5100 switches. EX4300 switches can also be interconnected into a Virtual Chassis with QFX3500, QFX3600, and QFX5100 switches. You can connect up to ten standalone EX4300, QFX3500, QFX3600, or QFX5100 switches into a QFX Series Virtual Chassis and manage the interconnected switches as a single chassis. The advantages of connecting multiple switches into a Virtual Chassis include better-managed bandwidth at a network layer, simplified configuration and maintenance because multiple devices can be managed as a single device, increased fault tolerance and high availability (HA) because a Virtual Chassis can remain active and network traffic can be redirected to other member switches when a single member switch fails, and a flatter, simplified Layer 2 network topology that minimizes or eliminates the need for loop prevention protocols such as Spanning Tree Protocol (STP).

You configure a QFX Series Virtual Chassis by configuring 10-Gbps SFP+ or 40-Gbps QSFP+ interfaces into Virtual Chassis ports (VCPs). VCPs connect switches together to form a Virtual Chassis, and are responsible for passing all data and control traffic between member switches in the Virtual Chassis. All non-channelized 40-Gbps QSFP+ interfaces on QFX3500, QFX3600, and QFX5100 series switches can be configured into VCPs. All fixed 10-Gbps SFP+ interfaces, including 10-Gbps SFP+ uplink interfaces on EX4300 switches, can also be configured into VCPs.

You can increase VCP bandwidth between member switches by configuring multiple interfaces between the same two switches into VCPs. When multiple VCPs are interconnecting the same two member switches, a Link Aggregation Group (LAG) bundle is automatically formed when the VCPs are on interfaces supporting identical speeds. For instance, if you have two 40-Gbps QSFP+ interfaces configured as VCPs between member switches, a LAG with two member links with 80Gbps of total bandwidth is formed. 10-Gbps SFP+ and 40-Gbps QSFP+ interfaces configured as VCPs cannot be members of the same LAG, however.

## QFX5100 Switches in a Virtual Chassis

---

Virtual Chassis is supported on all QFX5100 switches starting in Junos OS Release 13.2X51-D20.

You can interconnect up to 10 switches into a Virtual Chassis. A QFX Series Virtual Chassis can contain up to ten total member switches, and the ten total member switches can include any combination of EX4300, QFX3500, QFX3600, and QFX5100 series switches.



**NOTE:** In Junos OS release 13.2X51-D20, you can interconnect up to ten QFX5100 switches into a Virtual Chassis with the exception of the QFX5100-96S switch, which you could configure into a non-mixed Virtual Chassis that included up to four QFX5100-96S switches only.

You can configure up to ten QFX5100-96S switches into a mixed or non-mixed Virtual Chassis starting in Junos OS release 13.2X51-D25.

---

### QFX3500 and QFX3600 Switches in a Virtual Chassis

Virtual Chassis is supported on QFX3500 and QFX3600 series switches. QFX3500 and QFX3600 series switches must be configured as standalone switches; the Virtual Chassis feature is not applicable to QFX devices in a QFabric.

QFX3500 and QFX3600 devices must be running a version of Junos OS for QFX devices that support Virtual Chassis. A QFX Series Virtual Chassis can contain up to ten total member switches and the ten total member switches can include any combination of EX4300, QFX3500, QFX3600, and QFX5100 series switches.

### EX4300 Switches in a QFX Series Virtual Chassis

Virtual Chassis is supported on EX4300 switches. Starting in Junos OS Release 13.2X51-D20, EX4300 switches can be interconnected into a Virtual Chassis with QFX3500 switches, QFX3600 switches, and QFX5100 switches.

A mixed or non-mixed Virtual Chassis that includes EX4300 switches can contain up to ten total member switches, and the ten total member switches can include any combination of EX4300, QFX3500, QFX3600, and QFX5100 series switches.

#### Related Documentation

- [Understanding QFX Series Virtual Chassis Components on page 6909](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)

## Understanding QFX Series Virtual Chassis Components

This topic describes the components of a QFX Series Virtual Chassis. A QFX Series Virtual Chassis is up to ten standalone QFX3500, QFX3600, or QFX5100 switches interconnected and managed as a single chassis. EX4300 switches can also be interconnected into a Virtual Chassis with QFX3500, QFX3600, and QFX5100 switches.

This topic does not discuss Virtual Chassis Fabric components. For information on Virtual Chassis Fabric components, see [“Understanding Virtual Chassis Fabric Components” on page 7035](#).

This topic covers:

- [Virtual Chassis Ports \(VCPs\) on page 6909](#)
- [Maximum Switch Support on page 6910](#)
- [Master Role on page 6910](#)
- [Backup Role on page 6911](#)
- [Linecard Role on page 6911](#)
- [Member Switch and Member ID on page 6912](#)
- [Mastership Priority on page 6913](#)

### Virtual Chassis Ports (VCPs)

You configure a QFX Series Virtual Chassis by configuring 10-Gbps SFP+ or 40-Gbps QSFP+ interfaces into Virtual Chassis ports (VCPs). VCPs connect switches together to

form a Virtual Chassis, and are responsible for passing all data and control traffic between member switches in the Virtual Chassis. All non-channelized 40-Gbps QSFP+ interfaces on QFX3500, QFX3600, and QFX5100 series switches can be configured into VCPs; 40-Gbps QSFP+ interfaces that have been channelized into SFP+ interfaces using a breakout cable cannot be configured into VCPs. All other SFP+ interfaces on QFX series switches can be configured into VCPs, and can also be used to interconnect EX4300 switches into a mixed Virtual Chassis.

You can increase VCP bandwidth between member switches by configuring multiple interfaces between the same two switches into VCPs. When multiple VCPs are interconnecting the same two member switches, a Link Aggregation Group (LAG) bundle is automatically formed when the VCPs are on interfaces supporting identical speeds. For instance, if you have two 40-Gbps QSFP+ interfaces configured as VCPs between member switches, a LAG with two member links with 80Gbps of total bandwidth is formed. 10-Gigabit SFP+ and 40-Gbps QSFP+ interfaces configured as VCPs cannot be members of the same LAG, however. See [“Understanding QFX Series Virtual Chassis Port Link Aggregation” on page 6921](#)

---

### Maximum Switch Support

You can interconnect up to 10 switches into a Virtual Chassis. The Virtual Chassis can contain up to ten total member switches and the ten total member switches can include any combination of EX4300, QFX3500, QFX3600, and QFX5100 series switches.



**NOTE:** In Junos OS release 13.2X51-D20, you can interconnect up to ten QFX5100 switches into a Virtual Chassis with the exception of the QFX5100-96S switch, which you could configure into a non-mixed Virtual Chassis that included up to four QFX5100-96S switches only.

You can configure up to ten QFX5100-96S switches into a mixed or non-mixed Virtual Chassis starting in Junos OS release 13.2X51-D25.

---

### Master Role

In a Virtual Chassis, each member switch is assigned one of three roles: master, backup, or linecard.

The member that functions in the master role in the Virtual Chassis:

- Manages the member switches.
- Runs Junos OS for the switches in a master role.
- Runs the chassis management processes and control protocols.
- Represents all the member switches interconnected within the Virtual Chassis configuration. (The hostname and other properties that you assign to this switch during setup apply to all members of the Virtual Chassis configuration.)

In a Virtual Chassis, one member functions as the master and a second member functions as the backup:



- In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions as the master member. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).
- In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

All switches that are not assigned the master or backup role function in the linecard role.

In a mixed Virtual Chassis, we recommend configuring the QFX5100 switches into the master and backup role. If the mixed Virtual Chassis does not contain QFX5100 switches, we recommend configuring QFX3500 or QFX3600 switches into the master and backup roles. You should only configure EX4300 switches into the master or backup role when your mixed Virtual Chassis contains one QFX series switch.

---

### Backup Role

The member that functions in the backup role in the Virtual Chassis:

- Maintains a state of readiness to take over the master role if the master fails.
- Runs Junos OS for switches in a backup role.
- Synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

You must have at least two member switches in the Virtual Chassis configuration in order to have a backup member.

- In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions in the backup role. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).
- In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

In a mixed Virtual Chassis, we recommend configuring the QFX5100 switches into the master and backup role. If the mixed Virtual Chassis does not contain QFX5100 switches, we recommend configuring QFX3500 or QFX3600 switches into the master and backup roles. You should only configure EX4300 switches into the master or backup role when your mixed Virtual Chassis contains one QFX series switch.

---

### Linecard Role

A member that functions in the linecard role in the Virtual Chassis:

- Runs only a subset of Junos OS.

- Does not run the chassis control protocols.
- Can detect certain error conditions (such as an unplugged cable) on any interfaces that have been configured on it through the master.

The Virtual Chassis configuration must have at least three members in order to include a linecard member.

- In a preprovisioned configuration, you can explicitly configure a member with the linecard role, which makes it ineligible for functioning as a master or backup.
- In a configuration that is not preprovisioned, the members that are not selected as master or backup function as linecard members of the Virtual Chassis configuration. The selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm. A switch with a mastership priority of 0 is always in the linecard role.

Any switch can function in the linecard role in a mixed or non-mixed Virtual Chassis.

In a mixed Virtual Chassis, we recommend configuring the QFX5100 switches into the master and backup role. If the mixed Virtual Chassis does not contain QFX5100 switches, we recommend configuring QFX3500 or QFX3600 switches into the master and backup roles. You should only configure EX4300 switches into the master or backup role when your mixed Virtual Chassis contains one QFX series switch.

---

### Member Switch and Member ID

Each standalone switch that supports Virtual Chassis is a potential member of a Virtual Chassis configuration. When one of those switches is powered on, it receives a member ID that can be seen by viewing the front-panel LCD or by entering the **show virtual-chassis** command. If the switch is powered on as a standalone switch, that member's member ID is always 0. When the switch is interconnected with other switches in a Virtual Chassis configuration, its member ID is assigned by the master based on various factors, such as the order in which the switch was added to the Virtual Chassis configuration or the member ID assigned by a preprovisioned configuration. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).

If the Virtual Chassis configuration previously included a member switch and that member was physically disconnected or removed from the Virtual Chassis configuration, its member ID is not available for assignment as part of the standard sequential assignment by the master. For example, you might have a Virtual Chassis configuration composed of member 0, member 2, and member 3, because member 1 was removed. When you add another member switch and power it on, the master assigns it as member 4.

The member ID distinguishes the member switches from one another. You use the member ID:

- To assign a mastership priority value to a member switch
- To configure interfaces for a member switch (The function is similar to that of a slot number on Juniper Networks routers.)

- To apply some operational commands to a member switch
- To display status or characteristics of a member switch

### Mastership Priority

In a configuration that is not preprovisioned, you can designate the role (master, backup, or linecard) that a member switch assumes by configuring its mastership priority (from **0** through **255**). The mastership priority value is the factor in the master election algorithm with the highest precedence for selecting the master of the Virtual Chassis configuration. A switch with a mastership priority of **0** never assumes the backup or master role.

The default value for mastership priority is **128**. When a standalone switch is powered on, it receives the default mastership priority value. Because it is the only member of the Virtual Chassis configuration, it is also the master. When you interconnect a standalone switch to an existing Virtual Chassis configuration (which implicitly includes its own master), we recommend that you explicitly configure the mastership priority of the members that you want to function as the master and backup.

In a preprovisioned configuration, you assign the role of each member switch.

#### Related Documentation

- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)

## Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric

This topic describes the requirements for a mixed Virtual Chassis or a mixed Virtual Chassis Fabric (VCF).

A mixed Virtual Chassis includes two or more types of EX Series switches, two or more types of QFX Series switches, or a mix of EX and QFX Series switches.

A mixed VCF is any VCF that includes two or more types of member switches. Because a VCF must use a QFX5100 switch as a spine device, a mixed VCF is any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switches.



**NOTE:** The optimal VCF topology is to use QFX5100 devices only. A VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces.

This topic covers:

- [Virtual Chassis Fabric Summary on page 6914](#)
- [Understanding Mixed Virtual Chassis Fabric on page 6914](#)
- [Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches on page 6915](#)

- [Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches on page 6915](#)
- [Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis on page 6916](#)
- [Understanding Mixed EX4300 and EX4600 Virtual Chassis on page 6916](#)
- [Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis on page 6916](#)

### Virtual Chassis Fabric Summary

[Table 665 on page 6914](#) provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed VCF. The table also includes license requirements and supported configuration methods.

**Table 665: Virtual Chassis Fabric Summary**

Category	Allowed Routing Engines	Allowed Line Cards	License Requirement	Configuration Methods
Non-mixed	QFX5100	QFX5100	Yes (on two QFX5100 switches operating in master and backup Routing Engine roles)	Autoprovisioning Preprovisioning Nonprovisioning (not recommended)
Mixed	QFX5100	QFX5100 QFX3600 QFX3500 EX4300	Yes (on two QFX5100 switches operating in master and backup Routing Engine roles)	Autoprovisioning Preprovisioning Nonprovisioning (not recommended)

### Understanding Mixed Virtual Chassis Fabric

A VCF must use a QFX5100 switch in the spine role. A mixed VCF is, therefore, any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switch.

The optimal method of configuring a VCF is to use QFX5100 devices only. A non-mixed VCF composed entirely of QFX5100 devices supports the largest breadth of features at

the highest scalability while also supporting the highest number of high-speed interfaces. You can, however, also configure a mixed VCF.

If you use QFX3600, QFX3500, or EX4300 devices as leaf devices in your VCF, you must configure all devices in your VCF into mixed mode. If you are turning a non-mixed VCF into a mixed VCF, you have to reboot the VCF to change the mixed mode setting.

### Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches

Table 666 on page 6915 provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed Virtual Chassis for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches. The table also includes license requirements and supported configuration methods.

**Table 666: Virtual Chassis Summary**

Category	Allowed Routing Engines	Allowed Line Cards	License Requirement	Configuration Methods
Non-mixed	QFX5100	QFX5100	No	Nonprovisioning Preprovisioning
	QFX3600 QFX3500	QFX3600 QFX3500	No	Nonprovisioning Preprovisioning
	EX4600	EX4600	No	Nonprovisioning Preprovisioning
	EX4300	EX4300	No	Nonprovisioning Preprovisioning
Mixed	QFX5100	QFX5100 QFX3600 QFX3500 EX4300	No	Nonprovisioning Preprovisioning
	QFX3600 QFX3500	QFX3600 QFX3500 EX4300	No	Nonprovisioning Preprovisioning
	EX4600	EX4600 EX4300	No	Nonprovisioning Preprovisioning

### Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches

In a mixed Virtual Chassis, the switch in the master Routing Engine role determines which switches are supported in the line card role of the mixed Virtual Chassis.

When a mixed Virtual Chassis is using a QFX5100 switch in the master Routing Engine role, you can use QFX5100, QFX3600, QFX3500, or EX4300 switches in the line card role.

When a mixed Virtual Chassis is using a QFX3600 or QFX3500 switch in the master Routing Engine role, you can use QFX3600, QFX3500, or EX4300 switches in the line card role.

In a mixed EX4300 and EX4600 Virtual Chassis, an EX4600 switch automatically assumes the Routing Engine role.

EX4600 switches can only be in a mixed Virtual Chassis with EX4300 switches. EX4600 switches cannot be in a mixed Virtual Chassis with QFX5100, QFX3600, or QFX3500 switches.

We recommend always configuring the same type of switch into the master and backup Routing Engine role, to ensure that the switch operating in the master role remains the same type of switch in the event of a switchover.

In most mixed Virtual Chassis, you must configure your Virtual Chassis to ensure a switch that supports the master Routing Engine assumes the master Routing Engine role. Without user configuration, any switch—with the exception of the EX4300 switch, which can never assume the master or backup Routing Engine role in a mixed Virtual Chassis or VCF—can assume the master or backup Routing Engine role.

### **Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis**

---

Up to ten EX4300 switches, QFX3500 switches, QFX3600 switches, and QFX5100 switches can be interconnected using Virtual Chassis ports (VCPs) to form a mixed or non-mixed Virtual Chassis. The mixed Virtual Chassis supports up to ten member switches regardless of the switches that compose the mixed Virtual Chassis.

EX4300 switches can also be interconnected into a mixed Virtual Chassis with EX4600 switches. See the following section for information on mixed EX4300 and EX4600 Virtual Chassis.

### **Understanding Mixed EX4300 and EX4600 Virtual Chassis**

---

EX4300 switches and EX4600 switches can be interconnected into the same Virtual Chassis. An EX4600 switch automatically assumes the master Routing Engine role in a mixed EX4300 and EX4600 Virtual Chassis, since EX4300 switches cannot assume the Routing Engine role in a mixed Virtual Chassis. EX4600 switches cannot be in a mixed Virtual Chassis with any other type of switch.

The mixed Virtual Chassis supports up to ten member switches.

### **Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis**

---

EX4200 switches, EX4500 switches, and EX4550 switches can be interconnected into the same Virtual Chassis to form a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4500 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, and EX4550 Virtual Chassis. The mixed Virtual Chassis supports up to 10 member switches regardless of whether the switches are EX4200 switches, EX4500 switches, or EX4550 switches. Any model of EX4200, EX4500, or EX4550 switch can be interconnected into the same mixed Virtual Chassis. The master election process that decides member switch roles in a mixed Virtual Chassis is identical to the

master election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the master, backup, or linecard role.

EX4200 switches, EX4500 switches, and EX4550 switches cannot be interconnected into a Virtual Chassis with any other switches.

#### Related Documentation

- [Virtual Chassis Fabric Overview on page 7033](#)
- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)
- [Understanding EX Series Virtual Chassis Components](#)

## Understanding How the Master in a Virtual Chassis Is Elected



**NOTE:** This topic does not apply to EX8200 Virtual Chassis. See *EX8200 Virtual Chassis Overview*.

All switches that are interconnected in a Virtual Chassis configuration are member switches of that Virtual Chassis. Each Virtual Chassis configuration has one member that functions as the *master* and controls the Virtual Chassis configuration.

When a Virtual Chassis configuration boots, the Juniper Networks Junos operating system (Junos OS) on the switches automatically runs a master election algorithm to determine which member switch assumes the role of master.

The algorithm proceeds from the top condition downward until the stated condition is satisfied:

1. Choose the member with the highest user-configured mastership priority (255 is the highest possible value). A switch with a mastership priority of 0 will always stay in the linecard role.
2. Choose the member that was master the last time the Virtual Chassis configuration booted.
3. Choose the member that has been included in the Virtual Chassis configuration for the longest period of time. (For this to be a deciding factor, there has to be a minimum time lapse of 1 minute between the power-ons of the individual interconnected member switches.)
4. Choose the member with the lowest MAC address.

The variations among switches and switch models do not impact the master election algorithm.

To ensure that a specific member is elected as the master:

1. Power on only the switch that you want to configure as master of the Virtual Chassis configuration.
2. Configure the mastership priority of that member to have the highest possible value (255).
3. Continue to configure other members through the master member.
4. Power on the other members.

You can also specify the switch roles by preprovisioning your Virtual Chassis. Preprovisioning a Virtual Chassis allows you to manually assign the member ID and role for each switch in the Virtual Chassis. See *Configuring an EX3300 Virtual Chassis (CLI Procedure)*, *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*, or [“Configuring a QFX Series Virtual Chassis \(CLI Procedure\)” on page 6931](#).

**Related  
Documentation**

- *EX8200 Virtual Chassis Overview*
- *EX Series Virtual Chassis Overview*
- [Understanding QFX Series Virtual Chassis on page 6907](#)
- *Understanding EX Series Virtual Chassis Components*
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)
- *Understanding EX Series Virtual Chassis Configuration*

## Understanding Software Upgrades in a QFX Series Virtual Chassis

This topic discusses software upgrades on a QFX Series Virtual Chassis. For information on software upgrades on a Virtual Chassis Fabric (VCF), see [“Understanding Software Upgrades in a Virtual Chassis Fabric” on page 7050](#)

In a Virtual Chassis, each member switch must be running the same version of Juniper Networks Junos operating system (Junos OS) that supports Virtual Chassis.

You can install a new Junos OS release on the entire Virtual Chassis or on a particular member in the Virtual Chassis by using the same CLI command that you use to install Junos OS on standalone switches—the [request system software add](#) command.

You can use the automatic software update feature to automatically update the Junos OS version on member switches as you add them to a Virtual Chassis. See [“Understanding Automatic Software Update on Virtual Chassis Member Switches” on page 6925](#). If you are not configuring the automatic software update feature, we recommend that you update the new member switch to the version of Junos OS running on the Virtual Chassis before adding the member switch to the Virtual Chassis.

**Related  
Documentation**

- [Understanding QFX Series Virtual Chassis Components on page 6909](#)



## Understanding Global Management of a Virtual Chassis




**NOTE:** This topic does not apply to EX8200 Virtual Chassis. See *Understanding Global Management of an EX8200 Virtual Chassis*.

A Virtual Chassis is composed of multiple switches, and it, therefore, has multiple console ports and multiple out-of-band management Ethernet ports located on the switches.

You can connect a PC or laptop directly to a console port of any member switch to set up and configure the Virtual Chassis. When you connect to the console port of any member switch, the console session is redirected to the master switch, as shown in [Figure 228 on page 6919](#).

**Figure 228: Console Session Redirection (EX4200 Virtual Chassis Pictured)**




If the master becomes unavailable, the console session is disconnected from the old master and a new session is established with the newly elected master.

An out-of-band management Ethernet port is often referred to simply as a management Ethernet port. It uses a dedicated management channel for device maintenance and allows a system administrator to monitor and manage the switch by remote control.

The Virtual Chassis configuration can be managed remotely through SSH or Telnet using a global management interface called the virtual management Ethernet (VME) interface. The VME interface is a logical interface representing all of the out-of-band management ports on the member switches. When you connect to the Virtual Chassis configuration using the VME interface's IP address, the connection is redirected to the master member as shown in [Figure 229 on page 6920](#).

**Figure 229: Management Ethernet Port Redirection to the VME Interface**



If the master management Ethernet link is unavailable, the session is redirected through the backup management Ethernet link. If there is no active management Ethernet link on the backup, the VME interface chooses a management Ethernet link on one of the linecard members, selecting the linecard member with the lowest member ID as its first choice.

You can configure an IP address for the VME global management interface at any time.

You can perform remote configuration and administration of all members of the Virtual Chassis configuration through the VME interface.

#### Related Documentation

- [Understanding Global Management of an EX8200 Virtual Chassis](#)
- [Understanding EX Series Virtual Chassis Components](#)
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)
- [Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)

## Understanding Nonvolatile Storage in a Virtual Chassis



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See [Understanding File Storage in an EX8200 Virtual Chassis](#) for information about EX8200 Virtual Chassis.

The EX Series or QFX Series switches store the Juniper Networks Junos operating system (Junos OS) system files in internal flash memory. In the Virtual Chassis configurations, both the master and the backup switch store the configuration information for all the member switches.

- [Nonvolatile Memory Features on page 6921](#)

### Nonvolatile Memory Features

Junos OS optimizes the way the Virtual Chassis stores its configuration if a member switch or the Virtual Chassis configuration is shut down improperly:

- If the master is not available, the backup switch takes on the role of the master and its internal flash memory takes over as the alternate location for maintaining nonvolatile configuration memory.
- If a member switch is taken offline for repair, the master stores the configuration of the member switch.

#### Related Documentation

- [Understanding File Storage in an EX8200 Virtual Chassis](#)
- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [Command Forwarding Usage with an EX Series Virtual Chassis](#)

## Understanding QFX Series Virtual Chassis Port Link Aggregation

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a logical point-to-point link, known as a *link*

*aggregation group (LAG) or bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links fails, the system automatically load-balances traffic across all remaining links.

Similarly, if a Virtual Chassis member switch that has LAG member interfaces on multiple member switches fails for any reason, the traffic traversing the LAG can be redirected through the active member switch. This setup has benefits for failover purposes and can be especially beneficial in cases when a member switch needs to be inactive for some time.

You can configure any optical uplink port that can be used to connect QFX devices configured as standalone switches together into a Virtual Chassis port (VCP). You can configure multiple optical uplink interfaces between two member switches in the same Virtual Chassis as VCPs. If you have configured two or more optical ports as VCPs connecting the same member switches, the optical uplink ports configured as VCPs automatically form a LAG provided the optical uplink ports are configured to operate at the same link speeds. Each LAG is assigned a positive-integer identifier called a *trunk ID*.

A LAG over uplink VCPs provides higher overall bandwidth for forwarding traffic between the member switches connected by the optical VCPs, faster management communications, and greater redundancy of operations among the members than would be available without the LAG. A LAG over optical VCPs provides an additional Virtual Chassis link throughput for the switches.



**NOTE:** The interfaces that are included within a bundle or LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.

---

**Related  
Documentation**

- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)

## Understanding Split and Merge in a Virtual Chassis

In a Virtual Chassis, two or more switches are connected together to form a unit that is managed as a single chassis. If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis. This situation could cause disruptions in the network if the two separate configurations share common resources, such as global IP addresses. The split and merge feature provides a method to prevent the separate Virtual Chassis configurations from adversely affecting the network and also allows the two parts to merge back into a single Virtual Chassis configuration.



**NOTE:** If a Virtual Chassis configuration splits into separate parts, we recommend that you resolve the problem that caused the Virtual Chassis configuration to split as soon as possible.

You can also use this feature to merge two active but separate Virtual Chassis that have not previously been part of the same configuration into one Virtual Chassis configuration.



**NOTE:** The split and merge feature is enabled by default on EX Series and QFX Series Virtual Chassis. You can disable the split and merge feature by using the `set virtual-chassis no-split-detection` command.

This topic describes:

- [What Happens When a Virtual Chassis Configuration Splits on page 6923](#)
- [Merging Virtual Chassis Configurations on page 6924](#)

### What Happens When a Virtual Chassis Configuration Splits

When a Virtual Chassis configuration splits into two separate Virtual Chassis configurations, the individual member switches detect this topology change and run the master election algorithm to select a new master for each of the two Virtual Chassis configurations. The new masters then determine whether their Virtual Chassis configuration remains active. One of the configurations remains active based on the following:

- It contains both the stable master and the stable backup (that is, the master and backup from the original Virtual Chassis configuration before the split).
- It contains the stable master and the configuration is greater than half the Virtual Chassis size.
- It contains the stable backup and is at least half the Virtual Chassis size.

In accordance with the rules given in the second and third list items, if the Virtual Chassis configuration splits into two equal parts and the stable master and stable backup are in different parts, then the part that contains the stable backup becomes active.



**NOTE:** The number of members in the Virtual Chassis configuration includes all member switches connected to date minus the number whose Virtual Chassis member IDs have been recycled (that is, made available for reassignment). Therefore, the size of the Virtual Chassis configuration increases when a new member switch is detected and decreases when a member switch's ID is recycled.

These rules ensure that only one of the two separate Virtual Chassis configurations created by the split remains active. The member switches in the inactive Virtual Chassis

configuration remain in a linecard role. For the inactive members to become active again, one of the following things must happen:

- The problem that caused the original Virtual Chassis configuration to split is resolved, allowing the two Virtual Chassis configurations to merge.
- You load the factory default configuration on the inactive members, which causes the inactive members to function as standalone switches or become part of a different Virtual Chassis configuration.



**NOTE:** When you remove a member switch from a Virtual Chassis configuration, we recommend that you recycle the member ID using the `request virtual-chassis recycle` command.

---

### Merging Virtual Chassis Configurations

---

There are two scenarios in which separate Virtual Chassis merge:

- A Virtual Chassis configuration that had split into two is now merging back into a single configuration because the problem that had caused it to split has been resolved.
- You want to merge two Virtual Chassis that had not previously been configured together.

Every Virtual Chassis configuration has a unique ID (VCID) that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a VCID using the `set virtual-chassis id` command. A VCID that you assign takes precedence over automatically assigned VCIDs.

When you reconnect the separate Virtual Chassis configurations or connect them for the first time, the members determine whether or not the separate Virtual Chassis configurations can merge. The members use the following rules to determine whether a merge is possible:

- If the Virtual Chassis configurations have the same VCID, then the configurations can merge. If the two Virtual Chassis were formed as the result of a split, they have the same VCID.
- If the VCIDs are different, then the two configurations can merge only if both are active (inactive configurations cannot merge, ensuring that members removed from one Virtual Chassis configuration do not become members of another Virtual Chassis configuration). If the configurations to merge are both active and one of them has a user-configured VCID, this ID becomes the ID of the merged Virtual Chassis. If neither Virtual Chassis has a user-configured VCID, then the VCID of the configuration with the highest mastership priority becomes the ID of the merged Virtual Chassis. The resulting merged Virtual Chassis configuration is active.

When you connect two Virtual Chassis configurations, the following events occur:

1. Connecting the two split Virtual Chassis configurations triggers the shortest-path-first (SPF) algorithm. The SPF algorithm computes the network topology and then triggers the master election algorithm. The master election algorithm waits for the members to synchronize the topology information before running.
2. The master election algorithm merges the VCIDs of all the members.
3. Each member runs the master election algorithm to select a master and a backup from among all members with the same VCIDs. For more information, see [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).
4. The master determines whether the Virtual Chassis configuration is active or inactive. (See [“What Happens When a Virtual Chassis Configuration Splits” on page 6923](#).)
5. If the Virtual Chassis configuration is active, the master assigns roles to all members. If the Virtual Chassis configuration is inactive, the master assigns all members the role of linecard.
6. When the other members receive their role from the master, they change their role to backup or linecard. They also use the active or inactive state information sent by the master to set their own state to active or inactive and to construct the Virtual Chassis member list from the information sent by the master.
7. If the Virtual Chassis state is active, the master waits for messages from the members indicating that they have changed their roles to the assigned roles, and then the master changes its own role to master.



**NOTE:** When you merge two Virtual Chassis that had not previously been part of the same Virtual Chassis configuration, any configuration settings (such as the settings for Telnet and FTP services, graceful Routing Engine switchover (GRES), fast failover, VLANs, and so on) that exist on the new master become the configuration settings for all members of the new Virtual Chassis, overwriting any other configuration settings.

#### Related Documentation

- [Disabling Split and Merge in a Virtual Chassis \(CLI Procedure\) on page 6944](#)
- [Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge \(CLI Procedure\) on page 6946](#)
- [Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge](#)
- [Understanding EX Series Virtual Chassis Configuration](#)
- [Understanding QFX Series Virtual Chassis on page 6907](#)

## Understanding Automatic Software Update on Virtual Chassis Member Switches

You can use the automatic software update feature to automatically update the Juniper Networks Junos operating system (Junos OS) version on prospective member switches as you add them to an EX Series or QFX Series Virtual Chassis.

This topic includes:

- [Automatic Software Update Basics on page 6926](#)
- [Automatic Software Update Restrictions on page 6926](#)

### [Automatic Software Update Basics](#)

---

When you have configured automatic software update on a Virtual Chassis, the Junos OS version is updated on the new member switch when you add it to the Virtual Chassis. The new member switch immediately joins the Virtual Chassis configuration and is put in the active state.

For a standalone switch to join an existing Virtual Chassis, it must be running the same version of Junos OS that is running on the Virtual Chassis master. When the master in a Virtual Chassis detects that a new switch has been added to the configuration, it checks the software version on the new switch. If the software version on the new switch is not the same as the version running on the master, the master keeps the new switch in the inactive state. If you have not enabled the automatic software update feature, you have to manually install the correct software version on each prospective member switch as it is added to the Virtual Chassis.

### [Automatic Software Update Restrictions](#)

---

You cannot use automatic software update in certain scenarios, and you must ensure that the software release version on the Virtual Chassis is supported by the release on the prospective member switch.

You cannot use the automatic software update feature to update software for a prospective member switch in the following scenarios:

- The Virtual Chassis was preprovisioned and is running Junos OS Release 10.4R2 or earlier.
- You configured the **mastership-priority** command to manually configure the mastership priority of at least one Virtual Chassis member switch and the Virtual Chassis was running Junos OS Release 10.4R2 or earlier when you committed this configuration.
- The Junos OS versions on the Virtual Chassis and the prospective member switch are different versions of the same major Junos OS release. For instance, if a Virtual Chassis is running Junos OS Release 10.4R1, the prospective member switch cannot be updated using automatic software update if it is running Junos OS Release 10.4R2, 10.4R3, or any other Junos OS Release 10.4 release version.

The automatic software update feature also has a Junos OS release dependency between the release that is already running on the Virtual Chassis and the release that is running on the prospective member switch.

[Table 667 on page 6927](#) summarizes automatic software update support for each Junos OS release combination.



Table 667: Automatic Software Update Support

Virtual Chassis Junos OS Release	Supported Junos OS Releases for Prospective Member Switches
All versions of Junos OS 9.0 through 9.6	All versions of Junos OS 9.0 through 9.6 Junos OS Releases 10.0R1 through 10.0R4 All versions of Junos OS Release 10.1 Junos OS Releases 10.2R1 through 10.2R3 Junos OS Releases 10.3R1 through 10.3R3
Junos OS Releases 10.0R1 through 10.0R4	All versions of Junos OS 9.0 through 9.6 All versions of Junos OS Release 10.1 Junos OS Releases 10.2R1 through 10.2R3 Junos OS Releases 10.3R1 through 10.3R3
Junos OS Release 10.0R5 and later 10.0 releases	Junos OS Release 10.2R4 and later 10.2 releases Junos OS Release 10.3R4 and later 10.3 releases All versions of Junos OS Release 10.4 All versions of Junos OS Release 11.1
All versions of Junos OS Release 10.1	All versions of Junos OS 9.0 through 9.6 Junos OS Releases 10.0R1 through 10.0R4 Junos OS Releases 10.2R1 through 10.2R3 Junos OS Releases 10.3R1 through 10.3R3
Junos OS Releases 10.2R1 through 10.2R3	All versions of Junos OS 9.0 through 9.6 Junos OS Releases 10.0R1 through 10.0R4 All versions of Junos OS Release 10.1 Junos OS Releases 10.3R1 through 10.3R3
Junos OS Release 10.2R4 and later 10.2 releases	Junos OS Release 10.0R5 Junos OS Release 10.3R4 and later 10.3 releases All versions of Junos OS Release 10.4 All versions of Junos OS Release 11.1
Junos OS Releases 10.3R1 through 10.3R3	All versions of Junos OS 9.0 through 9.6 Junos OS Releases 10.0R1 through 10.0R4 All versions of Junos OS Release 10.1 Junos OS Releases 10.2R1 through 10.2R3
Junos OS Release 10.3R4 and later 10.3 releases	Junos OS Release 10.0R5 All versions of Junos OS Release 10.4 All versions of Junos OS Release 11.1
Junos OS Releases 10.4R1 through 10.4R3	All versions of Junos OS 9.0 through 9.6 Junos OS Releases 10.0R1 through 10.0R4 All versions of Junos OS Release 10.1 Junos OS Releases 10.2R1 through 10.2R3 Junos OS Releases 10.3R1 through 10.3R3
Junos OS Release 10.4R4 and later 10.4 releases	Junos OS Release 10.0R5 Junos OS Release 10.2R4 and later 10.2 releases Junos OS Release 10.3R4 and later 10.3 releases All versions of Junos OS Release 11.1

Table 667: Automatic Software Update Support (*continued*)

Virtual Chassis Junos OS Release	Supported Junos OS Releases for Prospective Member Switches
Junos OS Release 11.1R1	All versions of Junos OS Release 10.4 Junos OS Release 11.2 and later Junos OS releases
Junos OS Release 11.1R2 and later Junos OS releases	Junos OS Release 10.0R5 Junos OS Release 10.2R4 and later 10.2 releases Junos OS Release 10.3R4 and later 10.3 releases Junos OS Release 11.2 and later Junos OS releases

**Related Documentation**

- [Understanding Software Upgrade in an EX Series Virtual Chassis](#)
- [Understanding Software Upgrades in a QFX Series Virtual Chassis on page 6918](#)
- [Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches](#)
- [Configuring Automatic Software Update on Virtual Chassis Member Switches \(CLI Procedure\) on page 6944](#)

**Understanding MAC Address Assignment on a Virtual Chassis**

In a Virtual Chassis, multiple switches—each with its own set of interfaces with unique MAC addresses—are connected together to form one chassis that can be managed as a single switch. The MAC address assigned to each network-facing interface on the switch changes when the switch joins a Virtual Chassis. Because all Layer 2 traffic decisions are based on an interface's MAC address, understanding MAC address assignment is important to understanding how network traffic is forwarded and received by the Virtual Chassis. For additional information about how a network uses MAC addresses to forward and receive traffic, see *Understanding Bridging and VLANs on EX Series Switches*.

When a Virtual Chassis is formed, the MAC address of the switch in the master role becomes the system MAC base address. The Virtual Chassis assigns the system MAC base address as the MAC address for all Layer 3 interfaces within the Virtual Chassis. The Virtual Chassis also assigns the system MAC base address to the virtual management Ethernet (VME) interface and to all of the virtual LANs (VLANs) in the Virtual Chassis.

The system MAC base address does not change in the event of a switchover if the switch that was originally configured in the master role remains a member of the Virtual Chassis. If the switch that was originally configured in the master role is removed from the Virtual Chassis, the MAC address of the current member switch in the master role is assigned as the system MAC base address after the MAC persistence timer interval has expired. You can configure the MAC persistence timer interval.

For Layer 2 and aggregated Ethernet interfaces, the Virtual Chassis assigns a unique MAC address that is derived from the member switch MAC address to each interface. The assignment of a unique MAC address to each network interface helps ensure that functions that require MAC address differentiation—such as redundant trunk groups

(RTGs), Link Aggregation Control Protocol (LACP), and general monitoring functions—can function properly.



**NOTE:** Unique MAC address assignment for Layer 2 and aggregated Ethernet interfaces in a Virtual Chassis was introduced in Junos OS Release 11.3. The same MAC address could be assigned to interfaces on different member switches in the same Virtual Chassis prior to this release.

If you reconfigure a Layer 2 interface into a Layer 3 interface, or the reverse, within a Virtual Chassis, the MAC address of that interface changes accordingly.

MAC addresses are assigned to interfaces in a Virtual Chassis automatically—no user configuration is possible or required. You can view the MAC addresses that are assigned to the interfaces by using the **show interfaces** command.

**Related  
Documentation**

- *Understanding MAC Address Assignment in an EX Series Switch*
- [Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis \(CLI Procedure\) on page 6943](#)
- *EX Series Virtual Chassis Overview*
- *EX8200 Virtual Chassis Overview*
- [Understanding QFX Series Virtual Chassis on page 6907](#)



## CHAPTER 81

# Configuration

- [Configuration Tasks on page 6931](#)
- [Configuration Statements on page 6946](#)

### Configuration Tasks

---

- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Adding a New Switch to an Existing QFX Series Virtual Chassis \(CLI Procedure\) on page 6936](#)
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 6938](#)
- [Configuring Mastership of a Virtual Chassis \(CLI Procedure\) on page 6941](#)
- [Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis \(CLI Procedure\) on page 6943](#)
- [Disabling Split and Merge in a Virtual Chassis \(CLI Procedure\) on page 6944](#)
- [Configuring Automatic Software Update on Virtual Chassis Member Switches \(CLI Procedure\) on page 6944](#)
- [Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge \(CLI Procedure\) on page 6946](#)

### Configuring a QFX Series Virtual Chassis (CLI Procedure)

This topic discusses configuring a QFX Series Virtual Chassis only. It does not apply to configuring a Virtual Chassis Fabric (VCF). For information on configuring a VCF, see [“Understanding Virtual Chassis Fabric Configuration” on page 7043](#).

You configure a QFX Series Virtual Chassis—a Virtual Chassis composed of QFX3500 standalone switches, QFX3600 standalone switches, QFX5100 standalone switches, or a combination of the three—by configuring 10-Gbps SFP+ or 40-Gbps QSFP+ interfaces connecting the member switches into Virtual Chassis ports (VCPs). EX4300 switches can also be included in a Virtual Chassis with QFX3500, QFX3600, and QFX5100 switches. All non-channelized QSFP+ uplink interfaces on standalone QFX series switches can be configured into VCPs. All fixed SFP+ and QSFP+ interfaces on can also be configured into VCPs.

A Virtual Chassis can only be configured on a QFX series device configured in standalone mode. The Junos OS image that supports Virtual Chassis—the Junos OS images that support Virtual Chassis include the “jinstall-qfx-3-” text in the filename when the Junos OS image is downloaded from the Software Center—for standalone switches must be downloaded on QFX3500 or QFX3600 switches in order for the member switches to be part of a Virtual Chassis. For QFX5100 switches, you must download the software image for the standalone switch.

A QFX Series Virtual Chassis can be configured with either:

- A preprovisioned configuration—You can deterministically control the member ID and role assigned to a member switch by tying it to its serial number.
- A nonprovisioned configuration—The master sequentially assigns a member ID to other member switches. The role is determined by the mastership priority value and other factors in the master election algorithm.



.....  
**NOTE:** A Virtual Chassis configuration has two Routing Engines—the master switch and the backup switch. Therefore, we recommend that you always use `commit synchronize` rather than simply `commit` to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved to both switches acting as Routing Engines.  
.....



.....  
**NOTE:** In Junos OS release 13.2X51-D20, you can interconnect up to ten QFX5100 switches into a Virtual Chassis with the exception of the QFX5100-96S switch, which you could configure into a non-mixed Virtual Chassis that included up to four QFX5100-96S switches only.  
.....

You can configure up to ten QFX5100-96S switches into a mixed or non-mixed Virtual Chassis starting in Junos OS release 13.2X51-D25.  
.....

Be sure that all switches that are interconnected into a Virtual Chassis are running the same version of Junos OS. See [“Upgrading Software” on page 134](#).

This topic includes:

- [Configuring a QFX Series Virtual Chassis with a Preprovisioned Configuration File on page 6932](#)
- [Configuring a QFX Series Virtual Chassis with a Nonprovisioned Configuration File on page 6934](#)

### [Configuring a QFX Series Virtual Chassis with a Preprovisioned Configuration File](#)

Preprovisioning a Virtual Chassis configuration allows you to assign the member ID and role for each switch in the Virtual Chassis.

To configure a Virtual Chassis using a preprovisioned configuration:



**NOTE:** You can configure a QFX Series Virtual Chassis while the cables are or are not physically connected.

1. Make a list of the serial numbers of all the switches to be connected in a Virtual Chassis configuration.
2. Note the desired role (**routing-engine** or **line-card**) of each switch. If you configure the member with a **routing-engine** role, it is eligible to function in the master or backup role. If you configure the member with a **line-card** role, it is not eligible to function in the master or backup role.
3. Power on only the switch that you plan to use as the master switch.
4. Specify the identification parameters for the switch by completing the initial configuration. See [“Configuring a QFX3500 Device as a Standalone Switch” on page 175](#) or *Configuring a QFX3600 Device as a Standalone Switch*, or *Configuring a QFX5100 Device*.



**NOTE:** The properties that you specify for the master switch apply to the entire Virtual Chassis configuration.

5. (Optional) Configure the master switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:  

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```
6. (Required for a mixed Virtual Chassis only) Set the master switch into mixed mode, and reboot the switch to complete the configuration:



**NOTE:** You do not have to complete this step if you are configuring a Virtual Chassis that includes QFX3500 and QFX3600 switches only.

You must complete this step if your Virtual Chassis includes a mix of QFX5100, EX4300, and QFX3500 or QFX3600 switches.

```
user@device> request virtual-chassis mode mixed reboot
```

7. After the reboot is complete, specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@switch# set preprovisioned
```

8. Specify all the members that you want included in the Virtual Chassis, listing each switch's serial number with the desired member ID and role:

```
[edit virtual-chassis]
user@switch# set member 0 serial-number abc123 role routing-engine
user@switch# set member 1 serial-number def456 role routing-engine
user@switch# set member 2 serial-number ghi789 role line-card
user@switch# set member 3 serial-number jkl012 role line-card
```

9. (Optional. Recommended for a two-member Virtual Chassis) Disable the split and merge feature:

```
[edit virtual-chassis]
```

user@switch# **set no-split-detection**

10. Power on the other member switches. The member IDs and roles have been determined by the configuration, so you can power on the member switches in any order.
11. (Required if you are configuring a mixed Virtual Chassis) Set each individual switch into mixed mode, and reboot the switch to complete the configuration:



**NOTE:** You do not have to complete this step if you are configuring a Virtual Chassis that includes QFX3500 and QFX3600 switches only.

You must complete this step if your Virtual Chassis includes a mix of QFX5100, EX4300, and QFX3500 or QFX3600 switches.

user@device> **request virtual-chassis mode mixed reboot**

12. (Optional) On each individual member switch, configure the ports that will be used to interconnect the member switches into VCPs using the following command:



**NOTE:** SFP+ and QSFP+ links are automatically turned into VCPs when the preprovisioned configuration is set.

This step is, therefore, optional and should only be used when a VCP is not automatically created.

user@switch> **request virtual-chassis vc-port set pic-slot *pic-slot-number* port *port-number***  
where *pic-slot-number* is the PIC slot number.

For instance, if you wanted to set port 0 on the QSFP+ interface on PIC slot 2 as a VCP:

user@switch> **request virtual-chassis vc-port set pic-slot 2 port 0**

The VCPs automatically bundle into a Link Aggregation Group when two or more interfaces of the same speed are configured into VCPs between the same two member switches. See *Understanding EX Series Virtual Chassis Port Link Aggregation*.



**NOTE:** You cannot modify the mastership priority when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two Routing Engines are assigned the same mastership priority value. However, the member that was powered on first has higher prioritization according to the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).

---

## Configuring a QFX Series Virtual Chassis with a Nonprovisioned Configuration File

You can use nonprovisioned configuration to configure a QFX Series Virtual Chassis.

To configure the Virtual Chassis using a nonprovisioned configuration:





**NOTE:** You can configure a QFX Series Virtual Chassis while the cables are or are not physically connected.

1. Power on only the switch that you plan to use as the master switch.
2. (Required for a mixed Virtual Chassis only) Set the master switch into mixed mode, and reboot the switch to complete the configuration:



**NOTE:** You do not have to complete this step if you are configuring a Virtual Chassis that includes QFX3500 and QFX3600 switches only.

You must complete this step if your Virtual Chassis includes a mix of QFX5100, EX4300, and QFX3500 or QFX3600 switches.

```
user@device> request virtual-chassis mode mixed reboot
```

3. After the master switch reboots, specify the identification parameters for the switch by completing the initial configuration. See [“Configuring a QFX3500 Device as a Standalone Switch” on page 175](#) or *Configuring a QFX3600 Device as a Standalone Switch* for details.



**NOTE:** The properties that you specify for the master switch apply to the entire Virtual Chassis configuration.

4. (Optional) Configure the master switch with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis:

```
user@switch# set interfaces vme unit 0 family inet address /ip-address/mask/
```

5. (Optional) Configure mastership priority for the other member switches:

```
[edit virtual-chassis]
user@switch# set member 0 mastership-priority 255
user@switch# set member 1 mastership-priority 255
```

6. (Optional. Recommended for a two-member Virtual Chassis) On the master switch, disable the split and merge feature:

```
[edit virtual-chassis]
user@switch# set no-split-detection
```

7. Power on the other member switches.
8. (Required for a mixed Virtual Chassis only) Set each individual switch into mixed mode, and reboot the switch to complete the configuration:



**NOTE:** You do not have to complete this step if you are configuring a Virtual Chassis that includes QFX3500 and QFX3600 switches only.

You must complete this step if your Virtual Chassis includes a mix of QFX5100, EX4300, and QFX3500 or QFX3600 switches.

```
user@device> request virtual-chassis mode mixed reboot
```

9. On each individual member switch, configure the ports that will be used to interconnect the member switches into VCPs using the following command:

```
user@switch> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
```

where *pic-slot-number* is the PIC slot number.

For instance, if you wanted to set port 0 on the QSFP+ interface on PIC slot 2 as a VCP:

```
user@switch> request virtual-chassis vc-port set pic-slot 2 port 0
```

The VCPs automatically bundle into a Link Aggregation Group when two or more interfaces of the same speed are configured into VCPs between the same two member switches. See *Understanding EX Series Virtual Chassis Port Link Aggregation*.



**NOTE:** If you do not edit the Virtual Chassis configuration file, a nonprovisioned configuration is generated by default. The mastership priority value for each member switch is 128. The master role is selected by default. You can change the role that is performed by the members by modifying the mastership priority. See [“Configuring Mastership of a Virtual Chassis \(CLI Procedure\)” on page 6941](#). We recommend that you specify the same mastership priority value for the desired master and backup members. In this example, the highest possible mastership priority has been assigned to two members. However, the member that was powered on first has higher prioritization according to the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#). The other members use the default mastership priority in this example, which configures them to function in the role of linecard.



**NOTE:** If you want to change the member ID that the master has assigned to a member switch, use the `request virtual-chassis renumber` command.

#### Related Documentation

- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [Configuring Mastership of a Virtual Chassis \(CLI Procedure\) on page 6941](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis](#)

## Adding a New Switch to an Existing QFX Series Virtual Chassis (CLI Procedure)

This procedure applies to QFX Series Virtual Chassis only. For the procedure on adding a switch to a Virtual Chassis Fabric (VCF), see [“Adding a Device to a Virtual Chassis Fabric” on page 7062](#).

You can use this procedure to add an EX4300, QFX3500, QFX3600, or QFX5100 device to an existing QFX Series Virtual Chassis. A QFX Series Virtual Chassis is a Virtual Chassis composed of QFX3500 series switches, QFX3600 series switches, QFX5100 switches, or a mix of QFX Series switches. A QFX Series Virtual Chassis can also include EX4300 switches.

Before you begin, be sure you have:

- Mounted the new switch in a rack.
- Enabled automatic software update on the Virtual Chassis. See [“Configuring Automatic Software Update on Virtual Chassis Member Switches \(CLI Procedure\)” on page 6944](#).
- If you are expanding a preprovisioned configuration, made a note of the serial number (the number is on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.
- If you are expanding a preprovisioned configuration, edited the existing Virtual Chassis configuration to include the serial number of the new member switch. The parameters specified in the master Virtual Chassis configuration file are applied to the new switch after it has been interconnected to an existing member switch.



**NOTE:** If you are expanding a preprovisioned Virtual Chassis configuration, you can use the autoprovisioning feature to add member switches to that configuration.

- (Optional) Configured Ethernet interfaces on different member switches into the same LAG. See [“Configuring Link Aggregation” on page 2593](#).

An active member switch might temporarily go down before coming back up as part of this procedure. Having traffic load-balanced across member switches using a LAG helps alleviate traffic loss during this procedure.

To add a new member switch to an existing Virtual Chassis configuration:

1. If the new member switch has been previously configured, revert that switch's configuration to the factory defaults before interconnecting it into the Virtual Chassis. See [“Reverting to the Default Factory Configuration” on page 188](#).
2. (Required for a mixed Virtual Chassis only) Set the new switch into mixed mode, and reboot the switch to complete the configuration:



**NOTE:** You do not need to configure your Virtual Chassis into mixed mode if the Virtual Chassis is composed of QFX3500 and QFX3600 switches only.

```
user@device> request virtual-chassis mode mixed reboot
```

If you are adding a switch that converts a non-mixed Virtual Chassis into a mixed Virtual Chassis, you must also log onto the Virtual Chassis and enter the **request virtual-chassis mode mixed all-members reboot** command either before or after interconnecting the new switch into your Virtual Chassis.

3. Interconnect the new switch to one member of the existing Virtual Chassis. You interconnect the new member switch using a non-channelized QSFP+ interface or SFP+ interface.

Connect only one interface on the unpowered new switch to a VCP on a member switch in the existing Virtual Chassis at this point of the procedure.

4. Set the interface on the new member switch as a Virtual Chassis Port (VCP):  
`user@switch> request virtual-chassis vc-port set pic-slot slot-number port port-number`
5. Confirm that the new member switch is now included within the Virtual Chassis configuration by entering the `show virtual-chassis` command. The new member switch should be listed in the output and the **Status** is **Prsnt**.
6. Cable the next port into the Virtual Chassis, using Steps 2 through 5.



**CAUTION:** If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

7. If further Virtual Chassis configuration is needed, see [“Configuring a QFX Series Virtual Chassis \(CLI Procedure\)”](#) on page 6931.

#### Related Documentation

- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\)](#) on page 6931

## Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure)

---



**NOTE:** This topic does not apply to Virtual Chassis Fabric (VCF) or EX8200 Virtual Chassis. See [“Removing a Device From a Virtual Chassis Fabric”](#) on page 7069 or [Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis \(CLI Procedure\)](#).

You can replace a member switch in a Virtual Chassis without disrupting network service on the other members. You can retain the existing configuration of the member switch and apply it to a new member switch, or you can free up the member ID and make it available for assignment to a new member switch.

If you want to replace a member switch of a *mixed* Virtual Chassis that contains EX4200, EX4500, or EX4550 switches, see [Removing an EX4200, EX4500, or EX4550 Switch From a Mixed Virtual Chassis \(CLI Procedure\)](#).

To replace a member switch, use the procedure that matches what you need to accomplish:

- [Remove, Repair, and Reinstall the Same Switch on page 6939](#)
- [Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration on page 6939](#)
- [Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch on page 6940](#)

### Remove, Repair, and Reinstall the Same Switch

If you need to repair a member switch, you can remove it from the Virtual Chassis configuration without disrupting network service for the other members. The master stores the configuration for the member ID so that it can be reapplied when the member switch (with the same base MAC address) is reconnected.

To remove, repair, and reinstall the member switch:

1. Power off and disconnect the member switch to be repaired.
2. Repair, as necessary.
3. Reconnect the switch and power it on.

### Remove a Member Switch, Replace It with a Different Switch, and Reapply the Old Configuration

If you are unable to repair a member switch, you can replace it with a different member switch while retaining the previous configuration. The master stores the configuration of the member that was removed. When you connect a different member switch, the master assigns a new member ID. But the old configuration is still stored under the previous member ID of the previous member switch.



**NOTE:** If you have used a preprovisioned configuration, you can use the `replace` command to change the serial number in the Virtual Chassis configuration file. Substitute the serial number of the replacement member switch (on the back of the switch) for the serial number of the member switch that was removed.

To remove and replace a switch and reapply the old configuration:

1. Power off and disconnect the member switch to be replaced.
2. If the replacement member switch has been previously configured, revert that switch's configuration to the factory defaults. See *Reverting to the Default Factory Configuration for the EX Series Switch* for information on reverting to the factory default configuration on an EX Series switch or [“Reverting to the Default Factory Configuration” on page 188](#) for information on reverting to the factory default configuration on a QFX Series switch.
3. If you are interconnecting a switch using a dedicated VCP, connect one VCP on the replacement member switch to a VCP of another Virtual Chassis member switch.

If you are interconnecting a switch using an optical port configured as a VCP, cable the optical ports together then configure the port on the Virtual Chassis as a VCP:

```
user@switch> request virtual-chassis vc-port set pic-slot 1 port port-number
```

4. Power on the new member switch.
5. Confirm that the new member switch is now included in the Virtual Chassis configuration by checking the front-panel LCD or the for the member ID. It should display a member ID in the range from 0 through 9.

If you are using a switch that does not have an LCD interface, confirm the switch is part of the Virtual Chassis configuration by entering the **show virtual-chassis** and reviewing the output.

6. Cable the other VCP on the new member switch into the Virtual Chassis. Use the instruction in step 3 to complete this step.



**CAUTION:** If you immediately cable both VCPs on the new switch into the existing Virtual Chassis at the same time, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. Network traffic to this switch is dropped during the downtime.

The member switch will return to the normal operational state with no user intervention, and normal operation of the Virtual Chassis will resume after this downtime.

7. On the master switch, Issue the **request virtual-chassis renumber** command from the Virtual Chassis master to change the member switch's current member ID to the member ID of the member switch that was removed from the Virtual Chassis configuration.

### Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch

---

When you remove a member switch from the Virtual Chassis configuration, the master keeps that member switch's member ID in reserve. To make that member switch's member ID available for reassignment, issue the **request virtual-chassis recycle** command from the Virtual Chassis master.



**NOTE:** When you add or delete members in a Virtual Chassis configuration, internal routing changes might cause temporary traffic loss for a few seconds.

#### Related Documentation

- *Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)*
- *Adding a New Switch to an Existing QFX Series Virtual Chassis (CLI Procedure) on page 6936*
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*

- *Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4200 Switch to a Preprovisioned EX4500 Virtual Chassis or a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis (CLI Procedure)*

## Configuring Mastership of a Virtual Chassis (CLI Procedure)



**NOTE:** This topic applies to all EX Series Virtual Chassis except EX8200 Virtual Chassis. See *Configuring an EX8200 Virtual Chassis (CLI Procedure)* for information about EX8200 Virtual Chassis.

You can designate the role (master, backup, or linecard) that a member switch performs within any Virtual Chassis, whether or not you are using a preprovisioned configuration.



**NOTE:** A Virtual Chassis configuration has two Routing Engines—one is the switch in the master role and the other is the switch in the backup role. Therefore, we recommend that you always use `commit synchronize` rather than `commit` to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved in both Routing Engines.

This topic describes:

- [Configuring Mastership Using a Preprovisioned Configuration File on page 6941](#)
- [Configuring Mastership Using a Configuration File That Is Not Preprovisioned on page 6942](#)

### Configuring Mastership Using a Preprovisioned Configuration File

To configure mastership using a preprovisioned configuration:

1. Note the serial numbers of the switches that you want to function in the master role and backup role.
2. Power on only the switch that you want to function in the master role.
3. Edit the configuration to specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@switch# set preprovisioned
```

4. Specify the serial numbers of the member switches that you want to function as master and backup, specifying their role as **routing-engine**:

```
[edit]
user@switch# set virtual-chassis member 0 serial-number abc123 role routing-engine
user@switch# set virtual-chassis member 1 serial-number def456 role routing-engine
```



**NOTE:** You cannot directly modify the mastership priority value when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two members assigned the **routing-engine** role are assigned the same mastership priority value (128). However, the member that was powered on first has higher priority for the master role election according to the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#). Only two members can be configured with the **routing-engine** role.

5. Specify the serial numbers of any other member switches that you are including in the Virtual Chassis configuration. You can also explicitly configure their role as **line-card**.

### Configuring Mastership Using a Configuration File That Is Not Preprovisioned

To configure mastership of the Virtual Chassis through a configuration that is not preprovisioned:

1. Power on only the switch that you want to function in the master role.
2. Configure the highest possible mastership priority value (**255**) for the member that you want to function in the master role:  

```
[edit virtual-chassis]
user@switch# set member 0 mastership-priority 255
```
3. Configure the same mastership priority value (continue to edit the Virtual Chassis configuration on the master) for the member that you want to be in the backup role:  

```
[edit virtual-chassis]
user@switch# set member 1 mastership-priority 255
```



**NOTE:** We recommend that the master and backup have the same mastership priority value to prevent the master and backup status from switching back and forth between master and backup members in failover conditions.

4. Use the default mastership priority value (**128**) for the remaining member switches or configure the mastership priority to a value that is lower than the value specified for members functioning in the master and backup roles.

#### **Related Documentation**

- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis](#)
- [Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis \(CLI Procedure\)](#)
- [Adding an EX4200 Switch to a Preprovisioned EX4500 Virtual Chassis or a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis \(CLI Procedure\)](#)



- *Adding an EX4500 Switch to a Preprovisioned EX4200 Virtual Chassis (CLI Procedure)*
- *Adding an EX4500 Switch to a Nonprovisioned EX4200 Virtual Chassis (CLI Procedure)*

## Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure)

When a backup member takes control of a Virtual Chassis because of a reset or other temporary failure, the backup member uses the MAC address of the old master switch as the system MAC base address. This process helps ensure a smooth transition of mastership with no disruption to network connectivity.

The MAC persistence timer is used in situations in which the master switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed. If the old master switch does not rejoin the Virtual Chassis before the timer elapses, the new master switch starts using its own MAC address as the system's MAC base address. For information regarding how the system MAC base address is used to assign MAC addresses to ports in a Virtual Chassis, see ["Understanding MAC Address Assignment on a Virtual Chassis" on page 6928](#).

The default timer value is 10 minutes. The maximum timer value is 60 minutes.

You can disable the MAC persistence timer starting in Junos OS Release 12.2. When the MAC persistence timer is disabled, the MAC address of the old master switch is used as the system MAC base address; no MAC address changes occur within the Virtual Chassis even when the old master switch is no longer a member of the Virtual Chassis because it has been physically disconnected or removed.

To configure or modify the MAC persistence timer:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer minutes
```

To disable the MAC persistence timer:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer disable
```

### Related Documentation

- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- *Understanding EX Series Virtual Chassis Components*

## Disabling Split and Merge in a Virtual Chassis (CLI Procedure)

The split and merge feature is enabled by default on all EX Series switches and QFX Series devices in a Virtual Chassis. You can disable the split and merge feature. If you disable the split and merge feature and the Virtual Chassis splits, both parts of the split Virtual Chassis configuration remain active.

In a preprovisioned Virtual Chassis, if both of the Routing Engines end up in the same Virtual Chassis configuration after a split, the other part of the split Virtual Chassis configuration remains inactive. If the Routing Engines end up in different parts of the split Virtual Chassis configuration and the rest of the member switches are configured as having linecard roles, then a backup Routing Engine might not be selected for either part.

We recommend disabling split and merge on a Virtual Chassis with two member switches. A two-member switch Virtual Chassis that has disabled split and merge can reform more quickly and with less complications as a result of the feature being disabled.

To disable the split and merge feature in a Virtual Chassis:

```
[edit]
user@switch# set virtual-chassis no-split-detection
```

### Related Documentation

- [Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge](#)
- [Understanding Split and Merge in a Virtual Chassis on page 6922](#)

## Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)

The automatic software update feature allows you to automatically update the software version on prospective member switches as they are added so that they can join the Virtual Chassis.



**NOTE:** The version of Junos OS running on the Virtual Chassis must be compatible with the software running on the prospective member switch for an automatic software update to occur. For information on Junos OS compatibility and other automatic software update restrictions, see [“Understanding Automatic Software Update on Virtual Chassis Member Switches” on page 6925](#).

Before you begin, ensure that you know the name or the URL of the software package to be used by the automatic software update feature.

To configure the automatic software update feature for an EX Series or QFX Series Virtual Chassis with the exception of a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch:

[edit]

```
user@switch# set virtual-chassis auto-sw-update package-name package-name
```

To configure the automatic software update feature on a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch:

[edit]

```
user@switch# set virtual-chassis auto-sw-update ex-4200 package-name package-name
```

```
user@switch# set virtual-chassis auto-sw-update ex-4500 package-name package-name
```

If the software package is located on a local directory on the switch, use the following format for **package-name**:

***/pathname/package-name***

If the software package is to be downloaded and installed from a remote location, use one of the following formats:

***ftp://hostname/pathname/package-name***

***ftp://username:prompt@ftp.hostname.net/package-name***

***http://hostname/pathname/package-name***

If you are configuring a mixed Virtual Chassis containing at least one EX4200 switch and at least one EX4500 or EX4550 switch, use the **ex-4200** keyword when you are specifying a path to a package for the EX4200 switches and the **ex-4500** when you are specifying a path to a package for the EX4500 or EX4550 switches. You do not need to specify the **ex4500** keyword when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis, however, because the Junos OS package for an EX4500 switch updates the software for both EX4500 and EX4550 switches.

#### Related Documentation

- *Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches*
- [Understanding Automatic Software Update on Virtual Chassis Member Switches on page 6925](#)

## Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure)

Every Virtual Chassis has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a Virtual Chassis ID using the **set virtual-chassis id** command. When two Virtual Chassis configurations attempt to merge, the Virtual Chassis ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID for the newly merged Virtual Chassis configuration.

To configure the Virtual Chassis ID:

```
[edit]
user@switch# set virtual-chassis id id
```

### Related Documentation

- [Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge](#)
- [Understanding Split and Merge in a Virtual Chassis on page 6922](#)

## Configuration Statements

---

- [\[edit virtual-chassis\] Configuration Statement Hierarchy on page 6947](#)
- [aliases \(Virtual Chassis\) on page 6949](#)
- [alias-name \(Virtual Chassis aliases\) on page 6950](#)
- [auto-sw-update on page 6951](#)
- [id on page 6953](#)
- [location \(Virtual Chassis\) on page 6954](#)
- [mac-persistence-timer on page 6955](#)
- [mastership-priority on page 6956](#)
- [member on page 6958](#)
- [no-management-vlan on page 6959](#)
- [no-split-detection on page 6960](#)
- [package-name on page 6961](#)
- [preprovisioned on page 6962](#)
- [role on page 6963](#)
- [serial-number on page 6966](#)
- [serial-number \(Virtual Chassis aliases\) on page 6967](#)
- [traceoptions \(Virtual Chassis\) on page 6968](#)
- [vcp-no-hold-time on page 6971](#)
- [virtual-chassis on page 6973](#)

## [edit virtual-chassis] Configuration Statement Hierarchy

This topic lists supported and unsupported configuration statements in the **[edit virtual-chassis]** hierarchy level on EX Series and QFX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms.

For detailed information about feature support on specific EX Series or QFX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 6947](#)
- [Unsupported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 6948](#)

### Supported Statements in the [edit virtual-chassis] Hierarchy Level

The following hierarchy shows the **[edit virtual-chassis]** configuration statements supported on EX Series or QFX Series switches:

```
virtual-chassis {
  aliases {
    serial-number serial-number {
      alias-name alias-name;
    }
  }
  auto-provisioned;
  auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
  }
  fast-failover (ge | vcp disable | xe);
  graceful-restart {
    disable;
  }
  id id;
  mac-persistence-timer [minutes | disable];;
  member member-id {
    location location;
    mastership-priority number;
    no-management-vlan;
    role (line-card | routing-engine);
    serial-number;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
```

```
file filename <files number> <size size> <world-readable | no-world-readable> <match
  regex>;
flag flag ;
}
vc-port {
  lag-hash (packet-based | source-port-based);
}
vcp-no-hold-time;
}
```

---

### Unsupported Statements in the [edit virtual-chassis] Hierarchy Level

All statements in the [edit virtual-chassis] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### Related Documentation

- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX2200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)

## aliases (Virtual Chassis)

<b>Syntax</b>	<pre>aliases {   serial-number serial-number {     alias-name alias-name;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.
<b>Description</b>	<p>Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.</p> <p>An alias is not specified for a device until the alias name is specified using the <b>alias-name</b> keyword.</p> <p>The alias appears in the <b>Alias-Name</b> field in the <b>show virtual-chassis</b> command.</p> <p>Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 7035</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li> </ul>

## alias-name (Virtual Chassis aliases)

---

**Syntax** `alias-name alias-name;`

**Hierarchy Level** `[edit virtual-chassis aliases serial-number serial-number]`

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.

**Description** Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.

The alias appears in the **Alias-Name** field in the **show virtual-chassis** command.

Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.

In the following example, the **dc-floor-1** alias name is assigned to the member switch with the serial number AB0123456789.

### set serial-number

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

### show virtual-chassis

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 9d5d.5556.919a
Fabric Mode: Enabled

Member ID  Status   Serial No   Alias-Name   Model          Mstr  prio  Role
0 (FPC 0)  Prsnt    AB0123456789 dc-floor-1   qfx5100-48s-6q 129   Master
<additional output removed for brevity>
```

**Options** *alias-name*—The text label, or alias, assigned to the member switch by the user.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)



## auto-sw-update

<b>Syntax</b>	<pre> auto-sw-update {   (ex-4200   ex-4300   ex-4500   ex-4600   qfx-3   qfx-5)   package-name package-name; } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>The <b>ex-4200</b> and <b>ex-4500</b> options introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>The <b>ex-4300</b>, <b>qfx-3</b>, and <b>qfx-5</b> options introduced in Junos OS Release 13.2X51-D20.</p> <p>The <b>ex-4600</b> option introduced in Junos OS Release 13.2X51-D25.</p>
<b>Description</b>	<p>Enable the automatic software update feature for Virtual Chassis or Virtual Chassis Fabric (VCF) configurations.</p> <p>You should only use the keywords that specify a device—<b>ex-4300</b>, <b>ex-4600</b>, <b>qfx-3</b>, and <b>qfx-5</b>—when configuring automatic software update on a mixed Virtual Chassis or Virtual Chassis Fabric (VCF). You can simply specify the <i>package-name</i> without specifying the device keywords in non-mixed Virtual Chassis or VCF topologies.</p> <p>You must enter the <b>auto-sw-update</b> statement multiple times—once for each device family in your mixed Virtual Chassis or VCF—in most scenarios when enabling the automatic software update for a mixed Virtual Chassis or VCF.</p> <p>The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. You do not, therefore, need to specify the <b>ex-4500</b> keyword when configuring automatic software update for a mixed Virtual Chassis that include EX4500 and EX4550 switches only. You also only have to enter the <b>ex-4500</b> keyword once to configure automatic software update for all EX4500 and EX4550 member switches in the same mixed Virtual Chassis.</p> <p>The Junos OS package for a QFX3500 device updates the software for QFX3500 and QFX3600 devices. You do not, therefore, need to specify the <b>qfx-3</b> keyword when configuring automatic software update for a Virtual Chassis composed entirely of QFX3500 and QFX3600 devices. You also have to enter the <b>qfx-3</b> keyword only once to configure automatic software update for all QFX3500 and QFX3600 member devices in the same mixed Virtual Chassis.</p> <p>The remaining statement is explained separately.</p>
<b>Default</b>	The automatic software update feature is disabled.
<b>Options</b>	<p><b>package-name package-name</b>—Specify a path to a Junos OS software image.</p> <p><b>ex-4200</b>—Specify a path to a Junos OS image for an EX4200 switch when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed</p>

EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

**ex-4300**—Specify a path to a Junos OS image for an EX4300 switch when enabling automatic software update for a mixed Virtual Chassis or VCF.

**ex-4500**—Specify a path to a Junos OS image for an EX4500 switch, an EX4550 switch, or both types of switches when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. Therefore, you only enter this command once to upgrade the EX4500 and EX4550 member switches in the same mixed Virtual Chassis.

The **ex-4500** keyword also does not need to be specified when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis.

**ex-4600**—Specify a path to a Junos OS image for an EX4600 switch when enabling automatic software update for a mixed Virtual Chassis.

**qfx-3**—Specify a path to a Junos OS image for a QFX3500, QFX3600, or both types of devices when enabling automatic software update for a mixed VCF.

**qfx-5**—Specify a path to a Junos OS image for a QFX5100 device when enabling automatic software update for a mixed VCF.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</i></li><li>• <a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 6944</a></li><li>• <a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric on page 7050</a></li></ul>
------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## id

---

<b>Syntax</b>	<code>id id;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Configure the alphanumeric string that identifies a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.
<b>Options</b>	<i>id</i> —Virtual Chassis ID (VCID), which uses the ISO family address format—for example, <b>9622.6ac8.5345</b> .
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li> <li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 6946</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Understanding Virtual Chassis Member ID Numbering in an EX8200 Virtual Chassis</i></li> </ul>

## location (Virtual Chassis)

---

<b>Syntax</b>	<code>location location;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis member member-id</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Set a description of the location of the Virtual Chassis or VCF member switch or external Routing Engine.</p> <p>The <b>Location</b> field is visible to users who enter the <b>show virtual-chassis status detail</b> command.</p> <p>Setting this description has no effect on the operation of the member device.</p>
<b>Options</b>	<b>location</b> —Location of the current member switch or external Routing Engine. The <b>location</b> can be any single word.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <i>Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</i></li><li>• <i>Example: Configuring a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis</i></li><li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li><li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)</a></li><li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li></ul>

## mac-persistence-timer

<b>Syntax</b>	mac-persistence-timer [ <i>minutes</i>   <b>disable</b> ];
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>disable</b> introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>The maximum timer limit changed from no maximum timer limit to 60 minutes in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the master role as the system MAC base address after the original master switch is removed from the Virtual Chassis or VCF. The system MAC base address does not change in the event of a switchover provided the switch originally configured in the master role remains a member of the Virtual Chassis or VCF.</p> <p>The maximum timer limit is 60 minutes starting in Junos OS Release 12.2. There are no minimum or maximum timer limits in prior Junos OS releases.</p>
<b>Default</b>	The MAC persistence timer is set to 10 minutes by default.
<b>Options</b>	<p><b>minutes</b>—Time in minutes that the member switch in the backup role continues to use the system MAC base address of the old master before using its own system MAC base address after the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p> <p><b>disable</b>—Disable the MAC persistence timer. The system MAC base address never changes when the MAC persistence timer is disabled, even when the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure) on page 6943</a></li> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> </ul>

## mastership-priority

---

<b>Syntax</b>	<code>mastership-priority <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis member</a> <i>member-id</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Mastership priority option <b>0</b> introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>The mastership priority value is the most important factor in determining the role of the member switch within a nonprovisioned Virtual Chassis or VCF configuration. Other factors (see <a href="#">“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917</a>) also affect the election of the master.</p> <p>The mastership priority value takes the highest precedence in the master election algorithm. The member switch with highest mastership priority assumes the master Routing Engine role of the Virtual Chassis or VCF. Toggling back and forth between master and backup status in failover conditions is undesirable, so we recommend that you assign the same mastership priority value to both the master and the backup. Secondary factors in the master election algorithm determine which of these two members (that is, the two members that are assigned the highest mastership priority value) functions as the master of the Virtual Chassis or VCF.</p> <p>This statement is not used for the EX8200 Virtual Chassis, which determines mastership by external Routing Engine uptime. See <i>Understanding Virtual Chassis Roles in an EX8200 Virtual Chassis</i>.</p> <p>A switch with a mastership priority of <b>0</b> never takes the master or backup role.</p>
<b>Default</b>	128
<b>Options</b>	<p><i>number</i>—Mastership priority value.</p> <p><b>Range:</b> 0 through 255</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</a></li><li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</a></li></ul>

- *Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*

## member

---

<b>Syntax</b>	<pre>member <i>member-id</i> {     <i>location</i> <i>location</i>;     <i>mastership-priority</i> <i>number</i>;     <i>no-management-vlan</i>;     <i>serial-number</i> <i>serial-number</i>;     <i>role</i> <i>role</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Configure a switch or an XRE200 External Routing Engine as a member of a Virtual Chassis or a Virtual Chassis Fabric (VCF).
<b>Default</b>	<p>When an EX Series switch or a QFX Series devices configured in standalone mode is powered on but not interconnected through its Virtual Chassis ports (VCPs) with other member switches, its default member ID is 0.</p> <p>There is no default member ID in an EX8200 or EX9200 Virtual Chassis. An EX8200 or EX9200 Virtual Chassis must be preprovisioned, and that process configures the member IDs.</p>
<b>Options</b>	<p><b><i>member-id</i></b>—Identifies a specific member switch of a Virtual Chassis or VCF configuration.</p> <p>The exact range for a specific Virtual Chassis or VCF depends on the number of switches allowed in the Virtual Chassis or VCF.</p> <p>In an EX8200 Virtual Chassis, member IDs 0 through 7 are reserved for EX8200 member switches and member IDs 8 and 9 are reserved for the master and backup external Routing Engines.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</a></li><li>• <a href="#">Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</a></li></ul>



- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX9200 Virtual Chassis](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)

## no-management-vlan

<b>Syntax</b>	no-management-vlan;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis member member-id</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Remove the specified member's out-of-band management port from the virtual management Ethernet (VME) global management VLAN of the Virtual Chassis or VCF configuration.</p> <p>For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:</p> <pre>virtual-chassis {   member 2 {     no-management-vlan;   } }</pre> <p>You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell <b>ifconfig</b> command.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Setting Up a Multimember EX4200 Virtual Chassis Access Switch with a Default Configuration</a></li> <li>• <a href="#">Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 6919</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Configuration on page 7043</a></li> </ul>

## no-split-detection

---

<b>Syntax</b>	no-split-detection;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Disable the split and merge feature in a Virtual Chassis or VCF configuration.</p> <p>We recommend using this statement to disable the split and merge feature when configuring a two-member Virtual Chassis. Enabling this statement on a two-member Virtual Chassis ensures that both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split.</p>
<b>Default</b>	The split and merge feature is enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li><li>• <a href="#">Disabling Split and Merge in a Virtual Chassis (CLI Procedure) on page 6944</a></li><li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 6946</a></li><li>• <a href="#">Understanding Split and Merge in a Virtual Chassis on page 6922</a></li></ul>

## package-name

<b>Syntax</b>	<code>package-name <i>package-name</i>;</code>
<b>Hierarchy Level</b>	[edit virtual-chassis <a href="#">auto-sw-update</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Specify the software package name or location of the software package to be used by the automatic software update feature for Virtual Chassis or VCF.
<b>Default</b>	No package name is specified.
<b>Options</b>	<p><b><i>package-name</i></b>—Name of the software package or the URL to the software package to be used.</p> <ul style="list-style-type: none"> <li>If the software package is located on a local directory on the switch, use the following format for <b><i>package-name</i></b>:  <b><i>/pathname/package-name</i></b></li> <li>If the software package is to be downloaded and installed from a remote location, use one of the following formats:  <b><i>ftp://hostname/pathname/package-name</i></b> <b><i>ftp://username:prompt@ftp.hostname.net/package-name</i></b> <b><i>http://hostname/pathname/package-name</i></b></li> </ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</i></li> <li><a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)</a> on page 6944</li> <li><a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric</a> on page 7050</li> </ul>

## preprovisioned

---

<b>Syntax</b>	preprovisioned;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Enable the preprovisioned configuration mode for a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.</p> <p>When the preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the mastership priority or member ID of member switches.</p> <p>You must use this statement to configure an EX8200 Virtual Chassis. Nonprovisioned configuration of an EX8200 Virtual Chassis is not supported.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <i>Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</i></li><li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li><li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX9200 Virtual Chassis</i></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 6938</a></li></ul>

## role

<b>Syntax</b>	<code>role (line-card   routing-engine);</code>
<b>Hierarchy Level</b>	[edit <b>virtual-chassis</b> <b>preprovisioned member</b> <i>member-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Specify the roles of the members of the Virtual Chassis or a Virtual Chassis Fabric (VCF) in a preprovisioned Virtual Chassis.

### Virtual Chassis Fabric

Specify the role to be performed by each switch. In a VCF, the spine devices are configured into the Routing Engine role and the leaf devices are configured into the line card role. You can configure several devices into the Routine Engine role, but only two will operate in the Routing Engine role at a time. The role must be associated with the member's serial number.

### EX Series (except EX8200 Virtual Chassis) and QFX Series Virtual Chassis

Specify the role to be performed by each member switch. Associate the role with the member's serial number.

When you use a preprovisioned configuration, you cannot modify the mastership priority or member ID of member switches through the user interfaces. The mastership priority value is generated by the software, based on the assigned role:

- A member configured as **routing-engine** is assigned the mastership priority **129**.
- A member configured as **line-card** is assigned the mastership priority **0**.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the mastership priority **128**.

The configured role specifications are permanent. If both **routing-engine** members fail, a **line-card** member cannot take over as master of the Virtual Chassis configuration. You must delete the preprovisioned configuration to change the specified roles in a Virtual Chassis.

Explicitly configure two members as **routing-engine** and configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional members, they function in a linecard role by default. In that case, a member that is functioning in a linecard role can take over mastership if the members functioning as master and backup (**routing-engine** role) both fail.

### EX8200 Virtual Chassis

Specify the role to be performed by each XRE200 External Routing Engine and each EX8200 member switch. Associate the role with the member's serial number. An EX8200 Virtual Chassis cannot function when both external Routing Engines, which must be configured in the **routing-engine** role, have failed.

- Options**
- **line-card**—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis or VCF configuration other than the master or backup functions in the linecard role and runs only a subset of Junos OS for EX Series switches. A member functioning in the linecard role does not run the control protocols or the chassis management processes.

A Virtual Chassis must have at least three members for one member to function in the linecard role.

In an EX8200 Virtual Chassis configuration, all member switches must be in the linecard role.

- **routing-engine**—Enables the member to function as a master or backup of the Virtual Chassis or VCF configuration. The master manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

(All Virtual Chassis composed of EX Series switches, except EX8200 switches, or QFX Series devices) Specify two and only two members as **routing-engine**. The software determines which of the two members assigned the **routing-engine** role functions as master, based on the master election algorithm. See ["Understanding How the Master in a Virtual Chassis Is Elected" on page 6917](#). In these Virtual Chassis, the **routing-engine** role is associated with a switch.

(EX8200 Virtual Chassis) All XRE200 External Routing Engines must be in the **routing-engine** role.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------

**Related  
Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- *Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File*
- *Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines*
- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Configuring an EX9200 Virtual Chassis*
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- *Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)*
- *Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)*
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 6938](#)

## serial-number

---

<b>Syntax</b>	<code>serial-number serial-number;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis preprovisioned member member-id</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>In a preprovisioned Virtual Chassis or Virtual Chassis Fabric (VCF), specify the serial number of each member switch to be included in the configuration. If you do not include the serial number within the configuration, the switch cannot be recognized as a member of a preprovisioned configuration.</p> <p>In an EX8200 Virtual Chassis configuration, specify the serial number of each XRE200 External Routing Engine and each EX8200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the external Routing Engine or switch cannot be recognized as a member of the configuration.</p>
<b>Options</b>	<b>serial-number</b> —Permanent serial number for the external Routing Engine or for the member switch.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <i>Configuring an EX2200 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX3300 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX4300 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX9200 Virtual Chassis</i></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <i>Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)</i></li></ul>




## serial-number (Virtual Chassis aliases)

---

<b>Syntax</b>	<code>serial-number <i>serial-number</i> {     <i>alias-name alias-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis aliases</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series Virtual Chassis and Virtual Chassis Fabric (VCF).
<b>Description</b>	Specify the serial number that will be labeled with an alias in a Virtual Chassis or Virtual Chassis Fabric (VCF).  The remaining statements are explained separately.
<b>Options</b>	<b><i>serial-number</i></b> —Permanent serial number for the member switch in the Virtual Chassis or VCF.  You can retrieve the serial number for any device in your Virtual Chassis or VCF by entering the <b>show virtual-chassis</b> command and reviewing the output in the <b>Serial No</b> field.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 7035</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li> </ul>

## traceoptions (Virtual Chassis)

---

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> &lt;detail&gt; &lt;disable&gt; &lt;receive&gt; &lt;send&gt;; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>detail</b> added in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Define tracing operations for the Virtual Chassis or VCF.
<b>Default</b>	Tracing operations are disabled.
<b>Options</b>	<p><b>detail</b>—(Optional) Generate detailed trace information for a flag.</p> <p><b>disable</b>—(Optional) Disable a flag.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"><li>• <b>all</b>—All tracing operations.</li></ul> <div><p><b>TIP:</b> The <b>all</b> flag displays a subset of logs that are useful in debugging most issues. For more detailed information, use <b>all detail</b>.</p></div> <ul style="list-style-type: none"><li>• <b>auto-configuration</b>—Trace Virtual Chassis ports (VCPs) that have been automatically configured.</li><li>• <b>csn</b>—Trace Virtual Chassis complete sequence number (CSN) packets.</li><li>• <b>error</b>—Trace Virtual Chassis errored packets.</li></ul>

- **hello**—Trace Virtual Chassis hello packets.
- **krt**—Trace Virtual Chassis KRT events.
- **lsp**—Trace Virtual Chassis link-state packets.
- **lsp-generation**—Trace Virtual Chassis link-state packet generation.
- **me**—Trace Virtual Chassis ME events.
- **normal**—Trace normal events.
- **packets**—Trace Virtual Chassis packets.
- **parse**—Trace reading of the configuration.
- **psn**—Trace partial sequence number (PSN) packets.
- **route**—Trace Virtual Chassis routing information.
- **spf**—Trace Virtual Chassis SPF events.
- **state**—Trace Virtual Chassis state transitions.
- **task**—Trace Virtual Chassis task operations.

**no-stamp**—(Optional) Do not place a timestamp on any trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**receive**—(Optional) Trace received packets.

**replace**—(Optional) Replace a trace file rather than appending information to it.

**send**—(Optional) Trace transmitted packets.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB


**world-readable**—(Optional) Enable unrestricted file access.

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

**Related  
Documentation**

- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*
- [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 6975](#)
- *Verifying That Virtual Chassis Ports Are Operational*
- *Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis*
- *Troubleshooting an EX Series Virtual Chassis*

## vcp-no-hold-time

<b>Syntax</b>	vcp-no-hold-time;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
<b>Description</b>	<p>Disable the Virtual Chassis port (VCP) holddown timer for all VCPs in the Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The VCP holddown timer is an internal mechanism that delays a Virtual Chassis reconvergence for several seconds when a VCP becomes inactive. The purpose of this delay is to provide the VCP time to return online without having to reconverge the Virtual Chassis to adjust to the inactive VCP. All traffic to the VCP is dropped while the VCP is inactive. If the VCP remains down for a time that exceeds the VCP holddown timer, a Virtual Chassis reconvergence occurs.</p> <p>When this statement is enabled, the VCP holddown timer is disabled and the Virtual Chassis reconvergence occurs when a VCP becomes inactive. The period of time where traffic is dropped waiting for the VCP to return online is avoided.</p> <p>We recommend enabling this statement after a Virtual Chassis is operational. We recommend disabling this statement when you are adding or removing member switches from your Virtual Chassis.</p> <p>The VCP holddown timer cannot be viewed and is not user-configurable. You can only control whether the VCP holddown timer is enabled or disabled by configuring this statement.</p>
	<p> <b>NOTE:</b> For the EX4300 Virtual Chassis, you should enable the <code>vcp-no-hold-time</code> statement before performing a software upgrade using NSSU. If you do not enable the <code>vcp-no-hold-time</code> statement, the Virtual Chassis may split during the upgrade. A split Virtual Chassis can cause disruptions to your network, and you may have to manually reconfigure your Virtual Chassis after the NSSU if the split and merge feature was disabled. For more information about a split Virtual Chassis, see <a href="#">“Understanding Split and Merge in a Virtual Chassis” on page 6922</a></p>
<b>Default</b>	The VCP holddown timer is enabled by default on all devices that support this statement.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding EX4300 Virtual Chassis</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li> </ul>

- *Understanding EX Series Virtual Chassis Components*
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)

## virtual-chassis

<b>Syntax</b>	<pre> virtual-chassis {   aliases {     serial-number <i>serial-number</i> {       alias-name <i>alias-name</i>;     }   }   auto-provisioned   auto-sw-update {     (ex-4200   ex-4300   ex-4500   ex-4600   qfx-3   qfx-5)     package-name <i>package-name</i>;   }   fast-failover (ge   vcp disable   xe);   graceful-restart {     disable;   }   id <i>id</i>;   mac-persistence-timer [<i>minutes</i>   disable];;   member <i>member-id</i> {     location <i>location</i>;     mastership-priority <i>number</i>;     no-management-vlan;     serial-number;     role;   }   no-split-detection;   preprovisioned;   traceoptions (Virtual Chassis) {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt; &lt;match       <i>regex</i>&gt;;     flag <i>flag</i> ;   }   vc-port {     lag-hash (packet-based   source-port-based);   }   vcp-no-hold-time; } </pre>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Configure a Virtual Chassis or a Virtual Chassis Fabric (VCF).</p> <p>The remaining statements are explained separately.</p>
<b>Default</b>	A standalone EX Series switch is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A standalone XRE200 External Routing Engine or EX8200 switch is not part of an EX8200 Virtual Chassis until a Virtual Chassis configuration is set up.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
---------------------------------	-------------------------------------------------------------------------------------------------------------------

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <i>Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</i></li><li>• <i>Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</i></li><li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li><li>• <i>Configuring an EX3300 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX9200 Virtual Chassis</i></li></ul>
------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------



# Administration

- [Routine Monitoring on page 6975](#)
- [Operational Commands on page 6976](#)

## Routine Monitoring

---

- [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 6975](#)

### Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member

**Purpose** You can designate the role that a member performs within a Virtual Chassis or you can allow the role to be assigned by default. You can designate the member ID that is assigned to a specific switch by creating a permanent association between the switch's serial number and a member ID, using a preprovisioned configuration. Or you can let the member ID be assigned by the master, based on the sequence in which the member switch is powered on and on which member IDs are currently available.

The role and member ID of the member switch are displayed on the front-panel LCD.

Each member switch can be cabled to one or two other member switches, using either the dedicated Virtual Chassis ports (VCPs) on the rear panel, an uplink port that has been configured as a VCP, or an optical port that has been configured as a VCP. The members that are cabled together are considered neighbor members.

**Action** To display the role and member ID assignments using the CLI:

```
user@switch> show virtual-chassis
```

```
Virtual Chassis ID: 0000.e255.00e0
```

Member ID	Status	Serial No	Model	Mastership Priority	Role	Neighbor List ID, Interface
0 (FPC 0)	Prsnt	abc123	ex4200-48p	255	Master*	1 vcp-0 2 vcp-1
1 (FPC 1)	Prsnt	def456	ex4200-24t	255	Backup	2 vcp-0 0 vcp-1
2 (FPC 2)	Prsnt	abd231	ex4200-24p	128	Linecard	0 vcp-0 1 vcp-1

**Meaning** This output verifies that three EX4200 switches have been interconnected as a Virtual Chassis configuration through their dedicated VCPs to create an EX4200 Virtual Chassis. The display shows which of the VCPs is connected to which neighbor. The first port (**vcp-0**) of member **0** is connected to member **1** and the second port of member **0** (**vcp-1**) is connected to member **2**. The FPC slots for the switches are the same as the member IDs.

The **Mastership Priority** values indicate that the master and backup members have been explicitly configured, because they are not using the default value (**128**).



**NOTE:** This example uses output from an EX4200 Virtual Chassis. The output, with the exception of the **Model** column, would be identical on all other Virtual Chassis.

---

**Related  
Documentation**

- [Configuring Mastership of a Virtual Chassis \(CLI Procedure\) on page 6941](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Virtual Chassis on an EX Series Switch \(J-Web Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis](#)

---

## Operational Commands

- [clear virtual-chassis vc-port statistics](#)
- [request session member](#)
- [request virtual-chassis recycle](#)
- [request virtual-chassis renumber](#)
- [request virtual-chassis vc-port](#)
- [show virtual-chassis active-topology](#)
- [show virtual-chassis device-topology](#)
- [show virtual-chassis protocol adjacency](#)
- [show virtual-chassis protocol database](#)
- [show virtual-chassis protocol interface](#)
- [show virtual-chassis protocol route](#)
- [show virtual-chassis protocol statistics](#)
- [show virtual-chassis login](#)
- [show virtual-chassis](#)

- `show virtual-chassis vc-path`
- `show virtual-chassis vc-port`
- `show virtual-chassis vc-port statistics`

## clear virtual-chassis vc-port statistics

---

<b>Syntax</b>	<code>clear virtual-chassis vc-port statistics</code> <code>&lt;all-members&gt;</code> <code>&lt;interface-name&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member member-id&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. The options <b>all-members</b> and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric.
<b>Description</b>	Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).
<b>Options</b>	<b>none</b> —Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.  <b>all-members</b> —(Optional) Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.  <b>interface-name</b> —(Optional) Clear traffic statistics for the specified VCP.  <b>local</b> —(Optional) Clear traffic statistics for VCPs from the switch or external Routing Engine on which this command is entered.  <b>member member-id</b> —(Optional) Clear traffic statistics for VCPs from the specified member of a Virtual Chassis or VCF.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show virtual-chassis vc-port statistics on page 7024</a></li><li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li><li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li></ul>
<b>List of Sample Output</b>	<a href="#">clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis) on page 6978</a> <a href="#">clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 6979</a> <a href="#">clear virtual-chassis vc-port statistics member 3 on page 6979</a>

### Sample Output

#### clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis)

```
user@switch> clear virtual-chassis vc-port statistics
fpc0:
-----
Statistics cleared
```

**clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis)**

```
user@external-routing-engine> clear virtual-chassis vc-port statistics
```

```
member0:
```

```
-----  
Statistics cleared
```

```
member1:
```

```
-----  
Statistics cleared
```

```
member8:
```

```
-----  
Statistics cleared
```

```
member9:
```

```
-----  
Statistics cleared
```

**clear virtual-chassis vc-port statistics member 3**

```
user@switch> clear virtual-chassis vc-port statistics member 3
```


```
Cleared statistics on member 3
```

## request session member

---

<b>Syntax</b>	<code>request session member <i>member-id</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Start a session with the specified member of a Virtual Chassis or a VCF.
<b>Options</b>	<i>member-id</i> —Member ID for the specific member of the Virtual Chassis or VCF.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">member on page 6958</a></li><li>• <i>Understanding EX Series Virtual Chassis Components</i></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li></ul>

## request virtual-chassis recycle

<b>Syntax</b>	<code>request virtual-chassis recycle member-id <i>member-id</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
<b>Description</b>	<p>Make a previously used member ID available for reassignment.</p> <p>When you remove a member switch from the Virtual Chassis configuration, the master reserves that member ID. To make the member ID available for reassignment, you must use this command.</p>
	<div>  <p><b>NOTE:</b> You must run this command from the Virtual Chassis member in the master role.</p> </div>
<b>Options</b>	<code>member-id <i>member-id</i></code> —Specify the member ID that you want to make available for reassignment to a different member.
<b>Required Privilege Level</b>	system-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request virtual-chassis renumber on page 6982</a></li> <li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 6938</a></li> <li>• <a href="#">Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request virtual-chassis recycle member-id 3 on page 6981</a> <a href="#">request virtual-chassis recycle member-id 1 on page 6981</a>

### Sample Output

`request virtual-chassis recycle member-id 3`

```
user@switch> request virtual-chassis recycle member-id 3
```


### Sample Output

`request virtual-chassis recycle member-id 1`

```
user@external-routing-engine> request virtual-chassis recycle member-id 1
```

## request virtual-chassis renumber

---

<b>Syntax</b>	<code>request virtual-chassis renumber member-id <i>old-member-id</i> new-member-id <i>new-member-id</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
<b>Description</b>	Renumber a member of a Virtual Chassis configuration.
<div> <b>NOTE:</b> You must run this command from the Virtual Chassis member in the master role.</div>	
<b>Options</b>	<code>member-id <i>old-member-id</i></code> —Specify the ID of the member that you wish to renumber. <code>new-member-id <i>new-member-id</i></code> —Specify an unassigned member ID.
<b>Required Privilege Level</b>	system-control
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">request virtual-chassis recycle on page 6981</a></li><li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 6938</a></li><li>• <a href="#">Adding or Replacing a Member Switch or an External Routing Engine in an EX8200 Virtual Chassis (CLI Procedure)</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request virtual-chassis renumber member-id 5 new-member-id 4 on page 6982</a> <a href="#">request virtual-chassis renumber member-id 1 new-member-id 0 on page 6982</a>

### Sample Output

`request virtual-chassis renumber member-id 5 new-member-id 4`

```
user@switch> request virtual-chassis renumber member-id 5 new-member-id 4
```

`request virtual-chassis renumber member-id 1 new-member-id 0`

```
user@external-routing-engine> request virtual-chassis renumber member-id 1 new-member-id 0
```



## request virtual-chassis vc-port

<b>Syntax</b>	<code>request virtual-chassis vc-port set   delete &lt;fpc-slot <i>fpc-slot</i>&gt; pic-slot <i>pic-slot</i> port <i>port-number</i> &lt;member <i>member-id</i>&gt;</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>fpc-slot</b> introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Enable or disable an optical port as a Virtual Chassis port (VCP).</p> <p>If you omit <b>member <i>member-id</i></b>, this command defaults to enabling or disabling the uplink VCP or SFP network port configured as a VCP on the switch where the command is issued.</p> <p>On an EX3300 switch, uplink ports 2 and 3 are configured as VCPs by default. No other uplink ports on any other EX Series switches are configured as VCPs by default.</p> <p>You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.</p>
<b>Options</b>	<p><b>pic-slot <i>pic-slot</i></b>—Number of the PIC slot for the port on the switch.</p> <p><b>port <i>port-number</i></b>—Number of the port that is to be enabled or disabled as a VCP.</p> <p><b>member <i>member-id</i></b>—(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	system-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request virtual-chassis vc-port</a> (dedicated port)</li> <li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li> <li>• <a href="#">show virtual-chassis vc-port statistics on page 7024</a></li> <li>• <a href="#">clear virtual-chassis vc-port statistics on page 6978</a></li> <li>• <a href="#">Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Components</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">request virtual-chassis vc-port set pic-slot 1 port 0 on page 6984</a></p> <p><a href="#">request virtual-chassis vc-port set pic-slot 1 port 1 member 3 on page 6984</a></p> <p><a href="#">request virtual-chassis vc-port delete pic-slot 1 port 1 member 3 on page 6984</a></p>

## Sample Output

**request virtual-chassis vc-port set pic-slot 1 port 0**

user@switch> **request virtual-chassis vc-port set pic-slot 1 port 0**

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

**request virtual-chassis vc-port set pic-slot 1 port 1 member 3**

user@switch> **request virtual-chassis vc-port set pic-slot 1 port 1 member 3**

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

**request virtual-chassis vc-port delete pic-slot 1 port 1 member 3**

user@switch> **request virtual-chassis vc-port delete pic-slot 1 port 1 member 3**

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

## show virtual-chassis active-topology

<b>Syntax</b>	show virtual-chassis active-topology <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the active topology of the Virtual Chassis or VCF with next-hop reachability information.
<b>Options</b>	<p><b>none</b>—Display the active topology of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the active topology of all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the active topology of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> <li>• <i>Understanding EX Series Virtual Chassis Configuration</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis active-topology (EX4200 Virtual Chassis) on page 6986</a> <a href="#">show virtual-chassis active-topology (EX8200 Virtual Chassis) on page 6986</a> <a href="#">show virtual-chassis active-topology (Virtual Chassis Fabric) on page 6987</a>
<b>Output Fields</b>	<a href="#">Table 668 on page 6985</a> lists the output fields for the <b>show virtual-chassis active-topology</b> command. Output fields are listed in the approximate order in which they appear.

**Table 668: show virtual-chassis active-topology Output Fields**

Field Name	Field Description
<b>Destination ID</b>	Specifies the member ID of the destination.
<b>Next-hop</b>	<p>Specifies the member ID and Virtual Chassis port (VCP) of the next hop to which packets for the destination ID are forwarded.</p> <p>The next hop can be more than one device in a VCF.</p>

## Sample Output

### show virtual-chassis active-topology (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis active-topology
1                               1(vcp-1)

2                               1(vcp-1)

3                               1(vcp-1)

4                               1(vcp-1)

5                               8(vcp-0) 1(vcp-1)

6                               8(vcp-0)

7                               8(vcp-0)

8                               8(vcp-0)
```

### show virtual-chassis active-topology (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis active-topology
member0:
```

Destination ID	Next-hop
1	1(vcp-4/0/4.32768)
8	8(vcp-0/0.32768)
9	8(vcp-0/0.32768)

```
member1:
```

Destination ID	Next-hop
0	0(vcp-3/0/4.32768)
8	8(vcp-0/0.32768)
9	8(vcp-0/0.32768)

```
member8:
```

Destination ID	Next-hop
0	0(vcp-1/1.32768)
1	1(vcp-1/2.32768)
9	9(vcp-2/1.32768)

member9:

Destination ID	Next-hop
0	8(vcp-1/2.32768)
1	8(vcp-1/2.32768)
8	8(vcp-1/2.32768)

### show virtual-chassis active-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis active-topology  
fpc0:

Destination ID	Next-hop
1 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
2 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
3 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

fpc1:

Destination ID	Next-hop
0 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
2 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
3 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

fpc2:

Destination ID	Next-hop
0 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
1 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
3 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

## fpc3:

Destination ID	Next-hop
0 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
1 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
2 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

## fpc4:

Destination ID	Next-hop
0	0(vcp-255/0/48.32768)
1	1(vcp-255/0/49.32768)
2	2(vcp-255/0/50.32768)
3	3(vcp-255/0/51.32768)
5 0(vcp-255/0/48.32768)	3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768) 1(vcp-255/0/49.32768)
6 0(vcp-255/0/48.32768)	3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768) 1(vcp-255/0/49.32768)

## fpc5:

Destination ID	Next-hop
0	0(vcp-255/0/48.32768)

1	1(vcp-255/0/49.32768)	
2	2(vcp-255/0/50.32768)	
3	3(vcp-255/0/51.32768)	
4	3(vcp-255/0/51.32768)	2(vcp-255/0/50.32768)
0(vcp-255/0/48.32768)	1(vcp-255/0/49.32768)	
6	3(vcp-255/0/51.32768)	2(vcp-255/0/50.32768)
0(vcp-255/0/48.32768)	1(vcp-255/0/49.32768)	

fpc6:

Destination ID	Next-hop
0	0(vcp-255/0/0.32768)
1	1(vcp-255/0/1.32768)
2	2(vcp-255/0/2.32768)
3	3(vcp-255/0/3.32768)
4	3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768)
0(vcp-255/0/0.32768)	1(vcp-255/0/1.32768)
5	3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768)
0(vcp-255/0/0.32768)	1(vcp-255/0/1.32768)

## show virtual-chassis device-topology

<b>Syntax</b>	show virtual-chassis device-topology <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the device topology—the member and system IDs, the VCP numbers, and device status—for all hardware devices in the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the device topology for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding EX Series Virtual Chassis Port Link Aggregation</i></li> <li>• <i>Understanding EX8200 Virtual Chassis Topologies</i></li> </ul>
<b>Output Fields</b>	<a href="#">Table 669 on page 6990</a> lists the output fields for the <b>show virtual-chassis device-topology</b> command. Output fields are listed in the approximate order in which they appear.

**Table 669: show virtual-chassis device-topology Output Fields**

Field Name	Field Description
<b>Member</b>	Assigned member ID.
<b>Device</b>	Assigned device ID.  For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
<b>Status</b>	The status of the device within the Virtual Chassis or VCF. Outputs include: <ul style="list-style-type: none"> <li>• <b>Prsnt</b>—Device is currently connected to and participating in the Virtual Chassis or VCF.</li> <li>• <b>NotPrsnt</b>—Device is assigned but is not currently connected.</li> </ul>



Table 669: show virtual-chassis device-topology Output Fields (*continued*)

Field Name	Field Description
<b>System ID</b>	System ID of the device.  The system ID of the device is the device's MAC address.
<b>Member (Neighbor List)</b>	Assigned member ID of the neighbor device.
<b>Device (Neighbor List)</b>	Assigned device ID of the neighbor device.  For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
<b>Interface (Neighbor List)</b>	The interface connecting the device to the neighbor.

## Sample Output

### show virtual-chassis device-topology

```
user@switch> show virtual-chassis device-topology
```

```
member0:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0        0        Prsnt   0021.59f7.d000  8        8        vcp-0/0
                                1        1        vcp-4/0/1
1        1        Prsnt   0026.888d.6800  8        8        vcp-0/0
                                9        9        vcp-0/1
                                0        0        vcp-3/0/4
8        8        Prsnt   0000.4a75.9b7c  9        9        vcp-1/0
                                0        0        vcp-1/1
                                1        1        vcp-1/2
9        9        Prsnt   0000.73e9.9a57  8        8        vcp-1/0
                                1        1        vcp-1/1
```

```
member1:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0        0        Prsnt   0021.59f7.d000  8        8        vcp-0/0
                                1        1        vcp-4/0/1
1        1        Prsnt   0026.888d.6800  8        8        vcp-0/0
                                9        9        vcp-0/1
                                0        0        vcp-3/0/4
8        8        Prsnt   0000.4a75.9b7c  9        9        vcp-1/0
                                0        0        vcp-1/1
                                1        1        vcp-1/2
9        9        Prsnt   0000.73e9.9a57  8        8        vcp-1/0
                                1        1        vcp-1/1
```

```
member8:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
```

0	0	Prsnt	0021.59f7.d000	8	8	vcp-0/0
				1	1	vcp-4/0/1
1	1	Prsnt	0026.888d.6800	8	8	vcp-0/0
				9	9	vcp-0/1
				0	0	vcp-3/0/4
8	8	Prsnt	0000.4a75.9b7c	9	9	vcp-1/0
				0	0	vcp-1/1
				1	1	vcp-1/2
9	9	Prsnt	0000.73e9.9a57	8	8	vcp-1/0
				1	1	vcp-1/1

member9:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	0021.59f7.d000	8	8	vcp-0/0
				1	1	vcp-4/0/1
1	1	Prsnt	0026.888d.6800	8	8	vcp-0/0
				9	9	vcp-0/1
				0	0	vcp-3/0/4
8	8	Prsnt	0000.4a75.9b7c	9	9	vcp-1/0
				0	0	vcp-1/1
				1	1	vcp-1/2
9	9	Prsnt	0000.73e9.9a57	8	8	vcp-1/0
				1	1	vcp-1/1

#### show virtual-chassis device-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis device-topology  
fpc0:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
				1	1	vcp-255/0/49
5	5	Prsnt	100e.7eb5.80c0	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc1:

Neighbor List

Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
				1	1	vcp-255/0/49
5	5	Prsnt	100e.7eb5.80c0	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc2:

Neighbor List						
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
				1	1	vcp-255/0/49
5	5	Prsnt	100e.7eb5.80c0	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc3:

Neighbor List						
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3

1	1	Prsnt	100e.7eb8.3a40	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
2	2	Prsnt	100e.7eb5.d700	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
3	3	Prsnt	100e.7eb5.c440	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
4	4	Prsnt	100e.7eb5.7e40	6	6	vcp-255/0/1
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
5	5	Prsnt	100e.7eb5.80c0	0	0	vcp-255/0/48
				1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
6	6	Prsnt	100e.7eb6.3b00	1	1	vcp-255/0/49
				0	0	vcp-255/0/48
				3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc4:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
5	5	Prsnt	100e.7eb5.80c0	1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
6	6	Prsnt	100e.7eb6.3b00	0	0	vcp-255/0/48
				3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc5:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3

2	2	Prsnt	100e.7eb5.d700	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
3	3	Prsnt	100e.7eb5.c440	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
4	4	Prsnt	100e.7eb5.7e40	6	6	vcp-255/0/1
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
5	5	Prsnt	100e.7eb5.80c0	1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
6	6	Prsnt	100e.7eb6.3b00	0	0	vcp-255/0/48
				3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc6:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
				1	1	vcp-255/0/49
5	5	Prsnt	100e.7eb5.80c0	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

## show virtual-chassis protocol adjacency

---

<b>Syntax</b>	<code>show virtual-chassis protocol adjacency</code> <code>&lt;brief   detail   extensive&gt;</code> <code>&lt;all-members&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;system-id&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the Virtual Chassis Control Protocol (VCCP) adjacency statistics in the Virtual Chassis or VCF for all hardware devices.
<b>Options</b>	<p><b>none</b>—Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP adjacency statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p> <p><b>system-id</b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Understanding EX Series Virtual Chassis Port Link Aggregation</i></li><li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li></ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol adjacency on page 6997</a> <a href="#">show virtual-chassis protocol adjacency detail on page 6998</a>
<b>Output Fields</b>	<a href="#">Table 670 on page 6997</a> lists the output fields for the <b>show virtual-chassis protocol adjacency</b> command. Output fields are listed in the approximate order in which they appear.

Table 670: show virtual-chassis protocol adjacency Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the Virtual Chassis port (VCP) interface.	All levels
<b>System</b>	The MAC address of the device on the receiving side of the VCP link.	All levels
<b>State</b>	State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> <li>• <b>New</b>—The link is new.</li> <li>• <b>One-way</b>—The link is transmitting traffic in one direction.</li> <li>• <b>Initializing</b>—The link is initializing.</li> <li>• <b>Rejected</b>—The link is rejected.</li> </ul>	All levels
<b>Hold, Expires in</b>	Remaining holdtime of the adjacency.	All levels
<b>Priority</b>	Priority to become the designated intermediary system.	detail
<b>Up/Down Transitions</b>	Count of adjacency status transition changes from up to down or down to up.	detail
<b>Last transition</b>	Time of the last up/down transition.	detail

## Sample Output

### show virtual-chassis protocol adjacency

```
user@switch> show virtual-chassis protocol adjacency
```

```
member0:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          57
vcp-0/1.32768  0000.4a75.9b7c Up          59
vcp-4/0/1.32768 0026.888d.6800 Up          57
```

```
member1:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          58
vcp-0/1.32768  0000.73e9.9a57 Up          59
vcp-3/0/4.32768 0021.59f7.d000 Up          58
```

```
member8:
```

```
-----
Interface      System      State      Hold (secs)
vcp-1/0.32768  0000.73e9.9a57 Up          58
vcp-1/1.32768  0021.59f7.d000 Up          58
vcp-1/2.32768  0026.888d.6800 Up          59
vcp-2/0.32768  0021.59f7.d000 Up          59
```

```
member9:
```

```
-----
Interface      System      State      Hold (secs)
```

vcp-1/0.32768	0000.4a75.9b7c Up	58
vcp-1/1.32768	0026.888d.6800 Up	59

### show virtual-chassis protocol adjacency detail

```
user@switch> show virtual-chassis protocol adjacency detail
```

```
member0:
```

-----

```
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0000.4a75.9b7c
  interface-name: vcp-0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0026.888d.6800
  interface-name: vcp-4/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member1:
```

-----

```
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0000.73e9.9a57
  interface-name: vcp-0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

```
0021.59f7.d000
  interface-name: vcp-3/0/4.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member8:
```

-----

```
0000.73e9.9a57
  interface-name: vcp-1/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-1/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0026.888d.6800
  interface-name: vcp-1/2.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-2/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
member9:
```

-----

```
0000.4a75.9b7c
  interface-name: vcp-1/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```



```
0026.888d.6800
  interface-name: vcp-1/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

## show virtual-chassis protocol database

<b>Syntax</b>	show virtual-chassis protocol database <brief   detail   extensive> <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the Virtual Chassis Control Protocol (VCCP) database statistics for all hardware devices within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option. The <b>extensive</b> option provides all output and is most useful for customer support personnel.</p> <p><b>all-members</b>—(Optional) Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP database statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> <li>• <i>Understanding EX Series Virtual Chassis Components</i></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol database on page 7001</a> <a href="#">show virtual-chassis protocol database detail on page 7002</a>
<b>Output Fields</b>	<a href="#">Table 671 on page 7000</a> lists the output fields for the <b>show virtual-chassis protocol database</b> command. Output fields are listed in the approximate order in which they appear.

Table 671: show virtual-chassis protocol database Output Fields

Field Name	Field Description	Level of Output
LSP ID	Link-state protocol (LSP) data unit identifier.	All levels

Table 671: show virtual-chassis protocol database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Sequence</b>	Sequence number of the LSP.	All levels
<b>Checksum</b>	Checksum value of the LSP.	All levels
<b>Lifetime</b>	Remaining lifetime of the LSP, in seconds.	All levels
<b>Neighbor</b>	MAC address of the neighbor on the advertising system.	detail
<b>Interface</b>	Virtual Chassis port (VCP) interface name.	detail
<b>Metric</b>	Metric of the prefix or neighbor.	detail

The **extensive** output was omitted from this list. The **extensive** output is useful for customer support personnel only.

## Sample Output

### show virtual-chassis protocol database

```
user@switch> show virtual-chassis protocol database
```

```
member0:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   113
0021.59f7.d000.00-00  0x16882  0x3993   118
0026.888d.6800.00-00  0x1691f  0x82b7   116
  4 LSPs
```

```
member1:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x1691f  0x82b7   118
  4 LSPs
```

```
member8:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   118
0000.73e9.9a57.00-00  0xf361   0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x16920  0xa335   116
  4 LSPs
```

```
member9:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   116
0021.59f7.d000.00-00  0x16883  0x289    114
```

```
0026.888d.6800.00-00      0x16920   0xa335      116
4 LSPs
```

### show virtual-chassis protocol database detail

```
user@switch> show virtual-chassis protocol database detail
member0:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 114 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 118 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member1:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member8:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 118 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
```

```

Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15

```

member9:

```

-----

0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 116 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 113 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694f, Checksum: 0xa61a, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15

```

## show virtual-chassis protocol interface

---

<b>Syntax</b>	<code>show virtual-chassis protocol interface</code> <code>&lt;brief   detail&gt;</code> <code>&lt;all-members&gt;</code> <code>&lt;interface-name&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member member-id&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display information about Virtual Chassis Control Protocol (VCCP) statistics for VCCP-enabled interfaces within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display the VCCP interface statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail</b> —(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option.</p> <p><b>all-members</b>—(Optional) Display VCCP interface statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP interface statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP interface statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>EX Series Virtual Chassis Overview</i></li><li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li><li>• <i>Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li><li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li></ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol interface on page 7005</a>
<b>Output Fields</b>	<a href="#">Table 672 on page 7005</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 672: show virtual-chassis protocol interface Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the VCP.	All levels
<b>State</b>	State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> </ul>	All levels
<b>Metric</b>	Metric of the prefix or neighbor.	All levels

## Sample Output

### show virtual-chassis protocol interface

```
user@switch> show virtual-chassis protocol interface
```

```
member0:
```

```
-----
```

```
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Up	150
vcp-0/1.32768	Up	150
vcp-4/0/1.32768	Up	15
vcp-4/0/7.32768	Down	15

```
member1:
```

```
-----
```

```
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Up	150
vcp-0/1.32768	Up	150
vcp-3/0/4.32768	Up	15

```
member8:
```

```
-----
```

```
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Down	150
vcp-1/0.32768	Up	150
vcp-1/1.32768	Up	150
vcp-1/2.32768	Up	150
vcp-1/3.32768	Down	150
vcp-2/0.32768	Up	150
vcp-2/1.32768	Down	150
vcp-2/2.32768	Down	150
vcp-2/3.32768	Down	150

```
member9:
```

```
-----
```

```
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Down	150
vcp-1/0.32768	Up	150
vcp-1/1.32768	Up	150
vcp-1/2.32768	Down	150
vcp-1/3.32768	Down	150





## show virtual-chassis protocol route

<b>Syntax</b>	show virtual-chassis protocol route <all-members> <destination-id> <local> <member member-id>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the unicast and multicast Virtual Chassis Control Protocol (VCCP) routing tables within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display the unicast and multicast routing tables for all members of the Virtual Chassis.</p> <p><b>all-members</b>—(Optional) Display the unicast and multicast routing tables for all members of the Virtual Chassis or VCF.</p> <p><b>destination-id</b>—(Optional) Display the unicast and multicast routing tables to the specified destination member ID for each member of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the unicast and multicast routing tables on the device where this command is entered.</p> <p><b>member member-id</b>—(Optional) Display the unicast and multicast routing tables for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol route on page 7008</a>
<b>Output Fields</b>	<a href="#">Table 673 on page 7007</a> lists the output fields for the <b>show virtual-chassis protocol route</b> command. Output fields are listed in the approximate order in which they appear.

**Table 673: show virtual-chassis protocol route Output Fields**

Field Name	Field Description
<b>Dev</b>	MAC address of the member storing the VCCP routing table.
<b>Version</b>	Version of the shortest-path-first algorithm that generated the routing table.

Table 673: show virtual-chassis protocol route Output Fields (*continued*)

Field Name	Field Description
<b>System ID</b>	MAC address of the device.
<b>Version</b>	Version of the shortest-path-first (SPF) algorithm that generated the route.
<b>Metric</b>	The metric number to get to that device.
<b>Interface</b>	Name of the Virtual Chassis port (VCP) interface connecting the devices.
<b>Via</b>	MAC address of the next-hop device, if applicable.

## Sample Output

### show virtual-chassis protocol route

```

user@switch> show virtual-chassis protocol route
member0:
-----
Dev 0021.59f7.d000 ucast routing table           Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21      150 vcp-0/1.32768 0000.4a75.9b7c
0000.73e9.9a57    21      165 vcp-4/0/1.32768 0026.888d.6800
0021.59f7.d000    21         0
0026.888d.6800    21      15 vcp-4/0/1.32768 0026.888d.6800

Dev 0021.59f7.d000 mcast routing table           Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21
0000.73e9.9a57    21
0021.59f7.d000    21          vcp-4/0/1.32768
                   vcp-0/1.32768
0026.888d.6800    21

member1:
-----
Dev 0026.888d.6800 ucast routing table           Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25      150 vcp-0/0.32768 0000.4a75.9b7c
0000.73e9.9a57    25      150 vcp-0/1.32768 0000.73e9.9a57
0021.59f7.d000    25      15 vcp-3/0/4.32768 0021.59f7.d000
0026.888d.6800    25         0

Dev 0026.888d.6800 mcast routing table           Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25
0000.73e9.9a57    25          vcp-3/0/4.32768
0021.59f7.d000    25          vcp-0/1.32768
0026.888d.6800    25          vcp-3/0/4.32768
                   vcp-0/0.32768

```

vcp-0/1.32768

member8:

-----

Dev 0000.4a75.9b7c ucast routing table                      Current version: 39

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	39	0		
0000.73e9.9a57	39	150	vcp-1/0.32768	0000.73e9.9a57
0021.59f7.d000	39	150	vcp-2/0.32768	0021.59f7.d000
0026.888d.6800	39	150	vcp-1/2.32768	0026.888d.6800

Dev 0000.4a75.9b7c mcast routing table                      Current version: 39

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	39		vcp-1/0.32768	
			vcp-2/0.32768	
			vcp-1/2.32768	
0000.73e9.9a57	39			
0021.59f7.d000	39			
0026.888d.6800	39			

member9:

-----

Dev 0000.73e9.9a57 ucast routing table                      Current version: 31

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	31	150	vcp-1/0.32768	0000.4a75.9b7c
0000.73e9.9a57	31	0		
0021.59f7.d000	31	165	vcp-1/1.32768	0026.888d.6800
0026.888d.6800	31	150	vcp-1/1.32768	0026.888d.6800

Dev 0000.73e9.9a57 mcast routing table                      Current version: 31

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	31			
0000.73e9.9a57	31		vcp-1/0.32768	
			vcp-1/1.32768	
0021.59f7.d000	31			
0026.888d.6800	31			

## show virtual-chassis protocol statistics

<b>Syntax</b>	show virtual-chassis protocol statistics <all-members> <interface-name> <local> <member member-id>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the Virtual Chassis Control Protocol (VCCP) statistics for all hardware devices within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol statistics on page 7011</a>
<b>Output Fields</b>	<a href="#">Table 674 on page 7010</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 674: show virtual-chassis protocol statistics Output Fields**

Field Name	Field Description
<b>PDU type</b>	Protocol data unit type.
<b>Received</b>	Number of PDUs received since VCCP started or since the statistics were set to zero.
<b>Processed</b>	Number of PDUs received minus the number of PDUs dropped.

Table 674: show virtual-chassis protocol statistics Output Fields (*continued*)

Field Name	Field Description
<b>Drops</b>	Number of PDUs dropped.
<b>Sent</b>	Number of PDUs transmitted since VCCP started or since the statistics were set to zero.
<b>Rexmit</b>	Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.
<b>Total Packets Received</b>	Number of PDUs received since VCCP started or since the statistics were set to zero.
<b>Total Packets Sent</b>	Number of PDUs sent since VCCP started or since the statistics were set to zero.
<b>LSP queue length</b>	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
<b>SPF runs</b>	Number of shortest-path-first (SPF) calculations that have been performed.
<b>Fragments Rebuilt</b>	Number of link-state PDU fragments that the local system has computed.
<b>LSP Regenerations</b>	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
<b>Purges initiated</b>	Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.

## Sample Output

### show virtual-chassis protocol statistics

```

user@switch> show virtual-chassis protocol statistics
member0:
-----
IS-IS statistics for 0021.59f7.d000:
PDU type      Received    Processed      Drops      Sent      Rexmit
LSP            8166        8166           0         4551         0
HELLO          1659        1659           0         1693         0
CSNP             2            2             0            3         0
PSNP           1909        1909           0         2293         0
Unknown         0            0             0            0         0
Totals        11736       11736           0         8540         0

Total packets received: 11736 Sent: 8540

LSP queue length: 0 Drops: 0
SPF runs: 9
Fragments rebuilt: 1640
LSP regenerations: 1
Purges initiated: 0

member1:
-----
IS-IS statistics for 0026.888d.6800:

```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	10909	10909	0	12088	0
HELLO	1877	1877	0	2251	0
CSNP	3	3	0	3	0
PSNP	3846	3846	0	3732	0
Unknown	0	0	0	0	0
Totals	16635	16635	0	18074	0

Total packets received: 16635 Sent: 18074

LSP queue length: 0 Drops: 0  
SPF runs: 13  
Fragments rebuilt: 1871  
LSP regenerations: 2  
Purges initiated: 0

member8:

-----  
IS-IS statistics for 0000.4a75.9b7c:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	7935	7935	0	14865	0
HELLO	2695	2695	0	7124	0
CSNP	4	4	0	4	0
PSNP	4398	4398	0	3666	0
Unknown	0	0	0	0	0
Totals	15032	15032	0	25659	0

Total packets received: 15032 Sent: 25659

LSP queue length: 0 Drops: 0  
SPF runs: 26  
Fragments rebuilt: 2666  
LSP regenerations: 4  
Purges initiated: 0

member9:

-----  
IS-IS statistics for 0000.73e9.9a57:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	10800	10800	0	6327	0
HELLO	1492	1492	0	2356	0
CSNP	2	2	0	2	0
PSNP	2683	2683	0	3149	0
Unknown	0	0	0	0	0
Totals	14977	14977	0	11834	0

Total packets received: 14977 Sent: 11834

LSP queue length: 0 Drops: 0  
SPF runs: 19  
Fragments rebuilt: 1510  
LSP regenerations: 6  
Purges initiated: 0

## show virtual-chassis login

<b>Syntax</b>	<b>show virtual-chassis login</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Supply the address of the host that logged into the Virtual Chassis or VCF, or identify the location of the member switch that redirected the current session to a different member switch.</p> <p>You might need this information for tracing or troubleshooting purposes.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request session member on page 6980</a></li> <li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 6919</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis login (Direct Login to the Master Console Port) on page 7013</a></p> <p><a href="#">show virtual-chassis login (Backup Console Session Redirected to the Master Console Port) on page 7013</a></p>

### Sample Output

#### show virtual-chassis login (Direct Login to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host 248.1.2.3
```

#### show virtual-chassis login (Backup Console Session Redirected to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

## show virtual-chassis

<b>Syntax</b>	<b>show virtual-chassis</b> <b>&lt;status&gt;</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p> <p><b>Fabric ID</b>, <b>Fabric Mode</b>, and <b>Route Mode</b> output fields introduced in Junos OS Release 13.2X51-D20.</p> <p><b>Alias-Name</b> output field introduced in Junos OS Release 14.1X53-D10.</p>
<b>Description</b>	Display information about all members of the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display information about all Virtual Chassis or VCF member devices.</p> <p><b>status</b>—Same output as for <b>show virtual-chassis</b>.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis active-topology on page 6985</a></li> <li>• <a href="#">show virtual-chassis protocol adjacency on page 6996</a></li> <li>• <a href="#">show virtual-chassis vc-path on page 7018</a></li> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis (EX4200 Virtual Chassis) on page 7016</a></p> <p><a href="#">show virtual-chassis (EX8200 Virtual Chassis) on page 7016</a></p> <p><a href="#">show virtual-chassis (Virtual Chassis Fabric) on page 7017</a></p>
<b>Output Fields</b>	<p><a href="#">Table 675 on page 7014</a> lists the output fields for the <b>show virtual-chassis</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 675: show virtual-chassis Output Fields**

Field Name	Field Description
<b>Fabric ID</b>	Assigned ID used to identify the VCF.
<b>Fabric Mode</b>	Mode of the VCF: Enabled, Disabled, or Mixed.
<b>Virtual Chassis ID</b>	Assigned ID that applies to the entire Virtual Chassis or VCF.
<b>Virtual Chassis Mode</b>	Mode of the Virtual Chassis or VCF: Enabled, Disabled, or Mixed.



Table 675: show virtual-chassis Output Fields (*continued*)

Field Name	Field Description
<b>Member ID</b>	Assigned member ID and FPC: <ul style="list-style-type: none"> <li>On all EX Series Virtual Chassis except EX8200 Virtual Chassis, and on a VCF, the FPC number refers to the member ID assigned to the switch.</li> <li>On EX8200 Virtual Chassis, member IDs are numbered 0 through 9. The FPC number indicates the slot number of the line card within the Virtual Chassis. The FPC number on member 0 is always 0 through 15. The FPC number on member 1 is always 16 through 31. The FPC number on member 2 is always 32 through 47; and so on for the members.</li> </ul>
<b>Status</b>	For a nonprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is currently connected to the Virtual Chassis or VCF configuration.</li> <li><b>NotPrsnt</b> for a member ID that has been assigned but is not currently connected.</li> </ul> For a preprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis or VCF.</li> <li><b>Unprvsnd</b> for a member that is interconnected with the Virtual Chassis or VCF configuration but is not specified in the preprovisioned configuration file.</li> </ul>
<b>Serial No</b>	Serial number of the member device.
<b>Alias-Name</b>	The user-configured alias of the member device.  The <b>Alias-Name</b> field appears only if an alias has been configured for at least one device in the Virtual Chassis or VCF. Aliases are configured using the <b>alias-name</b> statement in the <code>[edit virtual-chassis aliases serial-number serial-number]</code> hierarchy.
<b>Model</b>	Model number of the member device.
<b>Mastership Priority</b>	Mastership priority value of the member device.
<b>Role</b>	Role of the member device: master, backup, or linecard.
<b>Mixed Mode</b>	Mixed mode configuration status: <ul style="list-style-type: none"> <li><b>Y</b> for a member device configured in mixed mode.</li> <li><b>N</b> for a member device not configured in mixed mode.</li> <li><b>NA</b> for a member device that cannot be configured in mixed mode.</li> </ul>
<b>Route Mode</b>	The route mode of the member device: fabric (F) or Virtual Chassis (V).
<b>Location</b>	Location of the member device.  If this field is empty, the location field was not set for the device.
<b>Neighbor List</b>	Member ID of the neighbor member to which this member's Virtual Chassis port (VCP) is connected.

## Sample Output

### show virtual-chassis (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis
Virtual Chassis ID: 0019.e250.47a0
Virtual Chassis Mode: Enabled

```

Member ID	Status	Serial No	Model	Mastership priority	Role	Mixed Mode	Neighbor List ID	Interface
0 (FPC 0)	Prsnt	AK0207360276	ex4200-24t	249	Master*	N	8	vcp-0
							1	vcp-1
1 (FPC 1)	Prsnt	AK0207360281	ex4200-24t	248	Backup	N	0	vcp-0
							2	vcp-1
2 (FPC 2)	Prsnt	AJ0207391130	ex4200-48p	247	Linecard	N	1	vcp-0
							3	vcp-1
3 (FPC 3)	Prsnt	AK0207360280	ex4200-24t	246	Linecard	N	2	vcp-0
							4	vcp-1
4 (FPC 4)	Prsnt	AJ0207391113	ex4200-48p	245	Linecard	N	3	vcp-0
							5	vcp-1
5 (FPC 5)	Prsnt	BP0207452204	ex4200-48t	244	Linecard	N	4	vcp-0
							6	vcp-1
6 (FPC 6)	Prsnt	BP0207452222	ex4200-48t	243	Linecard	N	5	vcp-0
							7	vcp-1
7 (FPC 7)	Prsnt	BR0207432028	ex4200-24f	242	Linecard	N	6	vcp-0
							8	vcp-1
8 (FPC 8)	Prsnt	BR0207431996	ex4200-24f	241	Linecard	N	7	vcp-0
							0	vcp-1

Member ID for next new member: 9 (FPC 9)

### show virtual-chassis (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis
Virtual Chassis ID: c806.0842.de51
Virtual Chassis Mode: Enabled

```

Member ID	Status	Serial No	Model	Mastership priority	Role	Neighbor List ID	Interface
0 (FPC 0-15)	Prsnt	BA0908380001	ex8216	0	Linecard	8	vcp-0/0
						8	vcp-0/1
						1	vcp-4/0/4
1 (FPC 16-31)	Prsnt	BT0909411634	ex8208	0	Linecard	8	vcp-0/0
						0	vcp-3/0/4
8 (FPC 128-143)	Prsnt	062009000021	ex-xre	128	Master	9	vcp-1/0
						1	vcp-1/2

```

9 (FPC 144-159) Prsnt 062009000022 ex-xre 128 Backup*
9 vcp-1/3
0 vcp-2/0
9 vcp-2/1
0 vcp-1/1
8 vcp-1/0
8 vcp-1/2
8 vcp-1/3
8 vcp-1/3

```

### show virtual-chassis (Virtual Chassis Fabric)

```

user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0282.5fa0.3f08
Fabric Mode: Enabled

```

List	Member ID	Status	Serial No	Model	Mstr prio	Role	Mixed	Route	Neighbor
Interface									
0 (FPC 0)	Prsnt	AB3112430001	qfx5100-48s	129	Master*	N	F	3	
vcp-255/1/0									2
vcp-255/1/1									4
vcp-255/1/2									4
vcp-255/1/3									4
1 (FPC 1)	Prsnt	AB3112230001	qfx5100-48s	129	Backup	N	F	3	
vcp-255/1/0									2
vcp-255/1/1									4
vcp-255/1/2									4
vcp-255/1/3									4
2 (FPC 2)	Prsnt	AB3112460011	qfx5100-48s	0	Linecard	N	F	1	
vcp-255/1/0									0
vcp-255/1/1									0
3 (FPC 3)	Prsnt	AB3112460011	qfx5100-48s	0	Linecard	N	F	1	
vcp-255/1/0									0
vcp-255/1/1									0
4 (FPC 4)	Prsnt	AB3112430011	qfx5100-48s	0	Linecard	N	F	1	
vcp-255/1/0									0
vcp-255/1/1									0

## show virtual-chassis vc-path

<b>Syntax</b>	<b>show virtual-chassis vc-path source-interface <i>interface-name</i> destination-interface <i>interface-name</i></b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.6 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
<b>Description</b>	Show the path a packet takes when going from a source interface to a destination interface in a Virtual Chassis configuration.
<b>Options</b>	<b>source-interface <i>interface-name</i></b> —Name of the interface from which the packet originates <b>destination-interface <i>interface-name</i></b> —Name of the interface to which the packet is delivered
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> <li>• <i>Understanding EX Series Virtual Chassis Configuration</i></li> <li>• <i>EX8200 Virtual Chassis Overview</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis vc-path source-interface destination-interface on page 7019</a>
<b>Output Fields</b>	<a href="#">Table 676 on page 7018</a> lists the output fields for the <b>show virtual-chassis vc-path</b> command. Output fields are listed in the approximate order in which they appear.

**Table 676: show virtual-chassis vc-path Output Fields**

Field Name	Field Description
<b>Hop</b>	The number of hops between the source and destination interfaces.
<b>Member</b>	The Virtual Chassis ID of the member switch that contains the Packet Forwarding Engine for each intermediate hop.
<b>PFE-Device</b>	The number of the Packet Forwarding Engine in each Virtual Chassis member through which a packet passes. Each Packet Forwarding Engine is the next hop of the preceding Packet Forwarding Engine.
<b>Interface</b>	The name of the interface through which the Packet Forwarding Engines are connected. The interface for the first hop is always the source interface and the interface for the last hop is always the destination interface. For intermediate hops, the <b>Interface</b> field denotes the Packet Forwarding Engines through which the packet passes on its way to the next hop.

## Sample Output

show virtual-chassis vc-path source-interface destination-interface

```
user@switch> show virtual-chassis vc-path source-interface ge-0/0/0 destination-interface
ge-1/0/1
vc-path from ge-0/0/0 to ge-1/0/1
Hop      Member    PFE-Device    Interface
0         0          1              ge-0/0/0
1         0          0              internal-1/24
2         1          3              vcp-0
3         1          4              ge-1/0/1
```

## show virtual-chassis vc-port

<b>Syntax</b>	show virtual-chassis vc-port <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the status of the Virtual Chassis ports (VCPs), including both the dedicated VCPs and the uplink ports configured as VCPs.
<b>Options</b>	<p><b>none</b>—Display the operational status of all VCPs of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the operational status of all VCPs on all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the operational status of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the operational status of all VCPs for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 7024</a></li> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> <li>• <i>Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis vc-port (EX4200 Virtual Chassis) on page 7022</a> <a href="#">show virtual-chassis vc-port (EX8200 Virtual Chassis) on page 7022</a> <a href="#">show virtual-chassis vc-port all-members on page 7023</a>
<b>Output Fields</b>	Table 677 on page 7020 lists the output fields for the <b>show virtual-chassis vc-port</b> command. Output fields are listed in the approximate order in which they appear.

Table 677: show virtual-chassis vc-port Output Fields

Field Name	Field Description
<i>fpcnumber</i>	The FPC number is the same as the member ID.

Table 677: show virtual-chassis vc-port Output Fields (*continued*)

Field Name	Field Description
Interface or PIC/Port	<p>VCP name.</p> <ul style="list-style-type: none"> <li>The dedicated VCPs in an EX4200 or EX4500 Virtual Chassis are <b>vcp-0</b> and <b>vcp-1</b>. The dedicated VCPs in an EX4550 Virtual Chassis are <b>VCP-1/0</b>, <b>VCP-1/1</b>, <b>VCP-2/0</b>, and <b>VCP-2/1</b>.</li> <li>Optical ports set as VCPs are named <b>1/0</b> and <b>1/1</b>, representing the PIC number and the port number.</li> <li>The native VCP (port 0) on an XRE200 External Routing Engine in an EX8200 Virtual Chassis is named <b>vcp-0</b>.</li> <li>The VCPs on each Virtual Chassis Control Interface (VCCI) module in an XRE200 External Routing Engine are named using the <b>vcp-slot-number/port-number</b> convention; for instance, <b>vcp-1/0</b>.</li> <li>The VCPs on EX8200 member switches are named using the <b>vcp-slot-number/pic-number/interface-number</b> convention; for instance, <b>vcp-3/0/2</b>.</li> <li>A <b>255</b> as the first number in your port number indicates that your VCP is part of a Link Aggregation group (LAG) bundle. For instance, a display of <b>vcp-255/1/0</b> indicates that the dedicated VCP named <b>vcp-1/0</b> is part of a LAG bundle. A display of <b>vcp-255/1/0</b> indicates that an uplink port that was previously named <b>xe-0/1/0</b> is now part of a VCP LAG bundle.</li> </ul>
Type	<p>Type of VCP:</p> <ul style="list-style-type: none"> <li><b>Dedicated</b>—The rear panel VCP on an EX4200, EX4500, or EX4550 switch, or any VCP link connected to an XRE200 External Routing Engine in an EX8200 Virtual Chassis.</li> <li><b>Configured</b>—Optical port configured as a VCP.</li> <li><b>Auto-Configured</b>—Optical port autoconfigured as a VCP.</li> </ul> <p>See <i>Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)</i> or <i>Setting a 10-Gigabit Ethernet Port as a Virtual Chassis Port in an EX8200 Virtual Chassis (CLI Procedure)</i> for information about configuring VCPs.</p>
Trunk ID	<p>A positive-number ID assigned to a link aggregation group (LAG) formed by the Virtual Chassis. The trunk ID value is –1 if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member in an EX4200 or EX4500 Virtual Chassis.</p> <p>Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater.</p> <p>The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.</p>
Status	<p>Interface status:</p> <ul style="list-style-type: none"> <li><b>absent</b>—Interface is not a VCP link.</li> <li><b>down</b>—VCP link is down.</li> <li><b>up</b>—VCP link is up.</li> </ul>
Speed (mbps)	Speed of the interface in megabits per second.
Neighbor ID/Interface	The Virtual Chassis member ID and interface of a VCP on a member that is connected to the interface or PIC/Port field in the same row as this interface.

## Sample Output

### show virtual-chassis vc-port (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis vc-port
```

```
fpc0:
```

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	1	Up	32000	1	vcp-1
vcp-1	Dedicated	2	Up	32000	0	vcp-0
1/0	Auto-Configured	3	Up	1000	2	vcp-255/1/0
1/0	Auto-Configured	3	Up	1000	2	vcp-255/1/1

### show virtual-chassis vc-port (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port
```

```
member0:
```

Interface or Slot/PIC/Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0/0	Dedicated	-1	Up	1000	8	vcp-1/1
vcp-0/1	Dedicated	-1	Up	1000	8	vcp-2/0
4/0/4	Configured	-1	Up	10000	1	vcp-3/0/4
4/0/7	Configured	-1	Down	10000		
4/0/3	Configured		Absent			
4/0/2	Configured		Absent			
4/0/5	Configured		Absent			
4/0/6	Configured		Absent			
4/0/1	Configured		Absent			
4/0/0	Configured		Absent			

```
member1:
```

Interface or Slot/PIC/Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0/0	Dedicated	-1	Up	1000	8	vcp-1/2
3/0/0	Configured	-1	Down	10000		
3/0/1	Configured	-1	Down	10000		
3/0/4	Configured	-1	Up	10000	0	vcp-4/0/4
3/0/5	Configured		Absent			
4/0/5	Configured		Absent			
4/0/4	Configured		Absent			

```
member8:
```

Interface or Slot/PIC/Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0/0	Dedicated	-1	Down	1000		
vcp-1/0	Dedicated	-1	Up	1000	9	vcp-1/0
vcp-1/1	Dedicated	-1	Up	1000	0	vcp-0/0
vcp-1/2	Dedicated	-1	Up	1000	1	vcp-0/0
vcp-1/3	Dedicated	-1	Up	1000	9	vcp-1/3
vcp-2/0	Dedicated	-1	Up	1000	0	vcp-0/1
vcp-2/1	Dedicated	-1	Up	1000	9	vcp-1/2
vcp-2/2	Dedicated	-1	Down	1000		



```
vcp-2/3          Dedicated          -1    Down          1000
```

```
member9:
```

```
-----
Interface      Type          Trunk  Status      Speed      Neighbor
or             or            ID      (mbps)      ID  Interface
Slot/PIC/Port
vcp-0/0        Dedicated     -1     Disabled    1000
vcp-1/0        Dedicated     -1     Up          1000        8    vcp-1/0
vcp-1/1        Dedicated     -1     Down        1000
vcp-1/2        Dedicated     -1     Up          1000        8    vcp-2/1
vcp-1/3        Dedicated     -1     Up          1000        8    vcp-1/3
```

### show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members
```

```
fpc0:
```

```
-----
Interface      Type          Trunk  Status      Speed      Neighbor
or             or            ID      (mbps)      ID  Interface
PIC / Port
vcp-0          Dedicated     1      Up          32000       1    vcp-1
vcp-1          Dedicated     2      Up          32000       0    vcp-0
1/0            Auto-Configured 3      Up          1000        2    vcp-255/1/0
1/1            Auto-Configured 3      Up          1000        2    vcp-255/1/1
```

```
fpc1:
```

```
-----
Interface      Type          Trunk  Status      Speed      Neighbor
or             or            ID      (mbps)      ID  Interface
PIC / Port
vcp-0          Dedicated     1      Up          32000       0    vcp-1
vcp-1          Dedicated     2      Up          32000       0    vcp-0
1/0            Auto-Configured -1     Up          1000        3    vcp-255/1/0
```

```
fpc2:
```

```
-----
Interface      Type          Trunk  Status      Speed      Neighbor
or             or            ID      (mbps)      ID  Interface
PIC / Port
vcp-0          Dedicated     1      Up          32000       3    vcp-1
vcp-1          Dedicated     2      Up          32000       3    vcp-0
1/0            Auto-Configured 3      Up          1000        0    vcp-255/1/0
1/1            Auto-Configured 3      Up          1000        0    vcp-255/1/1
```

```
fpc3:
```

```
-----
Interface      Type          Trunk  Status      Speed      Neighbor
or             or            ID      (mbps)      ID  Interface
PIC / Port
vcp-0          Dedicated     1      Up          32000       2    vcp-0
vcp-1          Dedicated     2      Up          32000       2    vcp-1
1/0            Auto-Configured -1     Up          1000        1    vcp-255/1/0
```

## show virtual-chassis vc-port statistics

---

<b>Syntax</b>	<pre>show virtual-chassis vc-port statistics &lt;all-members&gt; &lt;brief   detail   extensive &gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The options <b>all-members</b>, <b>brief</b>, <b>detail</b>, <b>extensive</b>, and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Display the traffic statistics collected on Virtual Chassis ports (VCPs).
<b>Options</b>	<p><b>none</b>—Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Display traffic statistics for VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">clear virtual-chassis vc-port statistics on page 6978</a></li><li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li><li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li><li>• <i>Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li></ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis vc-port statistics on page 7027</a></p> <p><a href="#">show virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 7028</a></p> <p><a href="#">show virtual-chassis vc-port statistics brief on page 7028</a></p> <p><a href="#">show virtual-chassis vc-port statistics extensive on page 7028</a></p> <p><a href="#">show virtual-chassis vc-port statistics member 0 on page 7030</a></p>

**Output Fields** Table 678 on page 7025 lists the output fields for the **show virtual-chassis vc-port statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 678: show virtual-chassis vc-port statistics Output Fields**

Field Name	Field Description	Level of Output
<b>fpcnumber</b>	(All Virtual Chassis except EX8200 Virtual Chassis. VCF) ID of the Virtual Chassis member. The FPC number is the same as the member ID.	All levels
<b>member number</b>	(EX8200 Virtual Chassis only) Member ID of the Virtual Chassis member.	All levels
<b>Interface</b>	VCP name.	<b>brief</b>
<b>Input Octets/Packets</b>	Number of octets and packets received on the VCP.	<b>brief, member, none</b>
<b>Output Octets/Packets</b>	Number of octets and packets transmitted on the VCP.	<b>brief, member, none</b>
<b>master: number</b>	Member ID of the master Routing Engine.	All levels
<b>Port</b>	VCP for which <b>RX</b> (Receive) statistics, <b>TX</b> (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared.	<b>detail, extensive</b>
<b>Total octets</b>	Total number of octets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Total packets</b>	Total number of packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Unicast packets</b>	Number of unicast packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Broadcast packets</b>	Number of broadcast packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Multicast packets</b>	Number of multicast packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>MAC control frames</b>	Number of media access control (MAC) control frames received and transmitted on the VCP.	<b>detail, extensive</b>

Table 678: show virtual-chassis vc-port statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>CRC alignment errors</b>	<p>Number of packets received on the VCP that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors:</p> <ul style="list-style-type: none"> <li>Invalid FCS with an integral number of octets (FCS error)</li> <li>Invalid FCS with a nonintegral number of octets (alignment error)</li> </ul>	<b>detail, extensive</b>
<b>Oversize packets</b>	Number of packets received on the VCP that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.	<b>detail, extensive</b>
<b>Undersize packets</b>	Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..	<b>detail, extensive</b>
<b>Jabber packets</b>	<p>Number of packets received on the VCP that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error.</p> <p><b>NOTE:</b> This definition of <i>jabber</i> is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define <i>jabber</i> as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.</p>	<b>detail, extensive</b>
<b>Fragments received</b>	<p>Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</p>	<b>detail, extensive</b>
<b>Ifout errors</b>	Number of outbound packets received on the VCP that could not be transmitted because of errors.	<b>detail, extensive</b>
<b>Packet drop events</b>	Number of outbound packets received on the VCP that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.	<b>detail, extensive</b>
<b>64 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>

Table 678: show virtual-chassis vc-port statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>65–127 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>
<b>128–255 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>
<b>256–511 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>
<b>512–1023 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>
<b>1024–1518 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>
<b>Rate packets per second</b>	Number of packets per second received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Rate bytes per second</b>	Number of bytes per second received and transmitted on the VCP.	<b>detail, extensive</b>

## Sample Output

### show virtual-chassis vc-port statistics

```
user@switch> show virtual-chassis vc-port statistics
fpc0:
```

```
-----
Interface          Input  Octets/Packets      Output  Octets/Packets
internal-0/24       0      / 0              0      / 0
internal-0/25       0      / 0              0      / 0
internal-1/26       0      / 0              0      / 0
internal-1/27       0      / 0              0      / 0
vcp-0               0      / 0              0      / 0
vcp-1               0      / 0              0      / 0
internal-0/26       0      / 0              0      / 0
internal-0/27       0      / 0              0      / 0
internal-1/24       0      / 0              0      / 0
internal-1/25       0      / 0              0      / 0
```

```
{master:0}
```

**show virtual-chassis vc-port statistics (EX8200 Virtual Chassis)**

```
user@external-routing-engine> show virtual-chassis vc-port statistics
```

```
member0:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-4/0/4           43171238 / 48152          47687133 / 51891
vcp-4/0/7           0 / 0                     0 / 0
```

```
member1:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-3/0/0           0 / 0                     0 / 0
vcp-3/0/1           0 / 0                     0 / 0
vcp-3/0/4           47695376 / 51899          43180556 / 48160
```

```
member8:
```

```
-----
```

```
member9:
```

```
-----
```

**show virtual-chassis vc-port statistics brief**

```
user@switch> show virtual-chassis vc-port statistics brief
```

```
fpc0:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
internal-0/24       0 / 0                     0 / 0
internal-0/25       0 / 0                     0 / 0
internal-1/26       0 / 0                     0 / 0
internal-1/27       0 / 0                     0 / 0
vcp-0               0 / 0                     0 / 0
vcp-1               0 / 0                     0 / 0
internal-0/26       0 / 0                     0 / 0
internal-0/27       0 / 0                     0 / 0
internal-1/24       0 / 0                     0 / 0
internal-1/25       0 / 0                     0 / 0
```

```
{master:0}
```

**show virtual-chassis vc-port statistics extensive**

```
user@switch> show virtual-chassis vc-port statistics extensive
```

```
fpc0:
```

```
-----
                                     RX               TX
Port: internal-0/24
Total octets:           0                   0
Total packets:          0                   0
Unicast packets:        0                   0
Broadcast packets:      0                   0
Multicast packets:      0                   0
MAC control frames:     0                   0
CRC alignment errors:   0
Oversize packets:       0
Undersize packets:      0
Jabber packets:         0
Fragments received:     0
```

```

Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

...

Port: vcp-0
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

Port: vcp-1
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

...

```

```
{master:0}
```

#### show virtual-chassis vc-port statistics member 0

```
user@switch>show virtual-chassis vc-port statistics member 0
fpc0:
```

```
-----
Interface           Input  Octets/Packets      Output  Octets/Packets
internal-0/24        0      / 0             0      / 0
internal-0/25        0      / 0             0      / 0
internal-1/26        0      / 0             0      / 0
internal-1/27        0      / 0             0      / 0
vcp-0                0      / 0             0      / 0
vcp-1                0      / 0             0      / 0
internal-0/26        0      / 0             0      / 0
internal-0/27        0      / 0             0      / 0
internal-1/24        0      / 0             0      / 0
internal-1/25        0      / 0             0      / 0
```

```
{master:0}
```



## PART 23

# Virtual Chassis Fabric

- [Overview on page 7033](#)
- [Configuration on page 7053](#)
- [Administration on page 7105](#)
- [Troubleshooting Procedures on page 7183](#)



# Overview

- [Virtual Chassis Fabric Overview on page 7033](#)

## Virtual Chassis Fabric Overview

---

- [Virtual Chassis Fabric Overview on page 7033](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric on page 7046](#)
- [Understanding Traffic Flow Through a Virtual Chassis Fabric on page 7050](#)
- [Understanding Software Upgrades in a Virtual Chassis Fabric on page 7050](#)

## Virtual Chassis Fabric Overview


The Juniper Networks Virtual Chassis Fabric (VCF) provides a low-latency, high-performance fabric architecture that can be managed as a single device. VCF is an evolution of the Virtual Chassis feature, which enables you to interconnect multiple devices into a single logical device, inside of a fabric architecture. The VCF architecture is optimized to support small and medium-sized data centers that contain a mix of 1-Gbps, 10-Gbps, and 40-Gbps Ethernet interfaces.



Video: [What is Virtual Chassis Fabric?](#)

A VCF is constructed using a spine-and-leaf architecture. In the spine-and-leaf architecture, each spine device is interconnected to each leaf device. A VCF supports up to twenty total devices, and up to four devices can be configured as spine devices. See [Figure 230 on page 7034](#) for an illustration of the VCF spine-and-leaf architecture.

Figure 230: VCF Spine-and-Leaf Architecture



Each spine device must be a QFX5100 device. In an optimal VCF configuration, the leaf devices are also QFX5100 devices. You can, however, also create a mixed VCF by configuring QFX3600, QFX3500, and EX4300 switches as leaf devices. See [“Understanding Virtual Chassis Fabric Components” on page 7035](#) for more information about the spine-and-leaf architecture.

A VCF provides the following benefits:

- **Latency**—VCF provides predictable low latency because it uses a fabric architecture that ensures each device is one or two hops away from every other device in the fabric. The weighted algorithm that makes traffic-forwarding decisions in a VCF is designed to avoid congestion and ensures low latency by intelligently forwarding traffic over all paths within the VCF to any destination device., ensuring predictable low latency for all traffic traversing the VCF.
- **Resiliency**—The VCF architecture provides a resilient framework because traffic has multiple paths across the fabric. Traffic is, therefore, easily diverted within the fabric when a device or link fails.
- **Flexibility**—You can easily expand the size of your VCF by adding devices to the fabric as your networking needs grow.
- **Investment protection**—In environments that need to expand because the capabilities of a traditional QFX5100, QFX3600, QFX3500, or EX4300 Virtual Chassis are maximized, a VCF is often a logical upgrade option because it enables the system to evolve without having to remove the existing, previously purchased devices from the network.
- **Manageability**—VCF provides multiple features that simplify configuration and management. VCF, for instance, has an autoprovisioning feature that enables you to plug and play devices into the fabric after minimal initial configuration. VCF leverages many of the existing configuration procedures from a Virtual Chassis, so that you can configure and maintain a VCF easily if you are already familiar with the procedures for configuring and maintaining a Virtual Chassis.

#### Related Documentation

- [Network Configuration Example: MetaFabric™ Architecture 1.1: Configuring Virtual Chassis Fabric and Network Director 1.6](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)

- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)

## Understanding Virtual Chassis Fabric Components

This topic describes the components of a Virtual Chassis Fabric (VCF).

This topic covers:

- [Spine-and-Leaf Topology on page 7035](#)
- [Spine Devices on page 7036](#)
- [Leaf Devices on page 7036](#)
- [Routing Engine Role on page 7037](#)
- [Linecard Role on page 7038](#)
- [Master Routing Engine Election Process on page 7038](#)
- [Virtual Chassis Ports \(VCPs\) on page 7039](#)
- [Automatic Virtual Chassis Port \(VCP\) Conversion on page 7039](#)
- [VCF Configuration Options on page 7040](#)
- [Fabric Mode on page 7040](#)
- [Mixed Mode on page 7041](#)
- [Virtual Management Ethernet Interface on page 7041](#)
- [Virtual Chassis Fabric Port Link Aggregation Group Bundles on page 7041](#)
- [Virtual Chassis Fabric License Requirements on page 7042](#)
- [Hardware Requirements for a Virtual Chassis Fabric on page 7042](#)
- [Software Requirements in a Virtual Chassis Fabric on page 7042](#)

### Spine-and-Leaf Topology


The VCF uses a spine-and-leaf architecture where each device in the fabric is either a spine device or a leaf device.

A VCF can have up to four spine devices, and up to twenty total devices. Each spine device has at least one direct Virtual Chassis port (VCP) connection to each leaf device in the VCF.

All traffic entering a leaf device can, therefore, be forwarded to any directly connected spine device and is always two hops away from any other leaf device—leaf device to leaf device traffic travels from the source leaf device to a spine device to the destination leaf device—within the VCF.

See [Figure 231 on page 7036](#) for an illustration of the VCF spine-and-leaf architecture:

Figure 231: VCF Spine-and-Leaf Architecture



Traffic is forwarded through a VCF using a weighted algorithm designed to avoid congestion. Traffic travelling across the VCF from one leaf device to another leaf device is forwarded using the best path available at the time, so any connection to a spine device can be used to transport traffic from one leaf device to another leaf device.

### Spine Devices

A spine device:

- Must be a QFX5100 device.
- Is configured into the Routing Engine role, but can operate in the Routing Engine role or the linecard role..



**NOTE:** Only two devices can simultaneously operate in the Routing Engine role within a VCF. A VCF, however, supports up to four spine devices. In scenarios where a VCF has three or more spine devices, the devices that are not operating in the Routing Engine role operate in the linecard role.

A spine device that is configured into the Routing Engine role but is operating in the linecard role assumes the Routing Engine role when an active Routing Engine fails.

- Has a direct connection to each leaf device.
- Typically connects a router, firewall, or other data center networking device to the VCF.

A VCF should always have at least two active spine devices. A VCF supports up to four spine devices.

You can configure any QFX5100 device as a spine device. In the most common VCF configurations, QFX5100-24Q devices are used as spine devices.

### Leaf Devices

A leaf device:

- Is optimally a QFX5100 device, but can also be a QFX3500, QFX3600, or EX4300 device.
- Has a direct connection to each spine device.
- Always operates in the linecard role.
- Typically connects an endpoint device—for instance, a server or other storage device in a data center—to the VCF.

A VCF can have up to twenty total devices and up to four devices can be configuring into spine devices. The devices that are not spine devices in a VCF operate as leaf devices.

In the most common VCF configurations, QFX5100-48S devices are used as leaf devices.

### Routing Engine Role

A VCF has two devices operating in the Routing Engine role—a master Routing Engine and a backup Routing Engine.

The device that functions as the master Routing Engine:

- Is a spine device.
- Manages the member devices.
- Runs the chassis management processes and control protocols.
- Represents all the member devices interconnected within the VCF configuration. (The hostname and other parameters that you assign to this device during setup apply to all members of the VCF.)

The device that functions as the backup Routing Engine:

- Is a spine device.
- Maintains a state of readiness to take over the master role if the master fails.
- Synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it preserves routing information and maintains network connectivity without disruption when the master is unavailable.

All spine devices in a VCF are configured into the Routing Engine role.



**NOTE:** Only two devices can simultaneously operate in the Routing Engine role within a VCF. A VCF, however, supports up to four spine devices. In scenarios where a VCF has three or more spine devices, the devices that are not operating in the Routing Engine role operate in the linecard role.

A spine device that is configured into the Routing Engine role but is operating in the linecard role assumes the Routing Engine role when an active Routing Engine fails.

The master and backup Routing Engines are selected by the master election algorithm. If a VCF has more than two spine devices, the spine devices that are not selected as the master or backup Routing Engine operate in the linecard role even though the devices are configured into the Routing Engine role.

A spine device operating in the linecard role can complete all spine-related functions with no limitations within a VCF. A spine device operating in the linecard role assumes the backup Routing Engine role when the master or backup Routing Engine fails.

### Linecard Role

---

All leaf devices in a VCF operate in the linecard role. In autoprovisioned configurations, leaf devices are assigned the linecard role when they are cabled into the VCF. In preprovisioned configurations, leaf devices are manually configured into the linecard role. In nonprovisioned configurations, leaf devices are assigned the linecard role according to the master election algorithm, which uses the mastership priority values to set the roles of each device in the VCF.

All spine devices in a VCF are configured into the Routing Engine role, but operate in the linecard role when they are not selected as the master or backup Routing Engine by the master election algorithm. A spine device operating in the linecard role can complete all spine-related functions with no limitations within a VCF. A spine device operating in the linecard role becomes the new backup Routing Engine when the master or backup Routing Engine fails.

A member that functions in the linecard role in a VCF:

- Runs only a subset of Junos OS.
- Detects certain error conditions (such as an unplugged cable) on any interfaces that have been configured on it through the device functioning as the master Routing Engine.

### Master Routing Engine Election Process

---

The device in the master Routing Engine role in a VCF is always a spine device.

In a preprovisioned or autoprovisioned VCF, up to four spine devices are assigned the Routing Engine role during the configuration process. The spine device that has been powered on the longest assumes the master Routing Engine role; the spine device that has been powered on the second longest assumes the backup Routing Engine role. The remaining spine devices assume the linecard role.

In a nonprovisioned VCF, the master and backup Routing Engines are selected using the following algorithm:

1. Choose the QFX5100 device with the highest user-configured mastership priority (255 is the highest possible value) as the master Routing Engine, and the QFX5100 switch with the second highest mastership priority value as the backup Routing Engine.  
A QFX5100 switch with a mastership priority of 0 will always stay in the linecard role.
2. Choose the QFX5100 device that was master the last time the VCF booted.



3. Choose the QFX5100 device that has been included in the VCF configuration for the longest period of time.
4. Choose the QFX5100 device with the lowest MAC address.

QFX3500, QFX3600, and EX4300 devices never assume the master or backup Routing Engine role in a VCF.

We strongly recommend that you configure the mastership priority of the QFX5100 devices in your VCF to ensure that the correct devices assume their intended roles when you configure your VCF using a nonprovisioned configuration.

### Virtual Chassis Ports (VCPs)

---

Virtual Chassis ports (VCPs) are used in a VCF to interconnect leaf devices to spine devices. All control and data traffic in a VCF is transported over VCPs.

VCPs in a VCF are either SFP+ connections that support 10-Gbps or QSFP+ connections that support 40-Gbps.

10-Gbps SFP+ and 40-Gbps QSFP+ links are automatically converted into VCPs in most scenarios when a device is added to an autoprovisioned or preprovisioned VCF. Automatic VCP conversion is discussed in more detail in the following section.

You can manually configure a 10-Gbps SFP+ and 40-Gbps QSFP+ link into a VCP.

Channelized interfaces cannot be configured into VCPs.

### Automatic Virtual Chassis Port (VCP) Conversion

---

10-Gbps SFP+ and 40-Gbps QSFP+ links are not configured into VCPs, by default.

10-Gbps SFP+ and 40-Gbps QSFP+ links are automatically converted into VCPs when:

- Link Layer Discovery Protocol (LLDP) is enabled on the interfaces on both ends of the link. LLDP is enabled by default.
- the device being added to the VCF is configured into fabric mode.
- The interfaces on both ends of the link are not configured as VCPs. The following interfaces are configured as VCPs:
  - The 40-Gbps QSFP+ port on an EX4300 switch, by default.
  - Any interface in the VCF that has been a VCP. If a device is removed from a VCF, the interface that was interconnected to the removed device remains configured as a VCP until it is configured into a network port using the **request virtual-chassis vc-port delete** command.
  - Any interface that has been configured into a VCP using the **request virtual-chassis vc-port set** command.

To change any of the above interfaces into a network interface so that the interface can become eligible for automatic VCP conversion, use the **request virtual-chassis vc-port delete** command.

- one of the devices is already part of a VCF that was autoprovisioned or preprovisioned.

Automatic VCP conversion does not work in nonprovisioned VCFs.

Automatic VCP conversion does not convert a VCP interface into a network interface when a device is removed from a VCF. If automatic VCP conversion has converted an interface into a VCP and you want the interface to function as a network interface, you must manually disable the VCP interface.

---

### VCF Configuration Options

You can configure a VCF using autoprovisioned, preprovisioned, or nonprovisioned configuration.

Autoprovisioned configuration allows you to *plug and play* leaf devices into a VCF after completing a minimal initial configuration procedure.

In a preprovisioned configuration, you deterministically control the devices in your VCF by associating each device's serial number to a member ID and role.

Nonprovisioned configuration is possible, but not recommended for most VCF installations. Nonprovisioned configuration is a highly manual procedure that should only be performed by expert users.

See [“Understanding Virtual Chassis Fabric Configuration” on page 7043](#) for additional information on the VCF configuration options.

---

### Fabric Mode

A device must be configured into fabric mode in order for it to join a VCF. You should always configure a device into fabric mode before interconnecting it into a VCF.

In preprovisioned and nonprovisioned configurations, a device is not participating as a VCF member until it is configured into fabric mode.

In autoprovisioned configurations, a spine device is not participating as a VCF member until it is configured into fabric mode. A spine device that is not configured into fabric mode is configured into fabric mode when it is interconnected into the VCF. The final step of the process of configuring the device into fabric mode is a device reboot. We strongly recommend configuring the spine device into fabric mode before interconnecting it into the VCF to eliminate this reboot.

A leaf device in an autoprovisioned configuration is also rebooted to complete the fabric mode configuration when it is interconnected into a VCF without being set into fabric mode. You can avoid the downtime that accompanies the reboot by setting the device into fabric mode before interconnecting it into the VCF.

A standalone device that is not part of a VCF should never be configured into fabric mode. A device is not in fabric mode, by default.

### Mixed Mode

---

The optimal method of configuring a VCF is to use QFX5100 devices only. A VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces.

You can, however, configure other devices as leaf devices in your VCF. QFX5100, QFX3600, QFX3500, or EX4300 devices can be used as leaf devices in a VCF.

If you use QFX3600, QFX3500, or EX4300 devices as leaf devices in your VCF, you must configure all devices in your VCF into mixed mode.

A device that is not part of a Virtual Chassis or a VCF with other devices should never be configured into mixed mode. A device is not configured into mixed mode, by default.

### Virtual Management Ethernet Interface

---

VCF configuration can be managed remotely using a global management interface called the virtual management Ethernet (VME) interface. The VME interface is a logical interface representing all of the out-of-band management ports on the member devices. When you connect to the VCF using the VME interface's IP address, the connection is always redirected to the device acting in the master Routing Engine role.

A VME interface should always be used to configure a VCF. The VME interface is not tied to a device, so it can always be used to log in to the VCF even after the master Routing Engine changes.

We strongly recommend cabling the management port on all spine devices to the network to ensure that you always have a direct connection to the master Routing Engine through the VME interface, regardless of which spine device assumes the master Routing Engine role. The management ports on leaf devices can also be used by the VME interface to access the VCF, so you can also cable leaf device management ports to the network, if desired.

### Virtual Chassis Fabric Port Link Aggregation Group Bundles

---

You can increase the bandwidth on links configured as VCPs within a VCF between two devices by configuring multiple same-speed links between two devices into VCPs. If, for instance, you configure two 40-Gbps QSFP+ links that are connecting the same devices in a VCF into VCPs, the two VCP links form one LAG bundle with two member links and 80-Gbps of total available bandwidth.

A VCP LAG bundle provides more bandwidth than a single VCP link can provide. A VCP LAG bundle also improves performance by load-sharing traffic across links within the bundle, and provides redundancy because traffic can be forwarded across another member link in the VCP LAG bundle when one VCP link fails.

VCP LAG bundling occurs automatically when same-speed VCP links are configured between two devices. No user configuration is required. VCP LAG bundling works only on same-speed VCP links; 10-Gbps and 40-Gbps links cannot be in the same VCP LAG bundle.

### Virtual Chassis Fabric License Requirements

---

A feature license is required to configure a VCF. The VCF feature license is an independent feature license; the enhanced feature licenses (EFLs) or advanced feature licenses (AFLs) that must be purchased to enable some features on some Juniper switches cannot be purchased to enable VCF.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master Routing Engine role and the other for the device in the backup Routing Engine role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show virtual-chassis** command.

### Hardware Requirements for a Virtual Chassis Fabric

---

A VCF can contain up to four devices configured as spines and up to twenty total devices.

All spine devices must be QFX5100 devices. We recommend optimizing the performance of your VCF by also configuring QFX5100 devices as your leaf devices. A non-mixed VCF has the highest port density and feature support for a VCF in addition to supporting more spine devices. Nevertheless, you can configure any combination of QFX5100, QFX3600, QFX3500, or EX4300 devices into leaf devices within your VCF.

You can configure any QFX5100 device as a spine device. In the most common VCF configurations, the QFX5100-24Q devices are used as spine devices and QFX5100-48S devices are used as leaf devices.

### Software Requirements in a Virtual Chassis Fabric

---

All devices in a VCF must be running the same version of Junos OS software that supports VCF. VCF is initially supported for twenty total member devices in Junos OS Release 13.2X51-D20 for QFX5100, QFX3600, QFX3500, and EX4300 devices.

The devices in the VCF must be using the version of software for standalone switches.

The flex software bundle is supported on non-mixed VCFs using QFX5100 member switches only. You cannot use the flex software bundle in mixed VCFs. The flex software bundle is the software that includes “jinstall-qfx-5-flex” text in the filename when it is downloaded from the Software Center.

We recommend configuring a device to the Junos OS release running on the VCF before interconnecting it into the VCF. For additional information on VCF software upgrades, see “[Understanding Software Upgrades in a Virtual Chassis Fabric](#)” on page 7050.

For information on software upgrade options for an operational VCF, see “[Understanding Software Upgrades in a Virtual Chassis Fabric](#)” on page 7050.

- Related Documentation**
- [Network Configuration Example: MetaFabric™ Architecture 1.1: Configuring Virtual Chassis Fabric and Network Director 1.6](#)
  - [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
  - [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
  - [Understanding Software Upgrades in a Virtual Chassis Fabric on page 7050](#)
  - [Virtual Chassis Fabric Overview on page 7033](#)

## Understanding Virtual Chassis Fabric Configuration

This topic describes the configuration options available for your Virtual Chassis Fabric (VCF).

This topic covers:

- [Virtual Chassis Fabric Setup on page 7043](#)
- [Configuration File Management in a VCF on page 7045](#)
- [Logging into a Virtual Chassis Fabric on page 7045](#)
- [Understanding Interface Numbering on page 7045](#)

### Virtual Chassis Fabric Setup

You must setup your VCF using one of the following options:

- [Autoprovisioned Virtual Chassis Fabric Configuration on page 7043](#)
- [Preprovisioned Virtual Chassis Fabric Configuration on page 7044](#)
- [Nonprovisioned Virtual Chassis Fabric Configuration on page 7044](#)

#### ***Autoprovisioned Virtual Chassis Fabric Configuration***

Autoprovisioned configuration allows you to “plug and play” leaf devices into a VCF after minimal initial configuration.

The minimal configuration requirements for autoprovioning a VCF include setting the configuration mode to autoprovioned and explicitly identifying the spine devices in your VCF by serial number. After this minimal configuration is complete, all supported devices—supported devices are either devices that have been zeroized or devices in factory default mode that have never been configured into a Virtual Chassis or VCF—are automatically added to the VCF as leaf devices when they are cabled to spine devices using supported 10-Gbps SFP+ ports or 40-Gbps QSFP+ ports. The Virtual Chassis ports (VCPs) are created automatically. Other parameters such as fabric and mixed mode are automatically detected and set.

A spine device in an autoprovioned configuration should be configured into fabric mode before being interconnected into a VCF. A spine device in an autoprovioned VCF must also have the same mixed mode setting as other member devices in the VCF. You should

configure your spine device into fabric mode and, if necessary, mixed mode before interconnecting it into the VCF.

A leaf device in an autoprovisioned configuration is rebooted to complete the fabric mode configuration when it is interconnected into a VCF without being set into fabric mode. The leaf device is also rebooted if the device needs to be configured into or out of mixed mode to participate in the VCF. You can avoid the downtime that accompanies the reboot of the leaf device by setting the leaf device into fabric mode and into or out of mixed mode before interconnecting it into the VCF.

### ***Preprovisioned Virtual Chassis Fabric Configuration***

In a preprovisioned configuration, you deterministically control the devices in your VCF by associating each device's serial number to a member ID and role.

The advantage of configuring a VCF using a preprovisioned configuration is that you can explicitly control which devices are added to your VCF, and in what roles. VCF configuration, notably, occurs automatically when two devices that have been configured into fabric mode (and mixed mode, if applicable) are interconnected by a supported 10-Gbps SFP+ port or a 40-Gbps QSFP+ port after the preprovisioned configuration is defined.

The disadvantage of using a preprovisioned configuration is that the configuration process is more manual than the autoprovisioned configuration process.

### ***Nonprovisioned Virtual Chassis Fabric Configuration***



**CAUTION:** We discourage nonprovisioned VCF configuration. You can configure all aspects of a VCF using autoprovisioned or preprovisioned configuration. Nonprovisioned VCF configuration should only be used by VCF experts in specialized scenarios.

A nonprovisioned VCF is the default method for creating a VCF; it is the configuration mode used when a VCF has not been configured into autoprovisioned or preprovisioned mode.

In a nonprovisioned VCF, member roles are determined by a mastership election algorithm. The first value checked by the mastership election algorithm is the mastership priority value. The devices with the highest mastership priority values assume the Routing Engine role, which is used by the spine devices in a VCF. All other devices assume the linecard role.

If two or more devices have the same mastership priority value and are candidates for the Routing Engine role, the mastership election algorithm uses other parameters to determine which device is elected as the Routing Engine. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).

The default mastership priority value for all devices is 128. You should always configure your spine devices with the highest mastership priority to ensure a spine device assumes the Routing Engine role.

In a nonprovisioned VCF, you must manually configure every VCP.

### Configuration File Management in a VCF

You configure a VCF by logging onto the master Routing Engine and making configuration changes. See the next section for information on logging into a VCF.

The configuration file that is modified when you are on the master Routing Engine is automatically shared with all other devices in the VCF when it is committed. Each device stores its own copy of the configuration file.

### Logging into a Virtual Chassis Fabric

The recommended method of logging into a VCF is through the use of a Virtual Management Ethernet (VME) interface. The VME interface is a logical interface representing all of the out-of-band management ports on the member devices. When you connect to the VCF configuration using the VME interface's IP address, the connection is always redirected to the management port on device in the master Routing Engine role. The VME interface is not tied to a device, so it can always be used to log in to the VCF even after the master Routing Engine changes. We recommend cabling the management ports—an *me* or *em* interface—on each spine device in your VCF to support the VME interface.

If you log in to the console port of any member device in a VCF, your session is automatically redirected to the device acting in the master Routing Engine role.

### Understanding Interface Numbering

Interfaces in Junos OS are specified as follows:

*type-fpc/pic/port*

A VCF applies this convention as follows:

- *type*—The interface type.
- *fpc*—Flexible PIC Concentrator. In a VCF, the *fpc* is the member ID of the switch. For instance, the *fpc* of member 16 in the VCF is 16.
- *pic*—the number of the PIC (Physical Interface Card) on the member device.
- *port*—the port number.

For more detailed information on interface numbering, see [“Understanding Interface Naming Conventions” on page 2401](#).

#### Related Documentation

- [Network Configuration Example: MetaFabric™ Architecture 1.1: Configuring Virtual Chassis Fabric and Network Director 1.6](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Configuring a Nonprovisioned Virtual Chassis Fabric on page 7059](#)

## Understanding Mixed EX Series and QFX Series Virtual Chassis or Virtual Chassis Fabric

This topic describes the requirements for a mixed Virtual Chassis or a mixed Virtual Chassis Fabric (VCF).

A mixed Virtual Chassis includes two or more types of EX Series switches, two or more types of QFX Series switches, or a mix of EX and QFX Series switches.

A mixed VCF is any VCF that includes two or more types of member switches. Because a VCF must use a QFX5100 switch as a spine device, a mixed VCF is any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switches.



**NOTE:** The optimal VCF topology is to use QFX5100 devices only. A VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces.

This topic covers:

- [Virtual Chassis Fabric Summary on page 7046](#)
- [Understanding Mixed Virtual Chassis Fabric on page 7047](#)
- [Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches on page 7047](#)
- [Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches on page 7048](#)
- [Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis on page 7049](#)
- [Understanding Mixed EX4300 and EX4600 Virtual Chassis on page 7049](#)
- [Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis on page 7049](#)

### Virtual Chassis Fabric Summary

---

[Table 665 on page 6914](#) provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed VCF. The table also includes license requirements and supported configuration methods.



Table 679: Virtual Chassis Fabric Summary

Category	Allowed Routing Engines	Allowed Line Cards	License Requirement	Configuration Methods
Non-mixed	QFX5100	QFX5100	Yes (on two QFX5100 switches operating in master and backup Routing Engine roles)	Autoprovisioning Preprovisioning Nonprovisioning (not recommended)
Mixed	QFX5100	QFX5100 QFX3600 QFX3500 EX4300	Yes (on two QFX5100 switches operating in master and backup Routing Engine roles)	Autoprovisioning Preprovisioning Nonprovisioning (not recommended)

### Understanding Mixed Virtual Chassis Fabric

A VCF must use a QFX5100 switch in the spine role. A mixed VCF is, therefore, any VCF that includes EX4300, QFX3500, or QFX3600 member switches in addition to the required QFX5100 switch.

The optimal method of configuring a VCF is to use QFX5100 devices only. A non-mixed VCF composed entirely of QFX5100 devices supports the largest breadth of features at the highest scalability while also supporting the highest number of high-speed interfaces. You can, however, also configure a mixed VCF.

If you use QFX3600, QFX3500, or EX4300 devices as leaf devices in your VCF, you must configure all devices in your VCF into mixed mode. If you are turning a non-mixed VCF into a mixed VCF, you have to reboot the VCF to change the mixed mode setting.

### Virtual Chassis Summary for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 Switches

Table 666 on page 6915 provides a high-level overview of the permitted hardware allowed in the routing engine and line card roles of a mixed and a non-mixed Virtual Chassis for QFX5100, QFX3600, QFX3500, EX4600, and EX4300 switches. The table also includes license requirements and supported configuration methods.

**Table 680: Virtual Chassis Summary**

Category	Allowed Routing Engines	Allowed Line Cards	License Requirement	Configuration Methods
Non-mixed	QFX5100	QFX5100	No	Nonprovisioning Preprovisioning
	QFX3600 QFX3500	QFX3600 QFX3500	No	Nonprovisioning Preprovisioning
	EX4600	EX4600	No	Nonprovisioning Preprovisioning
	EX4300	EX4300	No	Nonprovisioning Preprovisioning
Mixed	QFX5100	QFX5100 QFX3600 QFX3500 EX4300	No	Nonprovisioning Preprovisioning
	QFX3600 QFX3500	QFX3600 QFX3500 EX4300	No	Nonprovisioning Preprovisioning
	EX4600	EX4600 EX4300	No	Nonprovisioning Preprovisioning

#### **Understanding the Routing Engine Role in a Mixed Virtual Chassis Using EX4300, EX4600, QFX3500, QFX3600, or QFX5100 Member Switches**

In a mixed Virtual Chassis, the switch in the master Routing Engine role determines which switches are supported in the line card role of the mixed Virtual Chassis.

When a mixed Virtual Chassis is using a QFX5100 switch in the master Routing Engine role, you can use QFX5100, QFX3600, QFX3500, or EX4300 switches in the line card role.

When a mixed Virtual Chassis is using a QFX3600 or QFX3500 switch in the master Routing Engine role, you can use QFX3600, QFX3500, or EX4300 switches in the line card role.

In a mixed EX4300 and EX4600 Virtual Chassis, an EX4600 switch automatically assumes the Routing Engine role.

EX4600 switches can only be in a mixed Virtual Chassis with EX4300 switches. EX4600 switches cannot be in a mixed Virtual Chassis with QFX5100, QFX3600, or QFX3500 switches.

We recommend always configuring the same type of switch into the master and backup Routing Engine role, to ensure that the switch operating in the master role remains the same type of switch in the event of a switchover.

In most mixed Virtual Chassis, you must configure your Virtual Chassis to ensure a switch that supports the master Routing Engine assumes the master Routing Engine role. Without user configuration, any switch—with the exception of the EX4300 switch, which can never assume the master or backup Routing Engine role in a mixed Virtual Chassis or VCF—can assume the master or backup Routing Engine role.

### **Understanding EX4300, QFX3500, QFX3600, and QFX5100 Switches in a Virtual Chassis**

Up to ten EX4300 switches, QFX3500 switches, QFX3600 switches, and QFX5100 switches can be interconnected using Virtual Chassis ports (VCPs) to form a mixed or non-mixed Virtual Chassis. The mixed Virtual Chassis supports up to ten member switches regardless of the switches that compose the mixed Virtual Chassis.

EX4300 switches can also be interconnected into a mixed Virtual Chassis with EX4600 switches. See the following section for information on mixed EX4300 and EX4600 Virtual Chassis.

### **Understanding Mixed EX4300 and EX4600 Virtual Chassis**

EX4300 switches and EX4600 switches can be interconnected into the same Virtual Chassis. An EX4600 switch automatically assumes the master Routing Engine role in a mixed EX4300 and EX4600 Virtual Chassis, since EX4300 switches cannot assume the Routing Engine role in a mixed Virtual Chassis. EX4600 switches cannot be in a mixed Virtual Chassis with any other type of switch.

The mixed Virtual Chassis supports up to ten member switches.

### **Understanding EX4200, EX4500, and EX4550 Switches in a Mixed Virtual Chassis**

EX4200 switches, EX4500 switches, and EX4550 switches can be interconnected into the same Virtual Chassis to form a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, mixed EX4500 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, and EX4550 Virtual Chassis. The mixed Virtual Chassis supports up to 10 member switches regardless of whether the switches are EX4200 switches, EX4500 switches, or EX4550 switches. Any model of EX4200, EX4500, or EX4550 switch can be interconnected into the same mixed Virtual Chassis. The master election process that decides member switch roles in a mixed Virtual Chassis is identical to the master election process in a non-mixed Virtual Chassis, so any member switch in a mixed Virtual Chassis can assume the master, backup, or linecard role.

EX4200 switches, EX4500 switches, and EX4550 switches cannot be interconnected into a Virtual Chassis with any other switches.

#### **Related Documentation**

- [Virtual Chassis Fabric Overview on page 7033](#)
- [Understanding QFX Series Virtual Chassis on page 6907](#)
- [EX Series Virtual Chassis Overview](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)
- [Understanding EX Series Virtual Chassis Components](#)

## Understanding Traffic Flow Through a Virtual Chassis Fabric

A Virtual Chassis Fabric (VCF) forwards unicast traffic using a smart trunking algorithm that sends all traffic across multiple paths based on end-to-end available bandwidth. The smart trunking algorithm avoids unnecessary congestion due to improper traffic allocation while optimizing fabric bandwidth utilization because traffic is forwarded through the VCF relative to available bandwidth.

The smart trunking algorithm works by considering the overall available path bandwidth of each path in the VCF when making traffic-forwarding decisions, and then forwarding traffic across the paths relative to available path bandwidth. If a VCF with two spine devices, for instance, has one path from leaf device 1 to leaf device 4 that contains two 40-Gbps QSFP+ links and a second path from leaf device 1 to leaf device 4 that contains two 10-Gbps SFP+ links, the algorithm tries to balance traffic sent on the paths so that four times more packets are sent on the first path with 40 Gbps of available bandwidth across the entire path than are sent on the second path with 10 Gbps of total bandwidth.

You can optimize how traffic is forwarded through the VCF by adding spine devices to maximize the number of available paths between all leaf devices, and by using as many 40-Gbps QSFP+ interfaces as Virtual Chassis ports (VCPs) as possible.

VCF also supports adaptive load balancing (ALB). ALB enables the VCF trunking algorithm to use dynamic load information on interfaces and traffic queues to make forwarding decisions within the VCF. When ALB is implemented using flowlets, traffic flows that enter the VCF are spliced into smaller flows—flowlets—and individually forwarded across the VCF to the same destination device over different paths when the inactivity time between packet bursts on the sending interface exceeds the user-configurable inactivity interval. When ALB is implemented using per-packet mode, the sending interface actively monitors all paths available between two member devices and forwards traffic through the VCF using the best available path at the moment.

Implementing ALB using flowlets is effective in environments that periodically experience extremely large traffic flows—*elephant flows*—that are substantially larger than the majority of other traffic flowing through the VCF. The VCF is better able to manage the elephant flows by splicing them into smaller flowlets using ALB.

ALB is supported on a non-mixed VCF composed entirely of QFX5100 switches only. You should enable ALB using flowlets in non-mixed VCFs in environments where a small number of traffic flows are disproportionately larger than the majority of the other traffic flows.

**Related Documentation** • [Understanding Virtual Chassis Fabric Components on page 7035](#)

## Understanding Software Upgrades in a Virtual Chassis Fabric

This topic provides an overview of software upgrades on Virtual Chassis Fabric (VCF).

It contains the following sections:

- [Virtual Chassis Fabric Software Basics on page 7051](#)
- [Nonstop Software Upgrade \(NSSU\) on page 7051](#)
- [Automatic Software Update on page 7051](#)
- [Traditional Software Upgrade on page 7051](#)

---

### Virtual Chassis Fabric Software Basics

---

VCF is initially supported in Junos OS Release 13.2X51-D20. All devices in a VCF must be running the same version of Junos OS that supports VCF.

At initial VCF configuration, you should configure all devices to the same Junos OS release before interconnecting them into a VCF.

When you are adding a device to an existing VCF, you should update the Junos OS release on the new device to the Junos OS release running in the VCF before interconnecting it into the VCF. Updating the Junos OS on the device before interconnecting it helps ensure the device is gracefully added to the VCF, without the downtime that is required to reboot the device after an automatic software update or the troubleshooting that is required if the device isn't added to the VCF due to mismatched software releases.

Before you interconnect a device into a VCF, you should upgrade the software on the device being added to the VCF to the version of Junos OS running on the VCF.

---

### Nonstop Software Upgrade (NSSU)

---

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member devices in a VCF with minimal network traffic disruption during the upgrade.

NSSU upgrades the software on each device individually while all other devices continue normal operations.

For additional information on NSSU in a VCF, see [“Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric” on page 26](#).

---

### Automatic Software Update

---

Automatic software update automatically upgrades the Junos OS running on a device joining a VCF to the version of Junos OS running on the VCF at the moment the new device is cabled into the VCF.

Automatic software update is enabled using the **set virtual-chassis auto-sw-update** statement.

---

### Traditional Software Upgrade

---

You can upgrade software on a VCF using the traditional method of upgrading software for Junos OS by logging onto the master Routing Engine and using the **request system software add** command to initiate the upgrade on a non-mixed VCF or the **request system software add set [package-name package-name ...]** to initiate the upgrade on a mixed VCF, where *package-name* is the path to an image for one device family.

When you upgrade Junos OS on a VCF using the traditional software upgrade, the entire system is down until the upgrade is complete.

**Related  
Documentation**

- [Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade on page 139](#)
- [Upgrading Software for a Virtual Chassis Fabric on page 7070](#)
- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)

# Configuration

- [Configuration Tasks on page 7053](#)
- [Configuration Statements on page 7072](#)

## Configuration Tasks

---

- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Configuring a Nonprovisioned Virtual Chassis Fabric on page 7059](#)
- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)
- [Removing a Device From a Virtual Chassis Fabric on page 7069](#)
- [Upgrading Software for a Virtual Chassis Fabric on page 7070](#)

### Autoprovisioning a Virtual Chassis Fabric

Autoprovisioning a Virtual Chassis Fabric (VCF) enables you to “plug and play” devices into your VCF after minimal initial configuration.

Update all devices to the same version of Junos OS that supports VCF. See [“Upgrading Software” on page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To configure a VCF using autoprovisioning:

1. Make a list of the serial numbers of all the spine devices in the VCF. The spine devices must be QFX5100 devices. You can configure up to four spine devices in a VCF. You can get the device’s serial number in the **show virtual-chassis** output or by following the instructions in *Locating the Serial Number on a QFX5100 Device or Component*.
2. Configure each device into fabric mode. If needed, configure the devices into mixed mode.

Configure the device to reboot as part of the procedure to complete this configuration step.

Configure mixed mode if your VCF includes QFX3600, QFX3500, or EX4300 devices as leaf devices.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF composed of two or more types of devices:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** A spine device whose fabric or mixed mode setting is improperly set cannot join a VCF. You can check the mode settings by using the `show virtual-chassis mode` command.

We recommend that you set the fabric and mixed mode settings before you interconnect your spine devices into the VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a spine device into fabric or mixed mode after interconnecting your VCF.

The fabric and mixed mode settings are automatically updated for a leaf device when it is interconnected into an autoprovisioned VCF. If the fabric or mixed mode settings are changed when a leaf device is interconnected into a VCF, the leaf device reboots before joining the VCF.

---

3. When the reboot is complete, log in to one of the spine devices in your VCF.

4. Set the configuration mode to autoprovisioned:

```
[edit]
```

```
user@device# set virtual-chassis auto-provisioned
```

5. Configure your spine devices into the Routing Engine role:

```
[edit virtual-chassis]
```

```
user@device# set member member-id serial-number serial-number role routing-engine
```

For instance, to configure the four spine devices with the serial numbers “SERIALNUMB00”, “SERIALNUMB01”, “SERIALNUMB02”, and “SERIALNUMB03” into the Routing Engine role as members 0 through 3:

```
[edit virtual-chassis]
```

```
user@device# set member 0 serial-number SERIALNUMB00 role routing-engine
```

```
user@device# set member 1 serial-number SERIALNUMB01 role routing-engine
```

```
user@device# set member 2 serial-number SERIALNUMB02 role routing-engine
```

```
user@device# set member 3 serial-number SERIALNUMB03 role routing-engine
```

6. (Recommended) Configure a virtual management Ethernet (VME) interface for management of the VCF configuration:

```
[edit]
```

```
user@device# set interfaces vme unit 0 family inet address /ip-address/mask/
```





**NOTE:** A VME accesses the device in the master Routing Engine role using a management port, so cable management port em0 or em1 on each spine device in your VCF so the VME is available regardless of which spine device assumes the master Routing Engine role. See *Connecting a QFX Series Device to a Management Console*

7. Commit the configuration:

```
user@device# commit
```

8. Cable your VCF.

After your autoprovisioned VCF configuration is committed, you can cable any EX4300, QFX3500, QFX3600, or QFX5100 device that is zeroized or that has never been configured to a spine device using a supported SFP+ or QSFP+ interface. The device that is zeroized or in factory-default mode is added to the VCF as a leaf device. All VCPs are configured as part of this process.



**NOTE:** Mixed mode and fabric mode are checked and, if needed, set automatically on the device as part of this process. If the mixed or fabric mode has to be changed to become part of the VCF, the device reboots. The device participates in the VCF with no further user intervention after this reboot is complete.



**NOTE:** Automatic VCP conversion only works when the interfaces on both ends of the link are not configured into VCPs.

The 40-Gbps QSFP+ interfaces on EX4300 switches are configured as VCPs, by default. You must, therefore, delete the VCP on the 40-Gbps QSFP+ interface using the `request virtual-chassis vc-port delete` command before interconnecting it into the VCF in order for the link to be converted into a VCP. You can also manually configure the link into a VCP using the `request virtual-chassis vc-port set` command.

The device joins the VCF immediately without a reboot if the mixed or fabric mode setting does not need to be changed.

9. Install the VCF feature licenses.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the `show virtual-chassis` command.

After obtaining the licenses, follow the instructions in [“Generating License Keys” on page 75](#).

**Related Documentation**

- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)
- [Removing a Device From a Virtual Chassis Fabric on page 7069](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)

## Preprovisioning a Virtual Chassis Fabric

Preprovisioning a Virtual Chassis Fabric (VCF) configuration allows you to assign the member ID and role for each device in the VCF.

Update all devices to the same version of Junos OS that supports VCF. See [“Upgrading Software” on page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To preprovision a VCF:

1. Make a list of the serial numbers of all the devices to be connected in the VCF. You can get a device's serial number in the **show virtual-chassis** output or by following the instructions in *Locating the Serial Number on a QFX5100 Device or Component*, *Locating the Serial Number on a QFX3600 or QFX3600-I Device or Component*, *Locating the Serial Number on a QFX3500 Device or Component*, or *Locating the Serial Number on an EX4300 Switch or Component*.
2. Decide the desired role (**routing-engine** or **line-card**) for each device.

In a VCF, you configure up to four QFX5100 devices into the Routing Engine role as spine devices. All other devices are configured into the linecard role as leaf devices.



**NOTE:** Only two devices can simultaneously operate in the Routing Engine role within a VCF. A VCF, however, supports up to four spine devices. In scenarios where a preprovisioned VCF has three or more spine devices, the devices that are not operating in the Routing Engine role operate in the linecard role.

A spine device operating in the linecard role assumes the Routing Engine role when an active Routing Engine fails.

3. Configure each individual device into fabric mode. If needed, configure the devices into mixed mode.

Reboot each device to complete this configuration step.

Mixed mode must be configured if your VCF includes QFX3500, QFX3600, or EX4300 devices as leaf nodes.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF that includes EX4300, QFX3500, or QFX3600 devices as leaf nodes:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** A device whose fabric or mixed mode setting is improperly set cannot join a VCF. You can check the mode settings using the `show virtual-chassis mode` command.

We recommend that you set the fabric and mixed mode before you interconnect your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a device into fabric or mixed mode after interconnecting your VCF.

4. Log in to one of your spine devices after the reboot has completed.

5. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@device# set preprovisioned
```

6. Associate a member ID with a serial number for each device in your VCF, and configure the role for each device:

```
[edit virtual-chassis]
user@device# set member member-id serial-number serial-number role (line-card |
routing-engine)
```

Configure your spine devices into the Routing Engine role. You must use QFX5100 devices as your spine devices.

Configure your leaf devices into the linecard role.

For instance, if you wanted to preprovision a VCF with twenty member devices:

```
[edit virtual-chassis]
user@device# set member 0 serial-number SERIALNUMB00 role routing-engine
user@device# set member 1 serial-number SERIALNUMB01 role routing-engine
user@device# set member 2 serial-number SERIALNUMB02 role routing-engine
user@device# set member 3 serial-number SERIALNUMB03 role routing-engine
user@device# set member 4 serial-number SERIALNUMB04 role line-card
user@device# set member 5 serial-number SERIALNUMB05 role line-card
user@device# set member 6 serial-number SERIALNUMB06 role line-card
user@device# set member 7 serial-number SERIALNUMB07 role line-card
user@device# set member 8 serial-number SERIALNUMB08 role line-card
user@device# set member 9 serial-number SERIALNUMB09 role line-card
user@device# set member 10 serial-number SERIALNUMB10 role line-card
user@device# set member 11 serial-number SERIALNUMB11 role line-card
```

```

user@device# set member 12 serial-number SERIALNUMB12 role line-card
user@device# set member 13 serial-number SERIALNUMB13 role line-card
user@device# set member 14 serial-number SERIALNUMB14 role line-card
user@device# set member 15 serial-number SERIALNUMB15 role line-card
user@device# set member 16 serial-number SERIALNUMB16 role line-card
user@device# set member 17 serial-number SERIALNUMB17 role line-card
user@device# set member 18 serial-number SERIALNUMB18 role line-card
user@device# set member 19 serial-number SERIALNUMB19 role line-card

```

7. (Recommended) Configure a virtual management Ethernet (VME) interface for management of the VCF configuration:

```

[edit]
user@device# set interfaces vme unit 0 family inet address /ip-address/mask/

```



**NOTE:** A VME accesses the device in the master Routing Engine role using a management port, so cable management port em0 or em1 on each spine device in your VCF so the VME is available regardless of which spine device assumes the master Routing Engine role. See *Connecting a QFX Series Device to a Management Console*

8. Commit the configuration:

```

user@device# commit

```

9. Interconnect the spine device that you configured in the previous steps to all leaf devices by using the supported SFP+ and QSFP+ interfaces.



**NOTE:** The automatic Virtual Chassis port (VCP) conversion feature is enabled and automatically configures SFP+ and QSFP+ interfaces into VCPs when the VCF configuration mode is set to **preprovisioned**. You do not need to manually configure VCPs.

If you want to configure an SFP+ or QSFP+ interface into a network interface, disable LLDP on that interface. See [“Configuring LLDP” on page 1345](#).



**NOTE:** Automatic VCP conversion only works when the interfaces on both ends of the link are not configured into VCPs.

The 40-Gbps QSFP+ interfaces on EX4300 switches are configured as VCPs, by default. You must, therefore, delete the VCP on the 40-Gbps QSFP+ interface using the `request virtual-chassis vc-port delete` command before interconnecting it into the VCF in order for the link to be converted into a VCP. You can also manually configure the link into a VCP using the `request virtual-chassis vc-port set` command.

10. Interconnect all other spine devices to all other leaf devices using the supported SFP+ and QSFP+ interfaces.
11. Install the VCF feature licenses.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show virtual-chassis** command.

After obtaining the licenses, follow the instructions in “[Generating License Keys](#)” on [page 75](#).

#### Related Documentation

- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)
- [Removing a Device From a Virtual Chassis Fabric on page 7069](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)

## Configuring a Nonprovisioned Virtual Chassis Fabric



**CAUTION:** Configure your Virtual Chassis Fabric (VCF) using autoprovisioning or preprovisioning unless you have a compelling reason to use nonprovisioned configuration. You can configure all aspects of a VCF using autoprovisioned or preprovisioned configuration. The process for autoprovisioning your VCF is described in “[Autoprovisioning a Virtual Chassis Fabric](#)” on [page 7053](#) and the process for preprovisioning your VCF is described in “[Preprovisioning a Virtual Chassis Fabric](#)” on [page 7056](#).

Nonprovisioned VCF configuration is highly discouraged. Nonprovisioned VCF configuration should only be used by VCF experts in specialized scenarios.

A nonprovisioned VCF is the configuration mode used when a VCF has not been configured into autoprovisioned or preprovisioned mode.

In a nonprovisioned VCF, you configure the device roles by setting the mastership priority value of each device. If no mastership priority values are set, a master election algorithm process runs and selects the role for each device.

You must manually configure all Virtual Chassis ports (VCPs) in a nonprovisioned VCF. The automatic VCP conversion feature, which automatically configures supported 10-Gbps SFP+ links and 40-Gbps QSFP+ links into VCPs on autoprovisioned and preprovisioned VCFs, is not supported on nonprovisioned VCFs.

Update all devices to the same version of Junos OS that supports VCF. See “[Upgrading Software](#)” on [page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To configure a nonprovisioned VCF:

1. Power on the devices.
2. Configure each individual device into fabric mode. If needed, configure the devices into mixed mode.

Reboot each device to complete this configuration step.

A VCF must have QFX5100 devices in the spine role, and operates most efficiently when the leaf nodes are also QFX5100 devices. Mixed mode must be configured if your VCF also includes at least one QFX3600, QFX3500, or EX4300 device in the leaf role.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF using at least one QFX3600, QFX3500, or EX4300 device as a leaf device:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** A device whose fabric or mixed mode setting is improperly set cannot join a VCF. You can check the mode settings using the `show virtual-chassis mode` command.

We recommend setting the fabric and mixed mode settings before interconnecting your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

We strongly recommend configuring the mixed and fabric settings before you interconnect a device into a VCF. You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a device into fabric or mixed mode after you have interconnected your VCF.

- 
3. After the device reboots are complete, cable your spine devices to your leaf devices using supported SFP+ and QSFP+ interfaces.
  4. (Recommended) Configure a virtual management Ethernet (VME) interface for management of the VCF configuration:

[edit]

```
user@device# set interfaces vme unit 0 family inet address /ip-address/mask/
```



**NOTE:** A VME accesses the device in the master Routing Engine role using a management port, so cable management port em0 or em1 on each spine device in your VCF so the VME is available regardless of which spine device assumes the master Routing Engine role. See *Connecting a QFX Series Device to a Management Console*

5. Configure the desired SFP+ and QSFP+ interfaces into Virtual Chassis ports (VCPs):

```
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
```

The `show virtual-chassis vc-port` must be issued on the ports at both ends of the link in order for that link to be configured into a VCP.

6. Enter the `show virtual-chassis` command to confirm that the VCPs are operational and to learn the member ID of each member device in your VCF.

If you want to change the member ID that has been assigned to a member device, use the `request virtual-chassis renumber` command.

7. (Optional) Configure the mastership priority for each member device:

```
[edit virtual-chassis]
user@device# set member member-id mastership-priority number
```

In a nonprovisioned VCF, member roles are determined by a mastership election algorithm. The first value checked by the mastership election algorithm is the mastership priority value. The two QFX5100 devices with the highest mastership priority values assume the master and backup Routing Engine role, which must be used by the spine devices in a VCF. All other devices assume the linecard role.

QFX5100 devices assume the Routing Engine role, regardless of mastership priority settings. QFX5100 devices can also assume the linecard role.

QFX3600, QFX3500, and EX4300 devices always assume the linecard role in a VCF, regardless of the mastership priority settings.



**NOTE:** A spine device that isn't selected as master or backup Routing Engine assumes the linecard role. The spine devices should still be configured with a higher mastership priority value than the leaf devices to assure a spine device assumes the Routing Engine role when the master or backup Routing Engine fails.

If two or more devices have the same mastership priority value and are candidates for the Routing Engine role, the mastership election algorithm uses other parameters to determine which device is elected into the Routing Engine role. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#).

A device with a mastership priority of 0 never assumes the master or backup role.

For instance, to configure the mastership priority for member devices 0 through 19 in your VCF.

```
[edit virtual-chassis]
```

```
user@device# set member 0 mastership-priority 255
user@device# set member 1 mastership-priority 255
user@device# set member 2 mastership-priority 255
user@device# set member 3 mastership-priority 255
user@device# set member 4 mastership-priority 100
user@device# set member 5 mastership-priority 95
user@device# set member 6 mastership-priority 90
user@device# set member 7 mastership-priority 85
user@device# set member 8 mastership-priority 80
user@device# set member 9 mastership-priority 75
user@device# set member 10 mastership-priority 70
user@device# set member 11 mastership-priority 65
user@device# set member 12 mastership-priority 60
user@device# set member 13 mastership-priority 55
user@device# set member 14 mastership-priority 50
user@device# set member 15 mastership-priority 45
user@device# set member 16 mastership-priority 40
user@device# set member 17 mastership-priority 35
user@device# set member 18 mastership-priority 30
user@device# set member 19 mastership-priority 25
```

8. Install the VCF feature licenses.

For a VCF deployment, two license keys are recommended for redundancy—one for the device in the master role and the other for the device in the backup role.

To purchase a feature license for VCF, contact your Juniper Networks sales representative (<http://www.juniper.net/us/en/contact-us/sales-offices>). The Juniper sales representative will provide you with the feature license files and license keys. You will be asked to supply the chassis serial number of your switch; you can obtain the serial number by running the **show virtual-chassis** command.

After obtaining the licenses, follow the instructions in “[Generating License Keys](#)” on [page 75](#).

**Related  
Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)
- [Removing a Device From a Virtual Chassis Fabric on page 7069](#)
- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)

## Adding a Device to a Virtual Chassis Fabric

This topic describes how to add a device to a Virtual Chassis Fabric (VCF).

It contains the following sections:

- [Adding a Leaf Device to an Autoprovisioned Virtual Chassis Fabric on page 7063](#)
- [Adding a Spine Device to an Autoprovisioned Virtual Chassis Fabric on page 7064](#)
- [Adding a Spine or Leaf Device to a Preprovisioned Virtual Chassis Fabric on page 7065](#)
- [Adding a Spine or Leaf Device to a Nonprovisioned Virtual Chassis Fabric on page 7067](#)



### Adding a Leaf Device to an Autoprovisioned Virtual Chassis Fabric

Update your device to the same version of Junos OS running on the devices in the VCF. See [“Upgrading Software” on page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a leaf device to an autoprovisioned VCF:

1. Log onto the device that you are adding to the VCF.
2. (Optional) Perform this step if you want to avoid the downtime associated with an extra reboot when your device is interconnected into your VCF. If you do not perform this step, the VCF auto-detects the fabric and mixed mode settings and, if needed, reboots the device as part of the process of changing these settings.

Configure the leaf device into fabric mode. Configure your device into mixed mode if your VCF includes QFX3600, QFX3500, or EX4300 devices as leaf devices..

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local
```

If you are configuring a VCF composed of two or more types of devices:

```
user@device> request virtual-chassis mode fabric mixed local
```

3. If the leaf device that you are adding to the VCF has not previously been configured, proceed to the next step.

If your device has been configured, zeroize your device and reboot:

```
user@device> request system zeroize
```

```
warning: System will be rebooted and may not boot without configuration
Erase all data, including configuration and log files? [yes,no] (yes) yes
```



**NOTE:** You must zeroize your device if you have previously entered one or more configuration commands, including basic configuration commands.

Your device will not join the VCF if it contains any configuration until it has been zeroized.



**NOTE:** The `request virtual-chassis mode fabric local` and `request virtual-chassis mode fabric mixed local` commands are entered in operational mode, so those settings are maintained when the device is zeroized.

You cannot use other methods to set a device into factory default mode before inserting it into a VCF if it was previously configured in another Virtual Chassis or VCF. You must use **`request system zeroize`**.

For additional information on this procedure, see *Reverting to the Default Factory Configuration for the EX Series Switch* or [“Reverting to the Default Factory Configuration” on page 188](#).

4. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode. Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

5. Interconnect your leaf device into the existing spine devices, using at least one 10-Gbps SFP+ interface or 40-Gbps QSFP+ interface to connect to each spine device in the VCF.

An autoprovisioned VCF automatically adds a supported device in factory-default mode to the VCF when it is connected to a spine devices using a supported SFP+ or QSFP+ link. The SFP+ or QSFP+ link is automatically converted into a Virtual Chassis port (VCP) as part of this process.

No further configuration is required.

---

### Adding a Spine Device to an Autoprovisioned Virtual Chassis Fabric

---

Update your device to the same version of Junos OS running on the devices in the VCF before interconnecting it into the VCF. See [“Upgrading Software” on page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a spine device to an autoprovisioned VCF:

1. Log in to your VCF.
2. If you are replacing a spine device that is already part of the VCF, power off the spine device in the VCF.

Follow the steps in [“Removing a Device From a Virtual Chassis Fabric” on page 7069](#) to remove the device from the VCF.

3. Modify the configuration.

If your new spine device is replacing an existing spine, modify the configuration to remove the old spine.

You can skip this step if you are not replacing an existing spine device.

```
[edit virtual-chassis]
```

```
user@device# delete member member-id
```

where *member-id* is the member ID of the spine that is removed from this procedure.

Add the spine device to the configuration:

```
[edit virtual-chassis]
```

```
user@device# set member member-id serial-number serial-number role routing-engine
```

For instance, to configure a spine device with the serial number OU81234567890 as member 3:

```
[edit virtual-chassis]
```

```
user@device# set member 3 serial-number OU81234567890 role routing-engine
```

4. Commit the configuration.

```
[edit]
```

```
user@device# commit
```

5. Log in to the device that is going to be added to the VCF.
6. Configure the device into fabric mode. If needed, also configure the device into mixed mode.

Reboot the device to complete this configuration step.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF composed of QFX5100 devices and at least one other type of device:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** We recommend setting the fabric and mixed mode settings before interconnecting your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to set a device into fabric or mixed mode after interconnecting your VCF.

7. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode. Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

8. After the device reboots, interconnect the new device into the VCF by cabling the device to the leaf devices in the VCF using supported SFP+ or QSFP+ interfaces.

The SFP+ or QSFP+ links are converted into VCPs automatically.

The new spine device should be operational once the cabling is complete.

### Adding a Spine or Leaf Device to a Preprovisioned Virtual Chassis Fabric

Update your device to the same version of Junos OS running on the devices in the VCF before interconnecting it into the VCF. See [“Upgrading Software” on page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a spine or leaf device to a preprovisioned VCF:

1. Log in to your VCF.
2. If you are replacing a device that is already part of the VCF, power off the device in the VCF.

Follow the steps in [“Removing a Device From a Virtual Chassis Fabric”](#) on page 7069 to remove the device from the VCF.

3. Modify the configuration.

If your new device is replacing an existing device, modify the configuration to remove the old device.

You can skip this portion of the procedure if you are not replacing an existing device.

```
[edit virtual-chassis]
user@device# delete member member-id
```

where *member-id* is the member ID of the spine that is removed from this procedure.

Add the new device to the configuration:

```
[edit virtual-chassis]
user@device# set member member-id serial-number serial-number role routing-engine
```

For instance, to configure a device with the serial number OU81234567890 into the Routine Engine role as member 3:

```
[edit virtual-chassis]
user@device# set member 3 serial-number OU81234567890 role routing-engine
```

4. Commit the configuration.

```
[edit]
user@device# commit
```

5. Log in to the device that is going to be added to the VCF.

6. Configure the device into fabric mode. If needed, also configure the device into mixed mode.

Reboot the device to complete this configuration step.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF composed of two or more types of devices:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** If you are adding a QFX3600, QFX3500, or EX4300 device to a VCF that is composed entirely of QFX5100 devices, you must also log in to the VCF and set all of the devices in the VCF into mixed mode.

Log in to the VCF and enter the **request virtual-chassis mode mixed all-members reboot** command to perform this task.

The VCF reboots and incurs downtime to complete this procedure.

---



**NOTE:** We recommend that you set the fabric and mixed mode settings before you interconnect your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the `request virtual-chassis mode fabric local` or `request virtual-chassis mode mixed local` commands to recover a device that was not set into fabric or mixed mode before you interconnect it into your VCF.

7. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode. Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

8. After the device reboots, interconnect the new device into the VCF using supported SFP+ or QSFP+ interfaces.

The SFP+ or QSFP+ links are converted into VCPs automatically.

The new device should be operational shortly after the cabling is complete.

### Adding a Spine or Leaf Device to a Nonprovisioned Virtual Chassis Fabric



**CAUTION:** Configure your VCF using autoprovisioning or preprovisioning unless you have a compelling reason to use nonprovisioned configuration. You can configure all aspects of a VCF using autoprovisioned or preprovisioned configuration.

Nonprovisioned VCF configuration is highly discouraged. Nonprovisioned VCF configuration should only be used by VCF experts in specialized scenarios.

Update your device to the same version of Junos OS running on the devices in the VCF before interconnecting it into the VCF. See [“Upgrading Software” on page 134](#) or *Installing Software on an EX Series Switch with a Single Routing Engine (CLI Procedure)*.

To add a spine or leaf device to a nonprovisioned VCF:

1. Log in to your VCF.
2. If you are replacing a device that is already part of the VCF, power off the device in the VCF. Uncable the device once the power off is complete.

You can skip this step if you are adding a new device without replacing an existing device. You must skip this step if there is no configuration for the device that you are removing from the VCF.

If the device is configured, delete the device from the VCF configuration:

```
[edit virtual-chassis]
user@device# delete member member-id
```

where *member-id* is the member ID of the device that you are removing.

3. Log in to the device that you are going to add to the VCF.
4. Configure the device into fabric mode. If needed, also configure the device into mixed mode.

Reboot the device to complete this configuration step.

If you are configuring a VCF composed entirely of QFX5100 devices:

```
user@device> request virtual-chassis mode fabric local reboot
```

If you are configuring a VCF that includes at least one QFX3600, QFX3500, or EX4300 devices as a leaf device:

```
user@device> request virtual-chassis mode fabric mixed local reboot
```



**NOTE:** If you are adding a QFX3600, QFX3500, or EX4300 device to a VCF that is composed entirely of QFX5100 devices, you must also log in to the VCF and set all of the devices in the VCF into mixed mode.

Log in to the VCF and enter the **request virtual-chassis mode mixed all-members reboot** command to perform this task.

The VCF reboots and incurs downtime to complete this procedure.



**NOTE:** We recommend that you set the fabric and mixed mode settings before you interconnect your devices into a VCF to avoid the following issues:

- Incurring downtime as the devices reboot to commit the mixed mode or fabric settings.
- Manually correcting potential issues related to VCF formation because the device did not immediately join the VCF.

You can, however, use the **request virtual-chassis mode fabric local** or **request virtual-chassis mode mixed local** commands to set a device into fabric or mixed mode after interconnecting your VCF.

5. (Required only if you are adding a device that turns a non-mixed VCF into a mixed VCF) Log in to the VCF and set all devices in the VCF to mixed mode, Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed all-members reboot
```

The VCF experiences downtime as part of the reboot procedure.

6. After the device reboots, interconnect it into the VCF using supported SFP+ or QSFP+ interfaces.

7. Configure the SFP+ or QSFP+ interfaces into Virtual Chassis ports (VCPs):

```
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
user@device> request virtual-chassis vc-port set pic-slot pic-slot-number port port-number
```

The **request virtual-chassis vc-port** must be configured on the ports at both ends of the link in order for that link to be configured into a VCP.

8. (Optional) Log in to the VCF and set the mastership priority of the new device:

```
[edit virtual-chassis]
user@device# set member member-id mastership-priority number
```

If needed, enter the **show virtual-chassis** command to learn the member ID of the new member device in the VCF.

#### Related Documentation

- [Removing a Device From a Virtual Chassis Fabric on page 7069](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)

## Removing a Device From a Virtual Chassis Fabric

This topic describes how to remove a device from a Virtual Chassis Fabric (VCF):

To remove a device from a VCF:

1. Power off the device that you are removing from the VCF.
2. Uncable the device that you are removing from the VCF.
3. Log in to the Virtual Management ethernet (VME) interface. Remove the device from the VCF configuration.

You can skip this step if you are removing a device that was never configured.

```
[edit virtual-chassis]
user@device# delete member member-id
```

4. Delete the Virtual Chassis port (VCP) or ports on the devices that are still in the VCF but were connected to the removed device.

```
user@device> request virtual-chassis vc-port delete pic-slot pic-slot port port-number member member-id
```

When a device is removed from a VCF, the interface on the other end of the VCP link that was connected to the removed device remains configured as a VCP.

You can check the results of this command using the **show virtual-chassis vc-port** command.

5. (Required only if you are removing a device that turns a mixed VCF into a homogenous VCF) Log in to the VCF and disable mixed mode for all of the devices in the VCF, Configure all devices to reboot to complete this procedure.

```
user@device> request virtual-chassis mode mixed disable all-members reboot
```

This step should only be taken if you are removing a QFX3600, QFX3500, or EX4300 device from a mixed VCF and the only devices remaining in the VCF are QFX5100 devices.

The VCF experiences downtime as part of the reboot procedure.

6. Commit the configuration.

```
[edit]
user@device# commit
```

7. Power on the device that was removed from the VCF, and log in to it.
8. (Optional, but recommended) Delete the VCP or VCPs on the device that was removed:

```
user@device> request virtual-chassis vc-port delete pic-slot pic-slot port port-number member member-id
```

9. (Optional, but recommended) Reset the fabric and mixed mode settings.

If you are removing a device that was part of a VCF composed entirely of the same device:

```
user@device> request virtual-chassis mode fabric disable reboot
```

If you are removing a device that was part of a VCF composed of two or more device types:

```
user@device> request virtual-chassis mode fabric mixed disable reboot
```

Reboot the device to complete the process.

We recommend resetting the fabric and mixed mode settings immediately after removing it from the VCF to avoid any potential issues with your device if it is placed in your network in another role.

#### Related Documentation

- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Understanding Virtual Chassis Fabric Configuration on page 7043](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)

## Upgrading Software for a Virtual Chassis Fabric

This topic describes the processes that can be used to update software on an operational Virtual Chassis Fabric (VCF).

You should update the software on each device before initially interconnecting your VCF. This process describes the options that are available for upgrading software after a VCF is setup.



It contains the following sections:

- [NSSU on page 7071](#)
- [Automatic Software Update on page 7071](#)
- [Standard Upgrade on page 7071](#)

## NSSU

Nonstop software upgrade (NSSU) enables you to upgrade the software running on all member devices in a VCF with minimal network traffic disruption during the upgrade.

See “Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade” on page 139.

## Automatic Software Update

Automatic software update automatically upgrades the Junos OS running on a device joining a VCF to the version of Junos OS running on the VCF at the moment the new device is cabled into the VCF.

To configure the automatic software update feature for a VCF composed entirely of QFX5100 devices:

[edit]

```
user@device# set virtual-chassis auto-sw-update package-name package-name
```

To configure the automatic software update feature on a VCF composed of QFX5100 devices and at least one other type of device:

[edit]

```
user@device# set virtual-chassis auto-sw-update qfx-5 package-name package-name
```

```
user@device# set virtual-chassis auto-sw-update qfx-3 package-name package-name
```

```
user@device# set virtual-chassis auto-sw-update ex-4300 package-name package-name
```

where **qfx-5** specifies the path to the Junos OS used to run a QFX5100 devices, **qfx-3** specifies the path to the Junos OS used to run QFX3600 and QFX3500 devices, and **ex4300** specifies the path to the Junos OS used to run EX4300 switches.

If the software package is located on a local directory on the switch, use the following format for *package-name*:

***/pathname/package-name***

If the software package is to be downloaded and installed from a remote location, use one of the following formats:

***ftp://hostname/pathname/package-name***

***ftp://username:prompt@ftp.hostname.net/package-name***

***http://hostname/pathname/package-name***

## Standard Upgrade

You can upgrade software on a VCF using the traditional method of upgrading software for Junos OS by logging onto the master Routing Engine and using the **request system software add** command to initiate the upgrade on a non-mixed VCF or the **request system**

**software add set** [*package-name package-name ...*] to initiate the upgrade on a mixed VCF, where *package-name* is the path to an image for one device family.

When you upgrade Junos OS on a VCF using the traditional software upgrade, each device in the VCF must reboot. The entire system is down until the upgrade process is complete.

For information on performing this procedure, see [“Upgrading Software” on page 134](#).

**Related  
Documentation**

- [Adding a Device to a Virtual Chassis Fabric on page 7062](#)
- [Understanding Software Upgrades in a Virtual Chassis Fabric on page 7050](#)

---

## Configuration Statements

- [\[edit virtual-chassis\] Configuration Statement Hierarchy on page 7072](#)
- [aliases \(Virtual Chassis\) on page 7075](#)
- [alias-name \(Virtual Chassis aliases\) on page 7076](#)
- [auto-provisioned on page 7077](#)
- [auto-sw-update on page 7078](#)
- [enhanced-hash-key on page 7080](#)
- [fabric-load-balance on page 7082](#)
- [id on page 7083](#)
- [inactivity-interval \(Fabric Load Balance\) on page 7084](#)
- [location \(Virtual Chassis\) on page 7085](#)
- [mac-persistence-timer on page 7086](#)
- [mastership-priority on page 7087](#)
- [member on page 7089](#)
- [no-management-vlan on page 7090](#)
- [no-split-detection on page 7091](#)
- [package-name on page 7092](#)
- [preprovisioned on page 7093](#)
- [role on page 7094](#)
- [serial-number on page 7097](#)
- [serial-number \(Virtual Chassis aliases\) on page 7098](#)
- [traceoptions \(Virtual Chassis\) on page 7099](#)
- [virtual-chassis on page 7102](#)

### [edit virtual-chassis] Configuration Statement Hierarchy

This topic lists supported and unsupported configuration statements in the **[edit virtual-chassis]** hierarchy level on EX Series and QFX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.

- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms.

For detailed information about feature support on specific EX Series or QFX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 7073](#)
- [Unsupported Statements in the \[edit virtual-chassis\] Hierarchy Level on page 7074](#)

### Supported Statements in the [edit virtual-chassis] Hierarchy Level

The following hierarchy shows the **[edit virtual-chassis]** configuration statements supported on EX Series or QFX Series switches:

```
virtual-chassis {
  aliases {
    serial-number serial-number {
      alias-name alias-name;
    }
  }
  auto-provisioned;
  auto-sw-update {
    (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
    package-name package-name;
  }
  fast-failover (ge | vcp disable | xe);
  graceful-restart {
    disable;
  }
  id id;
  mac-persistence-timer [minutes | disable];;
  member member-id {
    location location;
    mastership-priority number;
    no-management-vlan;
    role (line-card | routing-engine);
    serial-number;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag ;
  }
  vc-port {
    lag-hash (packet-based | source-port-based);
  }
  vcp-no-hold-time;
}
```

### Unsupported Statements in the [edit virtual-chassis] Hierarchy Level

All statements in the **[edit virtual-chassis]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

#### **Related Documentation**

- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Configuring an EX4300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX2200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)

## aliases (Virtual Chassis)

<b>Syntax</b>	<pre>aliases {   serial-number serial-number {     alias-name alias-name;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.
<b>Description</b>	<p>Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.</p> <p>An alias is not specified for a device until the alias name is specified using the <b>alias-name</b> keyword.</p> <p>The alias appears in the <b>Alias-Name</b> field in the <b>show virtual-chassis</b> command.</p> <p>Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 7035</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li> </ul>

## alias-name (Virtual Chassis aliases)

---

**Syntax** `alias-name alias-name;`

**Hierarchy Level** `[edit virtual-chassis aliases serial-number serial-number]`

**Release Information** Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series switches.

**Description** Create an alias for a member switch in a Virtual Chassis or Virtual Chassis Fabric (VCF). An alias allows you to more clearly identify the member switches in your Virtual Chassis or VCF by assigning a text label to a member switch's serial number.

The alias appears in the **Alias-Name** field in the **show virtual-chassis** command.

Alias usage is optional and aliases are used for administrative purposes only. Setting an alias has no effect on the operation of the member switch.

In the following example, the **dc-floor-1** alias name is assigned to the member switch with the serial number AB0123456789.

### [set serial-number](#)

```
[edit virtual-chassis aliases]
user@switch# set serial-number AB0123456789 alias-name dc-floor-1
```

### [show virtual-chassis](#)

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 9d5d.5556.919a
Fabric Mode: Enabled

Member ID  Status  Serial No  Alias-Name  Model  Mstr  prio  Role
0 (FPC 0)  Prsnt    AB0123456789  dc-floor-1  qfx5100-48s-6q  129  Master
<additional output removed for brevity>
```

**Options** *alias-name*—The text label, or alias, assigned to the member switch by the user.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Understanding Virtual Chassis Fabric Components on page 7035](#)
- [Understanding QFX Series Virtual Chassis Components on page 6909](#)

## auto-provisioned

---

<b>Syntax</b>	auto-provisioned;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Enable the auto-provisioned configuration mode for a Virtual Chassis Fabric (VCF).</p> <p>When a VCF is autoprovisioned, you can plug and play leaf devices that have not been configured or are zeroized into your VCF without user configuration. The leaf devices are automatically configured into the linecard role and all other VCF configuration—configuring Virtual Chassis ports (VCPs), the member ID, fabric mode, mixed mode (if applicable), and other parameters—is completed without further user action when a supported spine device interconnects to the leaf device by using a 10-Gbps SFP+ or 40-Gbps QSFP+ link that can be converted into a VCP.</p> <p>A leaf device whose fabric or mixed mode setting is changed as part of the autoprovisioning process automatically reboots. You can avoid this reboot by configuring the fabric or mixed mode setting on the leaf device before interconnecting into the VCF.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Understanding Virtual Chassis Fabric Configuration on page 7043</a></li> </ul>

## auto-sw-update

---

<b>Syntax</b>	<pre>auto-sw-update {     (ex-4200   ex-4300   ex-4500   ex-4600   qfx-3   qfx-5)     package-name package-name; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>The <b>ex-4200</b> and <b>ex-4500</b> options introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>The <b>ex-4300</b>, <b>qfx-3</b>, and <b>qfx-5</b> options introduced in Junos OS Release 13.2X51-D20.</p> <p>The <b>ex-4600</b> option introduced in Junos OS Release 13.2X51-D25.</p>
<b>Description</b>	<p>Enable the automatic software update feature for Virtual Chassis or Virtual Chassis Fabric (VCF) configurations.</p> <p>You should only use the keywords that specify a device—<b>ex-4300</b>, <b>ex-4600</b>, <b>qfx-3</b>, and <b>qfx-5</b>—when configuring automatic software update on a mixed Virtual Chassis or Virtual Chassis Fabric (VCF). You can simply specify the <i>package-name</i> without specifying the device keywords in non-mixed Virtual Chassis or VCF topologies.</p> <p>You must enter the <b>auto-sw-update</b> statement multiple times—once for each device family in your mixed Virtual Chassis or VCF—in most scenarios when enabling the automatic software update for a mixed Virtual Chassis or VCF.</p> <p>The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. You do not, therefore, need to specify the <b>ex-4500</b> keyword when configuring automatic software update for a mixed Virtual Chassis that include EX4500 and EX4550 switches only. You also only have to enter the <b>ex-4500</b> keyword once to configure automatic software update for all EX4500 and EX4550 member switches in the same mixed Virtual Chassis.</p> <p>The Junos OS package for a QFX3500 device updates the software for QFX3500 and QFX3600 devices. You do not, therefore, need to specify the <b>qfx-3</b> keyword when configuring automatic software update for a Virtual Chassis composed entirely of QFX3500 and QFX3600 devices. You also have to enter the <b>qfx-3</b> keyword only once to configure automatic software update for all QFX3500 and QFX3600 member devices in the same mixed Virtual Chassis.</p> <p>The remaining statement is explained separately.</p>
<b>Default</b>	The automatic software update feature is disabled.
<b>Options</b>	<p><b>package-name package-name</b>—Specify a path to a Junos OS software image.</p> <p><b>ex-4200</b>—Specify a path to a Junos OS image for an EX4200 switch when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed</p>



EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

**ex-4300**—Specify a path to a Junos OS image for an EX4300 switch when enabling automatic software update for a mixed Virtual Chassis or VCF.

**ex-4500**—Specify a path to a Junos OS image for an EX4500 switch, an EX4550 switch, or both types of switches when enabling automatic software update for a mixed EX4200 and EX4500 Virtual Chassis, mixed EX4200 and EX4550 Virtual Chassis, or mixed EX4200, EX4500, or EX4550 Virtual Chassis.

The Junos OS package for an EX4500 switch updates the software for EX4500 and EX4550 switches. Therefore, you only enter this command once to upgrade the EX4500 and EX4550 member switches in the same mixed Virtual Chassis.

The **ex-4500** keyword also does not need to be specified when configuring automatic software update for a mixed EX4500 and EX4550 Virtual Chassis.

**ex-4600**—Specify a path to a Junos OS image for an EX4600 switch when enabling automatic software update for a mixed Virtual Chassis.

**qfx-3**—Specify a path to a Junos OS image for a QFX3500, QFX3600, or both types of devices when enabling automatic software update for a mixed VCF.

**qfx-5**—Specify a path to a Junos OS image for a QFX5100 device when enabling automatic software update for a mixed VCF.

<b>Required Privilege Level</b>	system—To view this statement in the configuration.
	system-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</i></li> <li>• <a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 6944</a></li> <li>• <a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric on page 7050</a></li> </ul>
------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## enhanced-hash-key

---

**Syntax**    enhanced-hash-key {  
              ecmp-resilient-hash;  
              fabric-load-balance {  
                  flowlet {  
                      inactivity-interval *interval*;  
                  }  
                  per-packet;  
              }  
              hash-mode {  
                  layer2-header;  
                  layer2-payload;  
              }  
              inet {  
                  no-ipv4-destination-address;  
                  no-ipv4-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-protocol;  
                  vlan-id;  
              }  
              inet6 {  
                  no-ipv6-destination-address;  
                  no-ipv6-source-address;  
                  no-l4-destination-port;  
                  no-l4-source-port;  
                  no-next-header;  
                  vlan-id;  
              }  
              layer2 {  
                  no-destination-mac-address;  
                  no-ether-type;  
                  no-source-mac-address;  
                  vlan-id;  
              }  
          }  
      }

**Hierarchy Level**    [edit forwarding-options]

**Release Information**    Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.  
                              Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.  
                              The **fabric-load-balance** statement introduced in Junos OS Release 14.1X53-D10.

**Description**    Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The remaining statements are explained separately.

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2590</a></li><li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396</a></li></ul>
------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## fabric-load-balance

---

<b>Syntax</b>	<pre>fabric-load-balance {     flowlet {         inactivity-interval interval;     }     per-packet; }</pre>
<b>Hierarchy Level</b>	[edit forwarding-options <b>enhanced-hash-key</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10.
<b>Description</b>	<p>Enable adaptive load balancing (ALB) for a VCF, and specify how ALB is implemented.</p> <p>When ALB is enabled, the Virtual Chassis ports (VCPs) are reset. Packets are dropped and might potentially arrive out of order for a brief period of time as a result of this VCP reset. Normal operation of the VCF resumes after the VCP reset with no further user action.</p>
<b>Default</b>	<p>ALB is disabled, by default.</p> <p>If you do not specify a mode when enabled ALB, ALB is enabled using flowlet mode with an inactivity timer of 16 microseconds.</p>
<b>Options</b>	<p><b>flowlet</b>—Implement ALB by using flowlets.</p> <p>When ALB is implemented using flowlets, traffic flows that enter the VCF are spliced into smaller flows—flowlets—and individually forwarded across the VCF to the same destination device over different paths when the inactivity time between packet bursts on the sending interface exceeds the user-configurable inactivity interval.</p> <p>The inactivity interval is 16 microseconds by default, and can be configured using the <b>inactivity-interval</b> statement. You should configure the inactivity interval to ensure in-order packet delivery, so that overall performance is not negatively impacted by the packet reordering process at the receiving device. To ensure in-order packet delivery, the inactivity interval should be larger than the largest latency skew among all the paths in the VCF from any node to any other node.</p> <p>Implementing ALB using flowlets is especially effective in environments that periodically experience extremely large traffic flows—<i>elephant flows</i>—that are substantially larger than the majority of other traffic flowing through the VCF. The VCF is better able to manage elephant flows by splicing them into smaller flowlets using ALB.</p> <p><b>per-packet</b>—Implement ALB using per-packet mode.</p> <p>When per-packet mode is enabled, the VCF forwarding algorithm dynamically monitors all paths in the VCF and forwards packets to destination devices using the best available path at that moment. Flows are reordered at the destination node when per-packet mode is used to enable ALB, so some performance impact due to packet reordering is experienced.</p>

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- [Understanding Traffic Flow Through a Virtual Chassis Fabric on page 7050](#)

## id

**Syntax** `id id;`

**Hierarchy Level** [edit [virtual-chassis](#)]

**Release Information** Statement introduced in Junos OS Release 9.3 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

**Description** Configure the alphanumeric string that identifies a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.

**Options** *id*—Virtual Chassis ID (VCID), which uses the ISO family address format—for example, **9622.6ac8.5345**.

**Required Privilege Level** system—To view this statement in the configuration.  
system-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge*
- [Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge \(CLI Procedure\) on page 6946](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Understanding Virtual Chassis Member ID Numbering in an EX8200 Virtual Chassis*

## inactivity-interval (Fabric Load Balance)

---

<b>Syntax</b>	<code>inactivity-interval <i>interval</i>;</code>
<b>Hierarchy Level</b>	[edit forwarding-options <a href="#">enhanced-hash-key fabric-load-balance</a> flowlet]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10.
<b>Description</b>	<p>Set the inactivity interval for adaptive load balancing (ALB) using flowlets within a VCF.</p> <p>The inactivity interval is the amount of time that occurs between packet bursts on a sending interface before a traffic flow is spliced into smaller traffic flows—flowlets—when ALB is implemented using flowlets. The flowlets are then individually forwarded across the VCF to the same destination device over different paths.</p> <p>You should configure the inactivity interval to ensure in-order packet delivery, so that overall performance is not negatively impacted by the packet re-ordering process at the receiving device. To ensure in-order packet delivery, the inactivity interval should be larger than the largest latency skew among all the paths in the VCF from any node to any other node.</p>
<b>Default</b>	<p>ALB is disabled, by default.</p> <p>If ALB is enabled without specifying a mode, ALB is enabled using flowlet mode with an inactivity interval of 16 microseconds.</p> <p>If ALB is enabled using flowlet mode without specifying an inactivity interval, the inactivity interval is set to 16 microseconds.</p>
<b>Options</b>	<p><b><i>interval</i></b>—The amount of time that occurs between packet bursts on a sending interface before a traffic flow is spliced into flowlets.</p> <p><b>Range:</b> 16 microseconds (<b>16us</b>) to 32 milliseconds(<b>32ms</b>).</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Traffic Flow Through a Virtual Chassis Fabric on page 7050</a></li></ul>

## location (Virtual Chassis)

<b>Syntax</b>	<code>location location;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis member member-id</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Set a description of the location of the Virtual Chassis or VCF member switch or external Routing Engine.</p> <p>The <b>Location</b> field is visible to users who enter the <b>show virtual-chassis status detail</b> command.</p> <p>Setting this description has no effect on the operation of the member device.</p>
<b>Options</b>	<b>location</b> —Location of the current member switch or external Routing Engine. The <b>location</b> can be any single word.
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <i>Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</i></li> <li>• <i>Example: Configuring a Preprovisioned Mixed EX4200 and EX4500 Virtual Chassis</i></li> <li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li> <li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li> </ul>

## mac-persistence-timer

---

<b>Syntax</b>	<code>mac-persistence-timer [<i>minutes</i>   <b>disable</b>];</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>disable</b> introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>The maximum timer limit changed from no maximum timer limit to 60 minutes in Junos OS Release 12.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Specify how long the Virtual Chassis or VCF continues to use the MAC address of the switch that was originally configured in the master role as the system MAC base address after the original master switch is removed from the Virtual Chassis or VCF. The system MAC base address does not change in the event of a switchover provided the switch originally configured in the master role remains a member of the Virtual Chassis or VCF.</p> <p>The maximum timer limit is 60 minutes starting in Junos OS Release 12.2. There are no minimum or maximum timer limits in prior Junos OS releases.</p>
<b>Default</b>	The MAC persistence timer is set to 10 minutes by default.
<b>Options</b>	<p><b>minutes</b>—Time in minutes that the member switch in the backup role continues to use the system MAC base address of the old master before using its own system MAC base address after the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p> <p><b>disable</b>—Disable the MAC persistence timer. The system MAC base address never changes when the MAC persistence timer is disabled, even when the switch in the master role is physically disconnected or removed from the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	<p><b>system</b>—To view this statement in the configuration.</p> <p><b>system-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of a Virtual Chassis (CLI Procedure) on page 6943</a></li><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li></ul>



## mastership-priority

<b>Syntax</b>	<code>mastership-priority <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis member</a> <i>member-id</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Mastership priority option <b>0</b> introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>The mastership priority value is the most important factor in determining the role of the member switch within a nonprovisioned Virtual Chassis or VCF configuration. Other factors (see <a href="#">“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917</a>) also affect the election of the master.</p> <p>The mastership priority value takes the highest precedence in the master election algorithm. The member switch with highest mastership priority assumes the master Routing Engine role of the Virtual Chassis or VCF. Toggling back and forth between master and backup status in failover conditions is undesirable, so we recommend that you assign the same mastership priority value to both the master and the backup. Secondary factors in the master election algorithm determine which of these two members (that is, the two members that are assigned the highest mastership priority value) functions as the master of the Virtual Chassis or VCF.</p> <p>This statement is not used for the EX8200 Virtual Chassis, which determines mastership by external Routing Engine uptime. See <i>Understanding Virtual Chassis Roles in an EX8200 Virtual Chassis</i>.</p> <p>A switch with a mastership priority of <b>0</b> never takes the master or backup role.</p>
<b>Default</b>	128
<b>Options</b>	<p><i>number</i>—Mastership priority value.</p> <p><b>Range:</b> 0 through 255</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</a></li> </ul>

- *Example: Configuring an EX4200 Virtual Chassis Interconnected Across Multiple Wiring Closets*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*

## member

<b>Syntax</b>	<pre>member <i>member-id</i> {   location <i>location</i>;   mastership-priority <i>number</i>;   no-management-vlan;   serial-number <i>serial-number</i>;   role <i>role</i>; }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Configure a switch or an XRE200 External Routing Engine as a member of a Virtual Chassis or a Virtual Chassis Fabric (VCF).
<b>Default</b>	<p>When an EX Series switch or a QFX Series devices configured in standalone mode is powered on but not interconnected through its Virtual Chassis ports (VCPs) with other member switches, its default member ID is 0.</p> <p>There is no default member ID in an EX8200 or EX9200 Virtual Chassis. An EX8200 or EX9200 Virtual Chassis must be preprovisioned, and that process configures the member IDs.</p>
<b>Options</b>	<p><b><i>member-id</i></b>—Identifies a specific member switch of a Virtual Chassis or VCF configuration.</p> <p>The exact range for a specific Virtual Chassis or VCF depends on the number of switches allowed in the Virtual Chassis or VCF.</p> <p>In an EX8200 Virtual Chassis, member IDs 0 through 7 are reserved for EX8200 member switches and member IDs 8 and 9 are reserved for the master and backup external Routing Engines.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</a></li> <li>• <a href="#">Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</a></li> </ul>

- [Configuring an EX3300 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX4200, EX4500, or EX4550 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX8200 Virtual Chassis \(CLI Procedure\)](#)
- [Configuring an EX9200 Virtual Chassis](#)
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)

---

## no-management-vlan

---

<b>Syntax</b>	no-management-vlan;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis member member-id</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Remove the specified member's out-of-band management port from the virtual management Ethernet (VME) global management VLAN of the Virtual Chassis or VCF configuration.</p> <p>For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:</p> <pre>virtual-chassis {   member 2 {     no-management-vlan;   } }</pre> <p>You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell <b>ifconfig</b> command.</p>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Setting Up a Multimember EX4200 Virtual Chassis Access Switch with a Default Configuration</a></li><li>• <a href="#">Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)</a></li><li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 6919</a></li><li>• <a href="#">Understanding Virtual Chassis Fabric Configuration on page 7043</a></li></ul>

## no-split-detection

---

<b>Syntax</b>	no-split-detection;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Disable the split and merge feature in a Virtual Chassis or VCF configuration.</p> <p>We recommend using this statement to disable the split and merge feature when configuring a two-member Virtual Chassis. Enabling this statement on a two-member Virtual Chassis ensures that both switches remain in the correct Virtual Chassis roles in the event of a Virtual Chassis split.</p>
<b>Default</b>	The split and merge feature is enabled.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Assigning the Virtual Chassis ID to Determine Precedence During an EX4200 Virtual Chassis Merge</i></li> <li>• <a href="#">Disabling Split and Merge in a Virtual Chassis (CLI Procedure) on page 6944</a></li> <li>• <a href="#">Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 6946</a></li> <li>• <a href="#">Understanding Split and Merge in a Virtual Chassis on page 6922</a></li> </ul>

## package-name

---

<b>Syntax</b>	<code>package-name <i>package-name</i>;</code>
<b>Hierarchy Level</b>	[edit virtual-chassis <a href="#">auto-sw-update</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Specify the software package name or location of the software package to be used by the automatic software update feature for Virtual Chassis or VCF.
<b>Default</b>	No package name is specified.
<b>Options</b>	<p><b><i>package-name</i></b>—Name of the software package or the URL to the software package to be used.</p> <ul style="list-style-type: none"><li>If the software package is located on a local directory on the switch, use the following format for <b><i>package-name</i></b>:  <b><i>/pathname/package-name</i></b></li><li>If the software package is to be downloaded and installed from a remote location, use one of the following formats:  <b><i>ftp://hostname/pathname/package-name</i></b> <b><i>ftp://username:prompt@ftp.hostname.net/package-name</i></b> <b><i>http://hostname/pathname/package-name</i></b></li></ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><a href="#">Example: Configuring Automatic Software Update on EX4200 Virtual Chassis Member Switches</a></li><li><a href="#">Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 6944</a></li><li><a href="#">Understanding Software Upgrades in a Virtual Chassis Fabric on page 7050</a></li></ul>

## preprovisioned

---

<b>Syntax</b>	preprovisioned;
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Enable the preprovisioned configuration mode for a Virtual Chassis or Virtual Chassis Fabric (VCF) configuration.</p> <p>When the preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the mastership priority or member ID of member switches.</p> <p>You must use this statement to configure an EX8200 Virtual Chassis. Nonprovisioned configuration of an EX8200 Virtual Chassis is not supported.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <i>Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File</i></li> <li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li> <li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li> <li>• <i>Configuring an EX9200 Virtual Chassis</i></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 6938</a></li> </ul>

## role

---

<b>Syntax</b>	<code>role (line-card   routing-engine);</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis preprovisioned member</a> <i>member-id</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Specify the roles of the members of the Virtual Chassis or a Virtual Chassis Fabric (VCF) in a preprovisioned Virtual Chassis.

### Virtual Chassis Fabric

Specify the role to be performed by each switch. In a VCF, the spine devices are configured into the Routing Engine role and the leaf devices are configured into the line card role. You can configure several devices into the Routine Engine role, but only two will operate in the Routing Engine role at a time. The role must be associated with the member's serial number.

### EX Series (except EX8200 Virtual Chassis) and QFX Series Virtual Chassis

Specify the role to be performed by each member switch. Associate the role with the member's serial number.

When you use a preprovisioned configuration, you cannot modify the mastership priority or member ID of member switches through the user interfaces. The mastership priority value is generated by the software, based on the assigned role:

- A member configured as **routing-engine** is assigned the mastership priority **129**.
- A member configured as **line-card** is assigned the mastership priority **0**.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the mastership priority **128**.

The configured role specifications are permanent. If both **routing-engine** members fail, a **line-card** member cannot take over as master of the Virtual Chassis configuration. You must delete the preprovisioned configuration to change the specified roles in a Virtual Chassis.

Explicitly configure two members as **routing-engine** and configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional members, they function in a linecard role by default. In that case, a member that is functioning in a linecard role can take over mastership if the members functioning as master and backup (**routing-engine** role) both fail.

### EX8200 Virtual Chassis



Specify the role to be performed by each XRE200 External Routing Engine and each EX8200 member switch. Associate the role with the member's serial number. An EX8200 Virtual Chassis cannot function when both external Routing Engines, which must be configured in the **routing-engine** role, have failed.

- Options**
- **line-card**—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis or VCF configuration other than the master or backup functions in the linecard role and runs only a subset of Junos OS for EX Series switches. A member functioning in the linecard role does not run the control protocols or the chassis management processes.

A Virtual Chassis must have at least three members for one member to function in the linecard role.

In an EX8200 Virtual Chassis configuration, all member switches must be in the linecard role.

- **routing-engine**—Enables the member to function as a master or backup of the Virtual Chassis or VCF configuration. The master manages all members and runs the chassis management processes and control protocols. The backup synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

(All Virtual Chassis composed of EX Series switches, except EX8200 switches, or QFX Series devices) Specify two and only two members as **routing-engine**. The software determines which of the two members assigned the **routing-engine** role functions as master, based on the master election algorithm. See [“Understanding How the Master in a Virtual Chassis Is Elected” on page 6917](#). In these Virtual Chassis, the **routing-engine** role is associated with a switch.

(EX8200 Virtual Chassis) All XRE200 External Routing Engines must be in the **routing-engine** role.

<b>Required Privilege</b>	system—To view this statement in the configuration.
<b>Level</b>	system-control—To add this statement to the configuration.

**Related  
Documentation**

- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
- [Preprovisioning a Virtual Chassis Fabric on page 7056](#)
- *Example: Configuring an EX4200 Virtual Chassis Using a Preprovisioned Configuration File*
- *Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines*
- *Configuring an EX3300 Virtual Chassis (CLI Procedure)*
- *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*
- *Configuring an EX8200 Virtual Chassis (CLI Procedure)*
- *Configuring an EX9200 Virtual Chassis*
- [Configuring a QFX Series Virtual Chassis \(CLI Procedure\) on page 6931](#)
- *Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)*
- *Adding a New EX4200 Switch to an Existing EX4200 Virtual Chassis (CLI Procedure)*
- [Replacing a Member Switch of a Virtual Chassis Configuration \(CLI Procedure\) on page 6938](#)

## serial-number


<b>Syntax</b>	<code>serial-number serial-number;</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis preprovisioned member member-id</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>In a preprovisioned Virtual Chassis or Virtual Chassis Fabric (VCF), specify the serial number of each member switch to be included in the configuration. If you do not include the serial number within the configuration, the switch cannot be recognized as a member of a preprovisioned configuration.</p> <p>In an EX8200 Virtual Chassis configuration, specify the serial number of each XRE200 External Routing Engine and each EX8200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the external Routing Engine or switch cannot be recognized as a member of the configuration.</p>
<b>Options</b>	<i>serial-number</i> —Permanent serial number for the external Routing Engine or for the member switch.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li> <li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li> <li>• <a href="#">Configuring an EX2200 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX3300 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX4300 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX8200 Virtual Chassis (CLI Procedure)</a></li> <li>• <a href="#">Configuring an EX9200 Virtual Chassis</a></li> <li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li> <li>• <a href="#">Configuring a Virtual Chassis on an EX Series Switch (J-Web Procedure)</a></li> </ul>

## serial-number (Virtual Chassis aliases)

---

<b>Syntax</b>	<code>serial-number <i>serial-number</i> {     <i>alias-name</i> <i>alias-name</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis aliases</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D10 for EX Series and QFX Series Virtual Chassis and Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Specify the serial number that will be labeled with an alias in a Virtual Chassis or Virtual Chassis Fabric (VCF).</p> <p>The remaining statements are explained separately.</p>
<b>Options</b>	<p><b><i>serial-number</i></b>—Permanent serial number for the member switch in the Virtual Chassis or VCF.</p> <p>You can retrieve the serial number for any device in your Virtual Chassis or VCF by entering the <b>show virtual-chassis</b> command and reviewing the output in the <b>Serial No</b> field.</p>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <a href="#">Understanding Virtual Chassis Fabric Components on page 7035</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li></ul>

## traceoptions (Virtual Chassis)

<b>Syntax</b>	<pre> traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;no-stamp&gt; &lt;replace&gt; &lt;size <i>size</i>&gt; &lt;world-readable       no-world-readable&gt;;     flag <i>flag</i> &lt;detail&gt; &lt;disable&gt; &lt;receive&gt; &lt;send&gt;; } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">virtual-chassis</a> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>detail</b> added in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Define tracing operations for the Virtual Chassis or VCF.
<b>Default</b>	Tracing operations are disabled.
<b>Options</b>	<p><b>detail</b>—(Optional) Generate detailed trace information for a flag.</p> <p><b>disable</b>—(Optional) Disable a flag.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 3 files</p> <p><b>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> <li><b>all</b>—All tracing operations.</li> </ul>
	<p> <b>TIP:</b> The <b>all</b> flag displays a subset of logs that are useful in debugging most issues. For more detailed information, use <b>all detail</b>.</p>
	<ul style="list-style-type: none"> <li><b>auto-configuration</b>—Trace Virtual Chassis ports (VCPs) that have been automatically configured.</li> <li><b>csn</b>—Trace Virtual Chassis complete sequence number (CSN) packets.</li> <li><b>error</b>—Trace Virtual Chassis errored packets.</li> </ul>

- **hello**—Trace Virtual Chassis hello packets.
- **krt**—Trace Virtual Chassis KRT events.
- **lsp**—Trace Virtual Chassis link-state packets.
- **lsp-generation**—Trace Virtual Chassis link-state packet generation.
- **me**—Trace Virtual Chassis ME events.
- **normal**—Trace normal events.
- **packets**—Trace Virtual Chassis packets.
- **parse**—Trace reading of the configuration.
- **psn**—Trace partial sequence number (PSN) packets.
- **route**—Trace Virtual Chassis routing information.
- **spf**—Trace Virtual Chassis SPF events.
- **state**—Trace Virtual Chassis state transitions.
- **task**—Trace Virtual Chassis task operations.

**no-stamp**—(Optional) Do not place a timestamp on any trace file.

**no-world-readable**—(Optional) Restrict file access to the user who created the file.

**receive**—(Optional) Trace received packets.

**replace**—(Optional) Replace a trace file rather than appending information to it.

**send**—(Optional) Trace transmitted packets.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through 1 GB

**Default:** 128 KB

**world-readable**—(Optional) Enable unrestricted file access.

<b>Required Privilege Level</b>	<b>system</b> —To view this statement in the configuration. <b>system-control</b> —To add this statement to the configuration.
---------------------------------	-----------------------------------------------------------------------------------------------------------------------------------

- Related Documentation**
- *Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis*
  - [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 6975](#)
  - *Verifying That Virtual Chassis Ports Are Operational*
  - *Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis*
  - *Troubleshooting an EX Series Virtual Chassis*

## virtual-chassis

```
Syntax  virtual-chassis {
        aliases {
            serial-number serial-number {
                alias-name alias-name;
            }
        }
        auto-provisioned
        auto-sw-update {
            (ex-4200 | ex-4300 | ex-4500 | ex-4600 | qfx-3 | qfx-5)
            package-name package-name;
        }
        fast-failover (ge | vcp disable | xe);
        graceful-restart {
            disable;
        }
        id id;
        mac-persistence-timer [minutes | disable];;
        member member-id {
            location location;
            mastership-priority number;
            no-management-vlan;
            serial-number;
            role;
        }
        no-split-detection;
        preprovisioned;
        traceoptions (Virtual Chassis) {
            file filename <files number> <size size> <world-readable | no-world-readable> <match
                regex>;
            flag flag ;
        }
        vc-port {
            lag-hash (packet-based | source-port-based);
        }
        vcp-no-hold-time;
    }
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 9.0 for EX Series switches.  
Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.  
Statement introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

**Description** Configure a Virtual Chassis or a Virtual Chassis Fabric (VCF).

The remaining statements are explained separately.

**Default** A standalone EX Series switch is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.



A QFX Series device configured in standalone mode is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

A standalone XRE200 External Routing Engine or EX8200 switch is not part of an EX8200 Virtual Chassis until a Virtual Chassis configuration is set up.

<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Autoprovisioning a Virtual Chassis Fabric on page 7053</a></li><li>• <a href="#">Preprovisioning a Virtual Chassis Fabric on page 7056</a></li><li>• <a href="#">Configuring a QFX Series Virtual Chassis (CLI Procedure) on page 6931</a></li><li>• <i>Example: Configuring an EX3300 Virtual Chassis with a Master and Backup</i></li><li>• <i>Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet</i></li><li>• <i>Example: Setting Up a Full Mesh EX8200 Virtual Chassis with Two EX8200 Switches and Redundant XRE200 External Routing Engines</i></li><li>• <i>Configuring an EX3300 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX8200 Virtual Chassis (CLI Procedure)</i></li><li>• <i>Configuring an EX9200 Virtual Chassis</i></li></ul>



# Administration

- [Routine Monitoring on page 7105](#)
- [Operational Commands on page 7107](#)

## Routine Monitoring

---

- [Verifying the Member ID, Role, Status, and Neighbor Member Connections of a Virtual Chassis Fabric Member Device on page 7105](#)
- [Verifying Virtual Chassis Port Connections in a Virtual Chassis Fabric on page 7106](#)
- [Verifying the Virtual Chassis Fabric Mode Settings on page 7107](#)

### Verifying the Member ID, Role, Status, and Neighbor Member Connections of a Virtual Chassis Fabric Member Device

**Purpose** Use this procedure to learn the current member ID, role, status, Virtual Chassis port (VCP) connections, and other information for the devices in your VCF.

Understanding the current member IDs, roles, device statuses, and VCP connections is required for routine monitoring of your VCF. You'll often need to identify this basic operational information to confirm a device or a VCP is working properly in the VCF, or how the VCF topology changed as a result of a configuration change or network error.

**Action** To display VCF status using the CLI:

#### `show virtual-chassis (Virtual Chassis Fabric)`

```
user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0282.5fa0.3f08
Fabric Mode: Enabled
```

				Mstr		Mixed Route Neighbor			
List				prio	Role	Mode	Mode	ID	
Member ID	Status	Serial No	Model						
Interface									
0 (FPC 0)	Prsnt	AB3112430001	qfx5100-48s	129	Master*	N	F	3	
vcp-255/1/0									2
vcp-255/1/1									4
vcp-255/1/2									4
vcp-255/1/3									4

1 (FPC 1)	Prsnt	AB3112230001	qfx5100-48s	129	Backup	N	F	3
vcp-255/1/0								2
vcp-255/1/1								4
vcp-255/1/2								4
vcp-255/1/3								1
2 (FPC 2)	Prsnt	AB3112460011	qfx5100-48s	0	Linecard	N	F	0
vcp-255/1/0								1
vcp-255/1/1								0
3 (FPC 3)	Prsnt	AB3112460011	qfx5100-48s	0	Linecard	N	F	1
vcp-255/1/0								0
vcp-255/1/1								1
4 (FPC 4)	Prsnt	AB3112430011	qfx5100-48s	0	Linecard	N	F	0
vcp-255/1/0								1
vcp-255/1/1								0

**Meaning** This output verifies that fabric mode is enabled and that all devices in the VCF are participating in the fabric, as shown by the **Prsnt** status output for each device.

The Neighbor ID and Interface outputs show that all VCPs are operating correctly.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 7056](#)

## Verifying Virtual Chassis Port Connections in a Virtual Chassis Fabric

**Purpose** Verify the Virtual Chassis Ports (VCPs) in your Virtual Chassis Fabric (VCF).

You should use this command if you suspect a VCP link in your VCF is broken.

**Action** To display the VCPs of a device:

```
user@switch> show virtual-chassis vc-port member 4 fpc4:
```

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
0/48	Auto-Configured	-1	Up	40000	0	vcp-255/0/2
0/49	Auto-Configured	-1	Up	40000	1	vcp-255/0/2
0/50	Auto-Configured	-1	Up	40000	2	vcp-255/0/2
0/51	Auto-Configured	-1	Up	40000	3	vcp-255/0/2

**Meaning** All of the VCPs on this device are up and active.

If the **Status** of an interface is **Absent** or the interface that you thought was a VCP does not appear in the command output, you likely have a problem with a link that has not been converted into a VCP. In this scenario, configure the interface on the link into a VCP using the **request virtual-chassis vc-port** command.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 7056](#)

## Verifying the Virtual Chassis Fabric Mode Settings

**Purpose** Verify the Virtual Chassis Fabric (VCF) mode settings on a device.

You must configure all devices in a VCF into fabric mode using the **request virtual-chassis mode fabric** command for the devices to operate in a VCF.

All VCFs use QFX5100 devices in the spine role. If a VCF uses a QFX3500, QFX3600, or EX4300 devices as a leaf node, you must also configure each device into mixed mode using the **request virtual-chassis mode mixed** command.

You must also configure a device out of mixed and fabric mode if it is removed from a VCF and placed into your network in a different role.

**Action** To display the current mode of a device:

```
user@switch> show virtual-chassis mode
fpc0:
```

```
-----
Current mode : Fabric with mixed devices
Future mode after reboot : Fabric with mixed devices
```

**Meaning** The output indicates that the switch is currently in mixed and fabric mode.

The output also indicates that the mode will not change when the device is rebooted without further configuration. You must reboot the device to change the fabric or mixed mode, so the **Future mode after reboot** output differs from the **Current mode** output when the mode has been changed but the device has not been rebooted.

- Related Documentation**
- [Autoprovisioning a Virtual Chassis Fabric on page 7053](#)
  - [Preprovisioning a Virtual Chassis Fabric on page 7056](#)

## Operational Commands

- [clear virtual-chassis vc-port statistics](#)
- [request session member](#)
- [request virtual-chassis mode](#)
- [request virtual-chassis reactivate](#)
- [request virtual-chassis vc-port](#)
- [request virtual-chassis vc-port diagnostics optics](#)
- [show forwarding-options enhanced-hash-key](#)
- [show virtual-chassis active-topology](#)
- [show virtual-chassis device-topology](#)

- [show virtual-chassis login](#)
- [show virtual-chassis mode](#)
- [show virtual-chassis protocol adjacency](#)
- [show virtual-chassis protocol database](#)
- [show virtual-chassis protocol interface](#)
- [show virtual-chassis protocol route](#)
- [show virtual-chassis protocol statistics](#)
- [show virtual-chassis](#)
- [show virtual-chassis vc-port](#)
- [show virtual-chassis vc-port diagnostics optics](#)
- [show virtual-chassis vc-port statistics](#)

## clear virtual-chassis vc-port statistics

<b>Syntax</b>	clear virtual-chassis vc-port statistics <all-members> <interface-name> <local> <member member-id>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. The options <b>all-members</b> and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric.
<b>Description</b>	Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).
<b>Options</b>	<p><b>none</b>—Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Clear traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Clear traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Clear traffic statistics for VCPs from the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Clear traffic statistics for VCPs from the specified member of a Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 7024</a></li> <li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis) on page 7109</a> <a href="#">clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 7110</a> <a href="#">clear virtual-chassis vc-port statistics member 3 on page 7110</a>

### Sample Output

#### clear virtual-chassis vc-port statistics (EX4200 Virtual Chassis)

```

user@switch> clear virtual-chassis vc-port statistics
fpc0:
-----
Statistics cleared

```

### clear virtual-chassis vc-port statistics (EX8200 Virtual Chassis)

```
user@external-routing-engine> clear virtual-chassis vc-port statistics
```

```
member0:
```

```
-----  
Statistics cleared
```

```
member1:
```

```
-----  
Statistics cleared
```

```
member8:
```

```
-----  
Statistics cleared
```

```
member9:
```

```
-----  
Statistics cleared
```

### clear virtual-chassis vc-port statistics member 3

```
user@switch> clear virtual-chassis vc-port statistics member 3
```

```
Cleared statistics on member 3
```



---

## request session member

---

<b>Syntax</b>	<code>request session member <i>member-id</i></code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Start a session with the specified member of a Virtual Chassis or a VCF.
<b>Options</b>	<i>member-id</i> —Member ID for the specific member of the Virtual Chassis or VCF.
<b>Required Privilege Level</b>	maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">member on page 6958</a></li><li>• <i>Understanding EX Series Virtual Chassis Components</i></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li></ul>

## request virtual-chassis mode

---

**Syntax**    request virtual-chassis mode  
              fabric  
              mixed  
              <disable>  
              <reboot>  
              <all-members>  
              <local>  
              <member *member-id*>

**Release Information**    Command introduced in Junos OS Release 11.1 for EX Series switches.  
                             Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.  
                             The **fabric** keyword introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices in a Virtual Chassis Fabric (VCF).  
                             Command introduced in Junos OS Release 13.2X51-D20 for VCF.

**Description**    Configure the mode for a device or multiple devices in a Virtual Chassis or a VCF.

A device must be configured in fabric mode to participate as a member device in a VCF.

A device must be configured in mixed mode when it is participating in a Virtual Chassis or a VCF with different types of devices.

Do not enable the **request virtual-chassis mode mixed** command for a standalone device or for a member switch that is intended to remain in a non-mixed Virtual Chassis or VCF. Enabling this command reduces the maximum scaling numbers for some features on the switch, Virtual Chassis, or VCF.

You do not need to configure mixed mode if the only devices in your Virtual Chassis are EX4500 and EX4550 switches.

To avoid potential traffic disruptions and configuration issues for a mixed Virtual Chassis, we recommend configuring mixed mode on your device before cabling it into your Virtual Chassis. We recommend rebooting your device to complete this configuration procedure before interconnecting your device into the Virtual Chassis.

To avoid potential traffic disruptions and configuration issues, we recommend configuring the fabric and, if applicable, the mixed mode settings on your device before cabling it into a VCF. We recommend rebooting your device to complete this configuration procedure before interconnecting your device into the VCF. You can change the fabric and mixed mode settings after the device has been added to a Virtual Chassis or VCF, however.

If you set some of the devices in a mixed Virtual Chassis or VCF to mixed mode using this command but not others, the mixed Virtual Chassis or VCF might not form. If you experience this issue, enter the **request virtual-chassis mode mixed all-members** command to set the Virtual Chassis mode to mixed for all devices in the Virtual Chassis or VCF. You then need to reboot the devices that have been set into mixed mode to complete the procedure. The Virtual Chassis or VCF forms after the devices have rebooted.

When you do not use this command to set any of the switches in a mixed EX4200 and EX4500 Virtual Chassis to mixed mode, a mixed EX4200 and EX4500 Virtual Chassis

forms with one of the switches assuming the master role if the switches are running Junos OS Release 11.4 or later. All other switches in the mixed EX4200 and EX4500 Virtual Chassis are placed into the linecard role. If you experience this behavior, enter the **request virtual-chassis mode mixed all-members** command to set the Virtual Chassis mode to mixed for all switches in the Virtual Chassis. You will then need to reboot the switches to complete the procedure. The Virtual Chassis will form after all of the switches have rebooted.

The Virtual Chassis mode setting is maintained through reboots even though it is set in operational mode.

- Options**
- none**—Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.
  - all-members**—(Optional) Set the Virtual Chassis mode for all members of the Virtual Chassis or VCF.
  - disable**—Disable the Virtual Chassis fabric or mixed mode setting if it was previously enabled.
  - fabric**—Set the device into fabric mode so that the device can participate in a VCF.
  - local**—(Optional) Set the Virtual Chassis mode on the member device where the command is issued.
  - member *member-id***—(Optional) Set the Virtual Chassis mode to mixed on the specified member of the Virtual Chassis or VCF.
  - mixed**—Set the device into mixed mode so that the device can participate in a mixed Virtual Chassis or mixed VCF.



**NOTE:** You do not need to set mixed mode if the only devices in your Virtual Chassis are QFX3500 and QFX3600 devices.

You do not need to configure mixed mode if the only devices in your Virtual Chassis are EX4500 and EX4550 switches.

**Required Privilege Level** system-control

- Related Documentation**
- [Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches \(CLI Procedure\)](#)
  - [Verifying the Virtual Chassis Fabric Mode Settings on page 7107](#)
  - [Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 6975](#)

**List of Sample Output** [request virtual-chassis mode mixed on page 7114](#)  
[request virtual-chassis mode fabric mixed reboot on page 7114](#)

## Sample Output

request virtual-chassis mode mixed

```
user@switch> request virtual-chassis mode mixed
```

## Sample Output

request virtual-chassis mode fabric mixed reboot

```
user@switch> request virtual-chassis mode fabric mixed reboot
```

## request virtual-chassis reactivate

---

<b>Syntax</b>	<code>request virtual-chassis reactivate</code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Reactivate a device that has been assigned a member ID but is not currently connected to the Virtual Chassis or VCF.</p> <p>You can use this command to reactivate a device that was previously part of the Virtual Chassis or VCF but whose status is no longer <b>Prsnt</b>.</p>
<b>Required Privilege Level</b>	system-control
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 6975</a></li> <li>• <i>Verifying the Member ID, Role, and Neighbor Member Connections of an EX8200 Virtual Chassis Member</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request virtual-chassis reactivate on page 7115</a>

### Sample Output

#### request virtual-chassis reactivate

```
user@switch> request virtual-chassis reactivate
```

## request virtual-chassis vc-port

---

<b>Syntax</b>	<b>request virtual-chassis vc-port set   delete</b> <fpc-slot <i>fpc-slot</i> > pic-slot <i>pic-slot</i> port <i>port-number</i> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Option <b>fpc-slot</b> introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Enable or disable an optical port as a Virtual Chassis port (VCP).</p> <p>If you omit <b>member <i>member-id</i></b>, this command defaults to enabling or disabling the uplink VCP or SFP network port configured as a VCP on the switch where the command is issued.</p> <p>On an EX3300 switch, uplink ports 2 and 3 are configured as VCPs by default. No other uplink ports on any other EX Series switches are configured as VCPs by default.</p> <p>You might experience a temporary traffic disruption immediately after creating or deleting a user-configured VCP in an EX8200 Virtual Chassis.</p>
<b>Options</b>	<p><b>pic-slot <i>pic-slot</i></b>—Number of the PIC slot for the port on the switch.</p> <p><b>port <i>port-number</i></b>—Number of the port that is to be enabled or disabled as a VCP.</p> <p><b>member <i>member-id</i></b>—(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	system-control
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">request virtual-chassis vc-port</a> (dedicated port)</li><li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li><li>• <a href="#">show virtual-chassis vc-port statistics on page 7024</a></li><li>• <a href="#">clear virtual-chassis vc-port statistics on page 6978</a></li><li>• <a href="#">Virtual Chassis Port (VCP) Interface Names in an EX8200 Virtual Chassis</a></li><li>• <a href="#">Understanding EX Series Virtual Chassis Components</a></li><li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request virtual-chassis vc-port set pic-slot 1 port 0 on page 7117</a> <a href="#">request virtual-chassis vc-port set pic-slot 1 port 1 member 3 on page 7117</a> <a href="#">request virtual-chassis vc-port delete pic-slot 1 port 1 member 3 on page 7117</a>

## Sample Output

**request virtual-chassis vc-port set pic-slot 1 port 0**

user@switch> request virtual-chassis vc-port set pic-slot 1 port 0

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

**request virtual-chassis vc-port set pic-slot 1 port 1 member 3**

user@switch> request virtual-chassis vc-port set pic-slot 1 port 1 member 3

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

**request virtual-chassis vc-port delete pic-slot 1 port 1 member 3**

user@switch> request virtual-chassis vc-port delete pic-slot 1 port 1 member 3

To check the results of this command, use the [show virtual-chassis vc-port](#) command.

## request virtual-chassis vc-port diagnostics optics

---

<b>Syntax</b>	<b>request virtual-chassis vc-port diagnostics optics</b>
<b>Release Information</b>	Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Run a digital optical monitoring (DOM) scan on the optical ports configured as Virtual Chassis ports (VCPs).</p> <p>Enter the <b>show virtual-chassis vc-port diagnostics optics</b> command to view the results of the diagnostic scan.</p> <p>On certain EX Series switches, the <b>request virtual-chassis vc-port diagnostics optics</b> command must be entered to run a diagnostic scan before you can gather the <b>show virtual-chassis vc-port diagnostics optics</b> output.</p>
<b>Required Privilege Level</b>	system-control
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show virtual-chassis vc-port diagnostics optics on page 7162</a></li></ul>

## Sample Output

### request virtual-chassis vc-port diagnostics optics

```
user@switch> request virtual-chassis vc-port diagnostics optics
fpc0:
-----
vc-port Diagnostics Optics Done
```



## show forwarding-options enhanced-hash-key

<b>Syntax</b>	<b>show forwarding-options enhanced-hash-key</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p><b>Fabric Load Balancing Options</b> output fields introduced in Junos OS Release 14.1X53-D10.</p>
<b>Description</b>	<p>Display information about which packet fields are used by the hashing algorithm to make hashing decisions.</p> <p>You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the <b>forwarding-options enhanced-hash-key</b> statement.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 2590</a></li> <li>• <a href="#">Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 2396</a></li> <li>• <a href="#">enhanced-hash-key on page 2644</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 7120</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 7121</a></p> <p><a href="#">show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 7121</a></p>
<b>Output Fields</b>	<p><a href="#">Table 250 on page 2763</a> lists the output fields for the <b>show forwarding-options enhanced-hash-key</b> command. Output fields are listed in the approximate order in which they first appear.</p>

**Table 681: show forwarding-options enhanced-hash-key Output Fields**

Field Name	Field Description
<b>Hash-Mode</b>	Current hash mode: Layer 2 header or Layer 2 payload.
<b>Protocol</b>	Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.
<b>Destination L4 Port</b>	Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.
<b>Source L4 Port</b>	Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.
<b>Destination IPv4 Addr</b>	Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.

Table 681: show forwarding-options enhanced-hash-key Output Fields (*continued*)

Field Name	Field Description
<b>Source IPv4 Addr</b>	Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
<b>Vlan id</b>	Indicates whether the Vlan id field is or is not used by the hashing algorithm: Yes or No.
<b>Next Hdr</b>	Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.
<b>Destination IPv6 Addr</b>	Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
<b>Source IPv6 Addr</b>	Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
<b>Ether Type</b>	Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.
<b>Destination MAC Address</b>	Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.
<b>Source MAC Address</b>	Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.
<b>Load Balancing Method</b>	Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet.  The load balancing method is flowlet by default, and can be configured using the <a href="#">fabric-load-balance</a> statement.
<b>Fabric Link Scale</b>	Indicates the fabric link scale, in mbps.
<b>Inactivity Interval</b>	Indicates the fabric load balance inactivity interval, in microseconds (us).  The inactivity interval is 16 microseconds by default, and can be configured using the <a href="#">inactivity-interval</a> statement.
<b>Hash Region Size/Trunk</b>	Indicates the hash region size, in buckets per fabric trunk.

## Sample Output

### show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
-----
```

```
Hash-Mode                               :layer2-payload
```

```
inet Hash settings-
```

```
-----
```

```
inet packet fields
```

```
Protocol                               : Yes
Destination L4 Port                    : Yes
Source L4 Port                         : Yes
Destination IPv4 Addr                  : Yes
Source IPv4 Addr                       : Yes
Vlan id                               : No
```

```
inet6 Hash settings-
```

```
-----
```

```
inet6 packet fields
```

```
Next Hdr                             : Yes
Destination L4 Port                    : Yes
Source L4 Port                         : Yes
Destination IPv6 Addr                  : Yes
Source IPv6 Addr                       : Yes
Vlan id                               : No
```

#### show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0
```

```
Current Hash Settings
```

```
-----
```

```
Hash-Mode                               : layer2-header
```

```
layer2 Hash settings-
```

```
-----
```

```
layer2 packet fields
```

```
Ether Type                           : Yes
Destination MAC Address                : Yes
Source MAC Address                     : Yes
VLAN ID                               : No
```

#### show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>
```

```
Fabric Load Balancing Options
```

```
-----
```

```
Load Balancing Method : Flowlet
Fabric Link Scale      : 40960 (mbps)
Inactivity Interval   : 16 (us)
Hash Region Size/Trunk : 1024 (buckets)
```

## show virtual-chassis active-topology

<b>Syntax</b>	show virtual-chassis active-topology <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the active topology of the Virtual Chassis or VCF with next-hop reachability information.
<b>Options</b>	<p><b>none</b>—Display the active topology of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the active topology of all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the active topology of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the active topology of the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> <li><i>Understanding EX Series Virtual Chassis Configuration</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis active-topology (EX4200 Virtual Chassis) on page 7123</a> <a href="#">show virtual-chassis active-topology (EX8200 Virtual Chassis) on page 7123</a> <a href="#">show virtual-chassis active-topology (Virtual Chassis Fabric) on page 7124</a>
<b>Output Fields</b>	<a href="#">Table 668 on page 6985</a> lists the output fields for the <b>show virtual-chassis active-topology</b> command. Output fields are listed in the approximate order in which they appear.

**Table 682: show virtual-chassis active-topology Output Fields**

Field Name	Field Description
<b>Destination ID</b>	Specifies the member ID of the destination.
<b>Next-hop</b>	<p>Specifies the member ID and Virtual Chassis port (VCP) of the next hop to which packets for the destination ID are forwarded.</p> <p>The next hop can be more than one device in a VCF.</p>

## Sample Output

### show virtual-chassis active-topology (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis active-topology
 1                      1(vcp-1)

 2                      1(vcp-1)

 3                      1(vcp-1)

 4                      1(vcp-1)

 5                      8(vcp-0) 1(vcp-1)

 6                      8(vcp-0)

 7                      8(vcp-0)

 8                      8(vcp-0)

```

### show virtual-chassis active-topology (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis active-topology
member0:

```

Destination ID	Next-hop
1	1(vcp-4/0/4.32768)
8	8(vcp-0/0.32768)
9	8(vcp-0/0.32768)

```

member1:

```

Destination ID	Next-hop
0	0(vcp-3/0/4.32768)
8	8(vcp-0/0.32768)
9	8(vcp-0/0.32768)

```

member8:

```

Destination ID	Next-hop
0	0(vcp-1/1.32768)
1	1(vcp-1/2.32768)
9	9(vcp-2/1.32768)

member9:

Destination ID	Next-hop
0	8(vcp-1/2.32768)
1	8(vcp-1/2.32768)
8	8(vcp-1/2.32768)

### show virtual-chassis active-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis active-topology  
fpc0:

Destination ID	Next-hop
1 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
2 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
3 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

fpc1:

Destination ID	Next-hop
0 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
2 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
3 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

fpc2:

Destination ID	Next-hop
0 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
1 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
3 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

fpc3:

Destination ID	Next-hop
0 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
1 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
2 6(vcp-255/0/1.32768)	4(vcp-255/0/2.32768) 5(vcp-255/0/3.32768)
4	4(vcp-255/0/2.32768)
5	5(vcp-255/0/3.32768)
6	6(vcp-255/0/1.32768)

fpc4:

Destination ID	Next-hop
0	0(vcp-255/0/48.32768)
1	1(vcp-255/0/49.32768)
2	2(vcp-255/0/50.32768)
3	3(vcp-255/0/51.32768)
5 0(vcp-255/0/48.32768)	3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768) 1(vcp-255/0/49.32768)
6 0(vcp-255/0/48.32768)	3(vcp-255/0/51.32768) 2(vcp-255/0/50.32768) 1(vcp-255/0/49.32768)

fpc5:

Destination ID	Next-hop
0	0(vcp-255/0/48.32768)

1	1(vcp-255/0/49.32768)	
2	2(vcp-255/0/50.32768)	
3	3(vcp-255/0/51.32768)	
4	3(vcp-255/0/51.32768)	2(vcp-255/0/50.32768)
0(vcp-255/0/48.32768)	1(vcp-255/0/49.32768)	
6	3(vcp-255/0/51.32768)	2(vcp-255/0/50.32768)
0(vcp-255/0/48.32768)	1(vcp-255/0/49.32768)	

fpc6:

Destination ID	Next-hop
0	0(vcp-255/0/0.32768)
1	1(vcp-255/0/1.32768)
2	2(vcp-255/0/2.32768)
3	3(vcp-255/0/3.32768)
4	3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768)
0(vcp-255/0/0.32768)	1(vcp-255/0/1.32768)
5	3(vcp-255/0/3.32768) 2(vcp-255/0/2.32768)
0(vcp-255/0/0.32768)	1(vcp-255/0/1.32768)



## show virtual-chassis device-topology

<b>Syntax</b>	show virtual-chassis device-topology <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the device topology—the member and system IDs, the VCP numbers, and device status—for all hardware devices in the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display the device topology for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the device topology for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the device topology for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding EX Series Virtual Chassis Port Link Aggregation</i></li> <li>• <i>Understanding EX8200 Virtual Chassis Topologies</i></li> </ul>
<b>Output Fields</b>	<a href="#">Table 669 on page 6990</a> lists the output fields for the <b>show virtual-chassis device-topology</b> command. Output fields are listed in the approximate order in which they appear.

**Table 683: show virtual-chassis device-topology Output Fields**

Field Name	Field Description
<b>Member</b>	Assigned member ID.
<b>Device</b>	Assigned device ID.  For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
<b>Status</b>	The status of the device within the Virtual Chassis or VCF. Outputs include: <ul style="list-style-type: none"> <li>• <b>Prsnt</b>—Device is currently connected to and participating in the Virtual Chassis or VCF.</li> <li>• <b>NotPrsnt</b>—Device is assigned but is not currently connected.</li> </ul>

Table 683: show virtual-chassis device-topology Output Fields (*continued*)

Field Name	Field Description
<b>System ID</b>	System ID of the device.  The system ID of the device is the device's MAC address.
<b>Member (Neighbor List)</b>	Assigned member ID of the neighbor device.
<b>Device (Neighbor List)</b>	Assigned device ID of the neighbor device.  For an EX8200 Virtual Chassis, the member ID and the device ID are always identical.
<b>Interface (Neighbor List)</b>	The interface connecting the device to the neighbor.

## Sample Output

### show virtual-chassis device-topology

```
user@switch> show virtual-chassis device-topology
```

```
member0:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0         0      Prsnt   0021.59f7.d000  8         8      vcp-0/0
                                1         1      vcp-4/0/1
1         1      Prsnt   0026.888d.6800  8         8      vcp-0/0
                                9         9      vcp-0/1
                                0         0      vcp-3/0/4
8         8      Prsnt   0000.4a75.9b7c  9         9      vcp-1/0
                                0         0      vcp-1/1
                                1         1      vcp-1/2
9         9      Prsnt   0000.73e9.9a57  8         8      vcp-1/0
                                1         1      vcp-1/1
```

```
member1:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
0         0      Prsnt   0021.59f7.d000  8         8      vcp-0/0
                                1         1      vcp-4/0/1
1         1      Prsnt   0026.888d.6800  8         8      vcp-0/0
                                9         9      vcp-0/1
                                0         0      vcp-3/0/4
8         8      Prsnt   0000.4a75.9b7c  9         9      vcp-1/0
                                0         0      vcp-1/1
                                1         1      vcp-1/2
9         9      Prsnt   0000.73e9.9a57  8         8      vcp-1/0
                                1         1      vcp-1/1
```

```
member8:
```

```
-----
Member  Device  Status  System ID      Neighbor List
                                Member  Device  Interface
```

0	0	Prsnt	0021.59f7.d000	8	8	vcp-0/0
				1	1	vcp-4/0/1
1	1	Prsnt	0026.888d.6800	8	8	vcp-0/0
				9	9	vcp-0/1
				0	0	vcp-3/0/4
8	8	Prsnt	0000.4a75.9b7c	9	9	vcp-1/0
				0	0	vcp-1/1
				1	1	vcp-1/2
9	9	Prsnt	0000.73e9.9a57	8	8	vcp-1/0
				1	1	vcp-1/1

member9:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	0021.59f7.d000	8	8	vcp-0/0
				1	1	vcp-4/0/1
1	1	Prsnt	0026.888d.6800	8	8	vcp-0/0
				9	9	vcp-0/1
				0	0	vcp-3/0/4
8	8	Prsnt	0000.4a75.9b7c	9	9	vcp-1/0
				0	0	vcp-1/1
				1	1	vcp-1/2
9	9	Prsnt	0000.73e9.9a57	8	8	vcp-1/0
				1	1	vcp-1/1

#### show virtual-chassis device-topology (Virtual Chassis Fabric)

user@device> show virtual-chassis device-topology  
fpc0:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
5	5	Prsnt	100e.7eb5.80c0	1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc1:

Neighbor List

Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
				1	1	vcp-255/0/49
5	5	Prsnt	100e.7eb5.80c0	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc2:

Neighbor List						
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
				1	1	vcp-255/0/49
5	5	Prsnt	100e.7eb5.80c0	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc3:

Neighbor List						
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3

1	1	Prsnt	100e.7eb8.3a40	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
2	2	Prsnt	100e.7eb5.d700	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
3	3	Prsnt	100e.7eb5.c440	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
4	4	Prsnt	100e.7eb5.7e40	6	6	vcp-255/0/1
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
5	5	Prsnt	100e.7eb5.80c0	0	0	vcp-255/0/48
				1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
6	6	Prsnt	100e.7eb6.3b00	0	0	vcp-255/0/48
				3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc4:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
5	5	Prsnt	100e.7eb5.80c0	1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
6	6	Prsnt	100e.7eb6.3b00	3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc5:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3

2	2	Prsnt	100e.7eb5.d700	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
3	3	Prsnt	100e.7eb5.c440	6	6	vcp-255/0/1
				4	4	vcp-255/0/2
				5	5	vcp-255/0/3
4	4	Prsnt	100e.7eb5.7e40	6	6	vcp-255/0/1
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
5	5	Prsnt	100e.7eb5.80c0	0	0	vcp-255/0/48
				1	1	vcp-255/0/49
				3	3	vcp-255/0/51
6	6	Prsnt	100e.7eb6.3b00	2	2	vcp-255/0/50
				1	1	vcp-255/0/49
				0	0	vcp-255/0/48
				3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

fpc6:

				Neighbor List		
Member	Device	Status	System ID	Member	Device	Interface
0	0	Prsnt	100e.7eb6.a900	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
1	1	Prsnt	100e.7eb8.3a40	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
2	2	Prsnt	100e.7eb5.d700	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
3	3	Prsnt	100e.7eb5.c440	4	4	vcp-255/0/2
				5	5	vcp-255/0/3
				6	6	vcp-255/0/1
4	4	Prsnt	100e.7eb5.7e40	3	3	vcp-255/0/51
				2	2	vcp-255/0/50
				0	0	vcp-255/0/48
5	5	Prsnt	100e.7eb5.80c0	1	1	vcp-255/0/49
				3	3	vcp-255/0/51
				2	2	vcp-255/0/50
6	6	Prsnt	100e.7eb6.3b00	1	1	vcp-255/0/49
				0	0	vcp-255/0/48
				3	3	vcp-255/0/3
				2	2	vcp-255/0/2
				0	0	vcp-255/0/0
				1	1	vcp-255/0/1

## show virtual-chassis login

<b>Syntax</b>	<b>show virtual-chassis login</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	<p>Supply the address of the host that logged into the Virtual Chassis or VCF, or identify the location of the member switch that redirected the current session to a different member switch.</p> <p>You might need this information for tracing or troubleshooting purposes.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request session member on page 6980</a></li> <li>• <a href="#">Understanding Global Management of a Virtual Chassis on page 6919</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis login (Direct Login to the Master Console Port) on page 7133</a></p> <p><a href="#">show virtual-chassis login (Backup Console Session Redirected to the Master Console Port) on page 7133</a></p>

### Sample Output

#### show virtual-chassis login (Direct Login to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host 248.1.2.3
```

#### show virtual-chassis login (Backup Console Session Redirected to the Master Console Port)

```
user@switch> show virtual-chassis login
Current login session initiated from host backup
```

## show virtual-chassis mode

<b>Syntax</b>	<b>show virtual-chassis mode</b> <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF). <b>Current mode</b> and <b>Future mode after reboot</b> fields introduced in Junos OS Release 13.2X51-D20.
<b>Description</b>	Display the Virtual Chassis or Virtual Chassis Fabric (VCF) mixed mode status.
<b>Options</b>	<p><b>none</b>—Display the Virtual Chassis or VCF mixed mode status for the device on which the command is entered.</p> <p><b>all-members</b>—(Optional) Display the Virtual Chassis or VCF mixed mode status for all member devices in the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the Virtual Chassis or VCF mixed mode status for the device on which the command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the Virtual Chassis or VCF mixed mode status for the specified member device..</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">request virtual-chassis mode on page 7112</a></li> <li>• <a href="#">Verifying the Virtual Chassis Fabric Mode Settings on page 7107</a></li> <li>• <a href="#">Configuring a Mixed Virtual Chassis with EX4200, EX4500, and EX4550 Member Switches (CLI Procedure)</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis mode (EX4200) on page 7135</a> <a href="#">show virtual-chassis mode (QFX5100) on page 7135</a>
<b>Output Fields</b>	<a href="#">Table 684 on page 7134</a> lists the output fields for the <b>show virtual-chassis mode</b> command.

Table 684: show virtual-chassis mode Output Fields

Field Name	Field Description
<b>Mixed Mode</b>	Specifies the mixed mode status of the member switch. Mixed mode is either <b>Enabled</b> or <b>Disabled</b> .



Table 684: show virtual-chassis mode Output Fields (*continued*)

Field Name	Field Description
<b>Current mode</b>	<p>Specifies the current mixed and fabric mode settings running on the member device or devices.</p> <p>A device reboot is required to change the fabric or mixed mode. The <b>Current mode</b> and <b>Future mode after reboot</b> are different when the mode has been changed but the device has not been rebooted.</p> <p>Outputs include:</p> <ul style="list-style-type: none"> <li>• <b>Fabric with mixed devices</b>—Fabric mode and mixed mode are enabled.</li> <li>• <b>Fabric with similar devices</b>—Fabric mode is enabled and mixed mode is disabled.</li> <li>• <b>Virtual Chassis with mixed devices</b>—Fabric mode is disabled and mixed mode is enabled.</li> <li>• <b>Virtual Chassis with similar devices</b>—Fabric mode is disabled and mixed mode is disabled.</li> </ul>
<b>Future mode after reboot</b>	<p>Specifies the mixed and fabric mode settings running on the member device or devices.</p> <p>A device reboot is required to change the fabric or mixed mode. The <b>Current mode</b> and <b>Future mode after reboot</b> are different when the mode has been changed but the device has not been rebooted.</p> <p>Outputs include:</p> <ul style="list-style-type: none"> <li>• <b>Fabric with mixed devices</b>—Fabric mode and mixed mode are enabled.</li> <li>• <b>Fabric with similar devices</b>—Fabric mode is enabled and mixed mode is disabled.</li> <li>• <b>Virtual Chassis with mixed devices</b>—Fabric mode is disabled and mixed mode is enabled.</li> <li>• <b>Virtual Chassis with similar devices</b>—Fabric mode is disabled and mixed mode is disabled.</li> </ul>

## Sample Output

### show virtual-chassis mode (EX4200)

```
user@switch>show virtual-chassis mode
fpc0:
-----
Mixed Mode: Disabled
```

## Sample Output

### show virtual-chassis mode (QFX5100)

```
user@switch>show virtual-chassis mode
fpc0:
-----
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc1:
-----
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc2:
-----
Current mode : Fabric with similar devices
Future mode after reboot : Fabric with similar devices

fpc3:
```

```
-----  
Current mode : Fabric with similar devices  
Future mode after reboot : Fabric with similar devices
```

fpc4:

```
-----  
Current mode : Fabric with similar devices  
Future mode after reboot : Fabric with similar devices
```

## show virtual-chassis protocol adjacency

<b>Syntax</b>	<pre>show virtual-chassis protocol adjacency &lt;brief   detail   extensive&gt; &lt;all-members&gt; &lt;local&gt; &lt;member member-id&gt; &lt;system-id&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Display the Virtual Chassis Control Protocol (VCCP) adjacency statistics in the Virtual Chassis or VCF for all hardware devices.
<b>Options</b>	<p><b>none</b>—Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display VCCP adjacency statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP adjacency statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p> <p><b>system-id</b>—(Optional) Display VCCP adjacency statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding EX Series Virtual Chassis Port Link Aggregation</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis protocol adjacency on page 7138</a></p> <p><a href="#">show virtual-chassis protocol adjacency detail on page 7139</a></p>
<b>Output Fields</b>	Table 670 on page 6997 lists the output fields for the <b>show virtual-chassis protocol adjacency</b> command. Output fields are listed in the approximate order in which they appear.

Table 685: show virtual-chassis protocol adjacency Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the Virtual Chassis port (VCP) interface.	All levels
<b>System</b>	The MAC address of the device on the receiving side of the VCP link.	All levels
<b>State</b>	State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> <li>• <b>New</b>—The link is new.</li> <li>• <b>One-way</b>—The link is transmitting traffic in one direction.</li> <li>• <b>Initializing</b>—The link is initializing.</li> <li>• <b>Rejected</b>—The link is rejected.</li> </ul>	All levels
<b>Hold, Expires in</b>	Remaining holdtime of the adjacency.	All levels
<b>Priority</b>	Priority to become the designated intermediary system.	detail
<b>Up/Down Transitions</b>	Count of adjacency status transition changes from up to down or down to up.	detail
<b>Last transition</b>	Time of the last up/down transition.	detail

## Sample Output

### show virtual-chassis protocol adjacency

```
user@switch> show virtual-chassis protocol adjacency
```

```
member0:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          57
vcp-0/1.32768  0000.4a75.9b7c Up          59
vcp-4/0/1.32768 0026.888d.6800 Up          57
```

```
member1:
```

```
-----
Interface      System      State      Hold (secs)
vcp-0/0.32768  0000.4a75.9b7c Up          58
vcp-0/1.32768  0000.73e9.9a57 Up          59
vcp-3/0/4.32768 0021.59f7.d000 Up          58
```

```
member8:
```

```
-----
Interface      System      State      Hold (secs)
vcp-1/0.32768  0000.73e9.9a57 Up          58
vcp-1/1.32768  0021.59f7.d000 Up          58
vcp-1/2.32768  0026.888d.6800 Up          59
vcp-2/0.32768  0021.59f7.d000 Up          59
```

```
member9:
```

```
-----
Interface      System      State      Hold (secs)
```

vcp-1/0.32768	0000.4a75.9b7c Up	58
vcp-1/1.32768	0026.888d.6800 Up	59

### show virtual-chassis protocol adjacency detail

```
user@switch> show virtual-chassis protocol adjacency detail
```

```
member0:
```

```
-----
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0000.4a75.9b7c
  interface-name: vcp-0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:37 ago
```

```
0026.888d.6800
  interface-name: vcp-4/0/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member1:
```

```
-----
0000.4a75.9b7c
  interface-name: vcp-0/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0000.73e9.9a57
  interface-name: vcp-0/1.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago
```

```
0021.59f7.d000
  interface-name: vcp-3/0/4.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 22:06:39 ago
```

```
member8:
```

```
-----
0000.73e9.9a57
  interface-name: vcp-1/0.32768, State: Up, Expires in 58 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-1/1.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0026.888d.6800
  interface-name: vcp-1/2.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
0021.59f7.d000
  interface-name: vcp-2/0.32768, State: Up, Expires in 57 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

```
member9:
```

```
-----
0000.4a75.9b7c
  interface-name: vcp-1/0.32768, State: Up, Expires in 59 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 19:26:38 ago
```

0026.888d.6800  
interface-name: vcp-1/1.32768, State: Up, Expires in 58 secs  
Priority: 0, Up/Down transitions: 1, Last transition: 22:17:36 ago

## show virtual-chassis protocol database

<b>Syntax</b>	show virtual-chassis protocol database <brief   detail   extensive> <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the Virtual Chassis Control Protocol (VCCP) database statistics for all hardware devices within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option. The <b>extensive</b> option provides all output and is most useful for customer support personnel.</p> <p><b>all-members</b>—(Optional) Display VCCP database statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display VCCP database statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display VCCP database statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</a></li> <li>• <a href="#">Understanding EX Series Virtual Chassis Components</a></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis Components on page 6909</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol database on page 7142</a> <a href="#">show virtual-chassis protocol database detail on page 7143</a>
<b>Output Fields</b>	<a href="#">Table 671 on page 7000</a> lists the output fields for the <b>show virtual-chassis protocol database</b> command. Output fields are listed in the approximate order in which they appear.

**Table 686: show virtual-chassis protocol database Output Fields**

Field Name	Field Description	Level of Output
LSP ID	Link-state protocol (LSP) data unit identifier.	All levels

Table 686: show virtual-chassis protocol database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Sequence</b>	Sequence number of the LSP.	All levels
<b>Checksum</b>	Checksum value of the LSP.	All levels
<b>Lifetime</b>	Remaining lifetime of the LSP, in seconds.	All levels
<b>Neighbor</b>	MAC address of the neighbor on the advertising system.	detail
<b>Interface</b>	Virtual Chassis port (VCP) interface name.	detail
<b>Metric</b>	Metric of the prefix or neighbor.	detail

The **extensive** output was omitted from this list. The **extensive** output is useful for customer support personnel only.

## Sample Output

### show virtual-chassis protocol database

```
user@switch> show virtual-chassis protocol database
```

```
member0:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   113
0021.59f7.d000.00-00  0x16882  0x3993   118
0026.888d.6800.00-00  0x1691f  0x82b7   116
  4 LSPs
```

```
member1:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x1691f  0x82b7   118
  4 LSPs
```

```
member8:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   118
0000.73e9.9a57.00-00  0xf361   0x27e8   114
0021.59f7.d000.00-00  0x16883  0x289    116
0026.888d.6800.00-00  0x16920  0xa335   116
  4 LSPs
```

```
member9:
```

```
-----
LSP ID          Sequence Checksum Lifetime
0000.4a75.9b7c.00-00  0x1dd80  0xc2e3   116
0000.73e9.9a57.00-00  0xf361   0x27e8   116
0021.59f7.d000.00-00  0x16883  0x289    114
```



```
0026.888d.6800.00-00      0x16920   0xa335      116
4 LSPs
```

### show virtual-chassis protocol database detail

```
user@switch> show virtual-chassis protocol database detail
member0:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 114 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 118 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member1:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbc, Checksum: 0x3111, Lifetime: 115 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15
```

```
0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

```
member8:
```

```
-----
0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 118 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150
```

```
0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150
```

```
0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 116 secs
```

```
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694e, Checksum: 0xca97, Lifetime: 115 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15

member9:
-----

0000.4a75.9b7c.00-00 Sequence: 0x1ddbd, Checksum: 0xfd83, Lifetime: 116 secs
Neighbor: 0000.73e9.9a57.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-1/1.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/2.32768 Metric: 150

0000.73e9.9a57.00-00 Sequence: 0xf381, Checksum: 0xe065, Lifetime: 117 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-1/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-1/1.32768 Metric: 150

0021.59f7.d000.00-00 Sequence: 0x168af, Checksum: 0x8b0b, Lifetime: 113 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0026.888d.6800.00 Interface: vcp-4/0/1.32768 Metric: 15

0026.888d.6800.00-00 Sequence: 0x1694f, Checksum: 0xa61a, Lifetime: 116 secs
Neighbor: 0000.4a75.9b7c.00 Interface: vcp-0/0.32768 Metric: 150
Neighbor: 0000.73e9.9a57.00 Interface: vcp-0/1.32768 Metric: 150
Neighbor: 0021.59f7.d000.00 Interface: vcp-3/0/4.32768 Metric: 15
```

## show virtual-chassis protocol interface

<b>Syntax</b>	<pre>show virtual-chassis protocol interface &lt;brief   detail&gt; &lt;all-members&gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Display information about Virtual Chassis Control Protocol (VCCP) statistics for VCCP-enabled interfaces within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display the VCCP interface statistics in brief form for all members of the Virtual Chassis or VCF.</p> <p><b>brief   detail</b> —(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> option provides more output than the <b>brief</b> option.</p> <p><b>all-members</b>—(Optional) Display VCCP interface statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP interface statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP interface statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP interface statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li> <li>• <i>Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol interface on page 7146</a>
<b>Output Fields</b>	<a href="#">Table 672 on page 7005</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 687: show virtual-chassis protocol interface Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Name of the VCP.	All levels
<b>State</b>	State of the link. Outputs include: <ul style="list-style-type: none"> <li>• <b>Up</b>—The link is up.</li> <li>• <b>Down</b>—The link is down.</li> </ul>	All levels
<b>Metric</b>	Metric of the prefix or neighbor.	All levels

## Sample Output

### show virtual-chassis protocol interface

```
user@switch> show virtual-chassis protocol interface
```

```
member0:
```

```
-----
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Up	150
vcp-0/1.32768	Up	150
vcp-4/0/1.32768	Up	15
vcp-4/0/7.32768	Down	15

```
member1:
```

```
-----
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Up	150
vcp-0/1.32768	Up	150
vcp-3/0/4.32768	Up	15

```
member8:
```

```
-----
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Down	150
vcp-1/0.32768	Up	150
vcp-1/1.32768	Up	150
vcp-1/2.32768	Up	150
vcp-1/3.32768	Down	150
vcp-2/0.32768	Up	150
vcp-2/1.32768	Down	150
vcp-2/2.32768	Down	150
vcp-2/3.32768	Down	150

```
member9:
```

```
-----
IS-IS interface database:
```

Interface	State	Metric
vcp-0/0.32768	Down	150
vcp-1/0.32768	Up	150
vcp-1/1.32768	Up	150
vcp-1/2.32768	Down	150
vcp-1/3.32768	Down	150



## show virtual-chassis protocol route

<b>Syntax</b>	show virtual-chassis protocol route <all-members> <destination-id> <local> <member member-id>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the unicast and multicast Virtual Chassis Control Protocol (VCCP) routing tables within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display the unicast and multicast routing tables for all members of the Virtual Chassis.</p> <p><b>all-members</b>—(Optional) Display the unicast and multicast routing tables for all members of the Virtual Chassis or VCF.</p> <p><b>destination-id</b>—(Optional) Display the unicast and multicast routing tables to the specified destination member ID for each member of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the unicast and multicast routing tables on the device where this command is entered.</p> <p><b>member member-id</b>—(Optional) Display the unicast and multicast routing tables for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol route on page 7149</a>
<b>Output Fields</b>	<a href="#">Table 673 on page 7007</a> lists the output fields for the <b>show virtual-chassis protocol route</b> command. Output fields are listed in the approximate order in which they appear.

**Table 688: show virtual-chassis protocol route Output Fields**

Field Name	Field Description
<b>Dev</b>	MAC address of the member storing the VCCP routing table.
<b>Version</b>	Version of the shortest-path-first algorithm that generated the routing table.

Table 688: show virtual-chassis protocol route Output Fields (*continued*)

Field Name	Field Description
<b>System ID</b>	MAC address of the device.
<b>Version</b>	Version of the shortest-path-first (SPF) algorithm that generated the route.
<b>Metric</b>	The metric number to get to that device.
<b>Interface</b>	Name of the Virtual Chassis port (VCP) interface connecting the devices.
<b>Via</b>	MAC address of the next-hop device, if applicable.

## Sample Output

### show virtual-chassis protocol route

```

user@switch> show virtual-chassis protocol route
member0:
-----
Dev 0021.59f7.d000 ucast routing table          Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21      150 vcp-0/1.32768 0000.4a75.9b7c
0000.73e9.9a57    21      165 vcp-4/0/1.32768 0026.888d.6800
0021.59f7.d000    21        0
0026.888d.6800    21      15 vcp-4/0/1.32768 0026.888d.6800

Dev 0021.59f7.d000 mcast routing table          Current version: 21
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    21
0000.73e9.9a57    21
0021.59f7.d000    21          vcp-4/0/1.32768
                   vcp-0/1.32768
0026.888d.6800    21

member1:
-----
Dev 0026.888d.6800 ucast routing table          Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25      150 vcp-0/0.32768 0000.4a75.9b7c
0000.73e9.9a57    25      150 vcp-0/1.32768 0000.73e9.9a57
0021.59f7.d000    25        15 vcp-3/0/4.32768 0021.59f7.d000
0026.888d.6800    25        0

Dev 0026.888d.6800 mcast routing table          Current version: 25
-----
System ID      Version  Metric Interface  Via
0000.4a75.9b7c    25
0000.73e9.9a57    25          vcp-3/0/4.32768
0021.59f7.d000    25          vcp-0/1.32768
0026.888d.6800    25          vcp-3/0/4.32768
                   vcp-0/0.32768

```

vcp-0/1.32768

member8:

-----

Dev 0000.4a75.9b7c ucast routing table                      Current version: 39

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	39	0		
0000.73e9.9a57	39	150	vcp-1/0.32768	0000.73e9.9a57
0021.59f7.d000	39	150	vcp-2/0.32768	0021.59f7.d000
0026.888d.6800	39	150	vcp-1/2.32768	0026.888d.6800

Dev 0000.4a75.9b7c mcast routing table                      Current version: 39

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	39		vcp-1/0.32768	
			vcp-2/0.32768	
			vcp-1/2.32768	
0000.73e9.9a57	39			
0021.59f7.d000	39			
0026.888d.6800	39			

member9:

-----

Dev 0000.73e9.9a57 ucast routing table                      Current version: 31

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	31	150	vcp-1/0.32768	0000.4a75.9b7c
0000.73e9.9a57	31	0		
0021.59f7.d000	31	165	vcp-1/1.32768	0026.888d.6800
0026.888d.6800	31	150	vcp-1/1.32768	0026.888d.6800

Dev 0000.73e9.9a57 mcast routing table                      Current version: 31

-----

System ID	Version	Metric	Interface	Via
0000.4a75.9b7c	31			
0000.73e9.9a57	31		vcp-1/0.32768	
			vcp-1/1.32768	
0021.59f7.d000	31			
0026.888d.6800	31			



## show virtual-chassis protocol statistics

<b>Syntax</b>	show virtual-chassis protocol statistics <all-members> <interface-name> <local> <member member-id>
<b>Release Information</b>	Command introduced in Junos OS Release 10.4 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the Virtual Chassis Control Protocol (VCCP) statistics for all hardware devices within the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display VCCP statistics for all members of the Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display VCCP statistics for the specified interface.</p> <p><b>local</b>—(Optional) Display VCCP statistics for the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display VCCP statistics for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>EX Series Virtual Chassis Overview</i></li> <li>• <a href="#">Understanding QFX Series Virtual Chassis on page 6907</a></li> <li>• <i>Understanding the Virtual Chassis Control Protocol in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis protocol statistics on page 7152</a>
<b>Output Fields</b>	<a href="#">Table 674 on page 7010</a> lists the output fields for the <b>show virtual-chassis protocol interface</b> command. Output fields are listed in the approximate order in which they appear.

**Table 689: show virtual-chassis protocol statistics Output Fields**

Field Name	Field Description
<b>PDU type</b>	Protocol data unit type.
<b>Received</b>	Number of PDUs received since VCCP started or since the statistics were set to zero.
<b>Processed</b>	Number of PDUs received minus the number of PDUs dropped.

Table 689: show virtual-chassis protocol statistics Output Fields (*continued*)

Field Name	Field Description
<b>Drops</b>	Number of PDUs dropped.
<b>Sent</b>	Number of PDUs transmitted since VCCP started or since the statistics were set to zero.
<b>Rexmit</b>	Number of PDUs retransmitted since VCCP started or since the statistics were set to zero.
<b>Total Packets Received</b>	Number of PDUs received since VCCP started or since the statistics were set to zero.
<b>Total Packets Sent</b>	Number of PDUs sent since VCCP started or since the statistics were set to zero.
<b>LSP queue length</b>	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
<b>SPF runs</b>	Number of shortest-path-first (SPF) calculations that have been performed.
<b>Fragments Rebuilt</b>	Number of link-state PDU fragments that the local system has computed.
<b>LSP Regenerations</b>	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
<b>Purges initiated</b>	Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.

## Sample Output

### show virtual-chassis protocol statistics

```

user@switch> show virtual-chassis protocol statistics
member0:
-----
IS-IS statistics for 0021.59f7.d000:
PDU type      Received    Processed      Drops      Sent      Rexmit
LSP            8166        8166           0         4551         0
HELLO          1659        1659           0         1693         0
CSNP             2            2             0            3         0
PSNP           1909        1909           0         2293         0
Unknown         0            0             0            0         0
Totals        11736       11736           0         8540         0

Total packets received: 11736 Sent: 8540

LSP queue length: 0 Drops: 0
SPF runs: 9
Fragments rebuilt: 1640
LSP regenerations: 1
Purges initiated: 0

member1:
-----
IS-IS statistics for 0026.888d.6800:

```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	10909	10909	0	12088	0
HELLO	1877	1877	0	2251	0
CSNP	3	3	0	3	0
PSNP	3846	3846	0	3732	0
Unknown	0	0	0	0	0
Totals	16635	16635	0	18074	0

Total packets received: 16635 Sent: 18074

LSP queue length: 0 Drops: 0  
 SPF runs: 13  
 Fragments rebuilt: 1871  
 LSP regenerations: 2  
 Purges initiated: 0

member8:

IS-IS statistics for 0000.4a75.9b7c:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	7935	7935	0	14865	0
HELLO	2695	2695	0	7124	0
CSNP	4	4	0	4	0
PSNP	4398	4398	0	3666	0
Unknown	0	0	0	0	0
Totals	15032	15032	0	25659	0

Total packets received: 15032 Sent: 25659

LSP queue length: 0 Drops: 0  
 SPF runs: 26  
 Fragments rebuilt: 2666  
 LSP regenerations: 4  
 Purges initiated: 0

member9:

IS-IS statistics for 0000.73e9.9a57:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	10800	10800	0	6327	0
HELLO	1492	1492	0	2356	0
CSNP	2	2	0	2	0
PSNP	2683	2683	0	3149	0
Unknown	0	0	0	0	0
Totals	14977	14977	0	11834	0

Total packets received: 14977 Sent: 11834

LSP queue length: 0 Drops: 0  
 SPF runs: 19  
 Fragments rebuilt: 1510  
 LSP regenerations: 6  
 Purges initiated: 0

## show virtual-chassis

<b>Syntax</b>	<b>show virtual-chassis</b> <b>&lt;status&gt;</b>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p> <p><b>Fabric ID</b>, <b>Fabric Mode</b>, and <b>Route Mode</b> output fields introduced in Junos OS Release 13.2X51-D20.</p> <p><b>Alias-Name</b> output field introduced in Junos OS Release 14.1X53-D10.</p>
<b>Description</b>	Display information about all members of the Virtual Chassis or VCF.
<b>Options</b>	<p><b>none</b>—Display information about all Virtual Chassis or VCF member devices.</p> <p><b>status</b>—Same output as for <b>show virtual-chassis</b>.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis active-topology on page 6985</a></li> <li>• <a href="#">show virtual-chassis protocol adjacency on page 6996</a></li> <li>• <a href="#">show virtual-chassis vc-path on page 7018</a></li> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis (EX4200 Virtual Chassis) on page 7156</a></p> <p><a href="#">show virtual-chassis (EX8200 Virtual Chassis) on page 7156</a></p> <p><a href="#">show virtual-chassis (Virtual Chassis Fabric) on page 7157</a></p>
<b>Output Fields</b>	<p><a href="#">Table 675 on page 7014</a> lists the output fields for the <b>show virtual-chassis</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 690: show virtual-chassis Output Fields**

Field Name	Field Description
<b>Fabric ID</b>	Assigned ID used to identify the VCF.
<b>Fabric Mode</b>	Mode of the VCF: Enabled, Disabled, or Mixed.
<b>Virtual Chassis ID</b>	Assigned ID that applies to the entire Virtual Chassis or VCF.
<b>Virtual Chassis Mode</b>	Mode of the Virtual Chassis or VCF: Enabled, Disabled, or Mixed.

Table 690: show virtual-chassis Output Fields (*continued*)

Field Name	Field Description
<b>Member ID</b>	Assigned member ID and FPC: <ul style="list-style-type: none"> <li>On all EX Series Virtual Chassis except EX8200 Virtual Chassis, and on a VCF, the FPC number refers to the member ID assigned to the switch.</li> <li>On EX8200 Virtual Chassis, member IDs are numbered 0 through 9. The FPC number indicates the slot number of the line card within the Virtual Chassis. The FPC number on member 0 is always 0 through 15. The FPC number on member 1 is always 16 through 31. The FPC number on member 2 is always 32 through 47; and so on for the members.</li> </ul>
<b>Status</b>	For a nonprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is currently connected to the Virtual Chassis or VCF configuration.</li> <li><b>NotPrsnt</b> for a member ID that has been assigned but is not currently connected.</li> </ul> For a preprovisioned configuration: <ul style="list-style-type: none"> <li><b>Prsnt</b> for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis or VCF.</li> <li><b>Unprvsnd</b> for a member that is interconnected with the Virtual Chassis or VCF configuration but is not specified in the preprovisioned configuration file.</li> </ul>
<b>Serial No</b>	Serial number of the member device.
<b>Alias-Name</b>	The user-configured alias of the member device.  The <b>Alias-Name</b> field appears only if an alias has been configured for at least one device in the Virtual Chassis or VCF. Aliases are configured using the <b>alias-name</b> statement in the <code>[edit virtual-chassis aliases serial-number serial-number]</code> hierarchy.
<b>Model</b>	Model number of the member device.
<b>Mastership Priority</b>	Mastership priority value of the member device.
<b>Role</b>	Role of the member device: master, backup, or linecard.
<b>Mixed Mode</b>	Mixed mode configuration status: <ul style="list-style-type: none"> <li><b>Y</b> for a member device configured in mixed mode.</li> <li><b>N</b> for a member device not configured in mixed mode.</li> <li><b>NA</b> for a member device that cannot be configured in mixed mode.</li> </ul>
<b>Route Mode</b>	The route mode of the member device: fabric (F) or Virtual Chassis (V).
<b>Location</b>	Location of the member device.  If this field is empty, the location field was not set for the device.
<b>Neighbor List</b>	Member ID of the neighbor member to which this member's Virtual Chassis port (VCP) is connected.

## Sample Output

### show virtual-chassis (EX4200 Virtual Chassis)

```

user@switch> show virtual-chassis
Virtual Chassis ID: 0019.e250.47a0
Virtual Chassis Mode: Enabled

```

Member ID	Status	Serial No	Model	Mastership priority	Role	Mixed Mode	Neighbor List ID	Interface
0 (FPC 0)	Prsnt	AK0207360276	ex4200-24t	249	Master*	N	8	vcp-0
							1	vcp-1
1 (FPC 1)	Prsnt	AK0207360281	ex4200-24t	248	Backup	N	0	vcp-0
							2	vcp-1
2 (FPC 2)	Prsnt	AJ0207391130	ex4200-48p	247	Linecard	N	1	vcp-0
							3	vcp-1
3 (FPC 3)	Prsnt	AK0207360280	ex4200-24t	246	Linecard	N	2	vcp-0
							4	vcp-1
4 (FPC 4)	Prsnt	AJ0207391113	ex4200-48p	245	Linecard	N	3	vcp-0
							5	vcp-1
5 (FPC 5)	Prsnt	BP0207452204	ex4200-48t	244	Linecard	N	4	vcp-0
							6	vcp-1
6 (FPC 6)	Prsnt	BP0207452222	ex4200-48t	243	Linecard	N	5	vcp-0
							7	vcp-1
7 (FPC 7)	Prsnt	BR0207432028	ex4200-24f	242	Linecard	N	6	vcp-0
							8	vcp-1
8 (FPC 8)	Prsnt	BR0207431996	ex4200-24f	241	Linecard	N	7	vcp-0
							0	vcp-1

Member ID for next new member: 9 (FPC 9)

### show virtual-chassis (EX8200 Virtual Chassis)

```

user@external-routing-engine> show virtual-chassis
Virtual Chassis ID: c806.0842.de51
Virtual Chassis Mode: Enabled

```

Member ID	Status	Serial No	Model	Mastership priority	Role	Neighbor List ID	Interface
0 (FPC 0-15)	Prsnt	BA0908380001	ex8216	0	Linecard	8	vcp-0/0
						8	vcp-0/1
						1	vcp-4/0/4
1 (FPC 16-31)	Prsnt	BT0909411634	ex8208	0	Linecard	8	vcp-0/0
						0	vcp-3/0/4
8 (FPC 128-143)	Prsnt	062009000021	ex-xre	128	Master	9	vcp-1/0
						1	vcp-1/2

```

9 (FPC 144-159) Prsnt 062009000022 ex-xre 128 Backup*
9 vcp-1/3
0 vcp-2/0
9 vcp-2/1
0 vcp-1/1
8 vcp-1/0
8 vcp-1/2
8 vcp-1/3
8 vcp-1/3

```

### show virtual-chassis (Virtual Chassis Fabric)

```

user@switch> show virtual-chassis
Preprovisioned Virtual Chassis Fabric
Fabric ID: 0282.5fa0.3f08
Fabric Mode: Enabled

```

List	Member ID	Status	Serial No	Model	Mstr prio	Role	Mixed	Route	Neighbor
Interface									
0 (FPC 0)	Prsnt	AB3112430001	qfx5100-48s	129	Master*	N	F	3	
vcp-255/1/0									2
vcp-255/1/1									4
vcp-255/1/2									4
vcp-255/1/3									4
1 (FPC 1)	Prsnt	AB3112230001	qfx5100-48s	129	Backup	N	F	3	
vcp-255/1/0									2
vcp-255/1/1									4
vcp-255/1/2									4
vcp-255/1/3									4
2 (FPC 2)	Prsnt	AB3112460011	qfx5100-48s	0	Linecard	N	F	1	
vcp-255/1/0									0
vcp-255/1/1									0
3 (FPC 3)	Prsnt	AB3112460011	qfx5100-48s	0	Linecard	N	F	1	
vcp-255/1/0									0
vcp-255/1/1									0
4 (FPC 4)	Prsnt	AB3112430011	qfx5100-48s	0	Linecard	N	F	1	
vcp-255/1/0									0
vcp-255/1/1									0

## show virtual-chassis vc-port

<b>Syntax</b>	show virtual-chassis vc-port <all-members> <local> <member <i>member-id</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	Display the status of the Virtual Chassis ports (VCPs), including both the dedicated VCPs and the uplink ports configured as VCPs.
<b>Options</b>	<p><b>none</b>—Display the operational status of all VCPs of the member switch where the command is issued.</p> <p><b>all-members</b>—(Optional) Display the operational status of all VCPs on all members of the Virtual Chassis or VCF.</p> <p><b>local</b>—(Optional) Display the operational status of the switch or external Routing Engine on which this command is entered.</p> <p><b>member <i>member-id</i></b>—(Optional) Display the operational status of all VCPs for the specified member of the Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show virtual-chassis vc-port statistics on page 7024</a></li> <li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li> <li>• <i>Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis vc-port (EX4200 Virtual Chassis) on page 7160</a> <a href="#">show virtual-chassis vc-port (EX8200 Virtual Chassis) on page 7160</a> <a href="#">show virtual-chassis vc-port all-members on page 7161</a>
<b>Output Fields</b>	Table 677 on page 7020 lists the output fields for the <b>show virtual-chassis vc-port</b> command. Output fields are listed in the approximate order in which they appear.

Table 691: show virtual-chassis vc-port Output Fields

Field Name	Field Description
<b>fpcnumber</b>	The FPC number is the same as the member ID.



Table 691: show virtual-chassis vc-port Output Fields (*continued*)

Field Name	Field Description
Interface or PIC/Port	<p>VCP name.</p> <ul style="list-style-type: none"> <li>The dedicated VCPs in an EX4200 or EX4500 Virtual Chassis are <b>vcp-0</b> and <b>vcp-1</b>. The dedicated VCPs in an EX4550 Virtual Chassis are <b>VCP-1/0</b>, <b>VCP-1/1</b>, <b>VCP-2/0</b>, and <b>VCP-2/1</b>.</li> <li>Optical ports set as VCPs are named <b>1/0</b> and <b>1/1</b>, representing the PIC number and the port number.</li> <li>The native VCP (port 0) on an XRE200 External Routing Engine in an EX8200 Virtual Chassis is named <b>vcp-0</b>.</li> <li>The VCPs on each Virtual Chassis Control Interface (VCCI) module in an XRE200 External Routing Engine are named using the <b>vcp-slot-number/port-number</b> convention; for instance, <b>vcp-1/0</b>.</li> <li>The VCPs on EX8200 member switches are named using the <b>vcp-slot-number/pic-number/interface-number</b> convention; for instance, <b>vcp-3/0/2</b>.</li> <li>A <b>255</b> as the first number in your port number indicates that your VCP is part of a Link Aggregation group (LAG) bundle. For instance, a display of <b>vcp-255/1/0</b> indicates that the dedicated VCP named <b>vcp-1/0</b> is part of a LAG bundle. A display of <b>vcp-255/1/0</b> indicates that an uplink port that was previously named <b>xe-0/1/0</b> is now part of a VCP LAG bundle.</li> </ul>
Type	<p>Type of VCP:</p> <ul style="list-style-type: none"> <li><b>Dedicated</b>—The rear panel VCP on an EX4200, EX4500, or EX4550 switch, or any VCP link connected to an XRE200 External Routing Engine in an EX8200 Virtual Chassis.</li> <li><b>Configured</b>—Optical port configured as a VCP.</li> <li><b>Auto-Configured</b>—Optical port autoconfigured as a VCP.</li> </ul> <p>See <i>Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure)</i> or <i>Setting a 10-Gigabit Ethernet Port as a Virtual Chassis Port in an EX8200 Virtual Chassis (CLI Procedure)</i> for information about configuring VCPs.</p>
Trunk ID	<p>A positive-number ID assigned to a link aggregation group (LAG) formed by the Virtual Chassis. The trunk ID value is <b>-1</b> if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member in an EX4200 or EX4500 Virtual Chassis.</p> <p>Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater.</p> <p>The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.</p>
Status	<p>Interface status:</p> <ul style="list-style-type: none"> <li><b>absent</b>—Interface is not a VCP link.</li> <li><b>down</b>—VCP link is down.</li> <li><b>up</b>—VCP link is up.</li> </ul>
Speed (mbps)	Speed of the interface in megabits per second.
Neighbor ID/Interface	The Virtual Chassis member ID and interface of a VCP on a member that is connected to the interface or PIC/Port field in the same row as this interface.

## Sample Output

### show virtual-chassis vc-port (EX4200 Virtual Chassis)

```
user@switch> show virtual-chassis vc-port
```

```
fpc0:
```

Interface or PIC / Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0	Dedicated	1	Up	32000	1	vcp-1
vcp-1	Dedicated	2	Up	32000	0	vcp-0
1/0	Auto-Configured	3	Up	1000	2	vcp-255/1/0
1/0	Auto-Configured	3	Up	1000	2	vcp-255/1/1

### show virtual-chassis vc-port (EX8200 Virtual Chassis)

```
user@external-routing-engine> show virtual-chassis vc-port
```

```
member0:
```

Interface or Slot/PIC/Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0/0	Dedicated	-1	Up	1000	8	vcp-1/1
vcp-0/1	Dedicated	-1	Up	1000	8	vcp-2/0
4/0/4	Configured	-1	Up	10000	1	vcp-3/0/4
4/0/7	Configured	-1	Down	10000		
4/0/3	Configured		Absent			
4/0/2	Configured		Absent			
4/0/5	Configured		Absent			
4/0/6	Configured		Absent			
4/0/1	Configured		Absent			
4/0/0	Configured		Absent			

```
member1:
```

Interface or Slot/PIC/Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0/0	Dedicated	-1	Up	1000	8	vcp-1/2
3/0/0	Configured	-1	Down	10000		
3/0/1	Configured	-1	Down	10000		
3/0/4	Configured	-1	Up	10000	0	vcp-4/0/4
3/0/5	Configured		Absent			
4/0/5	Configured		Absent			
4/0/4	Configured		Absent			

```
member8:
```

Interface or Slot/PIC/Port	Type	Trunk ID	Status	Speed (mbps)	Neighbor ID	Interface
vcp-0/0	Dedicated	-1	Down	1000		
vcp-1/0	Dedicated	-1	Up	1000	9	vcp-1/0
vcp-1/1	Dedicated	-1	Up	1000	0	vcp-0/0
vcp-1/2	Dedicated	-1	Up	1000	1	vcp-0/0
vcp-1/3	Dedicated	-1	Up	1000	9	vcp-1/3
vcp-2/0	Dedicated	-1	Up	1000	0	vcp-0/1
vcp-2/1	Dedicated	-1	Up	1000	9	vcp-1/2
vcp-2/2	Dedicated	-1	Down	1000		

```
vcp-2/3      Dedicated      -1   Down      1000
```

```
member9:
```

```
-----
Interface    Type            Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
Slot/PIC/Port
vcp-0/0      Dedicated       -1     Disabled  1000
vcp-1/0      Dedicated       -1     Up        1000      8   vcp-1/0
vcp-1/1      Dedicated       -1     Down      1000
vcp-1/2      Dedicated       -1     Up        1000      8   vcp-2/1
vcp-1/3      Dedicated       -1     Up        1000      8   vcp-1/3
```

### show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members
```

```
fpc0:
```

```
-----
Interface    Type            Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated       1      Up        32000    1   vcp-1
vcp-1        Dedicated       2      Up        32000    0   vcp-0
1/0          Auto-Configured 3      Up        1000     2   vcp-255/1/0
1/1          Auto-Configured 3      Up        1000     2   vcp-255/1/1
```

```
fpc1:
```

```
-----
Interface    Type            Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated       1      Up        32000    0   vcp-1
vcp-1        Dedicated       2      Up        32000    0   vcp-0
1/0          Auto-Configured -1     Up        1000     3   vcp-255/1/0
```

```
fpc2:
```

```
-----
Interface    Type            Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated       1      Up        32000    3   vcp-1
vcp-1        Dedicated       2      Up        32000    3   vcp-0
1/0          Auto-Configured 3      Up        1000     0   vcp-255/1/0
1/1          Auto-Configured 3      Up        1000     0   vcp-255/1/1
```

```
fpc3:
```

```
-----
Interface    Type            Trunk  Status    Speed    Neighbor
or           or              ID     Status    (mbps)   ID  Interface
PIC / Port
vcp-0        Dedicated       1      Up        32000    2   vcp-0
vcp-1        Dedicated       2      Up        32000    2   vcp-1
1/0          Auto-Configured -1     Up        1000     1   vcp-255/1/0
```

## show virtual-chassis vc-port diagnostics optics

---

<b>Syntax</b>	<code>show virtual-chassis vc-port diagnostics optics</code> <code>&lt;all-members&gt;</code> <code>&lt;interface-name&gt;</code> <code>&lt;local&gt;</code> <code>&lt;member member-id&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 12.2 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).
<b>Description</b>	<p>Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.</p> <p>On some EX Series switches, the <b>request virtual-chassis vc-port diagnostics optics</b> command must be entered to run a diagnostic scan before you can gather the <b>show virtual-chassis vc-port diagnostics optics</b> output.</p>
<b>Options</b>	<p><b>none</b>—Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>all-members</b>—(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display diagnostics information for the transceiver installed in a specified VCP.</p> <p><b>local</b>—(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li><li>• <i>Installing a Transceiver in an EX Series Switch</i></li><li>• <i>Removing a Transceiver from an EX Series Switch</i></li><li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show virtual-chassis vc-port diagnostics optics on page 7165</a> <a href="#">show virtual-chassis vc-port diagnostics optics (interface-name) on page 7170</a>

[show virtual-chassis vc-port diagnostics optics local on page 7172](#)

[show virtual-chassis vc-port diagnostics optics \(member member-id\) on page 7174](#)

**Output Fields** [Table 678 on page 7025](#) lists the output fields for the **show virtual-chassis vc-port diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

**Table 692: show virtual-chassis vc-port diagnostics optics Output Fields**

Field Name	Field Description
FPC	Displays the FPC slot number.
Virtual chassis port	Displays the name of the VCP.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is <i>On</i> or <i>Off</i> .
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is <i>On</i> or <i>Off</i> .
Laser bias current high warning	Displays whether the laser bias power setting high warning is <i>On</i> or <i>Off</i> .
Laser bias current low warning	Displays whether the laser bias power setting low warning is <i>On</i> or <i>Off</i> .
Laser output power high alarm	Displays whether the laser output power high alarm is <i>On</i> or <i>Off</i> .
Laser output power low alarm	Displays whether the laser output power low alarm is <i>On</i> or <i>Off</i> .
Laser output power high warning	Displays whether the laser output power high warning is <i>On</i> or <i>Off</i> .
Laser output power low warning	Displays whether the laser output power low warning is <i>On</i> or <i>Off</i> .
Module temperature high alarm	Displays whether the module temperature high alarm is <i>On</i> or <i>Off</i> .
Module temperature low alarm	Displays whether the module temperature low alarm is <i>On</i> or <i>Off</i> .
Module temperature high warning	Displays whether the module temperature high warning is <i>On</i> or <i>Off</i> .
Module temperature low warning	Displays whether the module temperature low warning is <i>On</i> or <i>Off</i> .

Table 692: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage high alarm	Displays whether the module voltage high alarm is <i>On</i> or <i>Off</i> .
Module voltage low alarm	Displays whether the module voltage low alarm is <i>On</i> or <i>Off</i> .
Module voltage high warning	Displays whether the module voltage high warning is <i>On</i> or <i>Off</i> .
Module voltage low warning	Displays whether the module voltage low warning is <i>On</i> or <i>Off</i> .
Laser rx power high alarm	Displays whether the receive laser power high alarm is <i>On</i> or <i>Off</i> .
Laser rx power low alarm	Displays whether the receive laser power low alarm is <i>On</i> or <i>Off</i> .
Laser rx power high warning	Displays whether the receive laser power high warning is <i>On</i> or <i>Off</i> .
Laser rx power low warning	Displays whether the receive laser power low warning is <i>On</i> or <i>Off</i> .
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.

Table 692: show virtual-chassis vc-port diagnostics optics Output Fields (*continued*)

Field Name	Field Description
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

## Sample Output

### show virtual-chassis vc-port diagnostics optics

```

user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc1:
-----
Virtual chassis port: vcp-0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-1
  Optical diagnostics                : N/A

fpc2:
-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                 : 4.130 mA
  Laser output power                  : 0.2450 mW / -6.11 dBm
  Module temperature                 : 32 degrees C / 90 degrees F
  Module voltage                     : 3.3530 V
  Receiver signal average optical power : 0.0971 mW / -10.13 dBm
  Laser bias current high alarm      : Off
  Laser bias current low alarm       : Off

```

```

Laser bias current high warning      : Off
Laser bias current low warning       : Off
Laser output power high alarm        : Off
Laser output power low alarm         : Off
Laser output power high warning      : Off
Laser output power low warning       : Off
Module temperature high alarm        : Off
Module temperature low alarm         : Off
Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 14.998 mA
Laser bias current low alarm threshold : 0.998 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 1.198 mA
Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.600 V
Module voltage low alarm threshold : 3.000 V
Module voltage high warning threshold : 3.499 V
Module voltage low warning threshold : 3.099 V
Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                  : 5.428 mA
Laser output power                  : 0.4760 mW / -3.22 dBm
Module temperature                  : 28 degrees C / 83 degrees F
Module voltage                      : 3.3440 V
Receiver signal average optical power : 0.4002 mW / -3.98 dBm
Laser bias current high alarm       : Off
Laser bias current low alarm        : Off
Laser bias current high warning     : Off
Laser bias current low warning      : Off
Laser output power high alarm       : Off
Laser output power low alarm        : Off
Laser output power high warning     : Off
Laser output power low warning      : Off
Module temperature high alarm       : Off
Module temperature low alarm        : Off
Module temperature high warning     : Off
Module temperature low warning      : Off
Module voltage high alarm           : Off
Module voltage low alarm            : Off
Module voltage high warning         : Off
Module voltage low warning          : Off
Laser rx power high alarm           : Off

```



```

Laser rx power low alarm           : Off
Laser rx power high warning        : Off
Laser rx power low warning         : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc3:

-----  
Virtual chassis port: vcp-255/0/2

```

Laser bias current           : 7.876 mA
Laser output power           : 0.5330 mW / -2.73 dBm
Module temperature           : 26 degrees C / 78 degrees F
Module voltage               : 3.3060 V
Receiver signal average optical power : 0.4885 mW / -3.11 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm  : Off
Laser output power low alarm   : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm  : Off
Module temperature low alarm   : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm      : Off
Module voltage low alarm       : Off
Module voltage high warning    : Off
Module voltage low warning     : Off
Laser rx power high alarm      : Off
Laser rx power low alarm       : Off
Laser rx power high warning    : Off
Laser rx power low warning     : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F

```

```
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current : 5.052 mA
Laser output power : 0.5030 mW / -2.98 dBm
Module temperature : 24 degrees C / 75 degrees F
Module voltage : 3.2890 V
Receiver signal average optical power : 0.5028 mW / -2.99 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm
Virtual chassis port: vcp-255/0/4
Laser bias current : 7.978 mA
Laser output power : 0.5460 mW / -2.63 dBm
Module temperature : 24 degrees C / 76 degrees F
```

```

Module voltage : 3.3060 V
Receiver signal average optical power : 0.6305 mW / -2.00 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 14.500 mA
Laser bias current low alarm threshold : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

fpc4:

```

-----
Virtual chassis port: vcp-0
  Optical diagnostics : N/A
Virtual chassis port: vcp-1
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/4
  Laser bias current : 7.860 mA
  Laser output power : 0.5370 mW / -2.70 dBm
  Module temperature : 24 degrees C / 75 degrees F
  Module voltage : 3.2920 V
  Receiver signal average optical power : 0.6271 mW / -2.03 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm : Off
  Laser output power high warning : Off

```

```

Laser output power low warning           : Off
Module temperature high alarm            : Off
Module temperature low alarm             : Off
Module temperature high warning          : Off
Module temperature low warning           : Off
Module voltage high alarm                : Off
Module voltage low alarm                 : Off
Module voltage high warning              : Off
Module voltage low warning               : Off
Laser rx power high alarm                : Off
Laser rx power low alarm                 : Off
Laser rx power high warning              : Off
Laser rx power low warning               : Off
Laser bias current high alarm threshold  : 14.500 mA
Laser bias current low alarm threshold   : 3.500 mA
Laser bias current high warning threshold : 14.500 mA
Laser bias current low warning threshold : 3.500 mA
Laser output power high alarm threshold  : 1.8620 mW / 2.70 dBm
Laser output power low alarm threshold   : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold  : 75 degrees C / 167 degrees F
Module temperature low alarm threshold   : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold      : 3.135 V
Laser rx power high alarm threshold       : 1.9952 mW / 3.00 dBm
Laser rx power low alarm threshold        : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold      : 0.1023 mW / -9.90 dBm

```

#### show virtual-chassis vc-port diagnostics optics (interface-name)

```

user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:

```

```

fpc1:

```

```

fpc2:

```

```

Virtual chassis port: vcp-255/0/3
Laser bias current           : 5.448 mA
Laser output power           : 0.4770 mW / -3.21 dBm
Module temperature           : 28 degrees C / 82 degrees F
Module voltage               : 3.3450 V
Receiver signal average optical power : 0.3973 mW / -4.01 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm  : Off
Laser output power low alarm   : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm  : Off
Module temperature low alarm   : Off

```

```

Module temperature high warning      : Off
Module temperature low warning       : Off
Module voltage high alarm            : Off
Module voltage low alarm             : Off
Module voltage high warning          : Off
Module voltage low warning           : Off
Laser rx power high alarm            : Off
Laser rx power low alarm             : Off
Laser rx power high warning          : Off
Laser rx power low warning           : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold   : 3.630 V
Module voltage low alarm threshold    : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold  : 3.135 V
Laser rx power high alarm threshold   : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold    : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold  : 0.1023 mW / -9.90 dBm

```

fpc3:

-----  
Virtual chassis port: vcp-255/0/3

```

Laser bias current      : 5.040 mA
Laser output power      : 0.5020 mW / -2.99 dBm
Module temperature      : 24 degrees C / 74 degrees F
Module voltage          : 3.2870 V
Receiver signal average optical power : 0.5073 mW / -2.95 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm     : Off
Module voltage low alarm      : Off
Module voltage high warning   : Off
Module voltage low warning    : Off
Laser rx power high alarm     : Off
Laser rx power low alarm      : Off
Laser rx power high warning   : Off
Laser rx power low warning    : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA

```

```

Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold  : 2.500 mA
Laser output power high alarm threshold   : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold    : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold  : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold   : 75 degrees C / 167 degrees F
Module temperature low alarm threshold    : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold  : 0 degrees C / 32 degrees F
Module voltage high alarm threshold       : 3.630 V
Module voltage low alarm threshold        : 2.970 V
Module voltage high warning threshold     : 3.465 V
Module voltage low warning threshold      : 3.135 V
Laser rx power high alarm threshold       : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold        : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold     : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold      : 0.1023 mW / -9.90 dBm

```

fpc4:

-----

#### show virtual-chassis vc-port diagnostics optics local

```

user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
  Optical diagnostics : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current : 4.130 mA
  Laser output power : 0.2450 mW / -6.11 dBm
  Module temperature : 32 degrees C / 90 degrees F
  Module voltage     : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm : Off
  Laser output power low alarm  : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm : Off
  Module temperature low alarm  : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm : Off
  Module voltage low alarm  : Off
  Module voltage high warning : Off
  Module voltage low warning : Off
  Laser rx power high alarm : Off
  Laser rx power low alarm  : Off
  Laser rx power high warning : Off
  Laser rx power low warning : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold  : 0.998 mA

```

```

Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 1.198 mA
Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
Module temperature high alarm threshold : 85 degrees C / 185 degrees F
Module temperature low alarm threshold : -10 degrees C / 14 degrees F
Module temperature high warning threshold : 80 degrees C / 176 degrees F
Module temperature low warning threshold : -5 degrees C / 23 degrees F
Module voltage high alarm threshold : 3.600 V
Module voltage low alarm threshold : 3.000 V
Module voltage high warning threshold : 3.499 V
Module voltage low warning threshold : 3.099 V
Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current : 5.426 mA
Laser output power : 0.4760 mW / -3.22 dBm
Module temperature : 28 degrees C / 83 degrees F
Module voltage : 3.3450 V
Receiver signal average optical power : 0.3955 mW / -4.03 dBm
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : Off
Laser rx power high warning : Off
Laser rx power low warning : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold : 3.630 V
Module voltage low alarm threshold : 2.970 V
Module voltage high warning threshold : 3.465 V
Module voltage low warning threshold : 3.135 V
Laser rx power high alarm threshold : 1.5849 mW / 2.00 dBm

```

```

Laser rx power low alarm threshold      : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold   : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold    : 0.1023 mW / -9.90 dBm

```

### show virtual-chassis vc-port diagnostics optics (member member-id)

```

user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:

```

```

-----
Virtual chassis port: vcp-2/0
  Optical diagnostics                : N/A
Virtual chassis port: vcp-2/1
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/14
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/15
  Optical diagnostics                : N/A
Virtual chassis port: vcp-255/0/24
  Laser bias current                  : 4.130 mA
  Laser output power                  : 0.2450 mW / -6.11 dBm
  Module temperature                  : 31 degrees C / 88 degrees F
  Module voltage                      : 3.3530 V
  Receiver signal average optical power : 0.0961 mW / -10.17 dBm
  Laser bias current high alarm       : Off
  Laser bias current low alarm        : Off
  Laser bias current high warning     : Off
  Laser bias current low warning      : Off
  Laser output power high alarm       : Off
  Laser output power low alarm        : Off
  Laser output power high warning     : Off
  Laser output power low warning      : Off
  Module temperature high alarm       : Off
  Module temperature low alarm        : Off
  Module temperature high warning     : Off
  Module temperature low warning      : Off
  Module voltage high alarm           : Off
  Module voltage low alarm            : Off
  Module voltage high warning         : Off
  Module voltage low warning          : Off
  Laser rx power high alarm           : Off
  Laser rx power low alarm            : Off
  Laser rx power high warning         : Off
  Laser rx power low warning          : Off
  Laser bias current high alarm threshold : 14.998 mA
  Laser bias current low alarm threshold : 0.998 mA
  Laser bias current high warning threshold : 14.000 mA
  Laser bias current low warning threshold : 1.198 mA
  Laser output power high alarm threshold : 0.7940 mW / -1.00 dBm
  Laser output power low alarm threshold : 0.0790 mW / -11.02 dBm
  Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
  Laser output power low warning threshold : 0.0990 mW / -10.04 dBm
  Module temperature high alarm threshold : 85 degrees C / 185 degrees F
  Module temperature low alarm threshold : -10 degrees C / 14 degrees F
  Module temperature high warning threshold : 80 degrees C / 176 degrees F
  Module temperature low warning threshold : -5 degrees C / 23 degrees F
  Module voltage high alarm threshold : 3.600 V
  Module voltage low alarm threshold : 3.000 V
  Module voltage high warning threshold : 3.499 V
  Module voltage low warning threshold : 3.099 V
  Laser rx power high alarm threshold : 1.5848 mW / 2.00 dBm

```



```

Laser rx power low alarm threshold      : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold   : 1.2589 mW / 1.00 dBm
Laser rx power low warning threshold    : 0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
Laser bias current                      : 5.418 mA
Laser output power                      : 0.4770 mW / -3.21 dBm
Module temperature                      : 28 degrees C / 83 degrees F
Module voltage                          : 3.3450 V
Receiver signal average optical power   : 0.3964 mW / -4.02 dBm
Laser bias current high alarm           : Off
Laser bias current low alarm            : Off
Laser bias current high warning         : Off
Laser bias current low warning          : Off
Laser output power high alarm           : Off
Laser output power low alarm            : Off
Laser output power high warning         : Off
Laser output power low warning          : Off
Module temperature high alarm           : Off
Module temperature low alarm            : Off
Module temperature high warning         : Off
Module temperature low warning          : Off
Module voltage high alarm               : Off
Module voltage low alarm                : Off
Module voltage high warning             : Off
Module voltage low warning              : Off
Laser rx power high alarm               : Off
Laser rx power low alarm                : Off
Laser rx power high warning             : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold  : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold  : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold  : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

## show virtual-chassis vc-port statistics

---

<b>Syntax</b>	<pre>show virtual-chassis vc-port statistics &lt;all-members&gt; &lt;brief   detail   extensive &gt; &lt;interface-name&gt; &lt;local&gt; &lt;member member-id&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The options <b>all-members</b>, <b>brief</b>, <b>detail</b>, <b>extensive</b>, and <b>local</b> were added in Junos OS Release 9.3 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).</p>
<b>Description</b>	Display the traffic statistics collected on Virtual Chassis ports (VCPs).
<b>Options</b>	<p><b>none</b>—Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output. Using the <b>brief</b> option is equivalent to entering the command with no options (the default). The <b>detail</b> and <b>extensive</b> options provide identical displays.</p> <p><b>all-members</b>—(Optional) Display traffic statistics for VCPs of all members of a Virtual Chassis or VCF.</p> <p><b>interface-name</b>—(Optional) Display traffic statistics for the specified VCP.</p> <p><b>local</b>—(Optional) Display traffic statistics for VCPs on the switch or external Routing Engine on which this command is entered.</p> <p><b>member member-id</b>—(Optional) Display traffic statistics for VCPs on the specified member of a Virtual Chassis or VCF.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">clear virtual-chassis vc-port statistics on page 6978</a></li><li>• <a href="#">show virtual-chassis vc-port on page 7020</a></li><li>• <i>Monitoring the Virtual Chassis Status and Statistics on EX Series Virtual Chassis</i></li><li>• <i>Verifying Virtual Chassis Ports in an EX8200 Virtual Chassis</i></li></ul>
<b>List of Sample Output</b>	<p><a href="#">show virtual-chassis vc-port statistics on page 7179</a></p> <p><a href="#">show virtual-chassis vc-port statistics (EX8200 Virtual Chassis) on page 7180</a></p> <p><a href="#">show virtual-chassis vc-port statistics brief on page 7180</a></p> <p><a href="#">show virtual-chassis vc-port statistics extensive on page 7180</a></p> <p><a href="#">show virtual-chassis vc-port statistics member 0 on page 7182</a></p>

**Output Fields** Table 678 on page 7025 lists the output fields for the **show virtual-chassis vc-port statistics** command. Output fields are listed in the approximate order in which they appear.

**Table 693: show virtual-chassis vc-port statistics Output Fields**

Field Name	Field Description	Level of Output
<b>fpcnumber</b>	(All Virtual Chassis except EX8200 Virtual Chassis. VCF) ID of the Virtual Chassis member. The FPC number is the same as the member ID.	All levels
<b>member number</b>	(EX8200 Virtual Chassis only) Member ID of the Virtual Chassis member.	All levels
<b>Interface</b>	VCP name.	<b>brief</b>
<b>Input Octets/Packets</b>	Number of octets and packets received on the VCP.	<b>brief, member, none</b>
<b>Output Octets/Packets</b>	Number of octets and packets transmitted on the VCP.	<b>brief, member, none</b>
<b>master: number</b>	Member ID of the master Routing Engine.	All levels
<b>Port</b>	VCP for which <b>RX</b> (Receive) statistics, <b>TX</b> (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared.	<b>detail, extensive</b>
<b>Total octets</b>	Total number of octets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Total packets</b>	Total number of packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Unicast packets</b>	Number of unicast packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Broadcast packets</b>	Number of broadcast packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>Multicast packets</b>	Number of multicast packets received and transmitted on the VCP.	<b>detail, extensive</b>
<b>MAC control frames</b>	Number of media access control (MAC) control frames received and transmitted on the VCP.	<b>detail, extensive</b>

Table 693: show virtual-chassis vc-port statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>CRC alignment errors</b>	<p>Number of packets received on the VCP that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors:</p> <ul style="list-style-type: none"> <li>Invalid FCS with an integral number of octets (FCS error)</li> <li>Invalid FCS with a nonintegral number of octets (alignment error)</li> </ul>	<b>detail, extensive</b>
<b>Oversize packets</b>	Number of packets received on the VCP that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.	<b>detail, extensive</b>
<b>Undersize packets</b>	Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..	<b>detail, extensive</b>
<b>Jabber packets</b>	<p>Number of packets received on the VCP that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error.</p> <p><b>NOTE:</b> This definition of <i>jabber</i> is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define <i>jabber</i> as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.</p>	<b>detail, extensive</b>
<b>Fragments received</b>	<p>Number of packets received on the VCP that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error.</p> <p>Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</p>	<b>detail, extensive</b>
<b>Ifout errors</b>	Number of outbound packets received on the VCP that could not be transmitted because of errors.	<b>detail, extensive</b>
<b>Packet drop events</b>	Number of outbound packets received on the VCP that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.	<b>detail, extensive</b>
<b>64 octets frames</b>	Number of packets received on the VCP (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).	<b>detail, extensive</b>

Table 693: show virtual-chassis vc-port statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
65–127 octets frames	Number of packets received on the VCP (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).	detail, extensive
128–255 octets frames	Number of packets received on the VCP (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).	detail, extensive
256–511 octets frames	Number of packets received on the VCP (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).	detail, extensive
512–1023 octets frames	Number of packets received on the VCP (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).	detail, extensive
1024–1518 octets frames	Number of packets received on the VCP (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets).	detail, extensive
Rate packets per second	Number of packets per second received and transmitted on the VCP.	detail, extensive
Rate bytes per second	Number of bytes per second received and transmitted on the VCP.	detail, extensive

## Sample Output

### show virtual-chassis vc-port statistics

```
user@switch> show virtual-chassis vc-port statistics
fpc0:
```

```
-----
Interface          Input  Octets/Packets      Output  Octets/Packets
internal-0/24       0      / 0                0      / 0
internal-0/25       0      / 0                0      / 0
internal-1/26       0      / 0                0      / 0
internal-1/27       0      / 0                0      / 0
vcp-0               0      / 0                0      / 0
vcp-1               0      / 0                0      / 0
internal-0/26       0      / 0                0      / 0
internal-0/27       0      / 0                0      / 0
internal-1/24       0      / 0                0      / 0
internal-1/25       0      / 0                0      / 0
```

```
{master:0}
```

**show virtual-chassis vc-port statistics (EX8200 Virtual Chassis)**

```
user@external-routing-engine> show virtual-chassis vc-port statistics
```

```
member0:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-4/0/4           43171238 / 48152          47687133 / 51891
vcp-4/0/7           0 / 0                     0 / 0
```

```
member1:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
vcp-3/0/0           0 / 0                     0 / 0
vcp-3/0/1           0 / 0                     0 / 0
vcp-3/0/4           47695376 / 51899          43180556 / 48160
```

```
member8:
```

```
-----
```

```
member9:
```

```
-----
```

**show virtual-chassis vc-port statistics brief**

```
user@switch> show virtual-chassis vc-port statistics brief
```

```
fpc0:
```

```
-----
Interface          Input Octets/Packets      Output Octets/Packets
internal-0/24       0 / 0                     0 / 0
internal-0/25       0 / 0                     0 / 0
internal-1/26       0 / 0                     0 / 0
internal-1/27       0 / 0                     0 / 0
vcp-0               0 / 0                     0 / 0
vcp-1               0 / 0                     0 / 0
internal-0/26       0 / 0                     0 / 0
internal-0/27       0 / 0                     0 / 0
internal-1/24       0 / 0                     0 / 0
internal-1/25       0 / 0                     0 / 0
```

```
{master:0}
```

**show virtual-chassis vc-port statistics extensive**

```
user@switch> show virtual-chassis vc-port statistics extensive
```

```
fpc0:
```

```
-----
RX TX
Port: internal-0/24
Total octets: 0 0
Total packets: 0 0
Unicast packets: 0 0
Broadcast packets: 0 0
Multicast packets: 0 0
MAC control frames: 0 0
CRC alignment errors: 0
Oversize packets: 0
Undersize packets: 0
Jabber packets: 0
Fragments received: 0
```

```

Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

...

Port: vcp-0
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

Port: vcp-1
Total octets:          0          0
Total packets:         0          0
Unicast packets:       0          0
Broadcast packets:     0          0
Multicast packets:     0          0
MAC control frames:    0          0
CRC alignment errors:  0
Oversize packets:     0
Undersize packets:     0
Jabber packets:        0
Fragments received:    0
Ifout errors:          0
Packet drop events:    0
64      octets frames: 0
65-127   octets frames: 0
128-255  octets frames: 0
256-511  octets frames: 0
512-1023 octets frames: 0
1024-1518 octets frames: 0
Rate packets per second: 0          0
Rate bytes per second:   0          0

...

```

```
{master:0}
```

#### show virtual-chassis vc-port statistics member 0

```
user@switch>show virtual-chassis vc-port statistics member 0  
fpc0:
```

```
-----  
Interface           Input  Octets/Packets      Output  Octets/Packets  
internal-0/24        0           / 0             0           / 0  
internal-0/25        0           / 0             0           / 0  
internal-1/26        0           / 0             0           / 0  
internal-1/27        0           / 0             0           / 0  
vcp-0                0           / 0             0           / 0  
vcp-1                0           / 0             0           / 0  
internal-0/26        0           / 0             0           / 0  
internal-0/27        0           / 0             0           / 0  
internal-1/24        0           / 0             0           / 0  
internal-1/25        0           / 0             0           / 0
```

```
{master:0}
```



# Troubleshooting Procedures

- [Troubleshooting Virtual Chassis Fabric on page 7183](#)

## Troubleshooting Virtual Chassis Fabric

---

This topic describes some of the following common troubleshooting issues for a Virtual Chassis Fabric (VCF):

- [Virtual Chassis Port Link Does Not Form on page 7183](#)
- [QFX5100 Leaf Device Assumes Routing Engine Role on page 7184](#)

### Virtual Chassis Port Link Does Not Form

**Problem** **Description:** You connect a 40-Gbps QSFP+ port or a 10-Gbps SFP+ port between a leaf device and a spine device in an autoprovisioned or preprovisioned VCF. You expect the automatic Virtual Chassis port (VCP) conversion feature to convert the link into a VCP link, but the conversion doesn't occur.

The [show virtual-chassis vc-port](#) output indicates that the status of the interface is **Absent** or one or both of interfaces don't appear in the [show virtual-chassis vc-port](#) output.

**Cause** If one end of a link is configured as a VCP and the other is not configured as a VCP, the VCP link does not form.

The automatic VCP conversion feature, therefore, does not work in the following situations:

- a 40-Gbps QSFP+ or 10-Gbps SFP+ interface on one end of the link is already configured as a VCP.

If you have previously removed a device from a VCF but haven't used the **request virtual-chassis vc-port delete** command to convert the interface that was connected to the removed device out of VCP mode, the interface is still configured as a VCP.

If you have removed a device from one Virtual Chassis or VCF and not changed the VCP port setting, the device being added to the VCF might also be configured as a VCP.

- a 40-Gbps QSFP+ port on an EX4300 switch, which is configured as a VCP by default, is interconnecting to a spine device.

**Solution** Manually configure the interface that is not configured as a VCP into a VCP using the **request virtual-chassis vc-port set** command.

### QFX5100 Leaf Device Assumes Routing Engine Role

**Problem** **Description:** A QFX5100 device configured as a leaf device assumes the Routing Engine role during VCF setup. The **show virtual-chassis** output confirms the role.

**Solution** The device can assume the Routing Engine role for several minutes during setup before it receives the configuration from the master Routing Engine, but eventually returns to the linecard role with no user intervention.

**Related Documentation**

- [Virtual Chassis Fabric Overview on page 7033](#)

## PART 24

# Troubleshooting

- [Overview on page 7187](#)
- [Administration on page 7197](#)
- [Troubleshooting on page 7209](#)



## CHAPTER 87

# Overview

- [General Troubleshooting on page 7187](#)
- [Alarms on page 7191](#)

## General Troubleshooting

---

- [Understanding Troubleshooting Resources on page 7187](#)
- [Troubleshooting Overview on page 7189](#)

## Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

[Table 663 on page 6893](#) provides a list of some of the troubleshooting resources.

**Table 694: Troubleshooting Resources on the QFX Series**

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	<a href="#">"Chassis Alarm Messages on a QFX3500 Device" on page 7192</a>
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	<i>Chassis Status LEDs on a QFX3500 Device</i>
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	<a href="#">"Interface Alarm Messages" on page 7195</a>
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	<a href="#">"Understanding Alarms" on page 7191</a>

Table 694: Troubleshooting Resources on the QFX Series (*continued*)

Troubleshooting Resource	Description	Documentation
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6561</a></li> <li>• <a href="#">Junos OS System Log Configuration Statements on page 6616</a></li> </ul>
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the <b>traceroute monitor</b> command to locate points of failure in a network.	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 333</a></li> <li>• <a href="#">Monitoring System Properties on page 334</a></li> <li>• <a href="#">traceroute monitor</a></li> </ul>
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	<i>Junos OS Automation Library</i>
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	<i>Junos XML API Operational Developer Reference</i>
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b> , <b>set</b> , and <b>commit</b> to perform those operations.	<i>NETCONF XML Management Protocol Developer Guide</i>
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	<ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs Support on page 6530</a></li> <li>• <a href="#">SNMP Traps Support on page 6546</a></li> <li>• <a href="#">Using the Traceroute MIB for SNMP Remote Operations</a></li> </ul>
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	<a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a>
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	<i>Service Automation</i>

Table 694: Troubleshooting Resources on the QFX Series (*continued*)

Troubleshooting Resource	Description	Documentation
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	<i>Service Automation</i>
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	<a href="http://kb.juniper.net">http://kb.juniper.net</a>

## Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series product.

[Table 664 on page 6895](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 695: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	See <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 7192</a> .
	Fan tray LED is blinking amber.	See <i>Fan Tray LED on a QFX3500 Device</i> .
	Chassis status LED for the power is blinking amber.	See <i>Chassis Status LEDs on a QFX3500 Device</i> .
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. See <i>Chassis Status LEDs on a QFX3500 Device</i> .

Table 695: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	<p>Check whether the port is a valid Gigabit Ethernet port (6 through 41).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a Fibre Channel port.	<p>Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a port as a 10-Gigabit Ethernet port.	<p>If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.</p> <p>If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
	Cannot configure a 40-Gbps QSFP+ interface.	<p>The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.</p> <p>See <i>QFX3500 Device Overview</i>.</p>
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.
Initial device configuration	Cannot configure management Ethernet ports.	<p>Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.</p> <p><b>NOTE:</b> The management ports are on the front panel of the QFX3500 switch. They are labeled <b>C0</b> and <b>C1</b> on the front panel. In the CLI they are referred to as <b>me0</b> and <b>me1</b>.</p> <p>See “<a href="#">Configuring a QFX3500 Device as a Standalone Switch</a>” on page 175.</p>



Table 695: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Software upgrade and configuration	Failed software upgrade.	See <a href="#">“Recovering from a Failed Software Installation” on page 121.</a>
	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: <ul style="list-style-type: none"> <li>• <a href="#">Loading a Previous Configuration File on page 1252</a></li> <li>• <a href="#">Reverting to the Default Factory Configuration on page 188</a></li> <li>• <a href="#">Reverting to the Rescue Configuration on page 189</a></li> <li>• <a href="#">Performing a Recovery Installation on page 116</a></li> </ul>
	Root password is lost or forgotten.	Recover the root password. See <a href="#">“Recovering the Root Password” on page 1233.</a>
Network interfaces	An aggregated Ethernet interface is down.	See <a href="#">“Troubleshooting an Aggregated Ethernet Interface” on page 1234.</a>
	Interface on built-in network port is down.	See <a href="#">“Troubleshooting Network Interfaces” on page 1234.</a>
	Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See <a href="#">“Troubleshooting Ethernet Switching” on page 1895.</a>
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See <a href="#">“Troubleshooting Firewall Filter Configuration” on page 5411.</a>

## Alarms

- [Understanding Alarms on page 7191](#)
- [Chassis Alarm Messages on a QFX3500 Device on page 7192](#)
- [Interface Alarm Messages on page 7195](#)
- [System Utilization Alarms on page 7195](#)

## Understanding Alarms

The QFX Series support different alarm types and severity levels. [Table 696 on page 7192](#) provides a list of alarm terms and definitions that may help you in monitoring the device.

Table 696: Alarm Terms and Definitions

Term	Definition
Alarm	Signal alerting you to conditions that might prevent normal operation. On the device, alarm indicators might include the LCD panel and LEDs on the device. The LCD panel (if present on the device) displays the chassis alarm message count. Blinking amber LEDs indicate yellow alarm conditions for chassis components.
Alarm condition	Failure event that triggers an alarm.
Alarm severity levels	<p>Seriousness of the alarm. The level of severity can be either major (red) or minor (yellow).</p> <ul style="list-style-type: none"> <li>Major (red)—Indicates a critical situation on the device that has resulted from one of the following conditions. A red alarm condition requires immediate action. <ul style="list-style-type: none"> <li>One or more hardware components have failed.</li> <li>One or more hardware components have exceeded temperature thresholds.</li> <li>An alarm condition configured on an interface has triggered a critical warning.</li> </ul> </li> <li>Minor (yellow or amber)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance. For example, a missing rescue configuration generates a yellow system alarm.</li> </ul>
Alarm types	<p>Alarms include the following types:</p> <ul style="list-style-type: none"> <li>Chassis alarm—Predefined alarm triggered by a physical condition on the device such as a power supply failure or excessive component temperature.</li> <li>Interface alarm—Alarm you configure to alert you when an interface link is down. Applies to <b>ethernet</b>, <b>fibre-channel</b>, and <b>management-ethernet</b> interfaces. You can configure a red (major) or yellow (minor) alarm for the link-down condition, or have the condition ignored.</li> <li>System alarm—Predefined alarm that might be triggered by a missing rescue configuration, failure to install a license for a licensed software feature, or high disk usage.</li> </ul>

#### Related Documentation

- *Chassis Alarm Messages on a QFX3008-I Interconnect Device*
- [Chassis Alarm Messages on a QFX3500 Device on page 7192](#)
- [Interface Alarm Messages on page 7195](#)
- [show chassis alarms on page 495](#)
- [show system alarms on page 974](#)

## Chassis Alarm Messages on a QFX3500 Device

Chassis alarms indicate a failure on the device or one of its components. Chassis alarms are preset and cannot be modified.

The chassis alarm message count is displayed on the LCD panel on the front of the device. To view the chassis alarm message text remotely, use the **show chassis lcd** CLI command.

Chassis alarms on QFX3500 devices have two severity levels:

- Major (red)—Indicates a critical situation on the device that has resulted from one of the conditions described in [Table 697 on page 7193](#). A red alarm condition requires immediate action.

- Minor (yellow or amber)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

Table 697 on page 7193 describes the chassis alarm messages on QFX3500 devices.

**Table 697: QFX3500 Chassis Alarm Messages**

Component	Alarm Type	CLI Message	Recommended Action
Fans	Major (red)	Fan/Blower Absent	The fan is missing. Install a fan.
		Fan Failure	Replace the fan and report the failure to customer support.
		Fan I2C Failure	Check the system log for one of the following messages and report the error message to customer support: <ul style="list-style-type: none"> <li>• CM ENV Monitor: Get fan speed failed.</li> <li>• CM ENV Monitor: Get fan speed failed <i>Fan-number</i> is NOT spinning @ correct speed, where <i>fan-number</i> may be 1, 2, or 3.</li> </ul>
		<i>fan-number</i> Not Spinning Fan	Remove and check the fan for obstructions, and then reinsert the fan. If the problem persists, replace the fan.
Power Supplies	Major (red)	PEM <i>pem-number</i> Airflow not matching Chassis Airflow	The power supply airflow direction is the opposite of the chassis airflow direction. Replace the power supply with a power supply that supports the same airflow direction as the chassis.
		PEM <i>pem-number</i> I2C Failure	Check the system log for one of the following messages and report the error message to customer support: <ul style="list-style-type: none"> <li>• I2C Read failed for device <i>number</i>, where <i>number</i> may be from 123 to 125.</li> <li>• PS <i>number</i>: Transitioning from online to offline, where power supply (PS) <i>number</i> may be 1 or 2.</li> </ul>
		PEM <i>pem-number</i> is not supported	Indicates a power supply problem, or the power supply is not supported on the device. Report the problem to customer support.
		PEM <i>pem-number</i> Not OK	Indicates a problem with the incoming AC or outgoing DC power. Replace the power supply.

Table 697: QFX3500 Chassis Alarm Messages (*continued*)

Component	Alarm Type	CLI Message	Recommended Action
	Minor (yellow)	<b>PEM <i>pem-number</i> Absent</b>	For information only. Indicates the device was powered on with two power supplies installed, but now one is missing. The device can continue to operate with a single power supply. If you wish to remove this alarm message, reboot the device with one power supply.
		<b>PEM <i>pem-number</i> is not powered</b>	For information only. Check the power cord connection and reconnect it if necessary.
		<b>PEM <i>pem-number</i> Power Supply Type Mismatch</b>	For information only. Indicates that an AC power supply and DC power supply have been installed in the same chassis. If you wish to remove this alarm message, reboot the device with two AC power supplies or two DC power supplies.
		<b>PEM <i>pem-number</i> Removed</b>	For information only. Indicates the device was powered on with two power supplies installed, but one has been removed. The device can continue to operate with a single power supply. If you wish to remove this alarm message, reboot the device with one power supply.
Temperature Sensors	Major (red)	<b><i>sensor-location</i> Temp Sensor Fail</b>	Check the system log for the following message and report it to customer support:  <b>Temp sensor <i>sensor-number</i> failed</b> , where <i>sensor-number</i> may range from 1 through 10.
		<b><i>sensor-location</i> Temp Sensor Too Hot</b>	Check environmental conditions and alarms on other devices. Ensure that environmental factors (such as hot air blowing around the equipment) are not affecting the temperature sensor. If the condition persists, the device may shut down.
	Minor (yellow)	<b><i>sensor-location</i> Temp Sensor Too Warm</b>	For information only. Check environmental conditions and alarms on other devices. Ensure that environmental factors (such as hot air blowing around the equipment) are not affecting the temperature sensor.

**Related Documentation**

- [Front Panel of a QFX3500 Device](#)

- [Configuring the Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types on page 161](#)
- [alarm on page 2630](#)

## Interface Alarm Messages

Interface alarms are alarms that you configure to alert you when an interface is down.

To configure an interface link-down condition to trigger a red or yellow alarm, or to configure the link-down condition to be ignored, use the **alarm** statement at the **[edit chassis]** hierarchy level. You can specify the **ethernet**, **fibre-channel**, or **management-ethernet** interface type.



**NOTE:** Fibre Channel alarms are only valid on QFX3500 devices.



**NOTE:** When red alarms or major alarms are issued on QFX5100 and EX4600 switches, the alarm LED glows amber instead of red.

By default, major alarms are configured for interface link-down conditions on the control plane and management network interfaces in a QFabric system. The link-down alarms indicate that connectivity to the control plane network is down. You can configure these alarms to be ignored using the **alarm** statement at the **[edit chassis]** hierarchy level.



**NOTE:** If you configure a yellow alarm on the QFX3008-I Interconnect device, it will be handled as a red alarm.

### Related Documentation

- [Understanding Alarms on page 7191](#)

## System Utilization Alarms

QFX Series devices provide system alarms that alert you when disk usage in the **/var** partition exceeds acceptable levels.

You can display the messages for these alarms by issuing the **show system alarms** operational mode command if the **/var** partition usage exceeds 75 percent. A usage level between 76 and 90 percent indicates high usage and raises a minor alarm condition, whereas a usage level above 90 percent indicates that the partition is full and raises a major alarm condition.

The following sample output from the **show system alarms** command shows system alarm messages that are displayed when disk usage is exceeded on the switch.

```
user@host> show system alarms
4 alarms currently active
Alarm time          Class  Description
```

```
2013-10-08 20:08:20 UTC Minor RE 0 /var partition usage is high
2013-10-08 20:08:20 UTC Major RE 0 /var partition is full
2013-10-08 20:08:08 UTC Minor FPC 1 /var partition usage is high
2013-10-08 20:08:08 UTC Major FPC 1 /var partition is full
```



**BEST PRACTICE:** We recommend that you regularly request a system file storage cleanup to optimize the performance of the switch and prevent generating system alarms.

**Related  
Documentation**

- [Cleaning Up the System File Storage Space on page 7211](#)
- [Understanding Alarms on page 7191](#)
- [show system alarms on page 974](#)

## CHAPTER 88

# Administration

- [Routine Monitoring Using the CLI on page 7197](#)

### Routine Monitoring Using the CLI

---

- [Monitoring SNMP on page 7197](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 7199](#)
- [Monitoring RMON MIB Tables on page 7202](#)
- [Displaying a Log File from a Single-Chassis System on page 7203](#)
- [Monitoring System Log Messages on page 7204](#)
- [Monitoring Traffic Through the Router or Switch on page 7205](#)
- [Pinging Hosts on page 7207](#)

### Monitoring SNMP

There are several commands that you can access in Junos OS operational mode to monitor SNMP information. Some of the commands are:

- **show snmp health-monitor**, which displays the health monitor log and alarm information.
- **show snmp mib**, which displays information from the MIBs, such as device and system information.
- **show snmp statistics**, which displays SNMP statistics such as the number of packets, silent drops, and invalid output values.
- **show snmp rmon**, which displays the RMON alarm, event, history, and log information

The following example provides sample output from the **show snmp health-monitor** command:

```
user@switch> show snmp health-monitor
```

```
Alarm
```

Index	Variable description	Value	State
32768	Health Monitor: root file system utilization jnxHrStoragePercentUsed.1	58	active
32769	Health Monitor: /config file system utilization jnxHrStoragePercentUsed.2	0	active

```
32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0            35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon                          0 active
      Chassis daemon                      50 active
      Firewall daemon                     0 active
      Interface daemon                    5 active
      SNMP daemon                         11 active
      MIB2 daemon                         42 active
      ...
```

The following example provides sample output from the **show snmp mib** command:

```
user@switch> show snmp mib walk system
```

```
sysDescr.0    = Juniper Networks, Inc. qfx3500s internet router, kernel
JUNOS 11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC builder@abc.juniper.net:
/volume/build/junos/11.1/production/20100926.0/obj-xlr/bsd/sys/compile/JUNIPER-xxxxx
```

```
Build date: 2010-09-26 06:00:10 U
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0   = 24444184
sysContact.0  = J Smith
sysName.0     = Lab QFX3500
sysLocation.0 = Lab
sysServices.0 = 4
```

The following example provides sample output from the **show snmp statistics** command:

```
user@switch> show snmp statistics
```

```
SNMP statistics:
Input:
  Packets: 0, Bad versions: 0, Bad community names: 0,
  Bad community uses: 0, ASN parse errors: 0,
  Too bigs: 0, No such names: 0, Bad values: 0,
  Read onlys: 0, General errors: 0,
  Total request varbinds: 0, Total set varbinds: 0,
  Get requests: 0, Get nexts: 0, Set requests: 0,
  Get responses: 0, Traps: 0,
  Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
  Throttle drops: 0, Duplicate request drops: 0
Output:
  Packets: 0, Too bigs: 0, No such names: 0,
  Bad values: 0, General errors: 0,
  Get requests: 0, Get nexts: 0, Set requests: 0,
  Get responses: 0, Traps: 0
```

- Related Documentation
- [health-monitor on page 1416](#)
  - [show snmp mib on page 6874](#)
  - [show snmp statistics on page 1503](#)



## Tracing SNMP Activity on a Device Running Junos OS

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
  - chassisd
  - craftd
  - ilmid
  - mib2d
  - rmopd
  - serviced
  - snmpd
- When a trace file named **filename** reaches its maximum size, it is renamed **filename.0**, then **filename.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the *Junos OS System Log Messages Reference*.)
- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
traceoptions {
  file <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  memory-trace;
  no-remote-trace;
  no-default-memory-trace;
}
```

These statements are described in the following sections:

- [Configuring the Number and Size of SNMP Log Files on page 7200](#)
- [Configuring Access to the Log File on page 7200](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 7200](#)
- [Configuring the Trace Operations on page 7200](#)

### Configuring the Number and Size of SNMP Log Files

---

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed *filename.0*, then *filename.1*, and so on, until there are three trace files. Then the oldest trace file (*filename.2*) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

### Configuring Access to the Log File

---

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file no-world-readable;
```

### Configuring a Regular Expression for Lines to Be Logged

---

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file filename]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

### Configuring the Trace Operations

---

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
  all;
```

```

configuration;
database;
events;
general;
interface-stats;
nonvolatile-sets;
pdu;
policy;
protocol-timeouts;
routing-socket;
server;
subagent;
timer;
varbind-error;
}

```

Table 641 on page 6809 describes the meaning of the SNMP tracing flags.

**Table 698: SNMP Tracing Flags**

Flag	Description	Default Setting
<b>all</b>	Log all operations.	Off
<b>configuration</b>	Log reading of the configuration at the <b>[edit snmp]</b> hierarchy level.	Off
<b>database</b>	Log events involving storage and retrieval in the events database.	Off
<b>events</b>	Log important events.	Off
<b>general</b>	Log general events.	Off
<b>interface-stats</b>	Log physical and logical interface statistics.	Off
<b>nonvolatile-set</b>	Log nonvolatile SNMP set request handling.	Off
<b>pdu</b>	Log SNMP request and response packets.	Off
<b>policy</b>	Log policy processing.	Off
<b>protocol-timeouts</b>	Log SNMP response timeouts.	Off
<b>routing-socket</b>	Log routing socket calls.	Off
<b>server</b>	Log communication with processes that are generating events.	Off
<b>subagent</b>	Log subagent restarts.	Off
<b>timer</b>	Log internal timer events.	Off

Table 698: SNMP Tracing Flags (*continued*)

Flag	Description	Default Setting
<b>varbind-error</b>	Log variable binding errors.	Off

To display the end of the log for an agent, issue the **show log agentd | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where **agent** is the name of an SNMP agent.

- Related Documentation**
- [Configuring SNMP on a Device Running Junos OS](#)
  - [Configuration Statements at the \[edit snmp\] Hierarchy Level](#)
  - [Example: Tracing SNMP Activity](#)
  - [Configuring SNMP on page 1356](#)

## Monitoring RMON MIB Tables

**Purpose** Monitor remote monitoring (RMON) alarm, event, and log tables.

**Action** To display the RMON tables:

```
user@switch> show snmp rmon
Alarm
Index  Variable description                               Value State

      5 monitor
      jnxOperatingCPU.9.1.0.0                        5 falling threshold

Event
Index  Type                               Last Event
      1 log and trap                     2010-07-10 11:34:17 PDT
Event Index: 1
      Description: Event 1 triggered by Alarm 5, rising threshold (90) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 100)
      Time: 2010-07-10 11:34:07 PDT
      Description: Event 1 triggered by Alarm 5, falling threshold (75) crossed,
(variable: jnxOperatingCPU.9.1.0.0, value: 5)
      Time: 2010-07-10 11:34:17 PDT
```

**Meaning** The display shows that an alarm has been defined to monitor jnxRmon MIB object jnxOperatingCPU, which represents the CPU utilization of the Routing Engine. The alarm is configured to generate an event that sends an SNMP trap and adds an entry to the logTable in the RMON MIB. The log table shows that two occurrences of the event have been generated—one for rising above a threshold of 90 percent, and one for falling below a threshold of 75 percent.

- Related Documentation**
- [Configuring RMON Alarms and Events on page 6606](#)
  - [show snmp rmon on page 6877](#)

- [show snmp rmon history on page 6881](#)
- [clear snmp statistics on page 6857](#)
- [clear snmp history on page 6856](#)

## Displaying a Log File from a Single-Chassis System

To display a log file stored on a single-chassis system, enter Junos OS CLI operational mode and issue the following commands:

```
user@switch> show log log-filename
user@switch> file show log-file-pathname
```

By default, the commands display the file stored on the local Routing Engine.

The following example shows the output from the **show log messages** command:

```
user@switch1> show log messages
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Management
process): new instance detected (variable: sysAppElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon memory usage (Command-line
interface): new instance detected (variable: sysAppElmtRunMemory.5.8.2292)
...
Nov  4 12:08:30 switch1 rpdf[957]: task_connect: task BGP_100.10.10.1.6+179 addr
10.10.1.6+179: Can't assign requested
address
Nov  4 12:08:30 switch1 rpdf[957]: bgp_connect_start: connect 10.10.1.6 (Internal
AS 100): Can't assign requested address
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

The following example shows the output from the **file show** command. The file in the pathname **/var/log/processes** has been previously configured to include messages from the daemon facility.

```
user@switch1> file show /var/log/processes
Feb 22 08:58:24 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 22 20:35:07 switch1 snmpd[359]: SNMPD_THROTTLE_QUEUE_DRAINED:
trap_throttle_timer_handler: cleared all throttled traps
Feb 23 07:34:56 switch1 snmpd[359]: SNMPD_TRAP_WARM_START: trap_generate_warm:
SNMP trap: warm start
Feb 23 07:38:19 switch1 snmpd[359]: SNMPD_TRAP_COLD_START: trap_generate_cold:
```

```
SNMP trap: cold start
...
```

- Related Documentation**
- [Interpreting Messages Generated in Standard Format on page 6628](#)
  - [Interpreting Messages Generated in Structured-Data Format on page 6625](#)

## Monitoring System Log Messages

**Purpose** Display system log messages about the QFX Series. By looking through a system log file for any entries pertaining to the interface that you are interested in, you can further investigate a problem with an interface on the switch.

**Action** To view system log messages:

```
user@switch1> show log messages
```

## Sample Output

```
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:01 switch1 newsyslog[2283]: logfile turned over due to size>128K
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 1
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 2
Nov  4 11:30:06 switch1 chassism[952]: CM ENV Monitor: set fan speed is 65 percent
for Fan 3
...
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Management process): new instance detected (variable:
sysApp1ElmtRunMemory.5.6.2293)
Nov  4 11:52:53 switch1 snmpd[944]: SNMPD_HEALTH_MON_INSTANCE: Health Monitor:
jroute daemon
memory usage (Command-line interface): new instance detected (variable:
sysApp1ElmtRunMemory.5.8.2292)
...
Nov  4 12:10:24 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'exit '
Nov  4 12:10:27 switch1 mgd[2293]: UI_DBASE_LOGOUT_EVENT: User 'jsmith' exiting
configuration mode
Nov  4 12:10:31 switch1 mgd[2293]: UI_CMDLINE_READ_LINE: User 'jsmith', command
'show log messages'
```

**Meaning** The sample output shows the following entries in the **messages** file:

- A new log file was created when the previous file reached the maximum size of 128 kilobytes (KB).
- The fan speed for Fan 1, 2, and 3 is set at 65 percent.
- Health monitoring activity is detected.
- CLI commands were entered by the user jsmith.

- Related Documentation**
- [Overview of Junos OS System Log Messages on page 6560](#)
  - [Understanding the Implementation of System Log Messages on the QFabric System on page 6562](#)
  - [Example: Configuring System Log Messages on page 6568](#)
  - [clear log on page 350](#)
  - [show log on page 948](#)
  - [syslog on page 313](#)

## Monitoring Traffic Through the Router or Switch

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. [Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 7205](#)
2. [Displaying Real-Time Statistics About an Interface on the Router or Switch on page 7206](#)

### Displaying Real-Time Statistics About All Interfaces on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through all interfaces on the router or switch.

**Action** To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

## Sample Output

```
user@host> monitor interface traffic
host name                               Seconds: 15                               Time: 12:31:09
Interface  Link  Input packets  (pps)  Output packets  (pps)
so-1/0/0   Down    0              (0)      0              (0)
so-1/1/0   Down    0              (0)      0              (0)
so-1/1/1   Down    0              (0)      0              (0)
so-1/1/2   Down    0              (0)      0              (0)
so-1/1/3   Down    0              (0)      0              (0)
t3-1/2/0   Down    0              (0)      0              (0)
t3-1/2/1   Down    0              (0)      0              (0)
t3-1/2/2   Down    0              (0)      0              (0)
t3-1/2/3   Down    0              (0)      0              (0)
so-2/0/0   Up      211035         (1)      36778          (0)
so-2/0/1   Up      192753         (1)      36782          (0)
so-2/0/2   Up      211020         (1)      36779          (0)
so-2/0/3   Up      211029         (1)      36776          (0)
so-2/1/0   Up      189378         (1)      36349          (0)
so-2/1/1   Down    0              (0)      18747          (0)
so-2/1/2   Down    0              (0)      16078          (0)
so-2/1/3   Up      0              (0)      80338          (0)
at-2/3/0   Up      0              (0)      0              (0)
```

```
at-2/3/1    Down          0          (0)          0          (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

**Meaning** The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the **C** key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

### Displaying Real-Time Statistics About an Interface on the Router or Switch

**Purpose** Display real-time statistics about traffic passing through an interface on the router or switch.

**Action** To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

```
user@host> monitor interface interface-name
```

### Sample Output

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
  Input bytes:          5856541 (88 bps)
  Output bytes:         6271468 (96 bps)
  Input packets:        157629 (0 pps)
  Output packets:       157024 (0 pps)
Encapsulation statistics:
  Input keepalives:     42353
  Output keepalives:    42320
  LCP state: Opened
Error statistics:
  Input errors:         0
  Input drops:          0
  Input framing errors: 0
  Input runs:           0
  Input giants:         0
  Policed discards:     0
  L3 incompletes:       0
  L2 channel errors:    0
  L2 mismatch timeouts: 0
  Carrier transitions:  1
  Output errors:        0
  Output drops:         0
  Aged packets:         0
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count             1
  LOF count             1
  SEF count             1
  ES-S                  77
  SES-S                 77
SONET statistics:
  BIP-B1                0
  BIP-B2                0
```



```

REI-L                0
BIP-B3               0
REI-P                0
Received SONET overhead:  F1          : 0x00  J0          : 0xZ

```

**Meaning** The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in [Table 640 on page 6802](#).

**Table 699: Output Control Keys for the monitor interface Command**

Action	Key
Display information about the next interface. The <b>monitor interface</b> command scrolls through the physical or logical interfaces in the same order that they are displayed by the <b>show interfaces terse</b> command.	<b>N</b>
Display information about a different interface. The command prompts you for the name of a specific interface.	<b>I</b>
Freeze the display, halting the display of updated statistics.	<b>F</b>
Thaw the display, resuming the display of updated statistics.	<b>T</b>
Clear (zero) the current delta counters since <b>monitor interface</b> was started. It does not clear the accumulative counter.	<b>C</b>
Stop the <b>monitor interface</b> command.	<b>Q</b>

See the [CLI Explorer](#) for details on using match conditions with the **monitor traffic** command.

## Pinging Hosts

**Purpose** Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The switch sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to a specified host and receives ICMP echo responses.

**Action** To use the **ping** command to send four requests (ping count) to host3:  
**ping host count number**

## Sample Output

```

ping host3 count 4
user@switch> ping host3 count 4
PING host3.site.net (176.26.232.111): 56 data bytes
64 bytes from 176.26.232.111: icmp_seq=0 ttl=122 time=0.661 ms
64 bytes from 176.26.232.111: icmp_seq=1 ttl=122 time=0.619 ms

```

```
64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms
64 bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms

--- host3.site.net ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

**Meaning** • The **ping** results show the following information:

- Size of the ping response packet (in bytes).
- IP address of the host from which the response was sent.
- Sequence number of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
- Time-to-live (ttl) hop-count value of the ping response packet.
- Total time between the sending of the ping request packet and the receiving of the ping response packet, in milliseconds. This value is also called round-trip time.
- Number of ping requests (probes) sent to the host.
- Number of ping responses received from the host.
- Packet loss percentage.
- Round-trip time statistics: minimum, average, maximum, and standard deviation of the round-trip time.

**Related Documentation** • [Troubleshooting Overview on page 6895](#)  
• [Understanding Troubleshooting Resources on page 6893](#)

# Troubleshooting

- [Configuration and File Management on page 7209](#)
- [Ethernet Switching on page 7212](#)
- [Hardware on page 7217](#)
- [High Availability on page 7219](#)
- [Interfaces on page 7220](#)
- [Junos OS Basics on page 7226](#)
- [Layer 3 Protocols on page 7241](#)
- [MPLS on page 7242](#)
- [Network Management on page 7242](#)
- [Security on page 7253](#)
- [Services on page 7263](#)
- [Traffic Management on page 7266](#)
- [Virtual Chassis Fabric on page 7275](#)

## Configuration and File Management

---

- [Loading a Previous Configuration File on page 7209](#)
- [Reverting to the Default Factory Configuration on page 7210](#)
- [Reverting to the Rescue Configuration on page 7211](#)
- [Cleaning Up the System File Storage Space on page 7211](#)

### Loading a Previous Configuration File

You can use the **rollback** *<number>* command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

#### Syntax

**rollback** *<number>*

#### Options

- **none**— Return to the most recently saved configuration.
- **number**—Configuration to return to.
  - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
  - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

**Related  
Documentation**

- [Configuration File Terms on page 11](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

**Related  
Documentation**

- [Understanding Configuration Files on page 1242](#)
- [Loading a Previous Configuration File on page 1252](#)
- [Reverting to the Rescue Configuration on page 189](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```

2. Commit your changes.

```
[edit]
user@switch# commit filename
```

### Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1261](#)
- [Reverting to the Default Factory Configuration on page 188](#)
- [Configuration File Terms on page 11](#)

## Cleaning Up the System File Storage Space

**Problem**    **Description:** The system file storage space on the switch is full. Rebooting the switch does not solve the problem.

The following error message is displayed during a typical operation on the switch after the file storage space is full.

```
user@switch% cli
user@switch> configure
/var: write failed, filesystem is full
```

**Solution**    Clean up the file storage on the switch by deleting system files.

1. Request to delete system files on the switch.

```
user@switch> request system storage cleanup
```

The list of files to be deleted is displayed.

List of files to delete:

	Size	Date	Name
	11B	Jul 26 20:55	/var/jail/tmp/alarmd.ts
	124B	Aug 4 18:05	/var/log/default-log-messages.0.gz
	1301B	Jul 26 20:42	/var/log/install.0.gz
	387B	Jun 3 14:37	/var/log/install.1.gz
	4920B	Aug 4 18:05	/var/log/messages.0.gz
	20.0K	Jul 26 21:00	/var/log/messages.1.gz
	16.3K	Jun 25 13:45	/var/log/messages.2.gz
	804B	Aug 4 18:05	/var/log/security.0.gz
	16.8K	Aug 3 11:15	/var/log/security.1.gz

```
487B Aug  4 18:04 /var/log/wtmp.0.gz
855B Jul 29 22:54 /var/log/wtmp.1.gz
920B Jun 30 16:32 /var/log/wtmp.2.gz
94B Jun  3 14:36 /var/log/wtmp.3.gz
353.2K Jun  3 14:37 /var/sw/pkg/jloader-qfx-11.2I20110303_1117_dc-builder.tgz

124.0K Jun  3 14:30 /var/tmp/gres-tp/env.dat
0B Apr 14 16:20 /var/tmp/gres-tp/lock
0B Apr 14 17:37 /var/tmp/if-rtbdb/env.lock
12.0K Jul 26 20:55 /var/tmp/if-rtbdb/env.mem
2688.0K Jul 26 20:55 /var/tmp/if-rtbdb/shm_usr1.mem
132.0K Jul 26 20:55 /var/tmp/if-rtbdb/shm_usr2.mem
2048.0K Jul 26 20:55 /var/tmp/if-rtbdb/trace.mem
155B Jul 26 20:55 /var/tmp/krt_gencfg_filter.txt
0B Jul 26 20:55 /var/tmp/rtbdb/if-rtbdb
1400.6K Aug  3 10:13 /var/tmp/sfid.core.0.gz
1398.9K Aug  3 17:01 /var/tmp/sfid.core.1.gz
Delete these files ? [yes,no] (no)
```

2. Enter **yes** to delete the files.

3. Reboot the switch.



**BEST PRACTICE:** We recommend that you regularly request a system file storage cleanup to optimize the performance of the switch.

---

Related Documentation • [request system storage cleanup on page 466](#)

## Ethernet Switching

---

- [Troubleshooting Ethernet Switching on page 7212](#)
- [Troubleshooting Layer 2 Protocol Tunneling on page 7213](#)
- [Troubleshooting Private VLANs on page 7214](#)
- [Troubleshooting Q-in-Q and VLAN Translation Configuration on page 7217](#)

## Troubleshooting Ethernet Switching

**Problem**    **Description:** Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table.

Sometimes silent devices, such as syslog servers or SNMP trap receivers that receive UDP traffic but do not return acknowledgment (ACK) messages to the traffic source, fail to send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by

issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

**Solution** Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the ARP aging timer is from 1 through 240 minutes.)

```
[edit system arp]
user@switch# set aging-timer 3
```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 15 to 1,000,000 seconds.)

```
[edit vlans]
user@switch# set vlans sales mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP message, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table.

**Related Documentation**

- [arp](#)
- [mac-table-aging-time on page 1794](#)

## Troubleshooting Layer 2 Protocol Tunneling

- [Drop Threshold Statistics Might Be Incorrect on page 7213](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 7213](#)

### Drop Threshold Statistics Might Be Incorrect

**Problem** **Description:** L2PT processing is done by the CPU, and L2PT traffic to the CPU is rate limited to a maximum of 1000 pps. If traffic is received at a rate faster than this limit, the rate limit causes the traffic to be dropped before it hits the threshold and the dropped packets will not be reported in L2PT statistics. This can also occur if you configure a drop threshold that is less than 1000 pps but traffic is received at a faster rate. For example, if you configure a drop threshold of 900 pps and the VLAN receives traffic at rate of 1100 pps, L2PT statistics will show that 100 packets were dropped. The 100 packets dropped because of the rate limit are not reported. Similarly, if you do not configure a drop threshold and the VLAN receives traffic at rate of 1100 pps, the 100 packets dropped because of the rate limit are not reported.

**Solution** This is expected behavior.

### Egress Filtering of L2PT Traffic Not Supported

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a

firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

**Related Documentation**

- [Understanding Layer 2 Protocol Tunneling](#)
- [Configuring Layer 2 Protocol Tunneling](#)

## Troubleshooting Private VLANs

Use the following information to troubleshoot a private VLAN configuration.

- [Limitations of Private VLANs on page 7214](#)
- [Forwarding with Private VLANs on page 7214](#)
- [Egress Firewall Filters with Private VLANs on page 7215](#)
- [Egress Port Mirroring with Private VLANs on page 7216](#)

### Limitations of Private VLANs

---

The following constraints apply to private VLAN configurations:

- IGMP snooping is not supported with private VLANs.
- Routed VLAN interfaces are not supported on private VLANs
- Routing between secondary VLANs in the same primary VLAN is not supported.

### Forwarding with Private VLANs

---

**Problem Description:**

- When isolated VLAN or community VLAN tagged traffic is received on a PVLAN trunk port, MAC addresses are learned from the primary VLAN. This means that output from the [show ethernet-switching table](#) command shows that MAC addresses are learned from the primary VLAN and replicated to secondary VLANs. This behavior has no effect on forwarding decisions.
- If a packet with a secondary VLAN tag is received on a promiscuous port, it is accepted and forwarded.
- If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped.
  - The packet has a community VLAN tag.
  - The packet is destined to a unicast MAC address or multicast group MAC address that was learned on an isolated VLAN.
- If a packet is received on a PVLAN trunk port and meets both of the conditions listed below, it is dropped.



- The packet has an isolated VLAN tag.
- The packet is destined to a unicast MAC address or multicast group MAC address that was learned on a community VLAN.
- If a packet with a primary VLAN tag is received by a secondary (isolated or community) VLAN port, the secondary port forwards the packet.
- If you configure a community VLAN on one device and configure another community VLAN on a second device and both community VLANs use the same VLAN ID, traffic for one of the VLANs can be forwarded to the other VLAN. For example, assume the following configuration:
  - Community VLAN comm1 on switch 1 has VLAN ID 50 and is a member of primary VLAN pvlan100.
  - Community VLAN comm2 on switch 2 also has VLAN ID 50 and is a member of primary VLAN pvlan200.
  - Primary VLAN pvlan100 exists on both switches.

If traffic for comm1 is sent from switch 1 to switch 2, it will be sent to the ports participating in comm2. (The traffic will also be forwarded to the ports in comm1, as you would expect.)

**Solution** These are expected behaviors.

### Egress Firewall Filters with Private VLANs

**Problem Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port

- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).
- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

#### Egress Port Mirroring with Private VLANs

---

**Problem** **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

- Related Documentation**
- [Understanding Private VLANs](#)
  - [Understanding Secondary VLAN Trunk Ports and Promiscuous Access Ports on PVLANS](#)
  - [Creating a Private VLAN on a Single Switch](#)
  - [Creating a Private VLAN Spanning Multiple Switches](#)
  - [Example: Configuring PVLANS with Secondary VLAN Trunk Ports and Promiscuous Access Ports](#)

## Troubleshooting Q-in-Q and VLAN Translation Configuration

- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 7217](#)
- [Egress Port Mirroring with VLAN Translation on page 7217](#)

### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

**Problem** **Description:** If you create a firewall filter that includes a match condition of **dot1q-tag** or **dot1q-user-priority** and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the **set dot1q-tunneling ethertype 0x8100** statement at the **[edit ethernet-switching-options]** hierarchy level. You must also configure the other end of the link to use the same EtherType.

### Egress Port Mirroring with VLAN Translation

**Problem** **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

- Related Documentation**
- [Understanding Q-in-Q Tunneling and VLAN Translation](#)
  - [Example: Setting Up Q-in-Q Tunneling](#)

## Hardware

- [Troubleshooting QFX3100 Director Device Isolation on page 7218](#)

## Troubleshooting QFX3100 Director Device Isolation

**Problem**    **Description:** Both connections between the QFX3100 Director devices are broken so that one of the Director devices in a Director group becomes isolated from the group.

The redundant patch cables interconnecting the Director devices are critical links required for the operation of the Director group. The two inter-Director device links must remain connected when the Director devices are online. After the Director devices are installed and the Director group is active, if a single inter-Director device link loses and regains its connection, the operation of the Director group remains intact. However, the loss of both inter-Director device links causes one Director device to isolate itself from the Director group.



**WARNING:** Do not reconnect the inter-Director patch cables before properly restarting the isolated Director device. Restarting the active Director device instead of the isolated Director device can result in both Director devices rebooting, with a subsequent data loss.

**Environment:** This problem occurs between the two QFX3100 Director devices found in QFabric systems.

**Symptoms:** Symptoms of this problem include an unscheduled rebooting of one of the Director devices.

---

### Resolution    *Determine Which Director Device Is Isolated*

Before restoring the inter-Director device links, determine which one of the Director devices is in isolation.

To locate an isolated Director device, use one of the following methods:

- Review logs or management tools for standard SNMP traps issued from the Director group before the Director device became isolated.
  - If eth-2/6 links are down, the Director group cannot communicate. Normally, one of the devices reboots.
  - If both eth-2/6 and eth-7/8/9 links are down, the Director device is isolated from the control plane and is not providing fabric services.
  - Issue **show fabric session-host**.
- Use the CLI to determine the serial numbers of the active Director device.
  - Issue the **show fabric session-host** command.

```
root@qfabric>show fabric session-host
Identifier: 0281042010000013
```

- Issue the **show fabric administration inventory director-group status | grep “dg0|dg1”** command.

```
root@qfabrid> show fabric administration inventory director-group status | grep
“dg0|dg1”
```

```
dg0 online master 10.94.214.80 0% 13597976k 4 4 days, 22:36 hrs
dg1 online master 10.94.214.81 0% 18677380k 3 4 days, 22:25 hrs
dg0 0281042010000013 online master
dg1 0281042010000018 online backup
```

When the Director devices cannot communicate, the **show fabric administration inventory director-group** command only displays the Director device that is online.

### *Power Off the Isolated Director Device and Restore the Inter-Director Device Links*



**CAUTION:** Be sure you know which Director device is active and which is isolated. If you power off the active Director device, both Director devices reboot and cause potential data loss on the system.

To restore communication within the Director group:

1. Power off the isolated Director device.
2. Restore the inter-Director device links (port 3 to port 3) by firmly inserting the redundant patch cables.
3. Power on the previously isolated Director device. The Director device reboots.

#### Related Documentation

- [Connecting QFX3100 Director Devices in a Director Group](#)

## High Availability

- [Troubleshooting VRRP on page 7219](#)

## Troubleshooting VRRP

**Problem** **Description:** If you configure multiple VRRP groups on an interface (using multiple VLANs), traffic for some of the groups might be briefly dropped if a failover occurs. This can happen because the new master must send gratuitous ARP replies for each VRRP group to update the ARP tables in the connected devices, and there is a short delay between each gratuitous ARP reply. Traffic sent by devices that have not yet received the gratuitous ARP reply is dropped (until the device receives the reply and learns the MAC address of the new master).

**Solution** Configure a failover delay so that the new master delays sending gratuitous ARP replies for the period that you set. This allows the new master to send the ARP replies for all of the VRRP groups simultaneously.

**Related Documentation**

- [failover-delay on page 2324](#)

## Interfaces

---

- [Troubleshooting an Aggregated Ethernet Interface on page 7220](#)
- [Troubleshooting Network Interfaces on page 7220](#)
- [Troubleshooting Multichassis Link Aggregation on page 7221](#)

### Troubleshooting an Aggregated Ethernet Interface

**Problem**     **Description:** The **show interfaces terse** command shows that the LAG is down.

**Solution**     Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related Documentation**

- [Verifying the Status of a LAG Interface on page 2750](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)

### Troubleshooting Network Interfaces

**The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down**

---

**Problem**     **Description:** The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command **show interfaces interface-name**, the disabled port is not listed.

**Cause**     By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution**     Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

## Troubleshooting Multichassis Link Aggregation

Use the following information to troubleshoot multichassis link aggregation configuration.

- [MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table on page 7221](#)
- [MC-LAG Peer Does Not Go into Standby Mode on page 7222](#)
- [Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive on page 7222](#)
- [Redirect Filters Take Priority over User-Defined Filters on page 7222](#)
- [Operational Command Output Is Wrong on page 7223](#)
- [ICCP Connection Might Take Up to 60 Seconds to Become Active on page 7223](#)
- [MAC Address Age Learned on an MC-AE Interface Is Reset to Zero on page 7223](#)
- [MAC Address Is Not Learned Remotely in a Default VLAN on page 7224](#)
- [Snooping Entries Learned on MC-AE Interfaces Are Not Removed on page 7224](#)
- [ICCP Does Not Come Up After You Add or Delete an Authentication Key on page 7224](#)
- [Local Status Is Standby When It Should Be Active on page 7224](#)
- [Packets Loop on the Server When ICCP Fails on page 7224](#)
- [Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change on page 7225](#)
- [No Commit Checks Are Done for ICL-PL Interfaces on page 7225](#)
- [Double Failover Scenario on page 7225](#)
- [Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up on page 7225](#)
- [Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer on page 7226](#)
- [AE Interfaces Go Down on page 7226](#)
- [Flooding of Upstream Traffic on page 7226](#)

### MAC Addresses Learned on MC-AE Interfaces Are Not Removed from the MAC Address Table

**Problem Description:** When both of the multichassis aggregated Ethernet (MC-AE) interfaces on both connected multichassis link aggregation group (MC-LAG) peers are down, the MAC addresses learned on the MC-AE interfaces are not removed from the MAC address table. For example, if you disable the MC-AE interface (ae0) on both MC-LAG peers by issuing the **set interfaces ae0 disable** command and commit the configuration, the MAC table still shows the MAC addresses as being learned on the MC-AE interfaces of both MC-LAG peers:

```
user@switchA> show ethernet-switching table
Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries
VLAN      MAC address      Type      Age Interfaces
v10       *                Flood     - All-members
v10       00:10:94:00:00:01 Learn(L)    3:55 ae0.0 (MCAE)
```

v10	00:10:94:00:00:02	Learn(R)	0 xe-0/0/9.0
v20	*	Flood	- All-members
v30	*	Flood	- All-members
v30	84:18:88:de:b1:2e	Static	- Router

user@switchB> show ethernet-switching table

Ethernet-switching table: 6 entries, 2 learned, 0 persistent entries

VLAN	MAC address	Type	Age	Interfaces
v10	*	Flood		- All-members
v10	00:10:94:00:00:01	Learn(R)	0	ae0.0 (MCAE)
v10	00:10:94:00:00:02	Learn	40	xe-0/0/10.0
v20	*	Flood		- All-members
v30	*	Flood		- All-members
v30	84:18:88:df:83:0a	Static		- Router

**Solution** This is expected behavior.

---

#### MC-LAG Peer Does Not Go into Standby Mode

**Problem** **Description:** A multichassis link aggregation group (MC-LAG) peer does not go into standby mode if the MC-LAG peer IP address specified in the Interchassis Control Protocol (ICCP) configuration and the IP address specified in the multichassis protection configuration are different.

**Solution** To prevent failure to enter standby mode, make sure the peer IP address in the ICCP configurations and the IP address in multichassis protection configurations are the same.

---

#### Secondary MC-LAG Peer with Status Control Set to Standby Becomes Inactive

**Problem** **Description:** When the interchassis control link-protection link (ICL-PL) and multichassis aggregated Ethernet (MC-AE) interfaces go down on the primary multichassis link aggregation group (MC-LAG) peer, the secondary MC-LAG peer's MC-AE interfaces with status control set to standby become inactive instead of active.

**Solution** This is expected behavior.

---

#### Redirect Filters Take Priority over User-Defined Filters

**Problem** **Description:** Multichassis link aggregation group (MC-LAG) implicit failover redirection filters take precedence over user-configured explicit filters. This is expected behavior.

**Solution** This is expected behavior.



### Operational Command Output Is Wrong

**Problem Description:** After you deactivate the Interchassis Control Protocol (ICCP), the **show iccp** operational command output still shows registered client daemons, such as mcsnoopd, lacpd, and eswd.

For example:

```
user@switch> show iccp
Client Application: MCSNOOPD
Redundancy Group IDs Joined: None
```

```
Client Application: lacpd
Redundancy Group IDs Joined: 1
```

```
Client Application: eswd
Redundancy Group IDs Joined: 1
```

The **show iccp** command output always shows registered modules regardless of whether or not ICCP peers are configured.

**Solution** This is expected behavior.

### ICCP Connection Might Take Up to 60 Seconds to Become Active

**Problem Description:** When the Interchassis Control Protocol (ICCP) configuration and the routed VLAN interface (RVI) configuration are committed together, the ICCP connection might take up to 60 seconds to become active.

**Solution** This is expected behavior.

### MAC Address Age Learned on an MC-AE Interface Is Reset to Zero

**Problem Description:** When you activate and then deactivate an interchassis control link-protection link (ICL-PL), the MAC address age learned on the multichassis aggregated Ethernet (MC-AE) interface is reset to zero. The next-hop interface changes trigger MAC address updates in the hardware, which then triggers aging updates in the Packet Forwarding Engine (PFE). The result is that the MAC address age is updated to zero.

For example, the ICL-PL has been deactivated, and the **show ethernet-switching table** command output shows that the MAC addresses have an age of 0.

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 2 learned, 0 persistent entries
VLAN      MAC address      Type      Age Interfaces
v100      *                Flood     - All-members
v100      00:10:00:00:00:01 Learn(L)    0 ae0.0 (MCAE)
v100      00:10:00:00:00:02 Learn(L)    0 ae0.0 (MCAE)
```

**Solution** This is expected behavior.

### MAC Address Is Not Learned Remotely in a Default VLAN

**Problem**    **Description:** If a multichassis link aggregation group (MC-LAG) peer learns a MAC address in the default VLAN, the Interchassis Control Protocol (ICCP) does not synchronize the MAC address with the MAC address of the other MC-LAG peer.

**Solution**    This is expected behavior.

### Snooping Entries Learned on MC-AE Interfaces Are Not Removed

**Problem**    **Description:** When multichassis aggregated Ethernet (MC-AE) interfaces are configured on a VLAN that is enabled for multicast snooping, the membership entries learned on the MC-AE interfaces on the VLAN are not cleared when the MC-AE interfaces go down. This is done to speed up convergence time when the interfaces come up, or come up and go down.

**Solution**    This is expected behavior.

### ICCP Does Not Come Up After You Add or Delete an Authentication Key

**Problem**    **Description:** The Interchassis Control Protocol (ICCP) connection is not established when you add an authentication key and then delete it only at the global ICCP level. However, authentication works correctly at the ICCP peer level.

**Solution**    Delete the ICCP configuration , and then add the ICCP configuration.

### Local Status Is Standby When It Should Be Active

**Problem**    **Description:** If the multichassis aggregated Ethernet (MC-AE) interface is down when the state machine is in a synchronized state, the multichassis link aggregation group (MC-LAG) peer local status is standby. If the MC-AE interface goes down after the state machine is in an active state, then the local status remains active, and the local state indicates that the interface is down.

**Solution**    This is expected behavior.

### Packets Loop on the Server When ICCP Fails

**Problem**    **Description:** When you enable backup liveness detection for a multichassis link aggregation group (MC-LAG), and the backup liveness detection packets are lost because of a temporary failure on the MC-LAG, then both of the peers in the MC-LAG remain active. If this happens, both of the MC-LAG peers send packets to the connected server.

**Solution**    This is expected behavior.

### Both MC-LAG Peers Use the Default System ID After a Reboot or an ICCP Configuration Change

---

**Problem**    **Description:** After a reboot or after a new Interchassis Control Protocol (ICCP) configuration has been committed, and the ICCP connection does not become active, the Link Aggregation Control Protocol (LACP) messages transmitted over the multichassis aggregated Ethernet (MC-AE) interfaces use the default system ID. The configured system ID is used instead of the default system ID only after the MC-LAG peers synchronize with each other.

**Solution**    This is expected behavior.

### No Commit Checks Are Done for ICL-PL Interfaces

---

**Problem**    **Description:** There are no commit checks on the interface being configured as an interchassis control link-protection link (ICL-PL), so you must provide a valid interface name for the ICL-PL.

**Solution**    This is expected behavior.

### Double Failover Scenario

---

**Problem**    **Description:** If the following events happen in this exact order—the Interchassis Control Protocol (ICCP) goes down, and the multichassis aggregated Ethernet (MC-AE) interface on the multichassis link aggregation group (MC-LAG) peer in active mode goes down—a double failover occurs. In this scenario, the MC-LAG peer in standby mode does not detect what happens on the active MC-LAG peer. The MC-LAG peer in standby mode operates as if the MC-AE interface on the MC-LAG in active mode were up and blocks the interchassis control protocol-protection link (ICL-PL) traffic. The ICL-PL traffic is not forwarded.

**Solution**    This is expected behavior.

### Multicast Traffic Floods the VLAN When the ICL-PL Interface Goes Down and Up

---

**Problem**    **Description:** When the interchassis control link-protection link (ICL-PL) goes down and up, multicast traffic is flooded to all of the interfaces in the VLAN. The Packet Forwarding Engine (PFE) flag `Ip4McastFloodMode` for the VLAN is changed to `MCAST_FLOOD_ALL`. This problem only occurs when a multichassis link aggregation group (MC-LAG) is configured for Layer 2.

**Solution**    This is expected behavior.

### Layer 3 Traffic Sent to the Standby MC-LAG Peer Is Not Redirected to Active MC-LAG Peer

---

**Problem** **Description:** When the Interchassis Control Protocol (ICCP) is down, the status of a remote MC-LAG peer is unknown. Even if the MC-LAG peer is configured as standby, the traffic is not redirected to this peer because it is assumed that this peer is down.

**Solution** This is expected behavior.

### AE Interfaces Go Down

---

**Problem** **Description:** When a multichassis aggregated Ethernet (MC-AE) interface is converted to an aggregated Ethernet (AE) interface, it retains some MC-AE properties. For example, the AE interface might retain the administrative key of the MC-AE. When this happens, the AE interface goes down.

**Solution** Restart the Link Aggregation Control Protocol (LACP) on the multichassis link aggregation group (MC-LAG) peer hosting the AE interface to bring up the AE interface. Restarting LACP removes the MC-AE properties of the AE interface.

### Flooding of Upstream Traffic

---

**Problem** **Description:** When MAC synchronization is enabled, the multichassis link aggregation group (MC-LAG) peer can resolve Address Resolution Protocol (ARP) entries for the MC-LAG routed VLAN interface (RVI) with either of the MC-LAG peer MAC addresses. If the downstream traffic is sent with one MAC address (MAC1) but the peer has resolved the MAC address with a different MAC address (MAC2), the MAC2 address might not be learned by any of the access layer switches. Flooding of the upstream traffic for the MAC2 address might then occur.

**Solution** Make sure that downstream traffic is sent from the MC-LAG peers periodically to prevent the MAC addresses from aging out.

**Related Documentation**

- [Understanding Multichassis Link Aggregation on page 2411](#)
- [Example: Configuring Multichassis Link Aggregation on page 2471](#)
- [Configuring Multichassis Link Aggregation on page 2597](#)

## Junos OS Basics

---

- [Rebooting and Halting a Device on page 7227](#)
- [Recovering from a Failed Software Installation on page 7228](#)
- [Recovering the Root Password on page 7229](#)
- [Creating an Emergency Boot Device on page 7230](#)

- [Performing a Recovery Installation on page 7232](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 7233](#)
- [Troubleshooting Network Interfaces on page 7240](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 7240](#)

## Rebooting and Halting a Device

To reboot the switch, issue the **request system reboot** command.

```
user@switch> request system reboot ?
Possible completions:
<[Enter]>      Execute this command
all-members    Reboot all virtual chassis members
at             Time at which to perform the operation
both-routing-engines  Reboot both the Routing Engines
fast-boot      Enable fast reboot
in            Number of minutes to delay before operation
local         Reboot local virtual chassis member
member        Reboot specific virtual chassis member (0..9)
message       Message to display to all users
other-routing-engine  Reboot the other Routing Engine
|            Pipe through a command
{master:0}
```

```
user@switch> request system reboot
Reboot the system ? [yes,no] (no) yes
Rebooting switch
```



**NOTE:** Not all options shown in the preceding command output are available on all QFX Series and EX4600 devices. For example, the **fast-boot** option is available only on QFX5100. See the documentation for the [request system reboot](#) command for details about options.

Similarly, to halt the switch, issue the **request system halt** command.



**CAUTION:** Before entering this command, you must have access to the switch's console port in order to bring up the Routing Engine.

```
user@switch> request system halt ?
Possible completions:
<[Enter]>      Execute this command
all-members    Halt all virtual chassis members
at             Time at which to perform the operation
backup-routing-engine  Halt backup Routing Engine
both-routing-engines  Halt both Routing Engines
in            Number of minutes to delay before operation
local         Halt local virtual chassis member
member        Halt specific virtual chassis member (0..9)
message       Message to display to all users
other-routing-engine  Halt other Routing Engine
|            Pipe through a command
```



**NOTE:** When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member, use the `member` option. You cannot issue this command from the QFabric CLI.

Issuing the **request system halt** command on the switch halts the Routing Engine. To reboot a Routing Engine that has been halted, you must connect through the console.

**Related Documentation**

- [clear system reboot on page 355](#)
- [request system reboot on page 415](#)
- [request system halt on page 400](#)
- [request system power-off on page 410](#)
- *Connecting a QFX Series Device to a Management Console*

## Recovering from a Failed Software Installation

**Problem**    **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution**    If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [--format] [--external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.
- **external**—Installs the installation package onto external media (a USB stick, for example).

- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



**NOTE:** The root password cannot be recovered on a QFabric system.



**NOTE:** You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None

- Stop bits: 1
  - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.  
  
The terminal emulation screen on your management device displays the device's boot sequence.
  10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:  
  

```
Hit [Enter] to boot immediately, or space bar for command prompt.  
Booting [kernel] in 9 seconds...
```
  11. At the following prompt, enter **boot -s** to start up the system in single-user mode.  
  

```
ok boot -s
```
  12. At the following prompt, enter **recovery** to start the root password recovery procedure.  
  

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN  
for /bin/sh: recovery
```
  13. Enter configuration mode in the CLI.
  14. Set the root password. For example:  
  

```
user@switch# set system root-authentication plain-text-password
```
  15. At the following prompt, enter the new root password. For example:  
  

```
New password: juniper1  
Retype new password:
```
  16. At the second prompt, reenter the new root password.
  17. After you have finished configuring the password, commit the configuration.  
  

```
root@host# commit  
commit complete
```
  18. Exit configuration mode in the CLI.
  19. Exit operational mode in the CLI.
  20. At the prompt, enter **y** to reboot the device.  
  

```
Reboot the system? [y/n] y
```

**Related Documentation** • [Configuring the Root Password on page 1354](#)

## Creating an Emergency Boot Device

If Junos OS on the device is damaged in some way that prevents the software from loading properly, you can use an emergency boot device to repartition the primary disk and load a fresh installation of Junos OS. Use the following procedure to create an emergency boot device.

Before you begin, you need to download the installation media image for your device and Junos OS release from <http://www.juniper.net/customers/support/>.





**NOTE:** In the following procedure, we assume that you are creating the emergency boot device on a QFX device or EX4600 device. You can create the emergency boot device on another Juniper Networks switch or router, or any PC or laptop that supports Linux. The steps you take to create the emergency boot device vary, depending on the device.

To create an emergency boot device:

1. Use FTP to copy the installation media image into the `/var/tmp` directory on the device.
2. Insert a USB device into the USB port.
3. From the Junos OS command-line interface (CLI), start the shell:

```
user@device> start shell
%
```

4. Switch to the root account using the `su` command:

```
% su
Password: password
```



**NOTE:** The password is the root password for the device. If you logged in to the device as root, you do not need to perform this step.

5. Enter the following command on the QFX3500, QFX3600, and QFX3600-I devices:

```
root@device% dd if=/var/tmp/filename of=/dev/da1 bs=16k
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/install-media-qfx3500.junos_11.1 of=/dev/da1 bs=16k
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

6. Enter the following command on the QFX5100 and EX4600 devices:

```
root@device% dd if=/var/tmp/filename of=/dev/da0 bs=1048576
```

The device writes the installation media image to the USB device:

```
root@device% dd if=/var/tmp/jinstall-vjunos-usb-13.2.img of=/dev/da0 bs=1048576
11006+1 records in
11006+1 records out
180332544 bytes transferred in 71.764266 secs (2512846 bytes/sec)
```

7. Log out of the shell:

```
root@device% exit
% exit
user@device>
```

#### Related Documentation

- [USB Port Specifications for the QFX Series](#)
- [Performing a Recovery Installation on page 116](#)
- [Performing a QFabric System Recovery Installation on the Director Group on page 7233](#)
- [Performing a Recovery Installation on page 118](#)

## Performing a Recovery Installation

If Junos OS on your device is damaged in some way that prevents the software from loading correctly, you may need to perform a recovery installation using an emergency boot device (for example, a USB flash drive) to restore the default factory installation. Once you have recovered the software, you need to restore the device configuration. You can either create a new configuration as you did when the device was shipped from the factory, or if you saved the previous configuration, you can simply restore that file to the device.

If at all possible, you should try to perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device to use during the installation. See [“Creating an Emergency Boot Device” on page 176](#) for information on how to create an emergency boot device.
2. Copy the existing configuration in the file `/config/juniper.conf.gz` from the device to a remote system, such as a server, or to an emergency boot device. For extra safety, you can also copy the backup configurations (the files named `/config/juniper.conf.n`, where *n* is a number from 0 through 9) to a remote system or to an emergency boot device.



**WARNING:** The recovery installation process completely overwrites the entire contents of the internal flash storage.

3. Copy any other stored files to a remote system as desired.

To reinstall Junos OS:

1. Insert the emergency boot device into the QFX Series device.
2. Reboot the QFX Series device.



**NOTE:** Do not power off the device if it is already on.

```
[edit system]
user@device> request system reboot
```

If you do not have access to the CLI, power cycle the QFX Series device.

The emergency boot device (external USB install media) is detected. At this time, you can load the Junos OS from the emergency boot device onto the internal flash storage.

3. The software prompts you with the following options:

```
External USB install media detected.
You can load Junos from this media onto an internal drive.
Press 'y' to proceed, 'f' to format and install, or 'n' to abort.
Do you wish to continue ([y]/f/n)? f
```

4. Type **f** to format the internal flash storage and install the Junos OS on the emergency boot device onto the internal flash storage.

If you do not want to format the internal flash storage, type **y**.

The following messages are displayed:

```
Installing packages from external USB drive da1
Packages will be installed to da0, media size: 8G
```

```
Processing format options
Fri September  4 01:18:44 UTC 2012
```

```
-- IMPORTANT INFORMATION --
Installer has detected settings to format system boot media.
This operation will erase all data from your system.
```

```
Formatting installation disk .. this will take a while, please wait
Disabling platform watchdog - threshold 12 mins
```

```
Determining installation slice
Fri September  4 01:27:07 UTC 2012
```

5. The device copies the software from the emergency boot device, occasionally displaying status messages. Copying the software can take up to 12 minutes.

When the device is finished copying the software, you are presented with the following prompt:

```
*** Fri September  4 01:19:00 UTC 2012***
Installation successful..
Please select one of the following options:
Reboot to installed Junos after removing install media (default) ... 1
Reboot to installed Junos by disabling install media ..... 2
Exit to installer debug shell ..... 3
Install Junos to alternate slice ..... 4
Your choice: 4
NOTE: System installer will now install Junos to alternate slice
Do not power off or remove the external installer media or
interrupt the installation mechanism.
```

6. Select **4** to install Junos OS to the alternate slice of the partition, and then press Enter.
7. Remove the emergency boot device when prompted and then press Enter. The device then reboots from the internal flash storage on which the software was just installed. When the reboot is complete, the device displays the login prompt.
8. Create a new configuration as you did when the device was shipped from the factory, or restore the previously saved configuration file to the device.

**Related Documentation**

- [Creating an Emergency Boot Device on page 176](#)

## Performing a QFabric System Recovery Installation on the Director Group

If the software on your QFabric system is damaged in some way that prevents the software from loading correctly, or you need to upgrade the software on your QFabric system, you may need to perform a recovery installation on the Director group.

If possible, perform the following steps before you perform the recovery installation:

1. Ensure that you have an emergency boot device (for example, an external USB flash drive) for each of your Director devices to use during the recovery installation.

You can either use the external USB flash drive containing the software supplied by Juniper Networks, or you can use an external USB flash drive supplied by Juniper Networks on which you install the QFabric system install media.

2. Because the recovery installation process completely overwrites the entire contents of the Director device, make sure you back up any configuration files and initial setup information on a different external USB flash drive before you begin a recovery installation. You will need to restore this information as part of recovery process.

Use the **request system software configuration-backup** command to back up your configuration files and initial setup information:

```
user@switch> request system software configuration-backup path
```



**NOTE:** To recover the Director group, you must upgrade both Director devices in parallel. If you are recovering only one Director device in a Director group, and the software version will remain the same between the two Director devices, make sure that the other Director device is powered on and operational. If the software version of the Director device you are recovering will be different, make sure that the other Director device is powered off and is not operational.

- (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive on page 7234
- Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software on page 7236

#### (Optional) Creating an Emergency Boot Device Using a Juniper Networks External Blank USB Flash Drive

---

If you do not have an external USB flash drive preloaded with the software from Juniper Networks to use as an emergency boot device, you can create your own, using a blank external USB flash drive provided by Juniper Networks. Download the install media from the Juniper Networks Support website onto your UNIX workstation, uncompress and untar the software, and then burn the software image onto your Juniper Networks external USB (4-gigabyte) flash drive. Make sure you create two emergency boot devices, one for each Director device, so you can perform a recovery installation in parallel.

1. Using a Web browser, navigate to the <http://www.juniper.net/support>.
2. Click **Download Software**.
3. In the *Switchingbox*, click *Junos OS Platforms*.
4. In the *QFX Series* section, click the name of the platform for which you want to download software.

5. Click the *Software* tab and select the release number from the *Release* drop-down list.
6. Select the complete install media you want to download in the *QFabric System Install Media* section.  
A login screen appears.
7. Enter your name and password and press **Enter**.
8. Read the End User License Agreement, click the **I agree** radio button, and then click **Proceed**.
9. Log in and save the install media file to your UNIX workstation.
10. Use FTP to access the UNIX workstation where the install media resides.  
`ftp ftp://hostname/pathname install-media-qfabric-<version>.img.tgz`
11. When prompted, enter your username and password.
12. Make sure you are in binary mode by entering **binary** at the prompt.  
`binary`
13. Use the **get** command to transfer the installation package from the FTP host to your UNIX workstation.  
`get install-media-qfabric-<version>.img.tgz`
14. Close the FTP session:  
`bye`
15. Untar the *install-media-qfabric-<version>.img.tgz* file on your UNIX workstation.  
`tar -xvzf install-media-qfabric-11.3X30.6.img.tgz`
16. Insert a blank external USB (4-gigabyte) flash drive supplied by Juniper Networks into your UNIX workstation.
17. Burn the software image you just downloaded to your UNIX workstation onto your external USB flash drive using the **dd** command:  
`dd if=install-media-qfabric-11.3X30.6.img of=/dev/sdb bs=16k`  
250880+0 records in  
250880+0 records out  
4110417920 bytes (4.1 GB) copied, 5.10768 seconds, 805 MB/s
18. Perform the steps in [“Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software” on page 7236](#) to continue with the recovery installation.

## Performing a Recovery Installation Using a Juniper Networks External USB Flash Drive with Preloaded Software

---

This procedure describes how to perform a recovery installation using an external USB flash drive that contains Junos OS software.



**NOTE:** Since the recovery installation process completely overwrites the entire contents of the Director device, you will need to restore the required configuration files and initial setup information. The following procedure assumes you previously saved these backup files with the **request system software configuration-backup** command. Ensure that you have these backup files available on an external USB flash drive before you perform the following steps.

1. Insert the external USB flash drive into the Director device.
2. Perform one of the following tasks:
  - If you have access to the default partition, reboot the Director device by issuing the **request system reboot director-group** command.
  - If you do not have access to the default partition, power cycle the Director device.

The following menu appears on the Director device console when the Director device boots up:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

3. Type **install** and then press **Enter** to install the software on the Director device.

Once the installation process is complete, the Director device reboots, and the following menu appears on the Director device console:

```
Juniper Networks QFabric Director Install/Recovery Media
- To boot from the local disk, wait 10 seconds or press the Enter key.
- To reinstall the QFabric software on this Director device, type: install
```

4. Press **Enter**.

The Director device reboots from the local disk on which the software was just installed.

5. Log in as root on the Director device.

The following menu appears on the Director device console:

```
Before you can access the QFabric system, you must complete the initial setup
of the Director group by using the steps that follow.
If the initial setup procedure does not complete successfully, log out of the
Director device and then log back in to restart
this setup menu.
```

```
Continue?[y/n]
```

6. Enter **n** to bypass the initial setup script and enter the Director device root directory, where you can mount the external USB flash drive containing the configuration files and initial setup information.

7. Issue the **ls /mnt** command to list the *mount* directory.

```
root@dg0 ~]# ls /mnt
```

8. Issue the **mkdir** command to create a directory within the mount directory.

```
root@dg0 ~]# mkdir /mnt/myusb
```

9. Issue the **mount /dev/sdb2 /mnt/myusb/** command to mount the external USB flash drive to the local drive of the Director device.

```
root@dg0 ~]# mount /dev/sdb2 /mnt/myusb/
```

10. Issue the **ls -la /mnt/myusb/** command to verify the contents of your mounted external USB flashdrive.

```
root@dg0 ~]# ls -la /mnt/myusb/
total 1770884
drwxr-xr-x 2 root root      4096 Sep  7 05:16 .
drwxr-xr-x 3 root root      4096 Sep  7 10:15 ..
-rw-r--r-- 1 root root    4249 Sep  7 03:52 mybackup-20110907
```

11. Exit the Director device and log back in as root on the Director device.

The following menu appears:

Before you can access the QFabric system, you must complete the initial setup of the Director group by using the steps that follow.

If the initial setup procedure does not complete successfully, log out of the Director device and then log back in to restart this setup menu.

```
Continue?[y/n] y
Initial Configuration
```

You may enter the configuration manually or restore from a backup.

```
Specify a backup file? [y/n] : y
Please specify the full path of the configuration backup file. :
/mnt/myusb/mybackup-20110907
```

12. Enter **y** to continue.

13. Enter **y** and specify the path to the backup configuration file located on the external USB flash drive.

```
/mnt/myusb/mybackup-20110907
```

The following messages appear:

```
Saving temporary configuration...
Configuring peer...
connect error for 1.1.1.2:9001
Configuring local interfaces...
Configuring interface eth0 with [10.49.213.163/24:10.49.213.254]
Configured interface eth0 with [10.49.213.163/24:10.49.213.254]
Configuring QFabric software with initial pool of 4000 MAC addresses
[00:10:00:00:00:00 - 00:10:00:00:0f:3b]
Configuring QFabric address [10.49.213.50]
Reconfiguring QFabric software static configuration
Applying the new Director Device password
Applying the QFabric component password
```

```
First install initial configuration, generating and sharing SSH keys.  
First install initial configuration, generating SSH keys.  
connect error for 1.1.1.2:9001  
Shared SSH keys.  
Configuration complete. Director Group services will auto start within 30  
seconds.
```

The Director device reboots from the local disk on which the software was just installed.  
Exit the Director device session and log in to the QFabric default partition CLI.

14. Issue the **request system software configuration-restore** command and specify the path to the backup configuration file located on the external USB flash drive to load the previously saved QFabric system configuration.

15. From the default partition, issue the **request system reboot node-group all** command to reboot all of the Node groups in the QFabric system to ensure that all Node devices are running the same version of software as the Director-group.

```
user@switch> request system reboot node-group all
```

16. From the default partition, issue the **request system reboot fabric** command to reboot the Interconnect devices and the other components in the fabric in the QFabric system to ensure that Interconnect devices are running the same version of software as the Director group.

```
user@switch> request system reboot fabric
```

17. Log in to the default partition and issue the **show version component all** command to verify that all components are running the same version of software.

```
user@switch> show version component all  
dg1:  
-  
Hostname: qfabric  
Model: qfx3100  
JUNOS Base Version [11.3X30.6]  
  
dg0:  
-  
Hostname: qfabric  
Model: qfx3100  
JUNOS Base Version [11.3X30.6]  
  
NW-NG-0:  
-  
Hostname: qfabric  
Model: qfx-jvre  
JUNOS Base OS boot [11.3X30.6]  
JUNOS Base OS Software Suite [11.3X30.6]  
JUNOS Kernel Software Suite [11.3X30.6]  
JUNOS Crypto Software Suite [11.3X30.6]  
JUNOS Online Documentation [11.3X30.6]  
JUNOS Enterprise Software Suite [11.3X30.6]  
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]  
JUNOS Routing Software Suite [11.3X30.6]  
  
FC-0:  
-  
Hostname: qfabric  
Model: qfx-jvre  
JUNOS Base OS boot [11.3X30.6]  
JUNOS Base OS Software Suite [11.3X30.6]
```



```
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

FC-1:

```
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

DRE-0:

```
-
Hostname: dre-0
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

FM-0:

```
-
Hostname: qfabric
Model: qfx-jvre
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

nodedevice1:

```
-
Hostname: qfabric
Model: QFX3500
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
```

interconnectdevice1:

```
-
Hostname: qfabric
```

```
Model: QFX3108
JUNOS Base OS boot [11.3X30.6]
JUNOS Base OS Software Suite [11.3X30.6]
JUNOS Kernel Software Suite [11.3X30.6]
JUNOS Crypto Software Suite [11.3X30.6]
JUNOS Online Documentation [11.3X30.6]
JUNOS Enterprise Software Suite [11.3X30.6]
JUNOS Packet Forwarding Engine Support (QFX RE) [11.3X30.6]
JUNOS Routing Software Suite [11.3X30.6]
warning: from interconnectdevice0: Disconnected
```

- Related Documentation**
- *Performing the QFabric System Initial Setup on a QFX3100 Director Group*
  - *Upgrading Software on a QFabric System*
  - *request system software configuration-backup*
  - *request system software configuration-restore*

## Troubleshooting Network Interfaces

### The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

---

**Problem**    **Description:** The QFX Series has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command **show interfaces *interface-name***, the disabled port is not listed.

**Cause**    By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution**    Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

## Troubleshooting an Aggregated Ethernet Interface

**Problem**    **Description:** The **show interfaces terse** command shows that the LAG is down.

**Solution**    Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).

- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related  
Documentation**

- [Verifying the Status of a LAG Interface on page 2750](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 2462](#)

## Layer 3 Protocols

- [Troubleshooting Virtual Routing Instances on page 7241](#)

### Troubleshooting Virtual Routing Instances

- [Direct Routes Not Leaked Between Routing Instances on page 7241](#)

#### Direct Routes Not Leaked Between Routing Instances

**Problem** **Description:** Direct routes are not exported (leaked) between virtual routing instances. For example, consider the following scenario:

- Switch with two virtual routing instances:
  - Routing instance 1 connects to downstream device through interface xe-0/0/1.
  - Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the switch connects to the upstream device over a direct route and these routes are not leaked between instances.



**NOTE:** You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the switch over indirect routes.

**Solution** This is expected behavior.

**Related  
Documentation**

- [Understanding Virtual Router Routing Instances on page 2898](#)
- [Configuring Virtual Router Routing Instances on page 2908](#)
- [rib-group on page 3031](#)

## MPLS

---

- [Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch on page 7242](#)

### Issues and Limitations in Operation of MPLS Features on the QFX Series and on the EX4600 Switch

The following issues exist in the operation of MPLS features on QFX Series devices and on the EX4600 switch. In each case, the described behavior is the expected behavior.

- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.

#### Related Documentation

- [MPLS Feature Support on the QFX Series and EX4600 Switch Overview on page 4423](#)

## Network Management

---

- [Understanding Troubleshooting Resources on page 7242](#)
- [Troubleshooting Overview on page 7244](#)
- [QFX5100 Switch with Automation Enhancements Frequently Asked Questions on page 7247](#)
- [Recovering from a Failed Software Installation on page 7248](#)
- [Loading a Previous Configuration File on page 7249](#)
- [Reverting to the Default Factory Configuration on page 7250](#)
- [Reverting to the Rescue Configuration on page 7250](#)
- [Recovering the Root Password on page 7251](#)
- [Troubleshooting a Deprecated Network Analytics Configuration on page 7252](#)

### Understanding Troubleshooting Resources

This topic describes some of the troubleshooting resources available for the QFX Series. These resources include tools such as the Junos OS CLI, Junos Space applications, and the Advanced Insight Scripts (AI-Scripts).

[Table 663 on page 6893](#) provides a list of some of the troubleshooting resources.

Table 700: Troubleshooting Resources on the QFX Series

Troubleshooting Resource	Description	Documentation
Chassis alarms	Chassis alarms indicate a failure on the switch or one of its components. A chassis alarm count is displayed on the LCD panel on the front of the switch.	<a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 7192</a>
Chassis Status LEDs and Fan Tray LEDs	A blinking amber Power, Fan, or Fan Tray LED indicates a hardware component error. A blinking amber Status LED indicates a software error.	<i>Chassis Status LEDs on a QFX3500 Device</i>
Interface alarms	A predefined alarm (red or yellow) for an interface type is triggered when an interface of that type goes down.	<a href="#">“Interface Alarm Messages” on page 7195</a>
System alarms	A predefined alarm is triggered by a missing rescue configuration or problem with the software license.	<a href="#">“Understanding Alarms” on page 7191</a>
System log messages	The system log includes details of system and user events, including errors. Specify the severity and type of system log messages you wish to view or save, and configure the output to be sent to local or remote hosts.	<ul style="list-style-type: none"> <li>• <a href="#">Overview of Single-Chassis System Logging Configuration on page 6561</a></li> <li>• <a href="#">Junos OS System Log Configuration Statements on page 6616</a></li> </ul>
Junos OS operational mode commands	Operational mode commands can be used to monitor switch performance and current activity on the network. For example, use the <b>tracroute monitor</b> command to locate points of failure in a network.	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring System Process Information on page 333</a></li> <li>• <a href="#">Monitoring System Properties on page 334</a></li> <li>• <a href="#">tracroute monitor</a></li> </ul>
Junos OS automation scripts (event scripts)	Event scripts can be used to automate network troubleshooting and management tasks.	<i>Junos OS Automation Library</i>
Junos OS XML operational tags	XML operational tags are equivalent in function to operational mode commands in the CLI, which you can use to retrieve status information for a device.	<i>Junos XML API Operational Developer Reference</i>
NETCONF XML management protocol	The NETCONF XML management protocol defines basic operations that are equivalent to Junos OS CLI configuration mode commands. Client applications use the protocol operations to display, edit, and commit configuration statements (among other operations), just as administrators use CLI configuration mode commands such as <b>show</b> , <b>set</b> , and <b>commit</b> to perform those operations.	<i>NETCONF XML Management Protocol Developer Guide</i>

Table 700: Troubleshooting Resources on the QFX Series (*continued*)

Troubleshooting Resource	Description	Documentation
SNMP MIBs and traps	MIBs enable the monitoring of network devices from a central location. For example, use the Traceroute MIB to monitor devices remotely.	<ul style="list-style-type: none"> <li>• <a href="#">SNMP MIBs Support on page 6530</a></li> <li>• <a href="#">SNMP Traps Support on page 6546</a></li> <li>• <a href="#">Using the Traceroute MIB for SNMP Remote Operations</a></li> </ul>
AI-Scripts and Advanced Insight Manager (AIM)	AI-Scripts installed on the switch can automatically detect and monitor faults on the switch, and depending on the configuration on the AIM application, send notifications of potential problems and submit problem reports to Juniper Support Systems.	<a href="#">Advanced Insight Scripts (AI-Scripts) Release Notes</a>
Junos Space Service Now	This application enables you to display and manage information about problem events. When problems are detected on the switch by Advanced Insight Scripts (AI-Scripts) that are installed on the switch, the data is collected and sent to Service Now for your review and action.	<i>Service Automation</i>
Junos Space Service Insight	This application helps in accelerating operational analysis and managing the exposure to known issues. You can identify devices that are nearing their End Of Life (EOL) and also discover and prevent issues that could occur in your network. The functionality of Service Insight is dependent on the information sent from Service Now.	<i>Service Automation</i>
Juniper Networks Knowledge Base	You can search in this database for Juniper Networks product information, including alerts and troubleshooting tips.	<a href="http://kb.juniper.net">http://kb.juniper.net</a>

## Troubleshooting Overview

This topic provides a general guide to troubleshooting some typical problems you may encounter on your QFX Series product.

[Table 664 on page 6895](#) provides a list of problem categories, summary of the symptom or problem, and recommended actions with links to the troubleshooting documentation.

Table 701: Troubleshooting on the QFX Series

Problem Category	Symptom or Problem	Recommended Action
Switch hardware components	LCD panel shows a chassis alarm count.	See <a href="#">“Chassis Alarm Messages on a QFX3500 Device” on page 7192.</a>
	Fan tray LED is blinking amber.	See <a href="#">Fan Tray LED on a QFX3500 Device.</a>
	Chassis status LED for the power is blinking amber.	See <a href="#">Chassis Status LEDs on a QFX3500 Device.</a>
	Chassis status LED for the fan (on the management board) is blinking amber.	Replace the management board as soon as possible. See <a href="#">Chassis Status LEDs on a QFX3500 Device.</a>
Port configuration	Cannot configure a port as a Gigabit Ethernet port.	Check whether the port is a valid Gigabit Ethernet port (6 through 41).  See <a href="#">QFX3500 Device Overview.</a>
	Cannot configure a port as a Fibre Channel port.	Check whether the port is a valid Fibre Channel port (0 through 5 and 42 through 47).  See <a href="#">QFX3500 Device Overview.</a>
	Cannot configure a port as a 10-Gigabit Ethernet port.	If the port is not a 40-Gbps QSFP+ interface, check whether the port is in the range of 0 through 5 or 42 through 47. If one of the ports in that block (0 through 5 or 42 through 47) is configured as a Fibre Channel port, then all ports in that block must also be configured as Fibre Channel ports.  If the port is a 40-Gbps QSFP+ interface, make sure the configuration does not exceed the interface limit. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces, but because port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.  See <a href="#">QFX3500 Device Overview.</a>
	Cannot configure a 40-Gbps QSFP+ interface.	The 40-Gbps QSFP+ interfaces can only be used as 10-Gigabit Ethernet interfaces. Each 40-Gbps QSFP+ interface can be split into four 10-Gigabit Ethernet interfaces using a breakout cable. However, port 0 is reserved, so you can only configure an additional fifteen 10-Gigabit Ethernet interfaces.  See <a href="#">QFX3500 Device Overview.</a>
External devices (USB devices)	Upgrading software from a USB device results in an upgrade failure, and the system enters an invalid state.	Unplug the USB device and reboot the switch.

Table 701: Troubleshooting on the QFX Series (*continued*)

Problem Category	Symptom or Problem	Recommended Action
Initial device configuration	Cannot configure management Ethernet ports.	Configure the management ports from the console port. You cannot configure the management ports by directly connecting to them.  <b>NOTE:</b> The management ports are on the front panel of the QFX3500 switch. They are labeled <b>C0</b> and <b>C1</b> on the front panel. In the CLI they are referred to as <b>me0</b> and <b>me1</b> .  See <a href="#">“Configuring a QFX3500 Device as a Standalone Switch”</a> on page 175.
	Failed software upgrade.	See <a href="#">“Recovering from a Failed Software Installation”</a> on page 121.
Software upgrade and configuration	Active partition becomes inactive after upgrade.	
	Problem with the active configuration file.	See the following topics: <ul style="list-style-type: none"> <li>• <a href="#">Loading a Previous Configuration File</a> on page 1252</li> <li>• <a href="#">Reverting to the Default Factory Configuration</a> on page 188</li> <li>• <a href="#">Reverting to the Rescue Configuration</a> on page 189</li> <li>• <a href="#">Performing a Recovery Installation</a> on page 116</li> </ul>
	Root password is lost or forgotten.	Recover the root password. See <a href="#">“Recovering the Root Password”</a> on page 1233.
Network interfaces	An aggregated Ethernet interface is down.	See <a href="#">“Troubleshooting an Aggregated Ethernet Interface”</a> on page 1234.
	Interface on built-in network port is down.	See <a href="#">“Troubleshooting Network Interfaces”</a> on page 1234.
	Interface on port in which SFP or SFP+ transceiver is installed in an SFP+ uplink module is down.	
Ethernet switching	A MAC address entry in the Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch.	See <a href="#">“Troubleshooting Ethernet Switching”</a> on page 1895.
Firewall filter	Firewall configuration exceeded available Ternary Content Addressable Memory (TCAM) space.	See <a href="#">“Troubleshooting Firewall Filter Configuration”</a> on page 5411.



## QFX5100 Switch with Automation Enhancements Frequently Asked Questions

This FAQ addresses questions regarding using QFX5100 switches with automation enhancements, which were introduced at Junos OS Release 13.2X51-D15.

This FAQ covers the following questions:

- [Who Should You Contact If You Have Problems with Loading, Installing or Updating Libraries? on page 7247](#)
- [Who Should You Contact If You Have Problems with Puppet for Junos OS? on page 7247](#)
- [Who Should You Contact If You Have Problems with Chef for Junos OS? on page 7247](#)
- [What Happens to the User Partition If You Downgrade a QFX5100 Switch That Is Running the jinstall-qfx-5-flex-x.tgz Software Bundle to a QFX Switch That Is Running a Different QFX5100 Software Bundle? on page 7247](#)
- [How Do You Recover Junos OS Binaries That You Have Deleted? on page 7247](#)
- [How Do You Recover from a System Crash? on page 7247](#)
- [How Can You Verify That a QFX5100 Switch Is Running a jinstall-qfx-5-flex-x.tgz Software Bundle? on page 7248](#)

### Who Should You Contact If You Have Problems with Loading, Installing or Updating Libraries?

Contact Customer Support at <http://www.juniper.net/support>.

### Who Should You Contact If You Have Problems with Puppet for Junos OS?

You can obtain support for Puppet for Junos OS through the J-Net Forum for Puppet at [http://forums.juniper.net/t5/Puppet-for-JunOS/bd-p/puppet\\_junos](http://forums.juniper.net/t5/Puppet-for-JunOS/bd-p/puppet_junos).

### Who Should You Contact If You Have Problems with Chef for Junos OS?

You can obtain support for Chef for Junos OS through the J-Net Forum for Chef at [http://forums.juniper.net/t5/Chef-for-JunOS/bd-p/chef\\_junos](http://forums.juniper.net/t5/Chef-for-JunOS/bd-p/chef_junos).

### What Happens to the User Partition If You Downgrade a QFX5100 Switch That Is Running the jinstall-qfx-5-flex-x.tgz Software Bundle to a QFX Switch That Is Running a Different QFX5100 Software Bundle?

In this case, the user partition remains intact.



**NOTE:** If you make changes to the user partition while performing a unified in-service software upgrade (unified ISSU), the changes might be lost.

### How Do You Recover Junos OS Binaries That You Have Deleted?

You must reinstall the software package.

### How Do You Recover from a System Crash?

You must reinstall the software package.

### How Can You Verify That a QFX5100 Switch Is Running a jinstall-qfx-5-flex-x.tgz Software Bundle?

---

You cannot use the **show version** command to verify that a QFX5100 switch is running the jinstall-qfx-5-flex-x.tgz software bundle. However, there are two other ways to verify this.

- Use the **show configuration** command to check that you are running a Layer 3 configuration. See *Installing Junos OS Software with QFX5100 Switch Automation Enhancements*.
- Go to the shell and confirm that you can invoke Python. See “[Invoking the Python Interpreter](#)” on page 6587.

#### Related Documentation

- [Overview of QFX5100 Switch Automation Enhancements on page 6470](#)
- [Installing Junos OS Software with QFX5100 Switch Automation Enhancements](#)
- [Invoking the Python Interpreter on page 6587](#)
- [Chef for Junos Getting Started Guide](#)
- [Puppet for Junos OS Documentation](#)

### Recovering from a Failed Software Installation

**Problem**    **Description:** If the Junos OS appears to have been installed but the CLI does not work, or if the switch has no software installed, you can use this recovery installation procedure to install the Junos OS.

**Solution**    If a Junos OS image already exists on the switch, you can either install the new Junos OS package in a separate partition, in which case both Junos OS images remain on the switch, or you can remove the existing Junos OS image before you start the new installation process.

To perform a recovery installation:

1. Power on the switch. The loader script starts.
2. After the message **Loading /boot/defaults/loader.conf** appears, you are prompted with the following message:

**Hit [Enter] to boot immediately, or space bar for command prompt.**

Press the Spacebar to enter the manual loader. The **loader>** prompt appears.

3. Enter the following command:

```
loader> install [- --format] [- --external] source
```

where:

- **format**—Enables you to erase the installation media before installing the installation package. If you do not include this option, the system installs the new Junos OS in a different partition from that of the most recently installed Junos OS.

- **external**—Installs the installation package onto external media (a USB stick, for example).
- **source**—Represents the name and location of the Junos OS package, either on a server on the network or as a file on an external media, as shown in the following two examples:
  - Network address of the server and the path on the server; for example, **tftp://192.171.28/junos/jinstall-qfx-11.1R1.5-domestic-signed.tgz**
  - Junos OS package on a USB device (commonly stored in the root drive as the only file), for example, **file:///jinstall-qfx-11.1R1.5-domestic-signed.tgz**.

The installation now proceeds normally and ends with a login prompt.

## Loading a Previous Configuration File

You can use the **rollback <number>** command to return to a previously committed configuration file. A switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

### Syntax

**rollback <number>**

### Options

- **none**—Return to the most recently saved configuration.
- **number**—Configuration to return to.
  - **Range:** 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
  - **Default:** 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuration):

```
[edit]
user@switch# rollback 1
load complete
```

2. Activate the configuration you have loaded:

```
[edit]
user@switch# commit
```

**Related Documentation**

- [Configuration File Terms on page 11](#)

## Reverting to the Default Factory Configuration

If for any reason the current active configuration fails, you can revert to the default factory configuration. The default factory configuration contains the basic configuration settings. This is the first configuration of the switch, and it is loaded when the switch is first installed and powered on.

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

To revert the switch to the rescue configuration:

1. 

```
[edit]
user@switch# load factory-default
[edit]
user@switch# delete system commit factory-settings
[edit]
user@switch# commit
```

### Related Documentation

- [Understanding Configuration Files on page 1242](#)
- [Loading a Previous Configuration File on page 1252](#)
- [Reverting to the Rescue Configuration on page 189](#)

## Reverting to the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a device and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration. The rescue configuration is a previously committed, valid configuration.

To revert the switch to the rescue configuration:

1. Enter the **load override** command.

```
[edit]
user@switch# load override filename
```
2. Commit your changes.

```
[edit]
user@switch# commit filename
```

### Related Documentation

- [Setting or Deleting the Rescue Configuration on page 1261](#)
- [Reverting to the Default Factory Configuration on page 188](#)
- [Configuration File Terms on page 11](#)

## Recovering the Root Password

If you forget the root password, you can use the password recovery procedure to reset the root password.



**NOTE:** The root password cannot be recovered on a QFabric system.



**NOTE:** You need console access to the switch to recover the root password.

To recover the root password:

1. Power off the switch by switching off the AC power outlet of the device or, if necessary, by pulling the power cords out of the device's power supplies.
2. Turn off the power to the management device, such as a PC or laptop computer, that you want to use to access the CLI.
3. Plug one end of the Ethernet rollover cable supplied with the device into the RJ-45-to-DB-9 serial port adapter supplied with the device.
4. Plug the RJ-45-to-DB-9 serial port adapter into the serial port on the management device.
5. Connect the other end of the Ethernet rollover cable to the console port on the device.
6. Turn on the power to the management device.
7. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).
8. Configure the port settings as follows:
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
9. Power on the device by (if necessary) plugging the power cords into the device's power supply, or turning on the power to the device by switching on the AC power outlet the device is plugged into.

The terminal emulation screen on your management device displays the device's boot sequence.

10. When the following prompt appears, press the Spacebar to access the device's bootstrap loader command prompt:

```
Hit [Enter] to boot immediately, or space bar for command prompt.
Booting [kernel] in 9 seconds...
```

11. At the following prompt, enter **boot -s** to start up the system in single-user mode.  

```
ok boot -s
```
12. At the following prompt, enter **recovery** to start the root password recovery procedure.  

```
Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: recovery
```
13. Enter configuration mode in the CLI.
14. Set the root password. For example:  

```
user@switch# set system root-authentication plain-text-password
```
15. At the following prompt, enter the new root password. For example:  

```
New password: juniper1
Retype new password:
```
16. At the second prompt, reenter the new root password.
17. After you have finished configuring the password, commit the configuration.  

```
root@host# commit
commit complete
```
18. Exit configuration mode in the CLI.
19. Exit operational mode in the CLI.
20. At the prompt, enter **y** to reboot the device.  

```
Reboot the system? [y/n] y
```

**Related Documentation** • [Configuring the Root Password on page 1354](#)

## Troubleshooting a Deprecated Network Analytics Configuration

**Problem** **Description:** After a software upgrade to Junos OS Release 13.2X51-D15 from an earlier release, the network analytics configuration is no longer valid and the feature is disabled.

**Symptoms:** The network analytics configuration used in Junos OS Release 13.2X51-D10 has been deprecated in Release 13.2X51-D15. Issuing the **show services analytics** command results in the following output:

```
root@qfx5100# show services analytics

queue-statistics { ## Warning: 'queue-statistics' is deprecated
  interval 1;
}
```

**Cause** Junos OS Release 13.2X51-D15 added enhancements to the network analytics feature, resulting in significant changes in the CLI. The updated **[edit services analytics]** hierarchy level contains some statements that have replaced those that were previously released. As a result, the earlier configuration does not work in the new release.

**Solution** Use the new CLI statements to reconfigure the network analytics feature.

- Related Documentation**
- [Network Analytics Overview on page 6490](#)
  - [analytics on page 6667](#)

## Security

---

- [Troubleshooting Firewall Filter Configuration on page 7253](#)
- [Troubleshooting Policer Configuration on page 7259](#)

### Troubleshooting Firewall Filter Configuration

Use the following information to troubleshoot your firewall filter configuration.

- [Firewall Filter Configuration Returns a No Space Available in TCAM Message on page 7253](#)
- [Filter Counts Previously Dropped Packet on page 7255](#)
- [Matching Packets Not Counted on page 7255](#)
- [Counter Reset When Editing Filter on page 7256](#)
- [Cannot Include loss-priority and policer Actions in Same Term on page 7256](#)
- [Cannot Egress Filter Certain Traffic Originating on QFX Switch on page 7256](#)
- [Firewall Filter Match Condition Not Working with Q-in-Q Tunneling on page 7257](#)
- [Egress Firewall Filters with Private VLANs on page 7257](#)
- [Egress Filtering of L2PT Traffic Not Supported on page 7258](#)
- [Cannot Drop BGP Packets in Certain Circumstances on page 7258](#)
- [Invalid Statistics for Policer on page 7258](#)
- [Policers can Limit Egress Filters on page 7258](#)

#### Firewall Filter Configuration Returns a No Space Available in TCAM Message

**Problem** **Description:** When a firewall filter configuration exceeds the amount of available Ternary Content Addressable Memory (TCAM) space, the system returns the following **syslogd** message:

```
No space available in tcam.
Rules for filter filter-name will not be installed.
```

A switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of space available in the TCAM table. The filter is not applied, but the commit operation for the firewall filter configuration is completed in the CLI module.

**Solution** When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the filter and its binding and apply the new smaller firewall filter to the same binding:

1. Delete the filter and its binding to ports, VLANs, or Layer 3 interfaces. For example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter ingress-vlan-rogue-block
user@switch# delete vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# delete vlans employee-vlan filter input ingress-vlan-rogue-block
```

2. Commit the changes:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space. For example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block ...
```

4. Apply (bind) the new firewall filter to a port, VLAN, or Layer 3 interface. For example:

```
[edit]
user@switch# set vlans employee-vlan description "filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

5. Commit the changes:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing binding but not delete the original filter:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-ingress-vlan-rogue-block...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the binding of the original filter—for example:

```
[edit]
user@switch# set vlans employee-vlan description "smaller filter to block rogue devices on
employee-vlan"
user@switch# set vlans employee-vlan filter input new-ingress-vlan-rogue-block
```

Because you can apply no more than one firewall filter per VLAN per direction, the binding of the original firewall filter to the VLAN is overwritten with the new firewall filter **new-ingress-vlan-rogue-block**.

3. Commit the changes:

```
[edit]
user@switch# commit
```



**NOTE:** The original filter is not deleted and is still available in the configuration.

---



### Filter Counts Previously Dropped Packet

- Problem** **Description:** If you configure two or more filters in the same direction for a physical interface and one of the filters includes a counter, the counter will be incorrect if the following circumstances apply:
- You configure the filter that is applied to packets first to discard certain packets. For example, imagine that you have a VLAN filter that accepts packets sent to 10.10.1.0/24 addresses and implicitly discards packets sent to any other addresses. You apply the filter to the **admin** VLAN in the output direction, and interface xe-0/0/1 is a member of that VLAN.
  - You configure a subsequent filter to accept and count packets that are dropped by the first filter. In this example, you have a port filter that accepts and counts packets sent to 192.168.1.0/24 addresses that is also applied to xe-0/0/1 in the output direction.

The egress VLAN filter is applied first and correctly discards packets sent to 192.168.1.0/24 addresses. The egress port filter is applied next and counts the discarded packets as matched packets. The packets are not forwarded, but the counter displayed by the egress port filter is incorrect.

Remember that the order in which filters are applied depends on the direction in which they are applied, as indicated here:

Ingress filters:

1. Port (Layer 2) filter
2. VLAN filter
3. Router (Layer 3) filter

Egress filters:

1. Router (Layer 3) filter
2. VLAN filter
3. Port (Layer 2) filter

**Solution** This is expected behavior.

### Matching Packets Not Counted

**Problem** **Description:** If you configure two egress filters with counters for a physical interface and a packet matches both of the filters, only one of the counters includes that packet. For example:

- You configure an egress port filter with a counter for interface xe-0/0/1.
- You configure an egress VLAN filter with a counter for the **admin** VLAN, and interface xe-0/0/1 is a member of that VLAN.

- A packet matches both filters.

In this case, the packet is counted by only one of the counters even though it matched both filters.

**Solution** This is expected behavior.

---

#### Counter Reset When Editing Filter

**Problem Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution** This is expected behavior.

---

#### Cannot Include loss-priority and policer Actions in Same Term

**Problem Description:** You cannot include both of the following actions in the same firewall filter term in a QFX Series switch:

- **loss-priority**
- **policer**

If you do so, you see the following error message when you attempt to commit the configuration: "cannot support policer action if loss-priority is configured."

**Solution** This is expected behavior.

---

#### Cannot Egress Filter Certain Traffic Originating on QFX Switch

**Problem Description:** On a QFX Series switch, you cannot filter certain traffic with a firewall filter applied in the output direction if the traffic originates on the QFX switch. This limitation applies to control traffic for protocols such as ICMP (ping), STP, LACP, and so on.

**Solution** This is expected behavior.

### Firewall Filter Match Condition Not Working with Q-in-Q Tunneling

**Problem** **Description:** If you create a firewall filter that includes a match condition of `dot1q-tag` or `dot1q-user-priority` and apply the filter on input to a trunk port that participates in a service VLAN, the match condition does not work if the Q-in-Q EtherType is not 0x8100. (When Q-in-Q tunneling is enabled, trunk interfaces are assumed to be part of the service provider or data center network and therefore participate in service VLANs.)

**Solution** This is expected behavior. To set the Q-in-Q EtherType to 0x8100, enter the `set dot1q-tunneling ethertype 0x8100` statement at the `[edit ethernet-switching-options]` hierarchy level. You must also configure the other end of the link to use the same Ethertype.

### Egress Firewall Filters with Private VLANs

**Problem** **Description:** If you apply a firewall filter in the output direction to a primary VLAN, the filter also applies to the secondary VLANs that are members of the primary VLAN when the traffic egresses with the primary VLAN tag or isolated VLAN tag, as listed below:

- Traffic forwarded from a secondary VLAN trunk port to a promiscuous port (trunk or access)
- Traffic forwarded from a secondary VLAN trunk port that carries an isolated VLAN to a PVLAN trunk port.
- Traffic forwarded from a promiscuous port (trunk or access) to a secondary VLAN trunk port
- Traffic forwarded from a PVLAN trunk port. to a secondary VLAN trunk port
- Traffic forwarded from a community port to a promiscuous port (trunk or access)

If you apply a firewall filter in the output direction to a primary VLAN, the filter does *not* apply to traffic that egresses with a community VLAN tag, as listed below:

- Traffic forwarded from a community trunk port to a PVLAN trunk port
- Traffic forwarded from a secondary VLAN trunk port that carries a community VLAN to a PVLAN trunk port
- Traffic forwarded from a promiscuous port (trunk or access) to a community trunk port
- Traffic forwarded from a PVLAN trunk port. to a community trunk port

If you apply a firewall filter in the output direction to a community VLAN, the following behaviors apply:

- The filter is applied to traffic forwarded from a promiscuous port (trunk or access) to a community trunk port (because the traffic egresses with the community VLAN tag).
- The filter is applied to traffic forwarded from a community port to a PVLAN trunk port (because the traffic egresses with the community VLAN tag).

- The filter is *not* applied to traffic forwarded from a community port to a promiscuous port (because the traffic egresses with the primary VLAN tag or untagged).

**Solution** These are expected behaviors. They occur only if you apply a firewall filter to a private VLAN in the output direction and do not occur if you apply a firewall filter to a private VLAN in the input direction.

---

#### Egress Filtering of L2PT Traffic Not Supported

---

**Problem** **Description:** Egress filtering of L2PT traffic is not supported on the QFX3500 switch. That is, if you configure L2PT to tunnel a protocol on an interface, you cannot also use a firewall filter to filter traffic for that protocol on that interface in the output direction. If you commit a configuration for this purpose, the firewall filter is not applied to the L2PT-tunneled traffic.

**Solution** This is expected behavior.

---

#### Cannot Drop BGP Packets in Certain Circumstances

---

**Problem** **Description:** BGP packets with a time-to-live (TTL) value greater than 1 cannot be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface. BGP packets with TTL value of 1 or 0 can be discarded using a firewall filter applied to a loopback interface or applied on input to a Layer 3 interface.

**Solution** This is expected behavior.

---

#### Invalid Statistics for Policer

---

**Problem** **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution** This is expected behavior.

---

#### Policers can Limit Egress Filters

---

**Problem** **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional

egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

**Related Documentation**

- [Understanding FIP Snooping, FBF, and MVR Filter Scalability on page 5546](#)
- [Configuring Firewall Filters on page 5290](#)
- [Verifying That Firewall Filters Are Operational on page 5382](#)

## Troubleshooting Policer Configuration

- [Incomplete Count of Packet Drops on page 7260](#)
- [Counter Reset When Editing Filter on page 7260](#)
- [Invalid Statistics for Policer on page 7260](#)
- [Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 7260](#)

- [Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured on page 7261](#)
- [Policers Can Limit Egress Filters on page 7262](#)

---

### Incomplete Count of Packet Drops

---

**Problem**    **Description:** Under certain circumstances, Junos OS might display a misleading number of packets dropped by an ingress policer.

If packets are dropped because of ingress admission control, policer statistics might not show the number of packet drops you would expect by calculating the difference between ingress and egress packet counts. This might happen if you apply an ingress policer to multiple interfaces, and the aggregate ingress rate of those interfaces exceeds the line rate of a common egress interface. In this case, packets might be dropped from the ingress buffer. These drops are not included in the count of packets dropped by the policer, which causes policer statistics to underreport the total number of drops.

**Solution**    This is expected behavior.

---

### Counter Reset When Editing Filter

---

**Problem**    **Description:** If you edit a firewall filter term, the value of any counter associated with any term in the same filter is set to 0, including the implicit counter for any policer referenced by the filter. Consider the following examples:

- Assume that your filter has **term1**, **term2**, and **term3**, and each term has a counter that has already counted matching packets. If you edit any of the terms in any way, the counters for all the terms are reset to 0.
- Assume that your filter has **term1** and **term2**. Also assume that **term2** has a **policer** action modifier and the implicit counter of the policer has already counted 1000 matching packets. If you edit **term1** or **term2** in any way, the counter for the policer referenced by **term2** is reset to 0.

**Solution**    This is expected behavior.

---

### Invalid Statistics for Policer

---

**Problem**    **Description:** If you apply a single-rate two-color policer in more than 128 terms in a firewall filter, the output of the **show firewall** command displays incorrect data for the policer.

**Solution**    This is expected behavior.

---

### Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

---

**Problem**    **Description:** If you configure a policer to rate-limit throughput and apply it on egress to multiple interfaces on a QFX3500 switch or Node, the measured aggregate policed rate

might be twice the configured rate, depending on which interfaces you apply the policer to. The doubling of the policed rate occurs if you apply a policer to multiple interfaces and *both* of the following are true:

- There is at least one policed interface in the range xe-0/0/0 to xe-0/0/23 or the range xe-0/1/1 to xe-0/1/7.
- There is at least one policed interface in the range xe-0/0/24 to xe-0/0/47 or the range xe-0/1/8 to xe-0/1/15.

For example, if you configure a policer to rate-limit traffic at 1 Gbps and apply the policer (by using a firewall filter) to xe-0/0/0 and xe-0/0/24 in the output direction, each interface is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps. The same behavior occurs if you apply the policer to xe-0/1/1 and xe-0/0/24—each interface is rate-limited at 1 Gbps.

If you apply the same policer on egress to multiple interfaces in these groups, each *group* is rate-limited at 1 Gbps. For example, if you apply the policer to xe-0/0/0 through xe-0/0/4 (five interfaces) and xe-0/0/24 through xe-0/0/33 (ten interfaces), each group is rate-limited at 1 Gbps, for a total allowed throughput of 2 Gbps.

Here is another example: If you apply the policer to xe-0/0/0 through xe-0/0/4 and xe-0/1/1 through xe-0/1/5 (a total of ten interfaces), that group is rate-limited at 1 Gbps in aggregate. If you also apply the policer to xe-0/0/24, that one interface is rate-limited at 1 Gbps while the other ten are still rate-limited at 1 Gbps in aggregate.

Interfaces xe-0/1/1 through xe-0/1/15 are physically located on the QSFP+ uplink ports, according to the following scheme:

- xe-0/1/1 through xe-0/1/3 are on Q0.
- xe-0/1/4 through xe-0/1/7 are on Q1.
- xe-0/1/8 through xe-0/1/11 are on Q2.
- xe-0/1/12 through xe-0/1/15 are on Q3.

The doubling of the policed rate occurs only if the policer is applied in the output direction. If you configure a policer as described above but apply it in the input direction, the total allowed throughput for all interfaces is 1 Gbps.

**Solution** This is expected behavior.

### Filter-Specific Egress Policers on QFX3500 Devices Might Allow More Throughput Than Is Configured

**Problem** **Description:** You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in ternary content addressable memory (TCAM). If you create a filter-specific policer and

reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

**Solution** To prevent this unexpected behavior, use the information about TCAM slices presented in [“Planning the Number of Firewall Filters to Create” on page 5236](#) to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

### **Policers Can Limit Egress Filters**

---

**Problem** **Description:** The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters. Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.
- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
  - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
  - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

**Solution** You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.



- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

## Services

- [Troubleshooting Port Mirroring on page 7263](#)

### Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 7263](#)
- [Egress Port Mirroring with VLAN Translation on page 7265](#)
- [Egress Port Mirroring with Private VLANs on page 7265](#)

#### Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 7263](#)
- [Remote Port Mirroring Only on page 7265](#)

#### *Local and Remote Port Mirroring*

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:

- **interface**
- **ip-address**
- **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

**Remote Port Mirroring Only**

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

**Egress Port Mirroring with VLAN Translation**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

**Egress Port Mirroring with Private VLANs**

**Problem** **Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.
- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

- Related Documentation**
- [Understanding Port Mirroring on page 5425](#)
  - [Example: Configuring Port Mirroring for Local Analysis](#)
  - [Example: Configuring Port Mirroring for Remote Analysis](#)

---

## Traffic Management

- [Troubleshooting Dropped FCoE Traffic on page 7266](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth on page 7269](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth on page 7270](#)
- [Troubleshooting Egress Queue Bandwidth Impacted by Congestion on page 7271](#)
- [Troubleshooting an Unexpected Rewrite Value on page 7272](#)
- [Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic on page 7273](#)

## Troubleshooting Dropped FCoE Traffic

- Problem** **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.
- Cause** There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):
1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
  2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
  3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
  4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.

5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.
6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

---

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

**Solution** The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 5606 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

#### Related Documentation

- [show class-of-service congestion-notification on page 6305](#)
- [show class-of-service forwarding-class-set on page 6313](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 6174](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 5606](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 5559](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth

**Problem**    **Description:** The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth (shaping rate) configured for the queue.

**Cause**        When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.

The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

**Solution** When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

**Related  
Documentation**

- [shaping-rate on page 6278](#)
- [Example: Configuring Maximum Output Bandwidth on page 6101](#)
- [Example: Configuring Queue Schedulers on page 6081](#)
- [Understanding CoS Output Queue Schedulers on page 5868](#)

## Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

**Problem** **Description:** The minimum bandwidth of a queue or a priority group when measured at the egress port exceeds the minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).

**Cause** When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.



**NOTE:** The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

---

**Solution** When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit



rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

- Related Documentation**
- [guaranteed-rate on page 6247](#)
  - [transmit-rate on page 6285](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 6096](#)
  - [Example: Configuring Queue Schedulers on page 6081](#)
  - [Understanding CoS Output Queue Schedulers on page 5868](#)

## Troubleshooting Egress Queue Bandwidth Impacted by Congestion

**Problem** **Description:** Congestion on an egress port causes egress queues to receive less bandwidth than expected. Egress port congestion can impact the amount of bandwidth allocated to queues on the congested port and, in some cases, on ports that are not congested.

**Cause** Egress queue congestion can cause the ingress port buffer to fill above a certain threshold and affect the flow to the queues on the egress port. One queue receives its configured bandwidth, but the other queues on the egress port are affected and do not receive their configured share of bandwidth.

**Solution** The solution is to configure a drop profile to apply weighted random early detection (WRED) to the queue or queues on the congested ports.

Configure a drop profile on the queue that is receiving its configured bandwidth. This queue is preventing the other queues from receiving their expected bandwidth. The drop profile prevents the queue from affecting the other queues on the port.

To configure a WRED profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

- Related Documentation**
- [drop-profile on page 6227](#)
  - [Example: Configuring WRED Drop Profiles on page 6071](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 5966](#)
  - [Understanding CoS WRED Drop Profiles on page 5909](#)

## Troubleshooting an Unexpected Rewrite Value

**Problem**    **Description:** Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

**Cause**    If you configure a rewrite rule for a forwarding class on an egress port but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

**Solution**    If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.



**TIP:** If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.



**NOTE:** There are no default rewrite rules. You can bind one rewrite rule for each type (DSCP and IEEE 802.1) to a given interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value associations.

1. Assign a rewrite value to a forwarding class. Add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:  
  
[edit class-of-service rewrite-rules]

```
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **fcoe** is being randomly rewritten, and you want to use IEEE 802.1 code point **011** for the **fcoe** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class fcoe loss-priority high code-point
011
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1
custom-rw
```

#### Related Documentation

- [interfaces on page 6256](#)
- [rewrite-rules on page 6273](#)
- [Defining CoS Rewrite Rules on page 6182](#)
- [Monitoring CoS Rewrite Rules on page 6292](#)

## Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic

**Problem**    **Description:** In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets.

**Cause**    Failure of a queue to transmit packets for 12 consecutive seconds may be due to:

- A strict-high priority queue consuming all of the port bandwidth
- Several queues consuming all of the port bandwidth
- Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
- Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

**Solution**    If the cause is a strict-high priority queue or other queues consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the queues that are using all of the port bandwidth and preventing other queues from obtaining bandwidth on the port. You configure a maximum rate by creating a scheduler, using a scheduler map to apply it to a forwarding class (which maps to an output queue), and applying the scheduler map to the port using a forwarding class set and a traffic control profile.

To configure rate shaping using the CLI:

1. Name the existing scheduler or create a scheduler and define the maximum bandwidth as a rate or as a percentage:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name shaping-rate (rate | percent percentage)
```

2. Configure a scheduler map to associate the scheduler with the forwarding class (queue) that is consuming all of the port bandwidth:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

3. Associate the scheduler map with a traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name scheduler-map
scheduler-map-name
```

4. Associate the traffic control profile (and thus the scheduler map that contains the rate shaping queue scheduler) with a forwarding class set and apply them to the interface that is being reset:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set fc-set-name
output-traffic-control-profile traffic-control-profile-name
```

For example, a strict-high priority queue is using all of the bandwidth on interface **shpnode:xe-0/0/10** and preventing other queues from transmitting for 12 consecutive seconds. You decide to set a maximum rate of 7 Gbps on the strict-high priority queue to ensure that at least 3 Gbps of the port bandwidth is available to service other queues.

[Table 600 on page 6459](#) shows the topology for this example:

**Table 702: Components of the Rate Shaping Troubleshooting Example**

Component	Settings
Affected interface	<b>shpnode:xe-0/0/10</b>
Scheduler (strict-high priority scheduler)	Name: <b>shp-sched</b> Shaping rate: <b>7g</b> Priority: <b>strict-high</b>  <b>NOTE:</b> This example assumes that the scheduler already exists and has been configured as <b>strict-high</b> priority, but that rate shaping to prevent the strict-high priority traffic from using all of the port bandwidth has not been applied.
Scheduler map	Name: <b>shp-map</b> Forwarding class to associate with the <b>shp-sched</b> scheduler: <b>strict-high</b>  <b>NOTE:</b> This example assumes that a strict-high priority forwarding class has been configured and assigned the name <b>strict-high</b> .
Traffic control profile	Name: <b>shp-tcp</b>  <b>NOTE:</b> This example does not describe how to define a complete traffic control profile.

Table 702: Components of the Rate Shaping Troubleshooting Example (*continued*)

Component	Settings
Forwarding class set	<p>Name: <b>shp-pg</b></p> <p>To configure the scheduler, map it to the strict-high priority forwarding class, and apply it to interface <b>shpnode:xe-0/0/10</b> using the CLI:</p> <ol style="list-style-type: none"> <li>Specify the scheduler for the strict-high priority queue (<b>shp-sched</b>) with a maximum bandwidth of 7 Gbps: <pre>[edit class-of-service schedulers] user@switch# set shp-sched shaping-rate 7g</pre> </li> <li>Configure a scheduler map (<b>shp-map</b>) that associates the scheduler (<b>shp-sched</b>) with the forwarding class (<b>strict-high</b>): <pre>[edit class-of-service scheduler-maps] user@switch# set shp-map forwarding-class strict-high scheduler shp-sched</pre> </li> <li>Associate the scheduler map <b>shp-map</b> with a traffic control profile (<b>shp-tcp</b>): <pre>[edit class-of-service traffic-control-profiles] user@switch# set shp-tcp scheduler-map shp-map</pre> </li> <li>Associate the traffic control profile <b>shp-tcp</b> with a forwarding class set (<b>shp-pg</b>) and the affected interface (<b>shpnode:xe-0/0/10</b>): <pre>[edit class-of-service] user@switch# set interfaces shpnode:xe-0/0/10 forwarding-class-set shp-pg output-traffic-control-profile shp-tcp</pre> </li> </ol>
Related Documentation	<ul style="list-style-type: none"> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 5868</a></li> <li>• <a href="#">Defining CoS Queue Scheduling Priority on page 6171</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 6081</a></li> <li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 6094</a></li> <li>• <a href="#">Example: Configuring Forwarding Class Sets on page 6078</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 5966</a></li> </ul>

## Virtual Chassis Fabric

- [Troubleshooting Virtual Chassis Fabric on page 7275](#)

### Troubleshooting Virtual Chassis Fabric

This topic describes some of the following common troubleshooting issues for a Virtual Chassis Fabric (VCF):

- [Virtual Chassis Port Link Does Not Form on page 7276](#)
- [QFX5100 Leaf Device Assumes Routing Engine Role on page 7276](#)

### Virtual Chassis Port Link Does Not Form

---

**Problem**    **Description:** You connect a 40-Gbps QSFP+ port or a 10-Gbps SFP+ port between a leaf device and a spine device in an autoprovisioned or preprovisioned VCF. You expect the automatic Virtual Chassis port (VCP) conversion feature to convert the link into a VCP link, but the conversion doesn't occur.

The [show virtual-chassis vc-port](#) output indicates that the status of the interface is **Absent** or one or both of interfaces don't appear in the [show virtual-chassis vc-port](#) output.

**Cause**    If one end of a link is configured as a VCP and the other is not configured as a VCP, the VCP link does not form.

The automatic VCP conversion feature, therefore, does not work in the following situations:

- a 40-Gbps QSFP+ or 10-Gbps SFP+ interface on one end of the link is already configured as a VCP.

If you have previously removed a device from a VCF but haven't used the **request virtual-chassis vc-port delete** command to convert the interface that was connected to the removed device out of VCP mode, the interface is still configured as a VCP.

If you have removed a device from one Virtual Chassis or VCF and not changed the VCP port setting, the device being added to the VCF might also be configured as a VCP.

- a 40-Gbps QSFP+ port on an EX4300 switch, which is configured as a VCP by default, is interconnecting to a spine device.

**Solution**    Manually configure the interface that is not configured as a VCP into a VCP using the **request virtual-chassis vc-port set** command.

### QFX5100 Leaf Device Assumes Routing Engine Role

---

**Problem**    **Description:** A QFX5100 device configured as a leaf device assumes the Routing Engine role during VCF setup. The **show virtual-chassis** output confirms the role.

**Solution**    The device can assume the Routing Engine role for several minutes during setup before it receives the configuration from the master Routing Engine, but eventually returns to the linecard role with no user intervention.

**Related Documentation**

- [Virtual Chassis Fabric Overview on page 7033](#)